Actes de la conference sur le risque et l'assurance en l'honneur de Pierre Picard, organisée par le CREST (Center for Research in Economics and Statistics, ENSAE et IP Paris) qui s'est tenue les 14 et 15 septembre 2023 à l'Ecole Polytechnique (14 septembre) et à la Cité Internationale (15 Septembre).

Cette conférence était organisée autour de 5 thémes :

- Contract Theory : Communications de Jean-Charles Rochet (Toulouse School of Economics), "Dynamic Contracting with Many Agents" et François Salanié (INRAE & TSE), "Competitive Nonlinear Pricing under Adverse Selection".
- Public Policies : Communications de Enrico Biffis (Imperial College, London), "Short-lived gasses, carbon markets and climate risk mitigation" et Christian Gollier (Toulouse School of Economics), "Stress Discounting".
- Insurance : Communications de Pierre-Yves Geoffard (Paris School of Economics), "Road traffic accidents in France: compensation for body injury in France", Claude Fluet (Université Laval), "Consumer Protection in Retail Investments : Are Market Adjusted Damages Efficient?" et Michael Hoy (University of Guelph), "New Safety Technologies and Vehicle Safety".
- Risk Theory : Communications de Rachel J. Huang (National Central University, Taiwan), "A Simple Approach for Measuring Higher-Order Risk Attitudes", Arthur Snow (University of Georgia), "A Complete Characterization of Downside Risk Preference" et Richard Peter (University of Iowa), "The many faces of multivariate risk-taking: Risk apportionment for desirable and undesirable attributes".
- Empirical Insurance : Communications de Georges Dionne (HEC Montreal), "Consolidation of the US property and casualty insurance industry: Is climate risk a causal factor for mergers and acquisitions?" et Kili Wang (Tamkang University, Taiwan), "The case of insurance claim manipulation--induced by market powered distribution channel offering multi-faceted services".

Les articles de Enrico Biffis et Pierre-Yves Geoffard ne sont pas disponibles car il sont encore à un stade préliminaire.

Money and Taxes Implement Optimal Dynamic Mechanisms

Bruno Biais (HEC), Hans Gersbach (ETHZ), Jean-Charles Rochet (TSE),
Ernst Ludwig von Thadden (Mannheim), and Stéphane Villeneuve (TSE)[1]

August 2023

Abstract

We analyze capital allocation and risk sharing between a principal and many agents, who privately observe their output. Incentive compatibility requires that agents bear part of their idiosyncratic risk. The larger the agents' risk exposure, the larger the rents the principal can extract from them. The optimal dynamic mechanism can be implemented by a market equilibrium with money and taxes. Inflation affects agents' portfolio choice between risky capital and safe money. To implement the optimal mechanism, the principal sets the inflation rate so that agents' risk exposure is the same in equilibrium as in the optimal mechanism.

[1] We gratefully acknowledge comments from seminar participants at Mannheim University, Bonn University, Berkeley, ACPR, HEC, WFA 2023 and the Lemma-Rice conference on Money, especially Guillaume Rocheteau, Jacques Olivier, Noémie Pinardon-Touati, Stefan Ruenzi, and David Sraer.

Money and Taxes Implement Optimal Dynamic Mechanisms

1 Introduction

How should capital be allocated and risks shared in a dynamic economy without aggregate risk? In the absence of informational frictions, the answer is clear: capital should be allocated according to expected individual productivities and risks should be eliminated by diversification. However, when information about individual outputs is private, one must also take into account incentive compatibility constraints. This paper studies how these constraints affect capital accumulation and risk sharing.

To address these issues, we consider an infinite horizon economy with a continuum of risk averse agents and a single good that can be consumed or invested as capital, similar to the economy studied in Angeletos (2007). Each agent operates a project whose output is proportional to the amount of capital under her management and subject to idiosyncratic shocks. Individual unit outputs are i.i.d. so that a version of the law of large numbers applies, implying that aggregate output is deterministic.

We first show that, under symmetric information, in the optimal allocation agents fully mutualize idiosyncratic risks, and consumption is deterministic. Moreover, since agents' productivities are i.i.d. across agents and periods, capital allocation is not influenced by past performance.

Next we consider the case in which agents privately observe their individual output and can secretly consume some of it, as in Bolton and Scharfstein (1990). In contrast with output, capital is observable. Applying the revelation principle, we study truthful revelation mechanisms, in which agents truthfully report their output to the principal, who then allocates consumption and capital according to the reports. Thus the dynamic optimal mechanism allocates capital and consumption to maximize the principal's utility, subject to the participation and incentive constraints of the agents and the aggregate resource constraint.

To provide agents with incentives not to divert output, the optimal contract specifies an increase (resp. decrease) of consumption **and** capital for agents whose output is larger (resp. smaller) than expected. Lucky agents (those that perform better in a given period) get more capital to manage in the next period, not because they are more skilled (performance is i.i.d. across agents and across periods) but because this provides incentives to report good performance instead of diverting output. In contrast with the symmetric information case, insurance is imperfect, because full insurance is not incentive compatible. So, the optimal mechanism exposes agents to a fraction of their idiosyncratic risk.

From a mathematical viewpoint, finding the optimal mechanism is challenging, as we need to extend to a continuum of agents the martingale techniques introduced by Sannikov (2008) in the one agent case. With only one agent, the Bellman equation that characterizes the optimal mechanism involves the partial derivatives of the value function with respect to two state variables: aggregate capital and the continuation utility promised to the (single) agent by the principal. In contrast, in our model with a continuum of agents, the state variables are aggregate capital and the **entire distribution** of continuation utilities across agents, which belongs to the space of probability distributions over $\mathbb{R}$, which we endow with the Wasserstein topology. So the value function of the principal solves a Bellman equation in an infinite dimensional space. We first determine the shape of this Bellman equation, which involves the Gateaux derivative of the value function with respect to the distribution of continuation utilities. Then, thanks to our log utility specification, we show that the dimension of states variables can be reduced to two: aggregate capital and the expectation of (a function of) agents' continuation utilities. These are sufficient statistics for the characterization of the optimal mechanism.[2] Thanks to the reduction of the dimension of the state space from infinity to two, we can fully characterize the dynamics of capital and consumption allocations as well as the distribution of continuation utilities across agents.

The optimal direct mechanism is remarkably simple: consumption and capital are allocated among agents proportionally to each agent's equivalent permanent consumption, defined as the constant lifetime stream of consumption giving the agent the same continuation utility as the mechanism. The equivalent permanent consumption of each agent grows at a constant rate in expectation, but is impacted by the agent's idiosyncatic output shock. The innovation in the growth rate of an agent's consumption or capital is proportional, by a constant $x \in (0, 1)$, to the agent's idiosyncratic output shock. x measures the extent to which the agent is exposed to the risk of her idiosyncratic output shock. By raising x the principal relaxes the incentive compatibility condition and can thus extract more rents from the agent, but this reduces allocative efficiency by reducing insurance. Thus there is a rent-efficiency tradeoff. We characterize the Pareto frontier of information-constrained Pareto optimal allocations, each point of which corresponds to a different value of x, i.e., a different compromise between rents and effciency. Because agents are exposed to their idiosyncratic shocks, inequality increases over time and agents become more and more heterogenous.[3] Moreover, while aggregate capital and output grow over time, growth is lower than under symmetric information. This is because incentive compatibility constrains how much new capital can be delegated to agents.

The above presented direct revelation mechanism is centralized as all agents report to the principal, who then reallocates consumption and capital among them. We show that a more decentralized implementation is possible, in which agents exchange goods against money in a market and the principal intervenes only via money issuance and taxation. When trading in the market, agents face a dynamic portfolio problem à la Merton. They choose how much to invest in capital and money, bearing in mind that the former has higher expected return but is riskier than the latter. The principal manipulates this portfolio choice by controlling the growth of money supply and thus inflation, which affects the excess rate of return of risky capital over money, so that agents' risk exposure is the same in equilibrium as in the optimal mechanism. Different monetary and fiscal policies implement different points on the incentive constrained Pareto frontier. In general the principal uses both taxation and seigneurage to raise revenue and extract rents from the agent.

Our results can be contrasted with the classical welfare theorems. In a convex economy with complete markets and no frictions, all competitive equilibria are efficient (first welfare theorem) and, converserly, all efficient allocations can be decentralized by a competitive equilibrium after appropriate lump sum transfers between agents (second welfare theorem). Other forms of taxations are distortive. The opposite is true in our economy with asymmetric information and endogenously incomplete markets. In particular, the first theorem does not hold: competitive equilibria without government intervention are con-

[2]Angeletos (2007) also avoids the "curse of dimensionality" with a log utility specification. A major difference is that in Angeletos (2007) institutions and market incompleteness constraints are exogenous while in our paper they are features of the endogenous optimal mechanism.

[3]More precisely, the coefficient of variation (standard deviation divided by the mean) of continuation utilities across agents increases over time.

strained inefficient. However, all constrained optimal allocations can be implemented as a market equilibrium provided the government chooses appropriate monetary and taxation policies. By setting the inflation rate and the tax rate in a particular way, the government can impact individual behaviour so that market imperfections are corrected and redistributive objectives are met. This can be viewed as an extension of the second welfare theorem to an economy with endogenously incomplete markets.

Literature: Our paper complements several strands of literature.

First, our analysis of dynamic contracting between one principal and many agents is related to the literature analyzing dynamic contracting between one principal and one agent, in particular the seminal work of DeMarzo and Fishman (2007a, 2007b) and Sannikov (2008), and the following analyses of Biais, Mariotti, Plantin, and Rochet (2007), DeMarzo and Sannikov (2006), Feng and Westerfield (2021), and Di Tella and Sannikov (2021). As in Biais, Mariotti, Rochet, and Villeneuve (2010) and DeMarzo, Fishman, He, and Wang (2012), firm size is determined by the optimal contract and is useful to provide incentives. However, in contrast with Biais, Mariotti, Rochet, Villeneuve (2010) and DeMarzo, Fishman, He and Wang (2012), in the present paper there are no capital adjustment costs. This enhances tractability, and gives rise to continuous reallocation of capital. He (2009) offers an interesting alternative approach in which firm size is affected by unobservable agent's effort. This differs from our model in which firm size is directly controlled by the principal, and what is unobservable is output.

The major contribution of the present paper relative to that literature is to embed the contracting problem in a general equilibrium context, with a population of agents and aggregate resource constraints. Thus we shed light on the impact of incentive constraints on the allocation of capital and consumption across agents. In particular, we show how incentive constraints generate increasing inequality in the population of agents. Moreover, we show how the dynamic optimal mechanism can be implemented by a market in which agents trade goods for money, and inequality is regulated by optimal taxes.

Second, our analysis is related to the dynamic macrofinance literature analyzing risk with exogenously incomplete markets (see Bewley, 1980, Aiyagari, 1994, Huggett, 1993, 1997, Krusell and Smith, 1998, Angeletos, 2007, Brunnermeier Sannikov, 2014, Gersbach, Rochet, and Von Thadden, 2022, Di Tella, 2020, and Achdou et al 2022.)[4]

The major contribution of the present paper relative to that literature is to provide microfoundations for market incompleteness.[5] Thus, the institutions and constraints we consider are endogenous features of the optimal dynamic mechanism. This helps clarify the consequences of informational frictions. For example, we reconcile two effects which, as explained by Angeletos (2007), had so far been viewed as distinct. While the literature in line with Bernanke and Gertler (1989) emphasizes how wealth affects the *ability* to invest in capital, Angeletos (2007) emphasizes how wealth affects the *willingness* to hold risky capital. Our mechanism design approach clarifies the common origin of these two forces: incentive compatibility constrains both how much capital agents are allocated and how much of the corresponding idiosyncratic risk they must bear. Consequently, in contrast with Angeletos (2007), in our analysis frictions unambiguously lower capital accumulation.

[4]Our focus on the distribution across agents and our reliance on mean field techniques are in line with Achdou et al (2022).

[5]Another difference is that, while most of that literature studies labor income risk, our paper, like Angeletos (2007) considers capital return risk.

Third, our focus on money in the implementation of the optimal mechanism links our paper to the new monetarist literature initiated by the seminal papers of Kiyotaki and Wright (1989, 1993) and presented by Williamson and Wright (2011). A common theme with that literature is that money arises endogenously, as a useful instrument, instead of being a constraint as in cash in advance models or exogenous as in money-in-the-utility-function models. Money in our implementation encodes the memory of past performance in line with Kocherlakota (1998) and provides consumption insurance in line with Berentsen and Rocheteau (2004).

There are important differences, however, between our analysis and the new monetarist literature. First, instead of starting from a characterization of optimal allocations in a setting with money, we characterize the optimal mechanism in a real economy with only goods and no money, and then we introduce money as a tool to implement the optimal mechanism. Second, while the new monetarist literature assumes large households (Shi, 1997) or the alternation of decentralized and centralized markets (Lagos and Wright, 2005) so that at the beginning of each period all agents start with the same amount of money, in our framework agents have endogenously heterogeneous money holdings, and we characterize the dynamics of this heterogeneity. Third, in the new monetarist literature, agents are homogenous at the beginning of each period, so that the optimal allocation is pinned down by a static mechanism. In contrast, in our setting agents' continuation utilities vary stochastically over time and the optimal allocation is set by a dynamic mechanism.

Finally, we complement the mechanism design approach to optimal taxation pioneered by Mirrlees (1971), Diamond and Mirrlees (1978), and Diamond (1998), and further developed by the new dynamic public finance literature (e.g., Golosov, Kocherlakota, Tsyvinski (2003), Golosov, Tsyvinski (2007), and Fahri, Werning (2010)). A major difference is that, in these papers, risk and information asymmetry are about wage earners' skills, while, in our paper, risk and information asymmetry are about managers' capital returns. Correspondingly, unlike in these papers, the dynamic of capital allocation plays a key role in our analysis. Another major difference is that the optimal taxation literature focuses on one policy tool, namely the tax system, while in our set-up, the government chooses also budgetary policy (the consumption of the principal) and monetary policy (how much money is issued).

Structure of the paper: The complete analysis of the continuous time model under asymmetric information is difficult and mathematically complex. In order to build intuition, the next section presents a simple two period model, which illustrates some (but not all) of the economic forces at play in our full model. Then, Section 3 introduces the continuous time model, and solves the symmetric information case, which provides a useful benchmark for the analysis of asymmetric information. Section 4 determines the Bellman equation that characterizes the principal value function under asymmetric information. Then we make a guess on the shape of the optimal policy, qualitatively close to that obtained under symmetric information, and we show that this candidate policy is indeed the full solution of our problem. Section 5 shows that the optimal direct mechanism can be implemented with money and taxes. Section 6 points to the link between our results and optimal tax theory. Section 7 offers a brief conclusion. Proofs not in the text are in Appendix A.

2 The two-period case

To build intuition, we first present a simple version of the model with three dates and no discounting. There is a single good which can be used for consumption or investment. Agents can invest the good and generate returns whose distribution is i.i.d across agents. We first characterize the symmetric first best allocation. Then we turn to the case in which agents privately observe their returns. While the first best allocation is not incentive compatible, we characterize the symmetric second best allocation and show it can be implemented with money and taxes.

2.1 A simple two-period model

The investment technology has constant returns to scale. There is a mass 1 continuum of agents, that are identical at time 0: each of them is endowed with one unit of capital good. Time 1 unit returns, which are i.i.d. across agents, can be high ($R^H = 1 + \sigma$) or low ($R^L = 1 - \sigma$) with equal probability 1/2. [6] To ensure that returns are positive we assume $\sigma < 1$. For simplicity we assume time 2 returns are deterministic and equal to 1.

Agents with high time 1 return are referred to as type $s = H$ and agents with low time 1 return as type $s = L$. After time 1 output is realized, agents consume and are allocated capital which, invested, generates output at time 2. This output is consumed by the agents at time 2. The ex-ante utility of each agent is

$$E\left[U(C_1^s) + U(C_2^s)\right],$$

where the function U is strictly concave, increasing and differentiable.

2.2 The Symmetric First Best Allocation

Since agents are ex ante identical, we focus on the symmetric first best allocation, which is characterized by the consumption profile $\{C_1^s, C_2^s)\}$, $s \in \{L, H\}$ maximizing

$$E\left[U(C_1^s) + U(C_2^s)\right],$$

subject to the intertemporal resource constraint:[7]

$$E\left[C_1^s + C_2^s\right] \leq 1.$$

We thus obtain our first proposition:

Proposition 1 *In the symmetric first best allocation,* $C_t^s = \frac{1}{2}, \forall (t, s)$.

The intuition is straightforward. Because returns are i.i.d across agents, by the law of large numbers there is no aggregate risk. Agents are only exposed to idiosyncratic shocks and it is optimal to mutualize this idiosyncratic risk. Initially there is one unit of good, which can be consumed or invested. Since the rate of return and the discount rate are equal, half the endowment is consumed at time 1 ($C_1^s = \frac{1}{2}$) while the other half is invested at time 1 and consumed at time 2 ($C_2^s = \frac{1}{2}$). Note that this allocation is characterized by complete consumption smoothing over time, and complete insurance against output shocks.

[6] Thus, net average rates of returns are 0, like the discount rate. This is just for the sake of simplicity. In the continuous time analysis below, average rates of returns and discount rates are strictly positive.

[7] The intertemporal resource constraint obtains by adding the time 1 resource constraint $E\left[C_1^s + k_1^s\right] \leq 1$, and the time 2 resource constraint $E\left[C_2^s\right] \leq E[k_1^s]$.

2.3 The Symmetric Second Best Allocation

Next turn to the case in which agents privately observe their time 1 return. By the revelation principle, we can restrict attention to direct mechanisms, where agents report and transfer their return to the principal, who then sets their consumption and capital allocation. When an agent's return is high she can pretend it is low, and secretly consume the difference (2σ). When an agent's return is low, she cannot pretend it is high: to do so she would have to transfer high output to the principal, but such high output is not available to the agent. Thus, there is only one incentive compatibility constraint:

$$U(C_1^H) + U(C_2^H) \geq U(C_1^L + 2\sigma) + U(C_2^L).$$

Clearly, the incentive compatibility constraint does not hold for the first best allocation, in which $C_1^H = C_2^H = C_1^L = C_2^L$. Maximising agents' expected utility under resource and incentive constraints yields the properties of the second best allocation, stated in the next proposition.

Proposition 2 *In the symmetric second best allocation, consumptions are such that* $C_1^H = C_2^H$, $C_2^H > C_2^L$ *and* $C_1^L > C_2^L$.

As in the first best allocation, high types have the same consumption at time 1 and time 2. This is the standard "no distortion at the top" result. By contrast, consumption smoothing is imperfect for low types: $C_1^L > C_2^L$. An agent with high return pretending to have had low return secretly consumes an additional 2σ at time 1. This lowers his time 1 marginal utility relative to his time 2 marginal utility. Consequently, giving time 1 consumption to the agent with low returns tightens the incentive constraint less than giving time 2 consumption to that agent. Hence, it is optimal to set $C_1^L > C_2^L$, although this is a distortion relative to the first best, in which agents have the same consumption at the two periods. Moreover risk sharing is also imperfect: The intertemporal utility of high types is strictly higher than that of low types. Low consumption after low output, although suboptimal under symmetric information, is constrained-optimal under asymmetric information, because it incentivizes truthful reporting. Thus with information asymmetry there is imperfect insurance. Figure 1 plots, for $\sigma = 0.3$ and log utility, the time 1 and time 2 consumptions in the first and second best. In the logarithmic utility case, the second best allocation can be computed (almost) explicitly.

Proposition 3 *When* $U(c) = \log c$, *the symmetric second best allocation is such that*

$$C_1^H = C_2^H = (\frac{1+\sigma}{2})(1-\mu^2),$$

$$C_2^L = (\frac{1+\sigma}{2})(1-\mu)^2,$$

$$C_1^L = (\frac{1+\sigma}{2})(1+\mu)^2 - 2\sigma,$$

where μ *is the unique positive solution of the equation:*

$$\frac{\mu(1+\mu)^2}{1+3\mu} = \frac{\sigma}{1+\sigma}.$$

Moreover $\frac{\sigma}{1+\sigma} < \mu < 1$. *Correspondingly, in the second best:*

1. *At date 1, successful agents get more consumption than unsuccessful agents:* $C_1^H > C_1^L$.
2. *At date 2, successful agents get more consumption than unsuccessful agents:* $C_2^H > C_2^L$.
3. *Informational frictions reduce aggregate investment:* $C_2^H + C_2^L < 1$.
4. *Unsuccessful agents are partially insured by successful agents:* $(C_2^L + C_1^L) > (C_2^H + C_1^H) - 2\sigma$.

Properties 1 and 2 imply that incentives are optimally provided when $C_1^H > C_1^L$ and $C_2^H > C_2^L$: higher consumption at both dates for successful agents. Property 3 shows that informational frictions reduce investment. This will also be the case in the full model presented below. Finally, Property 4 shows that some insurance can be achieved in spite of informational frictions: agents with low output obtain larger consumption in the optimal mechanism than in autarky.

2.4 Implementation with Money and Taxes

The optimal direct mechanism that we have characterized is completely centralized: all agents report to the principal, who then allocates goods across agents. However, a more decentralized implementation of the optimal allocation is possible, in which the good is allocated by a competitive market.[8] In this market the good is traded against fiat money issued by the principal. This money has value because the principal levies taxes at date 1 and requires the agents to pay these taxes in money. At date 0 the principal allocates to each agent m_0 units of money, which can be used at time 1 to buy or sell capital k_1^s at price $p = 1$,[9] and pay taxes contingent on wealth.[10] Note that the principal does not intervene in the good market. Since there is no market at date 2, time 2 consumption is equal to output, $C_2^s = k_1^s$.

At time 1, agent s has R^s units of good and m_0 units of money. The goods can be used for consumption C_1^s or invested as productive capital k_1^s, and a quantity S^s of goods can be sold for money. If $S^s < 0$, this means the agent is buying goods. So the budget constraint of the agent regarding goods is

$$C_1^s + k_1^s + S^s \leq R^s. \tag{1}$$

After trading, the amount of money held by the agent is equal to her initial endowment (m_0) plus or minus the proceeds of her time 1 sales (S^s). The total wealth of the agent is thus $e^s = k_1^s + m_0 + S^s = m_0 + R^s - C_1^s$. The agent uses her money to pay taxes $\tau(e^s)$. So the budget constraint of the agent regarding money is

$$\tau(m_0 + R^s - C_1^s) \leq m_0 + S^s. \tag{2}$$

[8]Note that a pure market solution (no taxes) cannot implement the second best allocation since there are no gains from trade at date 1: successful agents are not willing to transfer resources to unsuccesful ones. In contrast with Diamond Dybvig (1981), if a bond market was created at $t = 1$, it would be inactive since in our model, agents have access to a storage technology at time 1.

[9]Price is indeterminate because taxes are in nominal terms. If the principal doubles the taxes and the money holdings, the price doubles but the real allocations are unchanged. So we normalize the price to 1.

[10]Taxes can be paid either at date 1 or 2.

This constraint must be binding for all agents, otherwise the value of money would be zero. Fiat money has positive value in this finite horizon model because it is required to pay taxes. At time 1, after observing her type, agent s chooses C_1^s and C_2^s to maximize

$$U(C_1^s) + U(C_2^s),$$

subject to the two constraints (1) and (2), which can be combined as

$$C_1^s + C_2^s + \tau(m_0 + R^s - C_1^s) \leq R^s + m_0.$$

That is expenses, equal to the sum of consumption at both dates and taxes, must be covered by resources, equal to the initial money endowment plus output. We want the solution to the agents' maximization problem to coincide with the constrained optimal allocation. This is satisfied when marginal tax rates for both agents are well chosen and when the goods market clears: $\mathbb{E}[S^s] = 0$. This second condition is satisfied when the aggregate money stock equals aggregate taxes, which could be interpreted as a form of equilibrium condition: money supply equals money demand. We obtain our next proposition.

Proposition 4 *1. The principal can implement the optimal mechanism* $(C_1^s, C_2^s)_{s\in\{L,H\}}$ *by distributing an amount* m_0 *of money to each agent and imposing a non linear wealth tax* $\tau(e)$ *such that*

$$\forall s, \tau'(e^s) = 1 - \frac{U'(C_1^s)}{U'(C_2^s)}, \tag{3}$$

which implies $\tau'(e^H) = 0$ *and* $\tau'(e^L) > 0$*.*

2. The aggregate money stock equals the expected value of taxes:

$$m_0 = \mathbb{E}[\tau(e^s)]. \tag{4}$$

Equation (3) states that the marginal tax rate is the wedge between the intertemporal marginal rate of substitution in the second best and the first best. $\tau'(e^H) = 0$ reflects that there is no distorsion at the top, while $\tau'(e^L) > 0$ reflects the distortion at the bottom. The intuition why money and taxes implement the optimal mechanism is the following:

First consider the agents with high time 1 output. They sell some of it, increasing their money holdings, which enables them to pay more taxes. Since taxes are increasing in capital, the ability to pay more taxes translates into the ability to hold more capital. And, since at time 2 agents consume the output from their capital, more capital translates into larger time 2 consumption, which implements the optimal mechanism. This is in line with theories of money as a record of good performance entitling money holders to consumption, i.e., "money as memory" (see Kocherlakota 1998).

Second consider the agents with low time 2 output. They can use some of their money to buy goods, and thus obtain some consumption smoothing. But since they have low money holdings, they cannot afford to pay large taxes, and therefore must have low capital investment and low time 2 consumption, again in line with the optimal mechanism. That unsuccessful agents use money to smooth the impact of shocks on consumption is in line with theories of money as a safe store of value in intertemporal consumption investment settings (see Merton 1969, 1971, and Berentsen and Rocheteau, 2002).

Taxation allows to create gains from trade between lucky (high types) and unlucky (low types) agents. Since lucky agents want to keep more wealth in order to consume more than unlucky agents at date 2, taxing wealth forces them to sell some of the good to unlucky agents, in order to get more money to pay their taxes. Unlucky agents buy the good because they know they will have to pay less taxes. This allows them to consume more at date 1.

There are interesting similarities between our model and the Diamond-Dybvig (1981) banking model: both models involve two periods, with consumption and investment. In both models, agents are ex-ante identical, and subject to privately observable independent shocks but there is no aggregate risk. However, the Diamond-Dybvig model involves time preference shocks, while ours involves output shocks. As a consequence, in Diamond-Dybvig there are gains from trade at the end of the first period, and opening a bond market improves the autarkic allocation. In our model there are no gains from trade at the interim date, and opening a bond market is useless. Taxation is used by the principal to create such gains from trade. Finally, in Diamond-Dybvig the Pareto optimal allocation can be implemented with private banks competing to offer demand deposits contracts while in our model some public intervention is needed, in the form of money and taxes.[11]

Finally note that, since the horizon is finite, the reason why money has value cannot be that it is a bubble. Here money has value because it is needed to pay taxes, in line with chartalism (Knapp, 1924). In the infinite horizon analysis below, the above intuitions still hold, but additional effects come into play. For some parameter values money has a bubble component. Moreover, the inflation rate, which is controlled by the principal through money issuance plays an important role in the implementation of the optimal mechanism. Thus, in our continuous time infinite horizon model, there is an optimal level of inflation.

3 The infinite horizon case

We now extend the analysis to an infinite horizon model in continuous time. Idiosyncratic shocks are captured by independent Brownian motions, which are easy to define when there is a finite number N of agents, but more tricky with a continuum. We start therefore by describing the model with N agents and then take the limit as N tends to infinity.

3.1 The Model

The principal faces N ex-ante identical agents indexed by $i = 1, ...N$. In order to keep their total mass constant, we assume each of them has mass $1/N$: each agent becomes smaller as their number increases. All agents are infinitely lived with discount rate ρ and logarithmic utility. There is a single good, which can be used for consumption or as capital input in a constant return to scale technology operated by the agents. The total amount of capital K_t is allocated to the agents: agent i invests k_t^i/N units of the good in her production process. The feasibility constraint is

$$K_t = \frac{1}{N}\sum_i k_t^i. \tag{5}$$

The output of agent i is

$$dY_t^i = \frac{k_t^i}{N}[\mu dt + \sigma dB_t^i],$$

[11] Incidentally, when agents have log utilities, banks are not needed in Diamond Dybvig because the market allocation is already Pareto optimal. In our model, even with log utilities the market equilibrium allocation is not Pareto optimal: money and taxes are needed.

where μ is the expected rate of return (net of depreciation) of the technology and B^i, $i = 1, ...N$ are independent Brownian motions, which can be interpreted as idiosyncratic non persistent productivity shocks.

The law of motion of aggregate capital is

$$dK_t = \frac{1}{N} \sum_i \left(k_t^i [\mu dt + \sigma dB_t^i] - c_t^i dt \right) - c_t^P dt, \tag{6}$$

where c_t^i / N is the consumption flow of agent i, while c_t^P is the consumption flow of the principal. This law of motion is a resource constraint stating that investment (left hand side) is equal to total output net of depreciation minus consumption (right hand side). With a finite number of agents, there is some residual aggregate risk:

$$var(dK_t) = \frac{\sigma^2}{N^2} \sum_i [k_t^i]^2 dt. \tag{7}$$

However when N tends to infinity, if the capital allocation k_t^i is square Riemann integrable in i, we can determine the limit behavior of the economy. The aggregate amount of capital at date t converges to the Riemann integral[12] of k_t^i

$$K_t = \int_0^1 k_t^i di, \tag{8}$$

and its law of motion becomes deterministic:

$$dK_t = \left(\mu K_t - \int_0^1 c_t^i di - c_t^P \right) dt, \tag{9}$$

This is because $\frac{1}{N} \sum_i [k_t^i]^2$ has a finite limit ($\int_0^1 (k_t^i)^2 di$) and thus $var(dK_t)$ tends to zero when N goes to infinity.

3.2 Optimal allocations under symmetric information

We first consider the case in which idiosyncratic shocks are observable. This serves as a benchmark to which we then contrast the case in which agents privately observe shocks and can secretly divert output.

3.2.1 The maximization problem

The simplest way to characterize the Pareto frontier of the economy without frictions is to compute the maximum discounted expected utility that the principal can obtain, subject to the resource constraint and the constraint that each agent i gets a given level of utility ω^i.[13] When information is symmetric, since there is no aggregate risk, it is optimal not to expose the agents to any risk. As shown below, this contrasts with the asymmetric information case. Thus, under symmetric information, the consumption of agent i at date t is a deterministic function of t, denoted c_t^i. By construction, it satisfies

$$\omega^i = \int_0^\infty e^{-\rho t} \log c_t^i dt, \tag{10}$$

[12] Since the total mass of agents is 1, integrability of $(k_t^i)^2$ implies that (k_t^i) is also integrable.
[13] ω_i can be interpreted as a reservation utility, reflecting an outside option.

hereafter referred to as the promise keeping constraint. The objective of the principal is

$$\int_0^\infty e^{-\rho t} \log c_t^P dt, \tag{11}$$

to be maximized subject to the promise keeping condition (10) for all i, and the law of motion of capital:

$$\dot{K}_t = \mu K_t - \int_0^1 c_t^i di - c_t^P. \tag{12}$$

Integrating (12) over time and using the transversality condition ($\lim_{t\to\infty} e^{-\mu t} K_t = 0$), we obtain that the initial amount of capital is equal to the present value of future consumption, discounted at the rate of return on capital:[14]

$$K = \int_0^\infty exp(-\mu t)[\int_0^1 c_t^i di + c_t^P]dt. \tag{13}$$

The optimal mechanism maximizes the objective of the principal (11) under the promise keeping constraint (10) and the capital dynamics constraint (13).

3.2.2 Characterization of optimal allocations

The next proposition describes the solution of the maximization problem when information is symmetric:

Proposition 5 *Optimal allocations are such that:*

1. *Capital grows at constant rate $\mu - \rho$:*

$$K_t = Ke^{(\mu-\rho)t}.$$

2. *At each date t, the principal consumes a constant fraction γ^P of capital, i.e.,*

$$c_t^P = \gamma^P K_t.$$

3. *Agents' continuation utilities grow linearly:*

$$\omega_t^i = \omega^i + (\frac{\mu-\rho}{\rho})t.$$

4. *At each date t, agent i consumes a constant fraction of $exp(\rho\omega_t^i)$, i.e.,*

$$c_t^i = \gamma^A exp(\rho\omega_t^i),$$

where $\gamma^A = exp[-\frac{\mu-\rho}{\rho}]$.

5. *For all agents, the ratio $\frac{exp(\rho\omega_t^i)}{K_t}$ is constant over time.*

[14] This property also held in the simple two period model analyzed above. There the rate of return of capital was equal to 0, so the initial aggregate endowment of capital good was equal to the sum of the future aggregate consumptions.

Property 1 states that aggregate capital grows at a constant rate, equal to productivity μ minus the discount rate ρ. Correspondingly, the flow of aggregate consumption is a fraction ρ of aggregate capital.

Property 2 states that the principal consumes a constant fraction of capital. This arises because the principal has logarithmic utility. Properties 1 and 2, together with (12), imply that the aggregate consumption of the agents is a constant fraction of capital.

Property 3 states that, starting from its initial level ω, an agents' continuation utility grows linearly with time, the trend being equal to the growth rate of capital divided by the discount rate, which is the same for all agents. This implies that inequality across agents does not grow over time, which will not be the case with asymmetric information.

Property 4 states that, at time t, an agent's consumption is a constant fraction of $exp(\rho\omega)$, which can be interpreted as the "equivalent permanent consumption" namely the constant lifetime stream of consumption giving utility ω to an agent. Since the agent's utility function is logarithmic and her discount rate is ρ, the equivalent permanent consumption corresponding to ω is $exp(\rho\omega)$. Combined with properties 1 and 3, it implies that an agent's consumption grows at the same rate as aggregate capital.

This yields Property 5, which states that the ratio of an agent's equivalent permanent consumption to aggregate capital is a constant, equal for all agents, which we denote by z. Aggregating across agents, the ratio of aggregate equivalent permanent consumption to capital is constant and equal to z:

$$\frac{\int exp(\rho\omega_t)d\mathbb{P}(\omega)}{K_t} = z.$$

We can now compute the value function of the principal:

$$V = \int_0^\infty e^{-\rho t} \log c_t^P dt, \tag{14}$$

The above proposition implies that this value function only depends on two state variables: aggregate capital K and z, which summarizes all the necessary information on the probability distribution $\mathbb{P}$ of ω. This reduces the dimensionality of the problem from ∞ to 2. The value function of the principal can be computed explicitly:

$$\rho V = log(\rho exp\frac{\mu - \rho}{\rho} K - zK) \tag{15}$$

The first term in the log on the right hand side of (15) is the total amount of constant certainty equivalent consumption that can be allocated among the principal and the agents. It represents the present value of consuming a fraction ρ of capital K growing at rate $\mu - \rho$. The second term in the log on the right hand side of (15) is the aggregate equivalent permanent consumption $\int \exp \rho\omega d\mathbb{P}(\omega) = zK$ of the agents, which cannot exceed $[\rho exp\frac{\mu-\rho}{\rho}]K$. Thus the value function of the principal can be written

$$V(K, z) = \frac{logK}{\rho} + v(z),$$

where $v(z) = log(\rho exp\frac{\mu-\rho}{\rho} - z)$ is only defined for z in a bounded interval: $0 \leq z \leq \rho exp\frac{\mu-\rho}{\rho}$. Similar properties will also hold in the asymmetric information case. Finally, the Pareto frontier is linear in the space of equivalent permanent consumptions:

$$exp(\rho V) + \int exp(\rho\omega)d\mathbb{P}(\omega) = [\rho exp\frac{\mu - \rho}{\rho}]K, \tag{16}$$

where the left-hand side is the sum of the principal's equivalent permanent consumption and the aggregate agents' permanent consumption, while the right-hand side is the total amount of equivalent permanent consumption to be allocated among the principal and the agents. The Pareto frontier is depicted in Figure 2.

4 Optimal allocations under asymmetric information

We now turn to the case in which agents privately observe their individual output. By the revelation principle, we consider revelation mechanisms. A mechanism is a mapping from the realized output dY_t^i, reported and delivered by agent i to the principal, into consumption and capital allocations for the agent. Since agents privately observe output, they can be tempted to divert a part of it and secretly consume it. To avoid this, the mechanism must induce truthful revelation, i.e., it must be incentive compatible.

4.1 Incentive compatibility

Consider an agent who would want to divert resources and consume secretly. Assuming the agent can only make absolutely continuous changes in the output process, the amount diverted is denoted by $\delta_t \, dt$. Defining

$$d\hat{B}_t^i = dB_t^i - \delta_t \, dt, \tag{17}$$

the dynamics of reported output writes as

$$d\hat{Y}_t^i = \mu k_t^i \, dt + \sigma k_t^i d\hat{B}_t^i.$$

Since the agent cannot secretly store, diversion cannot be negative: $\delta_t \geq 0$ for every t. The time 0 expected utility of an agent i who adopts a diversion strategy δ_t is

$$\omega_0^i = \sup_{\delta} \; \mathbb{E}\left[\int_0^\infty e^{-\rho t} \log(c_t^i + \sigma k_t^i \delta_t) dt\right].$$

To provide incentives for truthful revelation, the principal changes the continuation utility of the agent as a function of her reports. Hence, by the martingale representation theorem, the dynamics of the continuation utility of agent i is

$$d\omega_t^i = (\rho\omega_t^i - \log(c_t^i)) \, dt + \sigma y_t^i d\hat{B}_t^i, \tag{18}$$

where y_t^i is a B_t^i-adapted process. On the equilibrium path we have

$$d\omega_t^i = (\rho\omega_t^i - \log(c_t^i)) \, dt + \sigma y_t^i dB_t^i. \tag{19}$$

Intuitively, y_t^i is the sensitivity of the agent's continuation utility with respect to her report. The principal must choose this sensitivity to incentivize the agent to report her output truthfully. The state variable for agent i is her continuation utility ω_t^i, so instead of denoting her consumption by c_t^i, we hereafter denote it by $c_t^A(\omega_t^i)$. An intuitive examination of the incentive compatibility condition is the following. The incentive compatibility condition states that the agent must be better off revealing dB_t truthfully, and getting

$$\log(c_t^A(\omega_t))dt + \sigma y_t(\omega_t) dB_t$$

than underreporting: $d\hat{B}_t = dB_t - \delta dt$ and getting

$$\begin{aligned} & \log(c_t^A(\omega_t) + \sigma\delta k_t(\omega_t))dt + \sigma y_t(\omega_t)d\hat{B}_t \\ = \quad & \log(c_t^A(\omega_t) + \sigma\delta k_t(\omega_t))dt + \sigma y_t(\omega_t)(dB_t - \delta dt). \end{aligned}$$

So the incentive compatibility condition is

$$\sigma y_t \geq \sup_{\delta \geq 0} \frac{\log((c_t^A(\omega_t) + \sigma\delta k_t(\omega_t)) - \log(c_t^A(\omega_t))}{\delta} = \frac{\sigma k_t(\omega_t)}{c_t^A(\omega_t)}.$$

This means that the sensitivity of continuation utility to performance has to be larger than the product of the capital $k_t(\omega_t)$ managed by the agent by her marginal utility of consumption. In the log utility case this writes as:

$$y_t(\omega_t) \geq \frac{k_t(\omega_t)}{c_t^A(\omega_t)}.$$

This leads to our next proposition.

Proposition 6 *The incentive compatibility condition is equivalent to the inequality*

$$\forall t, y_t(\omega_t) \geq \frac{k_t(\omega_t)}{c_t^A(\omega_t)}. \tag{20}$$

The incentive compatibility condition (20) implies that, in contrast with the symmetric information case, agents cannot fully mutualize the risk of their idiosyncratic shocks. Condition (20) also shows there is a tradeoff between risk-sharing and investment: providing more insurance to the agent, by reducing the sensitivity of her continuation value to output shocks is possible only at the cost of reducing capital relative to consumption. This is because increasing capital, and therefore output, increases the amount of resources the agent can divert, which tightens the incentive constraint. This tradeoff is similar to that arising in Biais, Mariotti, Rochet and Villeneuve (2010), where size of operation (similar to capital in the present context) was limited by incentive compatibility.

The agents being risk averse, it is never optimal for the principal to expose them to more risk than required by the incentive compatibility condition. In other words the incentive constraint (20) is always binding and we can eliminate the capital allocation variable by writing $k_t^i = y_t^i c_t^i$. Since $\mu > 0$, it is optimal to fully allocate the capital stock to the agents, implying that the aggregate capital constraint (8) writes as

$$\int y_t(\omega_t) c_t^A(\omega_t) c_t^A(\omega_t) d\mathbb{P}(\omega_t) = K_t$$

4.2 The Hamilton-Jacobi-Bellman equation

As in the first best, the value function of the principal does not depend on the specific value function of each individual, but on the distribution $\mathbb{P}$ of agents' continuation pay offs. It also does not depend on individual outputs or capital, but on aggregate capital, which is deterministic, and on aggregate output which is linear in aggregate capital. So K and $\mathbb{P}$ are the state variables of the principal's maximization problem. That is, the principal problem is a deterministic control problem in a space that is the product of $\mathbb{R}$ by the space of probability measures on $\mathbb{R}$, which we endow with the Wasserstein distance (see for example

Villani 2009). We characterize the principal's value function as the unique solution to the dynamic programming Hamilton-Jacobi-Bellman (HJB) equation in that space.

The main difficulty for exploiting the dynamic programming principle is to differentiate functionals defined on the Wasserstein space. There are various notions of derivatives with respect to measures which have been developed in connection with the theory of optimal transport and using Wasserstein metric on the space of probability measures, for details see Villani (2009) and Appendix B of the present paper. For our purpose, we use the notion of Gateaux differentiability that is presented in appendix. Following the traditional approach for control problems, we first determine the shape of the HJB equation that the value function of the principal must satisfy (necessary condition) and then establish a verification theorem showing that regulation solutions of this HJB equation solve our control problem (sufficient condition). To do so, consider the control problem of the principal

$$V(K, \mathbb{P}) = \sup_{(c_t^A(.), c_t^P, y_t(.)) \in \mathcal{K}} \int_0^\infty e^{-\rho t} \log c_t^P dt, \tag{21}$$

where the state equations are given by

$$\dot{K}_t = \mu K_t - c_t^P - \int c_t^A(\omega) d\mathbb{P}(\omega), \tag{22}$$

$$d\omega_t = (\rho \omega_t - \log c_t^A(\omega) dt + \sigma y_t dB_t, \tag{23}$$

and where the supremum is taken over the set $\mathcal{K}$ of admissible Markov controls (c^A, c^P, y) such that

$$\int y_t(\omega) c_t^A(\omega) d\mathbb{P}(\omega) = K_t. \tag{24}$$

Observe that the process K_t is deterministic. A second difficulty is that this control problem involves a constraint (24) that mixes control variables and state variables. To deal with this constraint, we introduce a related, unconstrained, problem as follows: for each function λ defined on the product of $\mathbb{R}$ by the space of probability measures on $\mathbb{R}$, which we will call from now on the *Lagrange multiplier*, consider the control problem

$$V_\lambda = \sup_{(c^A, c^P, y)} \int_0^\infty e^{-\rho t} \left(\log c_t^P + \lambda(K_t, \mathbb{P}) \left(K_t - \int y_t(\omega) c_t^A(\omega) d\mathbb{P}(\omega) \right) \right) dt.$$

We first state a result that establishes a link between the principal's value V and V_λ.

Proposition 7 *Suppose that for every Lagrange multiplier process, one can find an optimal control $u_\lambda = (c_\lambda^A, c_\lambda^P, y_\lambda)$ such that*

$$V_\lambda = \int_0^\infty e^{-\rho t} \left(\log c_{\lambda,t}^P + \lambda(K_t, \mathbb{P}) \left(K_t - \int y_{\lambda,t}(\omega) c_{\lambda,t}^A(\omega) d\mathbb{P}(\omega) \right) \right) dt.$$

Moreover, suppose that there exists $\lambda_0(.)$ such that $K_t = \int y_{\lambda_0,t}(\omega) c_{\lambda_0,t}^A(\omega) d\mathbb{P}(\omega)$, i.e. $u_{\lambda_0} \in \mathcal{K}$. Then, $V = V_{\lambda_0}$ and u_{λ_0} solves the constrained principal problem.

We are now in a position to derive the HJB equation associated with the unconstrained problem.

Proposition 8 *If the value function of the principal is sufficiently regular,*[15] *it satisfies the following HBJ equation:*

$$\rho V(K,\mathbb{P}) = \sup_{c^A(\dot{}),c^P,y(\dot{})} \left\{ \log c^P + \lambda(K,\mathbb{P}) \left(K - \int c^A(\omega) y(\omega) d\mathbb{P}(\omega) \right) \right. \tag{25}$$

$$+V_K(K,\mathbb{P}) \left(\mu K - c^P - \int c^A(\omega) d\mathbb{P}(\omega) \right)$$
$$\left. + \int \partial_\omega \delta V[K,\mathbb{P}](\omega)(\rho\omega - \log c^A(\omega)) d\mathbb{P}(\omega) + \int \partial_{\omega\omega} \delta V[K,\mathbb{P}](\omega) \frac{\sigma^2}{2} y^2(\omega) d\mathbb{P}(\omega) \right\},$$

where δV denotes the Gateaux gradient of V with respect to the measure $\mathbb{P}$ and ∂_ω (respectively $\partial_{\omega\omega}$) denote its first (respectively second) partial derivative in ω, while λ denotes the Lagrange multiplier associated with the capital allocation constraint.

Inspired by classical verification theorems for stochastic control of diffusion processes, we prove the following result, which is a consequence of the Itô formula given in appendix for functions defined on the Wasserstein space.

Proposition 9 *(Verification Theorem) Let $\lambda(.)$ be a Lagrange multiplier, and $v^\lambda(K)$ be C^1 with respect to K. Suppose that v^λ is a solution to (25) with the transversality condition $\lim_{t\to+\infty} e^{-\rho t} v^\lambda(K_t, P_{\omega_t^\mu}) = 0$ and there exists a control u^*_λ attaining the maximum in (25). Then $v^\lambda = V_\lambda$. Moreover, if there is a Lagrange multiplier λ_0 such that $u^*_{\lambda_0} \in \mathcal{K}$ then $v^{\lambda_0} = V$.*

4.3 A guess-and-verify approach

We now guess the form of the solution to the optimal control problem and show that the corresponding value function satisfies the Hamilton-Jacobi-Bellman equation (25), so that the guess is the actual solution of the problem.

4.3.1 A restricted control problem

Guided by the characterization of first-best allocations, we conjecture that optimal controls satisfy

$$C_t^P = \gamma^P K_t, C_t^A(\omega) = \gamma^A \exp(\rho\omega_t), \tag{26}$$

where γ^P and γ^A are positive constants. We also posit that $y_t(\omega) \equiv y$ is constant. This is what we call the restricted principal's problem. In the restricted problem, the feasibility constraint (24) gives for all $t \geq 0$,

$$K_t = y\gamma^A \int \exp(\rho\omega_t) d\mathbb{P}(\omega) = y\gamma^A Z_t,$$

where

$$Z_t = \mathbb{E}[\exp(\rho\omega_t)]. \tag{27}$$

[15]By this we mean that it is differentiable in K, Gateaux differentiable in $\mathbb{P}$ and that its Gateaux gradient in $\mathbb{P}$ is twice differentiable with respect to ω.

As a consequence, the ratio $\frac{Z_t}{K_t} \equiv z$ must be constant, and γ^A must be equal to the inverse of yz. Substituting (26) into (22) and using $\gamma^A = 1/(yz)$ and (27), we obtain the growth rate of capital

$$g := \mu - \gamma^P - \frac{1}{y}, \tag{28}$$

Since the ratio $\frac{Z_t}{K_t}$ is constant, the growth rates of K_t and Z_t must be equal. Thus

$$\frac{dZ_t}{Z_t} = \mathbb{E}[\rho d\omega_t] + \frac{\rho^2\sigma^2 y^2}{2}dt = gdt,$$

which implies the constraint:

$$\mu - \gamma^P - \frac{1}{y} = -\rho log\gamma^A + \frac{\rho^2\sigma^2 y^2}{2}.$$

Thus the value function of the restricted problem can be computed as

$$V(K, \mathbb{P}) = \frac{\log K}{\rho} + v(z). \tag{29}$$

where the function $v(z)$ satisfies

$$\rho v(z) = \sup_{y,\gamma^P}[\log\gamma^P + \frac{\mu - \gamma^P - \frac{1}{y}}{\rho}],$$

under the constraint:

$$\mu - \gamma^P - \frac{1}{y} = \rho logyz + \frac{\rho^2\sigma^2 y^2}{2}.$$

We obtain the next proposition, whose proof is in the appendix:

Proposition 10 *Let* $z = \frac{\int \exp(\rho\omega)d\mathbb{P}(\omega)}{K}$. *For* $0 < z < z_{max}$, *the value function of the restricted principal's problem writes as*

$$V(K, \mathbb{P}) = \frac{\log K}{\rho} + v(z). \tag{30}$$

where the function $v(z)$ *satisfies*

$$\rho v(z) = \sup_y[\log(\mu - \frac{1}{y} - \rho\log yz - \frac{\rho^2\sigma^2 y^2}{2}) + \log yz + \frac{\rho\sigma^2}{2}y^2]. \tag{31}$$

The solution to this problem is denoted $y(z)$. *The corresponding propensities to consume are*

$$\gamma^P(z) = \rho - \frac{1}{y(z) + \rho\sigma^2 y(z)^3}, \tag{32}$$

for the principal and $\gamma^A(z) = \frac{1}{zy(z)}$ *for the agent.*

In line with the incentive compatibility condition (20), which implies that y must be strictly positive as long as agents hold strictly positive capital, inspection of (30) reveals that the solution of the restricted principal's problem involves $y > 0$: in the optimal allocation, agents must bear some of their idiosyncratic risk.

4.3.2 The general case

We now show that the value function of the restricted problem satisfies the Bellman equation (25) and thus solves the complete problem. To do so, we substitute $V(K,\mathbb{P})$ from (30) in the HJB equation (25). We first compute the partial derivatives of order one:

$$V_K = \frac{1-\rho z v'(z)}{\rho K}, \delta V = \rho exp(\rho\omega)\frac{v'(z)}{K}$$

and then the derivatives of the Gateaux gradient of V:

$$\partial_\omega(\delta V) = \rho\delta V, \partial_{\omega\omega}(\delta V) = \rho^2\delta V$$

The Bellman equation becomes

$$\log K + \rho v(z) = \sup[log(\gamma^P K) + \lambda(K - \int \gamma^A(\omega)y(\omega)exp(\rho\omega)d\mathbb{P})$$
$$+[\frac{1}{\rho} - zv'(z)][\mu - \gamma_P - \frac{\int \gamma^A(\omega)exp(\rho\omega)d\mathbb{P}}{K}] + v'(z)\int \rho exp(\rho\omega)(-\log\gamma^A(\omega) + \frac{\rho\sigma^2}{2}y^2(\omega))d\mathbb{P}(\omega)].$$

Note that all the terms involving $\gamma^A(\omega)$ and $y(\omega)$ are multiplied by the same function of ω, namely the product of $exp(\rho\omega)$ by the density of $\mathbb{P}(\omega)$. Thus the pointwise maximum is attained for the same couple (y,γ^A), independently of ω. This implies that the solution is the same as that of the restricted problem, where we have assumed y and γ^A constant. Thus we can replace γ^A by $\frac{1}{yz}$ and γ^P by $(\mu - \frac{1}{y} - \rho\log yz - \frac{\rho^2\sigma^2y^2}{2})$ and the Bellman equation simplifies into:

$$\rho v(z) = \sup_y[\log(\mu - \frac{1}{y} - \rho\log yz - \frac{\rho^2\sigma^2y^2}{2}) + \log yz + \frac{\rho\sigma^2}{2}y^2],$$

which is the definition of the function $v(z)$. Thus we have established that the value function in (30) satisfies the Bellman equation of the full problem. Thus we have established the main result of our paper:

Proposition 11 *The value function of the full problem is*

$$V(K,\mathbb{P}) = \frac{logK}{\rho} + v(z),$$

where $z = \frac{\int exp(\rho\omega)d\mathbb{P}(\omega)}{K})$ *and the function* v *is defined by equation (82). The solution is such that:*

$$k(\omega) = \frac{exp(\rho\omega)}{z}, C^A(\omega) = \frac{exp(\rho\omega)}{zy(z)}, \gamma^P = \rho - \frac{1}{y(z) + \rho\sigma^2 y^3(z)},$$

where $y(z)$ *is the solution of (31).*

4.4 Properties of second best allocations

Taking stock of the analysis above, the next proposition summarizes the properties of optimal information constrained allocations. These properties are drastically simplified by the fact that date t allocations only depend on two state variables, namely the capital stock K_t and the ratio $z_t \equiv \frac{\int exp(\rho\omega_t)d\mathbb{P}(\omega)}{K_t}$. Moreover, along the optimal trajectories, this ratio is constant over time: $z_t \equiv z$, and optimal controls can all be expressed as functions of $y = y(z)$, the solution of (31).

Proposition 12 *Second best optimal allocations are such that:*

1. *Capital grows at a constant rate*
$$g = \mu - \rho - \frac{\rho\sigma^2 y}{1 + \rho\sigma^2 y^2}, \tag{33}$$
which is lower that the first best growth rate $\mu - \rho$.

2. *Agents' continuation utilities follow a drifted Brownian motion:*
$$\omega_t = \omega + \left(\frac{g}{\rho} - \frac{\rho\sigma^2 y^2}{2}\right) t + \sigma y B_t. \tag{34}$$

3. *At each date t, the principal consumes a constant fraction of the capital stock: $C_t^P = \gamma^P K_t$, where*
$$\gamma^P = \rho - \frac{1}{y + \rho\sigma^2 y^3}. \tag{35}$$

4. *At each date t, an agent consumes a constant fraction of $exp(\rho\omega_t)$: $C_t^A(\omega) = \gamma^A exp(\rho\omega_t)$, where*
$$\gamma^A = exp[-\frac{\mu - \rho}{\rho} + \frac{\rho^2\sigma^2 y}{1 + \rho\sigma^2 y^2} + \frac{\rho^2\sigma^2 y^2}{2}].$$

Property 1 shows that frictions reduce growth. This reflects incentive constraints, which restrict investment. When $\sigma = 0$, there is no incentive problem and the growth rate is equal to its first best level.

Property 2 implies that the cross section of agents' continuation payoffs gets more dispersed as time goes by. Even if all agents are ex ante identical, inequality necessarily increases over time, due to incentive compatibility constraints. Moreover, there is a simple relation between the continuation utility of an agent at date t and its performance over $(0, t)$. Indeed, the average productivity of the agent over $(0, t)$ is just $\mu + \sigma\frac{B_t}{t}$. Optimal compensation implies a simple, affine, relation between the continution utility ω_t and this performance measure, similarly to Holmstrom Milgrom (1985).

Finally, Properties 3 and 4 are similar to the first best case. This simplicity is due to our assumption that utilities are logarithmic and aggregate productity is constant. A characterization of optimal second best allocations in more general cases is probably much more difficult.

The above properties are parametrized by the sensitivity of agent's continuation utility to performance, y. Varying y does not qualitatively alter these properties, but it generates quantitative changes, e.g., in growth rates or principal's share of consumption. Below, we show how the information constrained Pareto frontier can be written as a function of y.

4.5 Information Constrained Pareto Frontier

The above analysis yields a characterization of the information constrained Pareto frontier in the space of equivalent permanent consumptions. To facilitate its representation, we focus on the case in which all agents start with the same continuation pay-off ω. We also take $K = 1$. In this case, $V(K, \mathbb{P})$ in (30) simplifies to $v(\exp(\rho\omega))$. The continuation utility of the agent is

$$\omega = E[\int_0^\infty e^{-\rho t}\log(C_t^A)dt] = \frac{\log\frac{1}{y}}{\rho} + \frac{\mu - \gamma^P - \frac{1}{y}}{\rho^2} - \frac{\sigma^2 y^2}{2},$$

while that of the principal is

$$v(\exp(\rho\omega)) = \int_0^\infty e^{-\rho t}\log(C_t^P)dt = \frac{\log\gamma^P}{\rho} + \frac{\mu - \gamma^P - \frac{1}{y}}{\rho^2}.$$

Substituting γ^P from (32) into ω and $v(\omega)$ enables us to parameterize the Pareto frontier as a function of y alone. We obtain that the equivalent permanent consumption of the agent is

$$\exp(\rho\omega) = \frac{1}{y}exp[\frac{g}{\rho} - \frac{\rho\sigma^2 y^2}{2}], \tag{36}$$

where g is the growth rate given in (33), while the equivalent permanent consumption of the principal is

$$\exp\rho v(\exp(\rho\omega)) = (\rho - \frac{1}{y + \rho\sigma^2 y^3})exp[\frac{g}{\rho}]. \tag{37}$$

(36) reflects that each agent consumes a fraction $\frac{1}{y}$ of its capital under management, which grows at average rate g, with volatility σy generating a risk premium, and is discounted at rate ρ. Similarly (37) reflects that the principal consumes a fraction $(\rho - \frac{1}{y+\rho\sigma^2 y^3})$ of the capital stock, but is not impacted by any risk, so that unlike in (36) there is no risk premium. As mentioned above, when $\sigma = 0$ there is no incentive problem. Correspondingly (36) and (37) reduce to

$$exp(\rho\omega) + exp(\rho v(\exp(\rho\omega))) = \rho exp[\frac{\mu - \rho}{\rho}],$$

the equation of the first best Pareto frontier (16), evaluated in the case in which all agents have the same utility ω and $K = 1$. It reflects that the total surplus $(\rho exp[\frac{\mu-\rho}{\rho}])$ must be shared between the principal and the agents.

5 Implementation by money and taxes

The direct revelation mechanism characterized above is completely centralized: all agents report to the principal, who then reallocates goods among agents. We now show that a more decentralized implementation is possible, in which the allocation of goods results from the equilibrium of a competitive market. In that implementation, the principal does not intervene in the reallocation of goods among agents, and relies only on monetary policy (via the inflation rate π) and fiscal policy (via the tax rate τ.)

Our analysis proceeds in two steps. First, we characterize the equilibrium allocation arising for a given policy (π, τ). There we show how the choice of π and τ determines the agents' and principal's equilibrium consumption processes, as well as the equilibrium growth of output and money supply. Second, we we show that, for any second best allocation, i.e., agents' and principal's second best consumption processes, there exists a policy (π, τ) for which this allocation is the equilibrium allocation. Thus, as explained below, we obtain a form of second welfare theorem.

5.1 Equilibrium

Our equilibrium analysis proceeds in three steps. First, we characterize the optimal consumption and investment of an agent for a given public policy (π, τ). Second, we spell out the market clearing condition, stating that, at each point in time, the supply of goods is equal to the demand for goods. Third, we derive the equilibrium growth rate induced by policy (τ, π) for output and money supply, as well as the equilibrium consumption share of the principal.

5.1.1 Agent's optimal policy

At $t = 0$, the principal endows each agent with money m_0 and commits to constant inflation rate π and tax rate τ. Normalizing p_0 to 1, the price of the good in money at time t is $p_t = \exp(\pi t)$. Agents hold capital (k_t) and money (m_t), so an agent's real wealth at time t is

$$e_t = k_t + \frac{m_t}{p_t}. \tag{38}$$

The dynamics of the capital holdings k_t of a given agent is given by:

$$dk_t = k_t(\mu dt + \sigma dB_t) - c_t dt - ds_t, \tag{39}$$

where ds_t denotes the agent's sales (purchases if negative) on the good market. Similarly, the dynamics of the agent's real money balances are

$$d(\frac{m_t}{p_t}) = ds_t - (\pi \frac{m_t}{p_t} + \tau e_t)dt, \tag{40}$$

Adding (39) and (40), ds_t cancels out and we obtain the dynamics of the agent's wealth

$$de_t = k_t(\mu dt + \sigma dB_t) - [c_t + \tau e_t + \pi(e_t - k_t)]dt. \tag{41}$$

Since there are no transaction costs, the agent can costlessly continuously rebalance her portfolio of money and capital and the only constraint is the wealth constraint. So e_t is the agent's state variable, while k_t and c_t can be viewed as the control variables. Equation (41) shows that the change in wealth of an agent is equal to output, minus consumption, taxes, and the decline in the real value of money holdings due to inflation. The latter can be interpreted as an inflation tax. Equation (41) and Ito's lemma imply that the value function $u(e)$ of the agents satisfies the following Bellman equation

$$\rho u(e) = Max_{k,c}[logc + u'(e)[\mu k - c - \tau e - \pi(e - k)] + \frac{\sigma^2 k^2}{2} u''(e)]. \tag{42}$$

The first order condition with respect to c is

$$\frac{1}{c} = u'(e).$$

The first order condition with respect to k is

$$k = \frac{\mu + \pi}{-\frac{u''(e)}{u'(e)}\sigma^2}.$$

Homogeneity implies that the value function is an affine transformation of $\log(e)$:

$$u(e) = \frac{\log(e)}{\rho} + u(1), \tag{43}$$

which implies

$$u'(e) = \frac{1}{\rho e}, u''(e) = -\frac{1}{\rho e^2}. \tag{44}$$

So the first order conditions yield

$$c = \rho e, \tag{45}$$

and

$$k = \frac{\mu + \pi}{\sigma^2} e. \tag{46}$$

That consumption and capital are constant fractions of wealth stems from the logarithmic utility specification. Denoting

$$x := \frac{\mu + \pi}{\sigma^2}, \tag{47}$$

the optimal portfolio choice of the agent is to invest a fraction x of her wealth in the risky asset and a fraction $1 - x$ in money, the safe asset. Condition (47) shows that the fraction of her wealth an agents invests in the risky asset is increasing in the inflation rate π, which determines the rate of return on money holdings.

5.1.2 Market clearing

Market clearing requires that the aggregate supply of goods by the agent be equal to the consumption of goods by the principal

$$\int_i \left(ds_t^i\right) di = c_t^P dt. \tag{48}$$

First, consider the left-hand side of (48). Since optimality requires a constant ratio of capital to wealth, each agent must buy or sell capital to equalize the growth rate of capital to that of wealth:

$$\frac{dk_t}{k_t} \equiv \frac{de_t}{e_t}.$$

The dynamics of an agent's capital holdings (39) $c_t = \rho e_t$, and $k_t = x e_t$ imply

$$\frac{dk_t}{k_t} = (\mu dt + \sigma dB_t) - \frac{\rho}{x} dt - \frac{ds_t}{k_t}.$$

The dynamics of an agent's wealth, (41), combined with $c_t = \rho e_t$, and $k_t = x e_t$ imply

$$\frac{de_t}{e_t} = x(\mu dt + \sigma dB_t) - [\rho + \tau + \pi(1 - x)]dt.$$

Equating the two yields

$$\frac{ds_t}{k_t} = (1 - x)(\mu dt + \sigma dB_t) - \frac{\rho}{x} dt + [\rho + \tau + \pi(1 - x)]dt,$$

which determines individual good sales ds_t

$$ds_t = [(\mu - \frac{\rho}{x} + \pi)(1 - x) + \tau]k_t dt + \sigma(1 - x)k_t dB_t,$$

and aggregate sales

$$\int \left(ds_t^i\right) di = \Big[(\mu - \frac{\rho}{x} + \pi)(1-x) + \tau\Big] K_t dt. \tag{49}$$

Second, turn to the right-hand side of (48), i.e., the consumption of the principal. By the budget constraint of the principal, this consumption is equal to the sum of the seigneurage and fiscal revenues, that is

$$c_t^P = g_M \int_i \left(\frac{m_t^i}{p_t}\right) di + \tau \int_i e_t^i di, \tag{50}$$

where g_M is the growth rate of the money supply. Now, by (38) and (46)

$$\frac{m_t}{p_t} = e_t - k_t = k_t \frac{1-x}{x}.$$

Substituting in (50) we have

$$c_t^P dt = \left(g_M \frac{1-x}{x} + \frac{\tau}{x}\right) K_t dt. \tag{51}$$

Equating (49) and (51), the market clearing condition is

$$(\mu - \frac{\rho}{x} + \pi)(1-x) + \tau = g_M \frac{1-x}{x} + \frac{\tau}{x}. \tag{52}$$

By (47), $\mu + \pi = x\sigma^2$. So (52) writes

$$(x\sigma^2 - \frac{\rho}{x})(1-x) = g_M \frac{1-x}{x} + \tau \frac{1-x}{x}.$$

Simplifying, this yields the rate of growth of money supply which must prevail in equilibrium when the government follows policy (τ, π).

$$g_M = \sigma^2 x^2 - \rho - \tau. \tag{53}$$

5.1.3 Equilibrium growth rate and principal's consumption

By definition, the growth rate of money is $g_M = g + \pi$. Equating this to (53) we obtain the equilibrium growth rate obtaining for policy (τ, π).

$$g = \sigma^2 x^2 - \rho - \tau - \pi. \tag{54}$$

By (51), the principal's consumption share of capital is

$$\gamma^P = g_M \frac{1-x}{x} + \frac{\tau}{x}.$$

Substituting (53), we have

$$\gamma^P = \left(\sigma^2 x^2 - \rho\right) \frac{1-x}{x} + \tau. \tag{55}$$

Summarizing the results derived above, we obtain the next proposition:

Proposition 13 *When the principal commits to a constant inflation rate π and a constant tax rate τ, equilibrium is as follows:*

- *Each agent consumes a constant fraction ρ of her wealth.*
- *Each agent holds a constant fraction $x = \frac{\mu+\pi}{\sigma^2}$ of her wealth in the risky asset and the complementary fraction in money.*
- *The growth rate of the money supply is $g_M = \sigma^2 x^2 - \rho - \tau$.*
- *The growth rate of output is $g = \sigma^2 x^2 - \rho - \tau - \pi$, and*
- *The principal's consumption share is $\gamma^P = \left(\sigma^2 x^2 - \rho\right) \frac{1-x}{x} + \tau$.*

The proposition clarifies that for any couple of policy variables π and τ, there is a unique stationary equilibrium allocation associated with the variables x, g_M, g, and γ^P characterized in the proposition.[16] However, we show below that only a subset of equilibrium allocations correspond to information constrained optimal allocations, in particular the laissez faire allocation with $\tau = 0$ is not information constrained optimal.

5.2 Implementation

To implement a second best allocation we need to find τ and π such that i) the dynamics of $u(e_t)$ in equilibrium is equal to that of ω_t in that second best allocation and ii) the consumption of the principal in equilibrium is equal to the consumption of the principal in that second best allocation. Let us look first at the identification of the utility of the agent in the second best and in equilibrium. Proposition 12 implies that in the second best the dynamics of an agent's utility is

$$d\omega_t = \left(\frac{g}{\rho} - \frac{\rho\sigma^2 y^2}{2}\right) dt + \sigma y dB_t, \tag{56}$$

where

$$g = \mu - \rho - \frac{\rho\sigma^2 y}{1 + \rho\sigma^2 y^2}. \tag{57}$$

Turning to the equilibrium, by Ito's Lemma the dynamics of an agent's utility is

$$du(e) = u'(e)de + \frac{1}{2}u''(e)\left(de\right)^2.$$

By (41), (44), $c = \rho e$, and $k = xe$, this is

$$du(e) = \frac{1}{\rho}\left[x\mu - (\rho + \tau + \pi(1-x))\right] dt - \frac{\sigma^2 x^2}{2\rho} dt + \frac{\sigma x}{\rho} dB_t. \tag{58}$$

For the equilibrium to implement the second best, we need to identify (56) and (58). For the Brownian term to be the same in the two equations, we need

$$x = \rho y. \tag{59}$$

[16] Instead of defining the principal's policy in terms of monetary (π) and fiscal (τ) policies, one could have equivalently defined it in terms of monetary (π) and budget (γ^P) policy. Because of the principal's budget constraint, stating the principal's consumption must be equal to the sum of seigneurage and tax revenues, setting τ is equivalent to setting γ^P

Substituting the value of x from (47) into (59) the equality becomes

$$\frac{\mu+\pi}{\sigma^2} = \rho y.$$

So, to ensure that the equilibrium implements the second best allocation parametrized by y, the principal must set inflation

$$\pi = \sigma^2 \rho y - \mu. \tag{60}$$

Once the two Brownian terms are equal, to identify (56) and (58) we need to identify the drifts, i.e., we must have

$$\frac{g}{\rho} - \frac{\rho\sigma^2 y^2}{2} = \frac{1}{\rho}\left[x\mu - (\rho + \tau + \pi(1-x))\right] - \frac{\sigma^2 x^2}{2\rho}. \tag{61}$$

After a few manipulations, explicited in the proof in the appendix, this is equivalent to

$$\tau = \sigma^2 \rho^2 y^2 \left(\frac{1-\sigma^2 y(1-\rho y)}{1+\sigma^2 \rho y^2}\right), \tag{62}$$

which, as shown in the proof in the appendix, also implies that the consumption of the principal is the same in equilibrium and in the second best. So we can state our next proposition:

Proposition 14 *Any second best allocation parametrized by y can be decentralized as the competitive equilibrium associated with public policy (π, τ), where*

$$\pi = \sigma^2 \rho y - \mu, \tag{63}$$

and

$$\tau = \sigma^2 \rho^2 y^2 \left(\frac{1-\sigma^2 y(1-\rho y)}{1+\sigma^2 \rho y^2}\right), \tag{64}$$

As noted after Proposition 10, the optimal allocation involves $y > 0$. By (64) this implies that tax rates implementing second best allocations must be strictly positive. Thus, a laissez-faire policy with $\tau = 0$ cannot implement a second best allocation. That is non zero taxes or subsidies are necessary to implement second best optimal allocations.

Proposition 14 is a form of second welfare theorem: For any Pareto optimal allocation, there exists a tax and monetary policy such that the competitive equilibrium yields the Pareto optimal allocation. But there are major differences between Proposition 14 and the classical second welfare theorem: First, in the classical welfare theorem, markets are perfect and complete. In contrast, in our analysis there are asymmetric information frictions, implying that markets are endogenously incomplete. Second, the classical second welfare theorem considers lumpsum taxes, which don't distort agents' behaviour. In contrast, in our analysis taxes are linear in wealth, and in conjunction with inflation, optimally affect agents' behaviour.

A key step to obtain Proposition 14 is equation (59) which states that $x = \rho y$. A priori, y and x are conceptually different objects. The former is the exposure of agents to their idiosyncratic risk in the optimal mechanism. The latter is the structure of agents' portfolio in market equilibrium, which is affected by π since inflation determines the relative attractiveness of the safe asset. To implement the optimal mechanism, inflation π must be set such that (59), because this ensures that agents have the same risk exposure in equilibrium and in the optimal mechanism.

When g_M is negative (monetary contraction), the principal uses taxes to finance his consumption and to "pump out" money from agents. This case is similar to our two-date model in that the value of money equals the sum of future taxes minus future public expenditures (primary surpluses).[17] However, for different parameter values, g_M can also be positive (monetary expansion), in which case the money stock grows without limit, and the primary surplus is negative. Then money can be viewed as a bubble: its value is positive, even though taxes are insufficient to cover the consumption of the principal. It is even possible that τ be negative, implying that the principal subsidizes the agents by distributing them part of the money he issues (helicopter money). This is sustainable when growth rate is sufficiently high.

6 Link to Optimal Tax Theory

The recent literature on the New Dynamic Public Finance (see Kocherlakota 2009 and Golosov et al. 2007) has shown that capital should be taxed in a dynamic economy when individual labor productivities are not publicly observable and are hit by idosyncratic shocks over time. We can show a similar result in our model, where it is capital income that is not publicly observable. Following the standard approach in optimal tax theory, consider the principal as a government selecting the tax scheme to maximizes intertemporal social welfare, taking as given that public expenditures are fixed to a given fraction γ_0^P of the capital stock. Note that here the consumption of the principal is exogenous while in our analysis it is endogenous. Using the direct mechanism approach, we can see that this problem amounts to finding the maximum ω_0 that is feasible for a given initial capital $K_0 = 1$ and a net expected productivity of capital $\mu - \gamma_0^P$. Note another difference with traditional tax theory: we allow the government to finance its deficits (or absorb its surpluses) by adjusting its monetary policy: money issuance (which can be negative) is determined by the difference between public expenditures and tax receipts. The problem then amounts to finding the highest continuation pay-off for agents on the second best Pareto frontier when μ is replaced by $\mu - \gamma_0^P$. Adapting the above analysis, one obtains the next proposition:

Proposition 15 *When government expenditures are a fixed fraction γ_0^P of the aggregate capital stock, the optimal policy mix (fiscal+monetary policies) consists of:*

1. *A linear wealth tax at rate: $\tau = \gamma_0^P - (1-x)(\sigma^2 x - \frac{\rho}{x})$, where x is the only positive solution of the cubic equation $\frac{\sigma^2}{\rho}x^3 + x = 1$,*
2. *A constant money issuance rate $g_M = \sigma^2 x - \frac{\rho}{x}$.*
3. *Moreover, the economy grows at a constant rate $g = \mu - \gamma_0^P - \frac{\rho}{x}$, and the inflation rate is also constant $\pi = \sigma^2 x - \mu + \gamma_0^P$.*

Note that taxes and inflation are an increasing function of government expenditures, while the money issuance rate only depends on σ and ρ.

[17] There is no discounting as the interest rate on money is 0.

7 Conclusion

In this paper we analyze capital allocation and risk sharing between a principal and many agents. We assume that agents privately observe their individual output and can secretly consume some of it, as in Bolton and Scharfstein (1990). To provide agents with incentives to reveal truthfully their output, the optimal dynamic mechanism allocates more capital and consumption to agents with better performance. Thus, while there is no aggregate risk, incentive compatibility precludes perfect insurance. Assuming logarithmic utility enables us to fully characterize the optimal dynamics of capital and consumption as well as the distribution of continuation utilities across agents.

Moreover, we show that the optimal dynamic mechanism can be implemented by market equilibrium with appropriatey chosen inflation and tax rates. Inflation determines the attractiveness of the safe asset relative to the risky asset, and thus agents' holdings of the latter. An appropriately chosen inflation rate gives agents the same risk exposure in equilibrium as in the optimal mechanism, so that the former implements the latter.

This implementation result is a form of second welfare theorem: For any Pareto optimal allocation, there exists a fiscal and monetary policy implementing that allocation in equilibrium. However, while in the classical welfare theorem, markets are perfect and complete, in our analysis markets are endogenously incomplete because of information asymmetry. Moreover, while in the classical second welfare theorem taxes are lumpsum so that they don't distort agents' behaviour, in our analysis taxes depend on wealth and optimally affect agents' behaviour. Finally note that we don't obtain a first theorem of welfare. Only a subset of the equilibria arising in our setting are information constrained Pareto optimum. In particular, the laissez-faire equilibrium, obtaining with zero taxation, is not information constrained Pareto optimal.

Appendix A: Proofs

Proof of Proposition 1: The Lagrangian is

$$E\left[U(C_1^s)+U(C_2^s)\right]+\lambda\left(1-E\left[C_1^s+C_2^s\right]\right),$$

where λ is the multiplier of the resource constraint. The first order condition with respect to C_t^s is $U'(C_t^s)=\lambda, \forall s,t$. So consumption is constant across types s and periods. Binding the resource constraint this yields $C_1^s=\frac{1}{2}$.

QED

Proof of Proposition 2: The Lagrangian is

$$\begin{aligned}&E\left[U(C_1^s)+U(C_2^s)\right]+\lambda\left(1-E\left[C_1^s+C_2^s\right]\right)\\&+\nu\left[U(C_1^H)+U(C_2^H)-U(C_1^L+2\sigma)-U(C_2^L)\right],\end{aligned}$$

where λ is the multiplier of the resource constraint and ν the multiplier of the incentive constraint. The first order condition with respect to C_t^H is:

$$U'(C_t^H)=\frac{\lambda}{1+\nu}, \forall t. \tag{65}$$

So $C_1^H=C_2^H$. The first order condition with respect to C_1^L is:

$$U'(C_1^L)-\nu U'(C_1^L+2\sigma)=\lambda. \tag{66}$$

The first order condition with respect to C_2^L is:

$$U'(C_2^L)=\frac{\lambda}{1-\nu}, \tag{67}$$

which, with (65), implies $C_2^L<C_2^H$. Now, (66) rewrites as

$$(1-\nu)U'(C_1^L)+\nu\left(U'(C_1^L)-U'(C_1^L+2\sigma)\right)=\lambda.$$

That is

$$U'(C_1^L)=\frac{\lambda}{1-\nu}-\frac{\nu}{1-\nu}\left(U'(C_1^L)-U'(C_1^L+2\sigma)\right),$$

which implies

$$U'(C_1^L)<\frac{\lambda}{1-\nu}.$$

Together with (67) this implies $C_2^L<C_1^L$.

QED

Proof of Proposition 3 Denoting by λ the multiplier of the resource constraint and μ the one of the IC constraint, the first order conditions give $C_1^H = C_2^L = \frac{1+\mu}{\lambda}, C_2^L = \frac{1-\mu}{\lambda}$, and

$$\frac{1}{C_1^L} - \frac{\mu}{C_1^L + 2\sigma} = \lambda.$$

Since the IC constraint is binding we can write $C_1^L + 2\sigma = \frac{C_1^H C_2^H}{C_2^L} = \frac{(1+\mu)^2)}{\lambda(1-\mu)}$. Similarly the resource constraint is binding, giving $C_1^L = 2 - \frac{3+\mu}{\lambda}$ By eliminating C_1^L between these two equations, we obtain $\frac{1}{\lambda} = \frac{(1+\sigma)(1-\mu)}{2}$ The expressions of $C_1^H = C_2^H, C_2^L$ are immediately deduced. Finally, the cubic equation in μ is obtained by plugging the expression of λ into the first order condition with respect to C_1^L.

Now we turn to the proof of the 4 properties stated in the proposition:

1. $C_1^H - C_1^L = 2\sigma - (1+\sigma)\mu(1+\mu)$ Using the equation defining μ we can write $\mu(1+\mu) = \frac{1+3\mu}{1+\mu}\frac{\sigma}{1+\sigma}$. Since $\mu < 1$, this is smaller than $\frac{2\sigma}{1+\sigma}$. This establishes property 1.

2. $C_2^H - C_2^L = (1+\sigma)\mu > 0$.

3. $C_2^H + C_2^L = (1+\sigma)(1-\mu) < 1$ since $\mu > \frac{\sigma}{1+\sigma}$.

4. $C_2^H + C_1^H = (1+\sigma)(1-\mu^2) < 1+\sigma$.

This ends the proof of the proposition.
QED

Proof of Proposition 4: The Lagrangian of the maximization problem faced by agent s is

$$U(C_1^s) + U(k_1^s) + \lambda^s \left[R^s + m_0 - (C_1^s + k_1^s + \tau(k_1^s))\right].$$

The first order condition with respect to time 1 consumption is

$$U'(C_1^s) = \lambda^s. \tag{68}$$

The first order condition with respect to investment is

$$U'(k_1^s) = \lambda^s \left[1 + \tau'(k_1^s)\right]. \tag{69}$$

Substituting (68) in (69), and noting that $C_2^s = k_1^s$, yields

$$\frac{U'(C_2^s)}{U'(C_1^s)} = \left[1 + \tau'(k_1^s)\right].$$

Since in the optimal mechanism $C_1^H = C_2^H$ and $C_1^L > C_2^L$, in the implementation we must have $\tau'(k_1^H) = 0$ and $\tau'(k_1^H) > 0$.

Binding the agent's goods budget constraint (1) and aggregating across agents yields

$$E\left[C_1^s + k_1^s + S^s\right] = E\left[R^s\right]. \tag{70}$$

Now the binding time 1 resource constraint faced by the planner is

$$E\left[C_1^s + k_1^s\right] = 1. \tag{71}$$

(70) and (71) imply

$$E\left[S^s\right] = 0,$$

which means that the goods market clears at time 1.

QED

Proof of Proposition 5: Denoting by β the Lagrange multiplier associated to the constraint on capital and λ^i the one associated to the promise keeping constraint for agent i, the Lagrangian writes, up to a constant:

$$L = \int_0^\infty e^{-\rho t} \log c_t^P dt + \int_0^1 \int_0^\infty \lambda^i e^{-\rho t} \log c_t^i dt di - \beta \int_0^\infty e^{-\mu t} [\int_0^1 c_t^i di + c_t^P] dt.$$

We can derive the first order conditions:

$$\frac{e^{-\rho t}}{c_t^P} = \beta e^{-\mu t},$$

with respect to c_t^P and

$$\frac{\lambda_t^i e^{-\rho t}}{c_t^i} = \beta e^{-\mu t},$$

with respect to c_t^i. This yields $c_t^P = \frac{exp(\mu-\rho)t}{\beta}$ and $c_t^i = \frac{\lambda^i exp(\mu-\rho)t}{\beta}$. Multiplying by ρ the promise keeping condition, we obtain

$$\rho\omega^i = \rho \int_0^\infty e^{-\rho t} \log c_t^i dt = \log(\frac{\lambda^i}{\beta}) + \frac{\mu - \rho}{\rho}.$$

Thus $\frac{\lambda^i}{\beta} = c_0^i = \gamma^A exp(\rho\omega^i)$. Now, we can multiply by ρ the constraint on capital, giving

$$\rho K = \gamma^P K + \gamma^A \int_0^1 exp(\rho\omega^i) di.$$

thus we can express γ^P as a function of the ratio of $\int_0^1 exp(\rho\omega^i)di$ and K, which we denote by z :

$$\gamma^P = \rho - \gamma^A \frac{\int_0^1 exp(\rho\omega^i)di}{K} = \rho - \gamma^A z.$$

Total consumption is thus $\rho K e^{(\mu-\rho)t}$. The dynamics of capital is:

$$\dot{K}_t = \mu K_t - \rho K e^{(\mu-\rho)t},$$

which gives after integration $K_t = K e^{(\mu-\rho)t}$. The optimal allocation is thus stationary: individual consumptions and aggregate capital all grow at rate $\mu - \rho$. Similarly $\rho \frac{d\omega_t^i}{dt} = \mu - \rho$.

QED

Proof of Proposition 7: Let $u = (c^A, c^P, y)$ be an admissible control, we will denote

$$J^u_\lambda = \int_0^\infty e^{-\rho t}\left(\log c^P_t + \lambda(K_t, \mathbb{P}_{\omega_t})\left(K_t - \int y_t(\omega) c^A_t(w) d\mathbb{P}_{\omega_t}(\omega)\right)\right) dt,$$

and

$$J^u = \mathbb{E}\left[\int_0^\infty e^{-\rho t}\log(c^P_t)\, dt\right] = \int_0^\infty e^{-\rho t}\hat{c}^P_t\, dt$$

For every Lagrange multiplier λ, we have $V_\lambda = J^{u_\lambda}_\lambda \geq J^u_\lambda$. In particular, for $\lambda = \lambda_0$ and $u \in \mathcal{K}$, we have $J^{u_{\lambda_0}} = J^{u_{\lambda_0}}_{\lambda_0} \geq J^u_{\lambda_0} = J^u$ and since $u_{\lambda_0} \in \mathcal{K}$, the proof is complete.
QED

Proof of Proposition 9: Fix $\mu \in \mathbb{P}_2(\mathbb{R})$ and a Lagrange multiplier λ and consider some arbitrary control $u(K_t, \mathbb{P}_{\omega_t}, \omega_t)$. We apply Itô's formula to $v^\lambda(K_t, \mathbb{P}_{\omega^\mu_t})$ between $s = 0$ and $s = t$ for $t > 0$.

$$\begin{aligned}
e^{-\rho t} v^\lambda(K_t, \mathbb{P}_{\omega^\mu_t}) = v(K,\mu) &+ \int_0^t e^{-\rho s}\left(-\rho v^\lambda(K_s, \mathbb{P}_{\omega^\mu_s}) + v_K(K_s, \mathbb{P}_{\omega^\mu_s})\left(\mu K - c^P - \int c^A(\omega) d\mathbb{P}_{\omega^\mu_s}(\omega)\right)\right) ds \\
&+ \int_0^T e^{-\rho s}\int \partial_\omega \delta v^\lambda[K_s, \mathbb{P}_{\omega^\mu_s}](\omega)(\rho\omega - \log c^A(\omega))\mathbb{P}_{\omega^\mu_s}(d\omega) \\
&+ \int_0^t e^{-\rho s}\int \partial_{\omega\omega}\delta v^\lambda[(K_s, \mathbb{P}_{\omega^\mu_s}](\omega)\frac{\sigma^2}{2} y^2(\omega)\mathbb{P}_{\omega^\mu_s}(d\omega)\, ds.
\end{aligned}$$

We deduce from the Bellman equation (25) satisfied by v^λ that

$$v^\lambda(K,\mu) \geq e^{-\rho t} v^\lambda(K_t, \mathbb{P}_{\omega^\mu_t}) + \int_0^t e^{-\rho s}\left(\log(c^P_{\lambda,s}) + \lambda(K_s, \mathbb{P}_{\omega^\mu_s})\left(K_s - \int y_{\lambda,s}(\omega) c^A_{\lambda,s}(\omega)\mathbb{P}_{\omega^\mu_s}(d\omega)\right)\right) ds.$$

Letting t tend to $+\infty$, we obtain using the transversality condition

$$v^\lambda(K,\mu) \geq \int_0^\infty e^{-\rho s}\left(\log(c^P_{\lambda,s}) + \lambda(K_s, \mathbb{P}_{\omega^\mu_s})\left(K_s - \int y_{\lambda,s}(\omega) c^A_{\lambda,s}(w) d\mathbb{P}_{\omega^\mu_s}(\omega)\right)\right) ds = J^u_\lambda.$$

Since the control is arbitrary, we obtain

$$v^\lambda(K,\mu) \geq V_\lambda.$$

On the other hand, let us apply the same Itô's argument with the control u^*_λ attaining the maximum in (25), we obtain

$$v^\lambda(K,\mu) = J^{u^*_\lambda}_\lambda \leq V_\lambda,$$

which yields that $v^\lambda = V_\lambda$. We conclude the proof by applying Proposition 7.
QED

Proof of Proposition (10): To obtain the dynamics of Z_t, we substitute $\gamma^A = 1/(yz)$ in $C^A_t(\omega) = \gamma^A \exp(\rho\omega_t)$, and then substitute the resulting expression into (23), which yields

$$d\omega_t = \log(yz)dt + \sigma y\, dB_t. \tag{72}$$

(72) and $Z_t = \mathbb{E}[\exp(\rho\omega_t)]$ yield

$$Z_t = Z_0\mathbb{E}\left[\exp(\rho\left(\log(yz)t + \sigma y B_t\right))\right] = Z_0\exp\left(\left(\rho\log(yz) + \frac{\rho^2\sigma^2 y^2}{2}\right)t\right), \tag{73}$$

which gives

$$\frac{dZ_t}{Z_t} = \left(\rho \log(yz) + \frac{\rho^2\sigma^2 y^2}{2}\right) dt. \tag{74}$$

By (28) and (74), equality of the growth rates of K_t and Z_t means that

$$\mu - \gamma^P - \frac{1}{y} = \rho \log(yz) + \frac{\rho^2\sigma^2 y^2}{2}. \tag{75}$$

The restricted principal's problem is thus characterized by the following maximization problem

$$V(K, \mathbb{P}) = \sup_{\gamma_P, y} \int_0^{+\infty} e^{-\rho t} \log(\gamma^P K_t) dt, \tag{76}$$

under the constraint (75) and the dynamics of capital

$$K_t = K exp((\mu - \gamma^P - \frac{1}{y})t). \tag{77}$$

Substituting K_t from (77) into (76), the latter writes

$$V(K, \mathbb{P}) = \sup_{\gamma_P, y} \int_0^{+\infty} [e^{-\rho t} \left(\log(\gamma_P K) + (\mu - \gamma^P - \frac{1}{y})t\right)]dt, \text{ s.t., (75)}. \tag{78}$$

Easy computations then show that (78) can be rewritten as

$$\rho V(K, \mathbb{P}) = \log K + \sup_{\gamma_P, y} \left(\log \gamma_P + \frac{\mu - \gamma^P - \frac{1}{y}}{\rho}\right), \text{ s.t., (75)}. \tag{79}$$

Using (75) we can express γ^P as a function of y and z

$$\gamma^P = \mu - \frac{1}{y} - \rho \left(\log(yz) + \frac{\rho\sigma^2 y^2}{2}\right).$$

Substituting the value of γ^P into (79), the latter writes as

$$\rho V(K, \mathbb{P}) = \log K + \sup_y \left(\log\left(\mu - \frac{1}{y} - \rho\left(\log(yz) + \frac{\rho\sigma^2 y^2}{2}\right)\right) + \log(yz) + \frac{\rho\sigma^2 y^2}{2}\right). \tag{80}$$

There exists a solution to (80) when the feasible set is non empty, i.e. when it is possible to find values of y for which the argument of the first log is positive. This is equivalent to

$$z < z_{max} := \max_y \frac{1}{y} exp[\frac{\mu}{\rho} - \frac{1}{\rho y} - \frac{\rho\sigma^2 y^2}{2}]. \tag{81}$$

Taking the first order condition in (80) and denoting

$$v(z) := \frac{1}{\rho} \sup_y \left(\log\left(\mu - \frac{1}{y} - \rho\left(\log(yz) + \frac{\rho\sigma^2 y^2}{2}\right)\right) + \log(yz) + \frac{\rho\sigma^2 y^2}{2}\right). \tag{82}$$

QED

Proof of Proposition 12: To prove Point 1 in Proposition 12 we start by observing that (22) states that the growth rate of capital is

$$g = \mu - \frac{\int c^A(\omega)d\mathbb{P}(\omega)}{K} - \frac{c^P}{K}$$

and that (26) states that

$$c^P = \gamma^P K, c^A(\omega) = \gamma^A \exp(\rho\omega).$$

Substituting the latter in the former, we have

$$g = \mu - \frac{\gamma^A \int \exp(\rho\omega)d\mathbb{P}(\omega)}{K} - \gamma^P.$$

By (27), this is

$$g = \mu - \gamma^A \frac{Z}{K} - \gamma^P. \tag{83}$$

As explained in the analysis of the restricted problem, (26) and (27) imply $\frac{Z_t}{K_t}$ is a constant, denoted by z, and $\gamma^A = \frac{1}{yz}$. Substituting in (83) yields

$$g = \mu - \frac{1}{y} - \gamma^P.$$

Substituting γ^P from (32), we obtain Point 1 in Proposition 12.

To prove Point 2 in Proposition 12, we start by recalling that (72) states

$$d\omega = \log(yz)dt + \sigma y dB_t$$

and that (75) implies

$$\log(yz) = \frac{\mu - \gamma^P - \frac{1}{y}}{\rho} - \frac{\rho\sigma^2 y^2}{2}.$$

Noting that the first term on the right-hand side is $\frac{g}{\rho}$, we obtain Point 2 in Proposition 12.

Point 3 in Proposition 12 is just a restatement of (32), while Points 4 and 5 are established at the beginning of the analysis of the restricted problem.

QED

Proof of Proposition 14: Replacing y by $\frac{x}{\rho}$ in the left-hand side, (61) becomes

$$\frac{g}{\rho} - \frac{\sigma^2 x^2}{2\rho} = \frac{1}{\rho}\left[x\mu - (\rho + \tau + \pi(1-x))\right] - \frac{\sigma^2 x^2}{2\rho}.$$

The second terms on both sides are the same and cancel out, so we are left with

$$g = x\mu - (\rho + \tau + \pi(1-x)). \tag{84}$$

Replacing y by $\frac{x}{\rho}$ in (57), the growth rate prevailing in the second best writes as

$$g = \mu - \rho - \frac{\sigma^2 x}{1 + \frac{\sigma^2}{\rho}x^2}. \tag{85}$$

Substituting $\pi = x\sigma^2 - \mu$ (from the definition of x), the right-hand side of (84) is

$$\mu - \rho - \tau - x\sigma^2(1-x).$$

So, (84) writes

$$\mu - \rho - \frac{\sigma^2 x}{1 + \frac{\sigma^2}{\rho}x^2} = \mu - \rho - \tau - \sigma^2 x(1-x),$$

which is equivalent to

$$\tau = \sigma^2 x^2 \left(\frac{1 - \frac{\sigma^2}{\rho}x(1-x)}{1 + \frac{\sigma^2}{\rho}x^2} \right),$$

which yields (62).

Implementation also requires that the equilibrium sale of goods by agents, which is equal to the principal's consumption in equilibrium, be equal to the principal's consumption in the second best

$$\mathbb{E}[ds_t] = \gamma^P K_t dt, \tag{86}$$

where γ^P is given by (35) which substituting $y = x/\rho$ is

$$\gamma^P = \rho - \frac{1}{x/\rho + \sigma^2 x^3/\rho^2}. \tag{87}$$

By (49), (86) is equivalent to

$$(\mu - \frac{\rho}{x} + \pi)(1-x) + \tau = \gamma^P,$$

that is

$$\tau = \gamma^P - (1-x)(\sigma^2 x - \frac{\rho}{x}). \tag{88}$$

Substituting in (88) the value of γ^P from Proposition 12, this yields

$$\tau = \rho - \frac{1}{x/\rho + \sigma^2 x^3/\rho^2} - (1-x)(\sigma^2 x - \frac{\rho}{x}),$$

which simplifies to

$$\tau = \sigma^2 x^2 \frac{1 - \frac{\sigma^2}{\rho}x(1-x)}{1 + \sigma^2 x/\rho}.$$

which, because $x = \rho y$, is equivalent to (62).

QED

Appendix B: Differential calculus in the Wasserstein space

Consider a real-valued function F defined on $\mathcal{P}_2(\mathbb{R})$ the set of probability measures on $\mathbb{R}$ with finite second moment. To apply a verification argument for the principal problem, we are interested in Itô's formula for F to describe the dynamic $t \to F(\mathbb{P}_{w_t})$. Itô's formula for F naturally requires differential calculus on the space of measures. We start by introducing the first variation of F, which is a standard notion of Gateaux differentiability for functions of measures relying on the convexity of $\mathcal{P}_2(\mathbb{R})$.

Definition 1 *A function F admits a linear derivative at $\mu \in \mathcal{P}_2(\mathbb{R})$ if there exists a real-valued and continuous function $\delta F[\mu] : \mathbb{R} \to \mathbb{R}$ such that for all ν in $\mathcal{P}_2(\mathbb{R})$, we have*

$$\lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \left(F((1-\varepsilon)\mu + \varepsilon\nu) - F(\mu) \right) = \int_{\mathbb{R}} \delta F[\mu](x)\, d(\nu - \mu)(x).$$

We will always assume that the linear derivative $\delta F[\mu]$ is twice continuously differentiable on $\mathbb{R}$ and we will denote $\partial_x \delta F[\mu]$ and $\partial_{xx} \delta F[\mu]$ its first and second derivatives. We will summarize these assumptions by saying that F is $C^2(\mathcal{P}_2)$. For a function F that is $C^2(\mathcal{P}_2)$, Itô's formula associated to the dynamic $t \to F(\mathbb{P}_{w_t})$ takes the following form, see [9], Chapter 5, Th. 5.99,

$$\begin{aligned} F(\mathbb{P}_{w_t}) = F(\mathbb{P}_{w_0}) &+ \int_0^t \mathbb{E}\left[\partial_x \delta F[\mathbb{P}_{w_s}](w_s)(\rho w_s - \log c^A(K_s, \mathbb{P}_{w_s}, w_s)) \right] ds \\ &+ \frac{1}{2} \int_0^t \mathbb{E}\left[\partial_{xx} \delta F[\mathbb{P}_{w_s}](w_s) \sigma^2 y^2(K_s, \mathbb{P}_{w_s}, w_s) \right] ds \end{aligned} \tag{89}$$

Example 16 *Let ϕ a twice continuously differentiable function on $\mathbb{R}$ and v a continuously differentiable function on $\mathbb{R}$. We consider the function F defined on $\mathcal{P}_2(\mathbb{R})$ by*

$$F(\mu) = v\left(\int_{\mathbb{R}} \phi(x)\mu(dx) \right).$$

Then, F is $C^2(\mathcal{P}_2)$ with

$$\delta F[\mu] = v'\left(\int_{\mathbb{R}} \phi(x)\mu(dx) \right) \phi,\ \partial_x \delta F[\mu] = v'\left(\int_{\mathbb{R}} \phi(x)\mu(dx) \right) \phi' \text{ and } \partial_{xx} \delta F[\mu] = v'\left(\int_{\mathbb{R}} \phi(x)\mu(dx) \right) \phi''.$$

References

[1] Aiyagari, S., 1994. "Uninsured Idiosyncratic Risk and Aggregate saving." *Quarterly Journal of Economics*, 659-684.

[2] Angeletos, G. M., 2007. "Uninsured Idiosyncratic Investment Risk and Aggregate saving." *Review of Economic Dynamics*, 1-30.

[3] Achdou, Y., J. Han, J.M. Lasry, P.L. Lions, and B. Moll, 2022. "Income and Wealth Distriution in Macroeconomics: A Continuous Time Approach." *Review of Economic Studies*, 45–86.

[4] Berentsen, A. and G. Rocheteau, 2004. "Money and Information." *Review of Economic Studies*, 915-944.

[5] Bewley, T., 1977. "The Permanent Income Hypothesis: A Theoretical Formulation." *Journal of Economic Theory*, 252-292.

[6] Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet, 2007. "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications." *Review of Economic Studies*, 345–390.

[7] Biais, B., T. Mariotti, J.-C. Rochet and S. Villeneuve. 2010. "Large Risks, Limited Liability and Dynamic Moral Hazard." *Econometrica*. 73-118.

[8] Bolton, P. and D. Sharfstein. 1990. "A theory of predation based on agency problems in financial contracting." *American Economic Review*. 93–106.

[9] Carmona, R. and F. Delarue, 2018. *Probabilistic Theory of Mean Field Games with Applications I*, vol 83, Probability Theory and Stochastic Modelling. Springer International Publishing.

[10] DeMarzo, P.M., and M.J. Fishman. 2007a. "Agency and Optimal Investment Dynamics." *Review of Financial Studies*. 151–188.

[11] DeMarzo, P.M., and M.J. Fishman. 2007b. "Optimal Long-Term Financial Contracting." *Review of Financial Studies*. 2079–2128.

[12] DeMarzo, P.M., and Y. Sannikov. 2006. "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model." *Journal of Finance*, 2681–2724.

[13] DeMarzo, P.M., M. Fishman, Z. He, and N. Wang, 2012, "Dynamic Agency and the q Theory of Investment," *Journal of Finance*, 2295-2340.

[14] Diamond, P., 1998, "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 88, 83-95.

[15] Diamond, P. and J. Mirrlees, 1978, "A model of Social Insurance with Variable Retirement," Journal of Public Economics, 10, 295-336.

[16] Di Tella, S., 2020. "Risk premia and the real effects of money." *American Economic Review*, 1995-2040.

[17] Di Tella, S., and Y. Sannikov, 2021. "Optimal Asset Management Contracts with Hidden Savings," *Econometrica*, 1099–1139.

[18] Farhi, E., and I. Werning, 2010, "Progressive Estate Taxation." *Quarterly Journal of Economics* 125 (2): 635-673.

[19] Feng, F., and M. Westerfield, 2021, "Dynamic Resource Allocation with Hidden Volatility," *Journal of Financial Economics*, 140 (2): 560-581.

[20] Gersbach, H., J.C. Rochet and E. von Thadden, 2022, "Public Debt and the Balance Sheet of the Private Sector," *discussion paper.*

[21] Golosov, M., N. Kocherlakota, and A. Tsyvinski, 2003, "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, 70 (3), 569-588.

[22] Golosov, M, and A. Tsyvinski, 2007, "Optimal Taxation with Endogenous Insurance markets," *Quarterly Journal of Economics*, 122 (2), 487-534.

[23] He, Z, 2009, "Optimal Executive Compensation when Firm Size follows Geometric Brownian Motion," *Review of Financial Studies*, 22, 859-892.

[24] Huguett, M., 1993. "The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies", *Journal of Economic Dynamics and Control*, 17, 953–969.

[25] Huguett, M., 1997, "The One-Sector Growth Model with Idiosyncratic Shocks: Steady States and Dynamics", *Journal of Monetary Economics*, 39, 385–403.

[26] Kiyotaki, N. and R. Wright, 1989. "On Money as a Medium of Exchange." *Journal of Political Economy.*

[27] Kiyotaki, N., and R. Wright, 1993. "A Search-Theoretic Approach to Monetary Economics." *The American Economic Review*, 63-77.

[28] Kocherlakota, N., 1998, "Money is memory." *Journal of Economic Theory*, 232-251.

[29] Krusell, P., and A. Smith, 1998, "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 106, 867–896.

[30] Lagos, R., and R. Wright, 2005, "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy*, 463-484.

[31] Mirrlees, J, 1971, "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies.* 38, 175-208.

[32] Sannikov, Y. 2008. "A Continuous-Time Version of the Principal-Agent Problem." *Review of Economic Studies.* 957–984.

[33] Shi, S, 1997, "A Divisible Model of Fiat Money." *Econometrica*, 75-102.

[34] Villani, C., 2009. *Optimal transport, old and new*, Springer.

[35] Williamson, S, and R. Wright, 2011, "New Monetarist Economics: Models." Handbook of Monetary Economics, Vol 2A. Elsevier.

Competitive Nonlinear Pricing under Adverse Selection*

Andrea Attar† Thomas Mariotti‡ François Salanié§

January 2, 2023

Abstract

This article surveys recent attempts at characterizing competitive allocations under adverse selection when each informed agent can privately trade with several uninformed parties: that is, trade is nonexclusive. We first show that requiring market outcomes to be robust to entry selects a unique candidate allocation, which involves cross-subsidies. We then study how to implement this allocation as the equilibrium outcome of a game in which the uninformed parties, acting as principals, compete by making offers to the informed agents. We show that equilibria typically fail to exist in competitive-screening games, in which these offers are simultaneous. We finally explore alternative extensive forms, and show that the candidate allocation can be implemented through a discriminatory ascending auction. These results yield sharp predictions for competitive nonexclusive markets.

Keywords: Adverse Selection, Entry-Proofness, Discriminatory Pricing, Nonexclusive Markets, Ascending Auctions.
JEL Classification: D43, D82, D86.

*The material in this article has been presented at the semi-plenary lecture "Constests and Contracts" of the 2020 World Congress of the Econometric Society (Università Bocconi, Milano, Italy). This research has benefited from financial support of the ANR (Programme d'Investissement d'Avenir ANR-17-EURE-0010), the MIUR (PRIN 2015), and the research foundation TSE-Partnership.

†Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, and Università degli Studi di Roma "Tor Vergata," Roma, Italy. Email: andrea.attar@tse-fr.eu.

‡Toulouse School of Economics, CNRS, University of Toulouse Capitole, Toulouse, France, CEPR, and CESifo. Email: thomas.mariotti@tse-fr.eu.

§Toulouse School of Economics, INRAE, University of Toulouse Capitole, Toulouse, France. Email: francois.salanie@tse-fr.eu.

1 Introduction

On the vast majority of markets, nonexclusivity is the rule: agents can privately trade with different parties, without having to inform each of them of these multiple relationships. Fortunately, detailed knowledge of individual trades is generally useless for the involved parties. Nor is it usually needed, moreover, to predict market outcomes: instead, by solely relying on the aggregate equality of supply and demand, general-equilibrium theory elegantly sidesteps the lack of such information to focus on the determination of equilibrium prices. These prices in turn convey all the information agents need to know in order to formulate their supplies and demands.

However, there are important cases in which information about the characteristics of the goods for trade is not symmetrically distributed, though it directly matters to the parties involved. The chief example is that of goods whose quality is privately known to a party but unknown to his trading partners, including, for instance, the sale of shares in a firm with different returns, the supply of labor services by workers with different productivities, or the design of insurance contracts for consumers with different riskiness.

These common-value situations raise important difficulties for the performance of markets. They also question the relevance of general-equilibrium theory, because, for instance, a buyer may now try to infer the missing information from the seller's behavior and not only from prices. This gives rise to new phenomena such as adverse selection, reflecting that agents who are more eager to trade are often endowed with goods of lower quality; moreover, these inferences are made even more difficult when trade is nonexclusive. Game-theoretical tools then become useful to precisely describe who trades with whom and how the distribution of information impacts behaviors and outcomes.

This article surveys recent developments in the theory of competitive markets under adverse selection and nonexclusivity. We are particularly interested in how price competition and the threat of entry lead to sharp predictions for equilibrium outcomes, in spite of the complexities associated with nonexclusive trading. We will show how different approaches often lead to the same prediction, though we also emphasize difficulties for the existence of an equilibrium, depending on the game studied and the solution concept adopted. These results may be of interest for many financial and insurance markets, including over-the-counter, life-insurance, and annuity markets; for labor markets such as the markets for professionals or freelance workers; and, more generally, for many markets in goods or services whose quality is the private information of agents on one side of the market.

Because this combination of adverse selection and nonexclusivity is the hallmark of the literature reviewed in this survey, it may be helpful to briefly recall how these two topics have been addressed in now classical works.

On the one hand, adverse selection has so far been typically studied under exclusive competition. This is by design in Akerlof's (1970) example of a market for an indivisible good. Subsequent works focusing on markets for perfectly divisible goods assume that exclusive contracts are enforceable. In Rothschild and Stiglitz (1976), this allows an insurance company to screen its customers by making low-risk consumers self-select into contracts with higher deductibles. In Leland and Pyle (1977), this allows an entrepreneur to signal the profitability of her project by retaining a greater or lesser equity share. In both cases, the observability of agents' aggregate trades is required to support equilibrium outcomes. Extensions by Prescott and Townsend (1984), Kehoe, Levine, and Prescott (2002), Bisin and Gottardi (2006), and Rustichini and Siconolfi (2008) of the standard existence and welfare theorems of general-equilibrium theory to private-information economies similarly restrict feasible allocations to those satisfying incentive-compatibility of individual trades on each market, which again requires strong observability assumptions.

On the other hand, nonexclusive competition has so far been mostly studied in private-values environments. As noticed above, in such cases, the functioning of Walrasian markets is unaffected by private information. The side-trading literature, from the early works of Hammond (1979, 1987), Allen (1985), and Jacklin (1987) to the more recent contributions of Cole and Kocherlakota (2001), Golosov and Tsyvinski (2007), and Farhi, Golosov, and Tsyvinski (2009), has accordingly investigated the limits that side trading on Walrasian markets, outside of the central planner's control, imposes on the set of allocations that can be achieved in standard Mirrlees (1971) or Diamond and Dybvig (1983) economies, which feature private information but not adverse selection.

By contrast, in common-values environments, assuming a priori that a Walrasian market exists and perfectly balances supply and demand is too much to ask for; in fact, we shall see in Section 2.2 that sellers on such a market, anticipating adverse selection, would like to reduce their competitive supply and thus choose to ration demand. In addition, nonexclusivity means that trades cannot be monitored, so that even a centralized market authority would not be able to ensure that all agents trade at the same price. Instead, the recent literature on nonexclusive markets under adverse selection assumes that contracting is bilateral, and allows for general contracts.

We now present the main findings of this literature. In line with Riley (2001), we adopt in this survey a strategic approach to the determination of market outcomes; in contrast with him, however, we exclusively focus on screening models.[1] A market is, therefore, described by a set of uninformed sellers competing through menus of contracts, or nonlinear tariffs, to serve the demand emanating from privately informed buyers. Nonexclusivity is captured by the assumption that, while each seller can monitor the trades each buyer conducts with him—which is what makes nonlinear pricing possible—he can monitor none of the trades this buyer makes with his competitors.

A useful entry point into the literature is to first abstract from the determination of individual tariffs and, in a reduced-form way, to directly impose properties on the market tariff that is obtained from them by aggregation. In line with Rothschild and Stiglitz (1976), who characterize the set of exclusive contracts preventing an entrant from making a profit, a desirable property of the market tariff is that it be entry-proof. Under nonexclusivity, this property means that no entrant can make a profit by offering a menu of contracts, given that each buyer is free to combine a contract offered by the entrant with a trade along the market tariff. Restricting attention to entry-proof market tariffs allows us to identify robust predictions for nonexclusive markets under adverse selection, which do not depend on the details of a specific extensive-form game. Section 3 of this article is devoted to the characterization of such tariffs.

The following insights emerge from the analysis (AMS (2020, 2021)).[2] First, neither the Rothschild and Stiglitz (1976) allocation, nor any of the second-best allocations characterized by Prescott and Townsend (1984) and Crocker and Snow (1985) can be implemented by an entry-proof market tariff. Second, entry-proofness, along with budget-feasibility, singles out a unique market tariff, which generally turns out to be nonlinear. The defining features of this tariff are that each marginal quantity is priced at the expected cost of serving the buyer types who optimally choose to trade it, and that gains from trade are exhausted subject to this constraint. For instance, in the two-type case, the low-cost type purchases her demand at a price equal to the expected cost, while the high-cost type in addition trades the quantity she prefers at a price that equals her own cost. The corresponding allocation, first described by Jaynes (1978), Hellwig (1988), and Glosten (1994)—and henceforth referred to as the JHG allocation—can thus be interpreted as a marginal version of the Akerlof (1970) competitive-equilibrium allocation, and stands out as a focal prediction for nonexclusive competitive markets under adverse selection. The entry-proofness criterion can also be seen as an external constraint imposed on the decisions of a central planner. AMS (2020)

[1] Extending the theory of signaling or, more generally, the informed-principal paradigm of Myerson (1983) to nonexclusive markets is a fascinating task for future research.

[2] To avoid repetitions, AMS hereafter stands for Attar, Mariotti, and F. Salanié.

emphasize that this constraint is so strong so as to limit the planner's feasible policies to a unique policy, which is not second-best. The social costs of side trading thus appear to be particularly severe under common values.

In line with the program outlined by Wilson (1989), a natural question is whether the JHG allocation can be implemented as an—ideally, unique—equilibrium outcome of a decentralized trading protocol. The contributions reviewed in Section 4 focus on competitive-screening games in which offers are simultaneously made by the uninformed sellers; they thus bring to bear insights and methods from the common-agency literature, in which several uninformed principals compete by posting menus of contracts to deal with a privately informed agent. While, as surveyed by Martimort (2006), common-agency games have mainly been used to tackle issues in industrial organization, public economics, and in the theory of organizations—from vertical contracting and the internal structure of the firm to lobbying and the relationships between governments and regulatory agencies—they have also provided, following the seminal contribution of Biais, Martimort, and Rochet (2000), a powerful tool for modeling competition under common values.

The results of this approach, however, are mixed. The only case in which the JHG allocation can be robustly implemented through a competitive-screening game is when the buyer's preferences are linear, subject to a capacity constraint. The JHG allocation then coincides with the Akerlof (1970) competitive-equilibrium allocation that maximizes the gains from trade, and it is the unique candidate-equilibrium allocation; moreover, there always exists an equilibrium in which sellers posts linear tariffs (AMS (2011)). These positive results thus extend the conclusions of Akerlof (1970) to the case of a divisible good. However, when the buyer has strictly convex preferences, the results crucially depend on fine modeling details such as the cardinality of the set of types. When the distribution of types is continuous, Biais, Martimort, and Rochet (2000) construct an equilibrium in which sellers posts strictly convex tariffs, and such that the resulting aggregate equilibrium allocation converges to the JHG allocation as the number of sellers grows large; however, as we show in Section 4.4, their existence result requires that some types be excluded from trade. When the distribution of types is discrete, exclusion is even more extreme, as an equilibrium exists only in the degenerate case where all types except possibly the highest-cost one do not trade in the JHG allocation (AMS (2014, 2019a)).

Section 5 provides more positive results, by exploring alternative extensive forms whereby uninformed sellers sequentially receive information about previously signed contracts or previously made offers. The bottom line is that transparency makes it easier for sellers to directly punish deviators, in contrast with competitive-screening games in which the burden of punishments entirely falls on the buyer. In this spirit, Beaudry and Poitevin (1995) and AMS (2021) implement the JHG allocation in a repeated game of signalling and in an ascending discriminatory auction with frequent offers, respectively. We also survey contributions by Jaynes (1978, 2011), Hellwig (1988), and Stiglitz, Yun, and Kosenko (2020) that allow for endogenous information disclosure.

Section 6 concludes on the empirical perspectives, in particular about tests for the presence of private information.

2 The Economy

We study a simple economy in which a single buyer (she) trades a divisible good with multiple, identical sellers (he). The buyer is endowed with private information about, for instance, the quality of the traded good. As usual, the buyer/sellers convention can be inverted thanks to a change of variables, so as to encompass a broad variety of situations:

- Insurance companies sell coverage to a consumer: Rothschild and Stiglitz (1976), Prescott and Townsend (1984), Crocker and Snow (1985), Hendren (2013).
- Market makers provide liquidity to an insider: Glosten (1989, 1994), Biais, Martimort, and Rochet (2000), Back and Baruch (2013).
- Investors purchase securities issued by a firm: Leland and Pyle (1977), Myers and Majluf (1984), DeMarzo and Duffie (1999), Biais and Mariotti (2005).
- Firms hire the services of a worker: Spence (1973), Miyazaki (1977).

In these situations, the private information of the buyer is directly relevant to the sellers, because it determines their production costs or their opportunity costs of selling. It is this common-value component that may generate adverse selection, as we now discuss in more formal terms.

2.1 The Model

Unless stated otherwise, the following assumptions are maintained throughout this article.

The Buyer The buyer's private information is represented by a type $i = 1, \ldots, I$ that takes a finite number of values with strictly positive probabilities m_i. Type i's preferences are represented by a utility function $u_i(Q, T)$ that is continuous and weakly quasiconcave in

(Q,T) and strictly decreasing in T, with the interpretation that Q is the nonnegative quantity of the good she purchases and T is the payment she makes in return. To define marginal rates of substitution without assuming differentiability, we let $\tau_i(Q,T)$ be the supremum of the set of prices p such that

$$u_i(Q,T) < \max\{u_i(Q+q, T+pq) : q \geq 0\}.$$

Thus $\tau_i(Q,T)$ is the slope of type i's indifference curve at the right of (Q,T). Quasiconcavity ensures that $\tau_i(Q,T)$ is finite, except possibly at $Q=0$, and that it is nonincreasing along an indifference curve of type i. For all i and $p>0$, we also define the demand $D_i(p)$ of type i as the set of quantities Q that maximize $u_i(Q,pQ)$. These demands are well-defined under the following Inada condition:

$$\text{For all } i,\ (Q,T), \text{ and } p>0,\ \arg\max\{u_i(Q+q,T+pq) : q \geq 0\} < \infty, \tag{1}$$

or if the domain of admissible quantities is compact. Types are ordered according to the weak single-crossing condition (Milgrom and Shannon (1994)), which states that higher types are at least as willing to increase their purchases as lower types are:

$$\text{For all } i<j,\ Q<Q',\ T, \text{ and } T',\ u_i(Q,T) \leq (<)\, u_i(Q',T') \text{ implies } u_j(Q,T) \leq (<)\, u_j(Q',T').$$

Weak single-crossing implies that $\tau_i(Q,T)$ and $D_i(p)$ are weakly increasing in i. For future reference, we also state the slightly stronger, strict single-crossing condition:

$$\text{For all } i<j,\ Q<Q',\ T, \text{ and } T',\ u_i(Q,T) \leq u_i(Q',T') \text{ implies } u_j(Q,T) < u_j(Q',T').$$

We shall occasionally make additional assumptions. Theorems 1 and 3, for instance, require that higher endowments of the good reduce the buyer's marginal rate of substitution:

Assumption 1 *For all i and T, $\tau_i(Q,T)$ is nonincreasing in Q.*

Our assumptions on the buyer's preferences hold, for instance, in a Rothschild and Stiglitz (1976) insurance economy in which the loss L is the same for all types: then i indexes the buyer's riskiness, Q is the amount of coverage she purchases, and T is the premium she pays in return. AMS (2021, Online Appendix C) show that these assumptions also hold in more general insurance economies, allowing for multiple loss levels or various forms of nonexpected utility. But our framework is relevant beyond insurance; in particular, first-best quantities may differ across types.

The Sellers On the supply side, sellers are identical, risk-neutral, and use the same linear technology. We denote by $c_i > 0$ the unit cost of serving type i, and by $\bar{c}_i$ the corresponding upper-tail conditional expectation of unit costs,

$$\bar{c}_i \equiv \mathbf{E}[c_j \,|\, j \geq i] = \frac{\sum_{j \geq i} m_j c_j}{\sum_{j \geq i} m_j}.$$

Adverse selection occurs if the unit cost c_i is nondecreasing in i. This case is the least conducive to trade, as types who are more willing to trade are also more costly to serve. In this article, we only rely on a slightly weaker assumption, namely, that $\bar{c}_i$ be nondecreasing in i. This weak adverse-selection condition is exactly equivalent to

$$\text{For all } j \leq i, \; c_j \leq \bar{c}_i. \tag{2}$$

Contracts A *contract* (q, t) between a seller and the buyer specifies a nonnegative quantity to be delivered by the seller and a transfer to be made in return by the buyer. Under nonexclusivity, two contracts (q, t) and (q', t') offered by different sellers can be added to form a trade $(q + q', t + t')$. In this article, we consider two different settings in turn.

- We first study when a nonexclusive market is entry-proof. We then only rely on the existence of a *market tariff* T, where $T(Q)$ is defined as the minimum transfer that allows the buyer to obtain a quantity Q. If an entrant proposes additional contracts, then the buyer can pick one of them, say (q, t), together with a trade $(Q, T(Q))$ along the market tariff, ending up with utility $u_i(Q + q, T(Q) + t)$.

- We next study competition in menus of nonexclusive contracts among $K \geq 2$ sellers. We then have to precisely specify the menus that are offered by the sellers $k = 1, \ldots, K$. We impose that each seller must always propose at least the null contract $(0, 0)$, so that the buyer may be seen as trading one contract (q^k, t^k) with each seller k, ending up with aggregate trade $(Q, T) \equiv (\sum_k q^k, \sum_k t^k)$ and utility $u_i(Q, T)$.

Let us emphasize that we focus on situations in which trade is not anonymous. Our view of nonexclusivity is thus that a seller can monitor the trades the buyer makes with him, though he cannot monitor the trades the buyer makes with the other sellers. This is a natural assumption to make in insurance markets, because an insurance contract must name a beneficiary. Preventing the buyer from making concealed repeat purchases from the same seller enables each seller to price the quantities he sells in a nonlinear way, charging different prices for different marginal units.

Our analysis and results extend to the case of multiple buyers, provided contracting is bilateral and the buyers' types are independent and identically distributed. Contracting is bilateral if trade between a seller and a buyer is only contingent on the information reported by the buyer to the seller, and not on the information this seller may obtain from other buyers.[3] Together with the linearity of costs, the independence of types across buyers then implies that the interactions between a seller and each of his potential customers can be studied separately. Finally, if the buyers' types are identically distributed, we can assume, using a symmetry argument, that each seller offers the same contracts to each buyer and that each type of each buyer facing the same choices behaves in the same way. In this way, the analysis of the multiple-buyer case can be reduced to that of the single-buyer case.

2.2 Benchmarks

The general question addressed in this article is how to define a notion of competitive allocation for the above-described economy. This section discusses a few benchmarks, with the purpose of introducing the main effects and difficulties.

The complete-information benchmark assumes that the buyer's type is made public before sellers make their supply decisions, so that information is symmetric about this collective risk (Malinvaud (1972)). The interpretation is that the good for trade comes in I observable varieties, each representing a different quality. When this quality i is revealed, the sellers learn their cost c_i; competition then implies that type i purchases her demand $D_i(c_i)$ at price c_i and that sellers make zero profits. Therefore, equilibria exist and are efficient.

When information becomes asymmetric, the now privately informed buyer optimally channels her demand to the market with the lowest price $p \equiv \min_i c_i$. But then aggregate profits $\mathbf{E}[(p-c_i)D_i(p)]$ are typically negative, except when unit costs are independent of the buyer's type, that is, $c_i \equiv c$ for all i. In this private-value case, sellers know their costs, while the tastes of the buyer are her private information. Nevertheless, a competitive market still plays its allocative role: in equilibrium, every type i purchases her demand $D_i(c)$ at price c, and sellers make zero profits. Once more, equilibria exist and are efficient.

We allow for private values as a limiting case, but our main focus is on the common-value case where the sellers' unit costs depend on the buyer's type. Proceeding as in Akerlof (1970), Pauly (1974) studies the case where linear pricing is imposed, independently of the quantity traded; that is, sellers stand ready to supply any quantity at the going price. Then an

[3]Under multilateral contracting, we enter the realm of multi-principal multi-agents models, in which unrestricted communication can be exploited to support many equilibrium allocations; see, for instance, Yamashita (2010).

equilibrium exists as soon as demand functions are continuous. Because equilibrium profits must be zero under constant returns to scale, the equilibrium price satisfies the equality

$$p = \mathbf{E}\left[c_i \, \frac{D_i(p)}{\mathbf{E}[D_j(p)]}\right]. \tag{3}$$

This formula is widely used in the annuity literature (Sheshinski (2008), Hosseini (2015), Rothschild (2015)). When the good is indivisible, each demand term in (3) is either zero or one, and we are back to the classical Akerlof (1970) formula that states that the price is equal to the expected unit cost of active types. With a divisible good, the formula in addition weighs the unit cost of serving each type by her demand. Because higher types have higher demands, it generally follows that the equilibrium price must lie above the expected unit cost of active types. Accordingly, the active types with the lowest costs subsidize the higher types, who are more costly to serve.

However, there exists a simple way to reduce the risk of having to sell too much to a high type at the going price: to this end, a seller need only post a *limit order* $(p, \overline{q})$ specifying the maximum quantity $\overline{q}$ he is ready to sell at price p. Such limit orders are commonly used on financial markets, and this may indeed be because they allow sellers to hedge against the risk of a high demand.[4] A well-chosen limit order, with a price just below p, is profitable because it reduces the loss-making sales to high types while preserving the profits from selling to low-cost types.[5] This, incidentally, shows that the Pauly (1974) outcome is not a competitive equilibrium: anticipating adverse selection, sellers would like to reduce their competitive supply, thereby collectively rationing demand.

We conclude that the linear-pricing construction is rather fragile under adverse selection. A natural step forward is to consider a competitive game in which sellers are allowed to post limit orders.[6] Notice that a collection of limit orders gives rise to a convex market tariff and, conversely, that any convex market tariff can be decomposed into a (possibly infinite) collection of (possibly infinitesimal) limit orders. In what follows, we sometimes impose that tariffs be convex, but we also explore cases where tariffs can be arbitrary.

[4]We follow the literature in using the term "limit order," although the maximum quantity $\overline{q}$ is here understood to apply to a single buyer, whereas a limit order on a financial market specifies the maximum aggregate quantity one is ready to sell to a set of traders.

[5]AMS (2019a, Lemma 4) formalize this intuition and show that, under adverse selection, such a limit order allows a seller to approximate the maximum profit he can earn when competing with a linear tariff.

[6]The game in which sellers compete by posting a single limit order implements the efficient competitive equilibria when information is symmetric and cost functions are weakly convex. In the present model with asymmetric information and linear cost functions, AMS (2018, pp. 1013–1014) show that this game has a pure-strategy equilibrium only in limiting cases where costs or demands do not depend on the buyer's type.

3 Entry-Proofness

The idea of using entry-proofness as a solution concept first originates in an attempt at simplicity, because this avoids the need to precisely describe the supply side of the economy or to fully specify the details of an extensive form. Moreover, entry-proof allocations, when they exist, are widely considered as capturing the idea of perfect competition.[7] We begin by studying when inactive markets are entry-proof. We then turn to active markets, for which the distinction between exclusive and nonexclusive competition becomes relevant, and we formulate a definition of entry-proofness consistent with nonexclusivity. We show in particular that entry-proofness selects a unique budget-balanced allocation, which exists under very general conditions. This requirement is thus more fruitful under nonexclusivity than under exclusivity.

3.1 Entry-Proofness in Inactive Markets

In this section, we describe the circumstances under which private information impedes trade altogether. We say that a market is *inactive* if the market tariff reduces to a single point, given by $T(0) = 0$, or, equivalently, if only the null contract $(0, 0)$ is available. Our goal is to find conditions ensuring that no entrant can make an offer leading to profitable trades. Accordingly, we say that an inactive market is *entry-proof* if, *for any menu of contracts offered by an entrant, the buyer has a best response such that the entrant earns at most zero expected profit.*

To characterize the inactive markets that are entry-proof, we first study the simplest case where the entrant offers a single contract. The key argument here is that, if this contract strictly attracts a type i, then it must also attract all types $j > i$: this is a simple consequence of weak single-crossing. Hence, from the entrant's viewpoint, the relevant unit cost is not the individual cost c_i of serving type i, but, rather, the expected cost $\bar{c}_i$ of serving types $j \geq i$. Notice that some other types $j < i$ may also be attracted by the entrant's offer, but the weak adverse-selection condition (2) ensures that this can only reduce the entrant's expected unit cost. This shows that the following condition is necessary for entry to be unprofitable.

Condition EP *For each i, $\tau_i(0, 0) \leq \bar{c}_i$.*

Notice that Condition EP does not rule out gains from trade, in the usual first-best sense of the term; that is, it may well be that $\tau_i(0, 0) > c_i$ for some i. AMS (2021, Theorem

[7] As argued by Rothschild and Stiglitz (1976, p. 642): "The basic idea underlying competitive markets involves free entry and noncollusive behavior among the participants in the market."

1) show that this necessary condition for entry-proofness is sufficient even when menus of contracts are allowed:

Theorem 1 *Under Assumption 1, an inactive market is entry-proof if and only if Condition EP is satisfied.*

The intuition is as follows. If the entrant offers an arbitrary menu of contracts, then, by weak single-crossing, the buyer has a best response with nondecreasing quantities, which we denote by $(q_i, t_i)_{i=1}^I$. Suppose that $q_i > q_{i-1}$ for some type i, and let us first locate the contract (q_{i-1}, t_{i-1}). We can safely assume that t_{i-1} is positive, as the entrant's profit can only be reduced by giving away costly production. We also know that type $i-1$ weakly prefers this contract to the null contract; by weak single-crossing again, so does type i. Therefore, the point (q_{i-1}, t_{i-1}) must lie in the north-east quadrant in Figure 1, at the right of the indifference curve of type i that goes through the origin.

Now, to be willing to trade the contract (q_i, t_i), it must be that type i, having already traded the contract (q_{i-1}, t_{i-1}), is willing to trade the additional *layer* $(q_i - q_{i-1}, t_i - t_{i-1})$. To evaluate her marginal rate of substitution at (q_{i-1}, t_{i-1}), we can use, in turn, the concavity of the indifference curve of type i, then Assumption 1, and finally Condition EP to obtain the following inequalities:

$$\tau_i(q_{i-1}, t_{i-1}) \leq \tau_i(\underline{q}_i, 0) \leq \tau_i(0, 0) \leq \overline{c}_i. \tag{4}$$

This implies that type i is not ready to pay more than $\overline{c}_i(q_i - q_{i-1})$ for the additional quantity $q_i - q_{i-1}$. Therefore,

$$t_i - t_{i-1} \leq \overline{c}_i(q_i - q_{i-1}).$$

Summing these inequalities over i with appropriate weights yields

$$\sum_i \left(\sum_{j \geq i} m_j \right) [t_i - t_{i-1} - \overline{c}_i(q_i - q_{i-1})] \leq 0.$$

Finally, rearranging terms in the spirit of Wilson (1993), we obtain

$$\sum_i m_i(t_i - c_i q_i) \leq 0,$$

which shows that entry cannot be profitable.

A noticeable feature of this proof is that it does not consider each contract (q_i, t_i) in isolation. Instead, the key role is played by layers of the form $(q_i - q_{i-1}, t_i - t_{i-1})$. Under

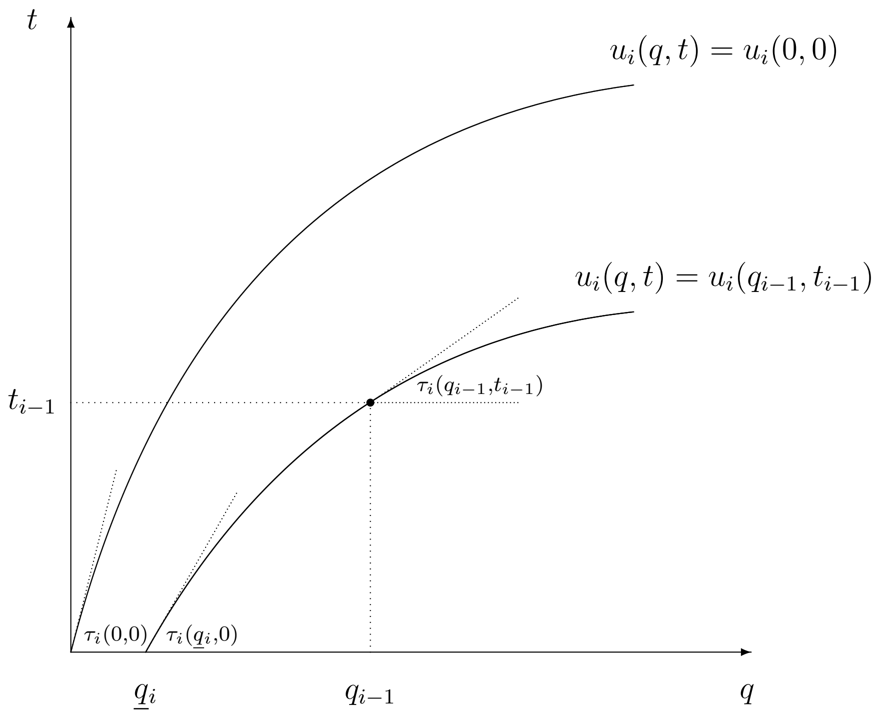

Figure 1: A graphical illustration of (4).

weak single-crossing, optimal quantities can be assumed to be nondecreasing in the buyer's type, so that the i^{th} layer can be thought of as traded by all types $j \geq i$, and thus has expected unit cost $\overline{c}_i$. Condition EP implies that, at this price, type i is not strictly willing to trade, so that each layer must yield a nonpositive expected profit. By contrast, some of the *contracts* proposed in a menu may yield positive profits. For instance, although the condition $t_1 \leq \overline{c}_1 q_1$ ensures that the expected profit on the first layer (q_1, t_1) is nonpositive, it may well be that $t_1 > c_1 q_1$.

AMS (2021) show that the assumptions of Theorem 1 can be weakened in several ways; however, the weak single-crossing condition and the seemingly innocuous Assumption 1 are tight. They also provide a result characterizing *market breakdown*, defined as a situation in which *any menu of contracts that strictly attracts at least some type yields a strictly negative expected profit, even if the buyer's best response is most favorable to the entrant.* Condition EP clearly remains necessary for this stronger concept, and it also remains sufficient under slightly stronger conditions on preferences. Earlier results were obtained by Mailath and Nöldeke (2008) for an economy in which the buyer has quadratic quasilinear preferences, and by Hendren (2013) for a Rothschild and Stiglitz (1976) insurance economy.

3.2 Entry-Proofness in Active Markets: The Two-Type Case

We now turn to *active* markets, on which nonnull contracts are available. In line with Rothschild and Stiglitz (1976), our goal is to characterize when entry on such a market is unprofitable, given the contracts available; in contrast with them, we suppose that the buyer can trade with several sellers. To this end, the proper object of study is the *market tariff*, which describes the frontier of the set of aggregate trades that can be achieved by trading on the market.

A market tariff specifies the minimum aggregate transfer $T(Q)$ required to purchase an aggregate quantity Q, with $T(Q) \equiv \infty$ if this is impossible; notice that we obviously have $T(0) = 0$. By assuming that T is lower semicontinuous, with a compact domain, we ensure that, for every type i, the problem of maximizing $u_i(Q, T(Q))$ admits a solution Q_i. We then say that the allocation $(Q_i, T(Q_i))_{i=1}^I$ is *implemented* by the tariff T. We assume that types are ordered according to the strict single-crossing condition, so that the optimal quantities Q_i are nondecreasing in i. Moreover, this allocation is *budget-feasible* if

$$\sum_i m_i[T(Q_i) - c_i Q_i] \geq 0. \tag{5}$$

Now, suppose an entrant can propose additional trades to the buyer, in the form of a menu of contracts that complement the market tariff. We say that the tariff T is *entry-proof* if, *for any menu of contracts offered by an entrant, the buyer has a best response such that the entrant earns at most zero expected profit, given that the buyer is free to combine any contract offered by the entrant with a trade along the tariff* T. The last clause of this definition is crucial, and captures the nonexclusivity of trade.

Our goal is to characterize the set of budget-feasible allocations that are implemented by entry-proof market tariffs. In this section, we focus on the two-type case $I = 2$, which is simple enough to allow for a precise discussion of the proof to the main result; the weak adverse-selection condition (2) then amounts to $c_1 \leq c_2$.

Thus consider an allocation $(Q_i, T_i)_{i=1}^2$ that is implemented by some market tariff T. Because this allocation is incentive-compatible, it satisfies $Q_2 \geq Q_1$ by strict single-crossing. Moreover, if T is entry-proof, then we must have

$$u_1(Q_1, T_1) \geq \max\{u_1(q, \bar{c}_1 q) : q \geq 0\}. \tag{6}$$

Otherwise, an entrant can offer a contract with unit price slightly above $\bar{c}_1$ that profitably attracts type 1, and remains profitable even if type 2 is attracted—recall that, by definition, $\bar{c}_1 = m_1 c_1 + m_2 c_2$. Similarly, we must have

$$u_2(Q_2, T_2) \geq \max\{u_2(Q_1 + q, T_1 + c_2 q) : q \geq 0\}. \tag{7}$$

Otherwise, an entrant can offer a contract with unit price slightly above c_2 that profitably attracts type 2 along with the contract (Q_1, T_1), and is even more profitable if type 1 is also attracted; notice that this second type of entry is specific to the nonexclusive case.[8] It follows from (6) that

$$T_1 \leq \bar{c}_1 Q_1. \tag{8}$$

Similarly, it follows from (7) that

$$T_2 \leq T_1 + c_2(Q_2 - Q_1). \tag{9}$$

However, rearranging terms in the spirit of Wilson (1993), the budget-feasibility constraint (5) can be rewritten as

$$T_1 - \bar{c}_1 Q_1 + m_2[T_2 - T_1 - c_2(Q_2 - Q_1)] \geq 0. \tag{10}$$

Thus, in light of (10), the equalities (8)–(9) are in fact equalities. That is, profits are zero on the first layer (Q_1, T_1), which is traded by both types; similarly, profits are zero on the second layer $(Q_2 - Q_1, T_2 - T_1)$, which is traded by type 2 only. This, in turn, implies that the inequalities (6)–(7) are also equalities. Overall, the four resulting equalities pin down the set of candidates for a budget-feasible allocation that is implemented by an entry-proof tariff. AMS (2020, Theorem 2) show that these necessary conditions for entry-proofness are sufficient even when menus of contracts are allowed:

Theorem 2 *Any budget-feasible allocation $(Q_i^*, T_i^*)_{i=1}^2$ that is implemented by an entry-proof market tariff satisfies*

$$Q_1^* \in \arg\max\{u_1(Q, \bar{c}_1 Q) : Q \geq 0\}, \tag{11}$$

$$T_1^* = \bar{c}_1 Q_1^*, \tag{12}$$

$$Q_2^* - Q_1^* \in \arg\max\{u_2(Q_1^* + q, T_1^* + c_2 q) : q \geq 0\}, \tag{13}$$

$$T_2^* - T_1^* = c_2(Q_2^* - Q_1^*). \tag{14}$$

Conversely, any allocation that satisfies (11)–(14) *can be implemented by the piecewise-linear convex market tariff*

$$T^*(Q) \equiv 1_{\{Q \leq Q_1^*\}} \bar{c}_1 Q + 1_{\{Q_1^* < Q \leq Q_2^*\}}[\bar{c}_1 Q_1^* + c_2(Q - Q_1^*)], \tag{15}$$

and this tariff is entry-proof.

[8]In a Rothschild and Stiglitz (1976) insurance economy, Stiglitz, Yun, and Kosenko (2020, Definition 1) base their analysis on the assumption that additional coverage is available without limits at price c_2, which implies an inequality similar to (7). They do not, however, state the first inequality (6).

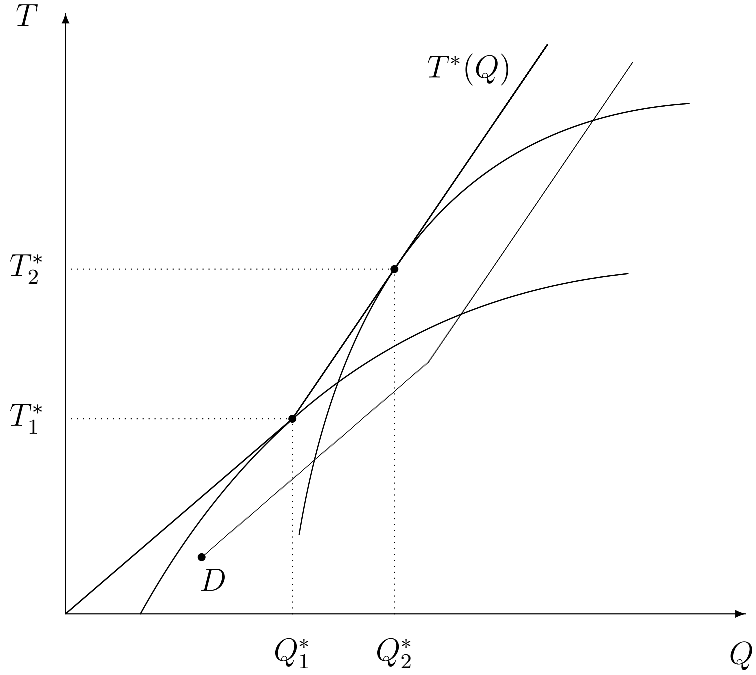

Figure 2: Blocking cream-skimming deviations.

When u_1 and u_2 are strictly quasiconcave—and also, generically, when they are only weakly quasiconcave—conditions (11)–(14) characterize a unique allocation. Notice that, because the low-cost type 1 obtains her demand at a price equal to the average cost $\overline{c}_1$, she subsidizes the high-cost type 2, though to a lesser degree than in the linear-pricing candidate with price (3) discussed in Section 2.2.

Concerning the second part of Theorem 2, it should be noted that the natural two-point tariff obtained by restricting the market tariff (15) to the quantities Q_1^* and Q_2^* does not generally resist entry, as an entrant may cream-skim type 1 and make a profit. To deter entry, we have to ensure that any such offer would also attract type 2. This is exactly what the convex tariff (15) achieves, by enabling type 2 to purchase any fraction of the first layer at price $\overline{c}_1$ and any additional quantity at price c_2. This ensures that any entrant's contract that would attract type 1, such as D in Figure 2, would also attract type 2, because type 2 can complement this contract by latent contracts made available by the market tariff and thereby reach aggregate trades that she strictly prefers to (Q_2^*, T_2^*). The need for latent contracts to block attempts at cream-skimming contrasts with the exclusive-competition case, where the revelation principle ensures that there is no need to distinguish a market tariff from the allocation it implements.

In summary, in the two-type case, entry-proofness singles out a generically unique budget-balanced allocation; moreover, the existence problem emphasized by Rothschild and Stiglitz (1976) under exclusive competition no longer arises, whatever the distribution of types.

The generality of this conclusion is striking: single-crossing is only used to ensure that the inequality $Q_2 \geq Q_1$ holds, Assumption 1 is not needed, and preferences and candidate tariffs can be arbitrary as long as the maximization problems in (6)–(7) admit a solution.

3.3 Entry-Proofness in Active Markets: The Convex-Tariff Case

We now extend these results to the case of an arbitrary number of types. We will see that this raises a subtle new difficulty; to deal with it, the key restriction we impose in this section is that the market tariff be *convex*. A case in point is when each seller k posts a convex tariff t^k such that $t^k(0) = 0$. An intuitive rationale is that this allows sellers to hedge against the risk of attracting high-cost types buying large quantities; for instance, in the market-microstructure literature, convex tariffs are often used to model collections of limit orders placed by strategic market makers and executed in order of price priority by an informed insider.[9] Then the market tariff $T(Q) \equiv \min\{\sum_k t^k(q^k) : \sum_k q^k = Q\}$, which incorporates the possibility of trading with several sellers on the market, is indeed a convex function of the aggregate quantity Q.[10]

For a seller contemplating entering on a market where existing trading opportunities are summarized by the market tariff T, everything is as if the market were inactive and every type i's preferences were represented by the indirect utility function

$$u_i^T(q,t) \equiv \max\{u_i(Q+q, T(Q)+t) : Q \geq 0\}. \tag{16}$$

Convexity of the market tariff ensures that, if the primitive utility functions $(u_i)_{i=1}^I$ satisfy the strict single-crossing property, then the indirect utility functions $(u_i^T)_{i=1}^I$ satisfy the weak single-crossing property. This allows AMS (2021, Theorem 2) to rely on Theorem 1, which deals with inactive markets, to tackle the case of an active market.

Theorem 3 *Under Assumption 1, an allocation $(Q_i^*, T(Q_i^*))_{i=1}^I$ is budget-feasible and is implemented by an entry-proof convex market tariff T^* with domain $[0, Q_I^*]$ if and only if they jointly satisfy the following recursive system:*

(i) $(Q_0^*, T^*(Q_0^*)) \equiv (0,0)$;

(ii) for each i, $Q_i^* - Q_{i-1}^* \in \arg\max\{u_i(Q_{i-1}^* + q, T^*(Q_{i-1}^*) + \bar{c}_i q) : q \geq 0\}$;

(iii) for each i, if $Q_{i-1}^* < Q_i^*$, then T^* is affine with slope $\bar{c}_i$ over the interval $[Q_{i-1}^*, Q_i^*]$.

[9] See, for instance, Biais, Martimort, and Rochet (2000, 2013), Back and Baruch (2013), AMS (2019a), and Baruch and Glosten (2019).

[10] This is because T is the *infimal convolution* of the convex tariffs t^k (Rockafellar (1970, Theorem 5.4)).

In particular, any such allocation is budget-balanced.

This result generalizes Theorem 2 to more than two types. While item (i) is merely a convention, (ii)–(iii) are substantial, and indicate how to recursively build a complete family of quantities, as well as the corresponding market tariff; by construction, this tariff is convex, because the upper-tail conditional expectation of unit costs is nondecreasing in the buyer's type. The proof parallels the argument provided in Section 3.2 for the two-type case: at each step, the entrant must be deterred from supplying a well-chosen quantity at a price slightly above $\bar{c}_i$. By single-crossing, if such an offer attracts type i, then it must also attract all types $j \geq i$, so that the offer is profitable as soon as type i is attracted. Therefore, entry-proofness implies the following inequalities:

$$\text{For each } i,\ u_i(Q_i^*, T^*(Q_i^*)) \geq \max\{u_i(Q_{i-1}^* + q, T(Q_{i-1}^*) + \bar{c}_i q) : q \geq 0\}. \tag{17}$$

It follows that no layer can be profitable,

$$\text{For each } i,\ T^*(Q_i^*) - T^*(Q_{i-1}^*) \leq \bar{c}_i(Q_i^* - Q_{i-1}^*). \tag{18}$$

Summing these inequalities as in Section 3.2, we obtain that the allocation $(Q_i^*, T(Q_i^*))_{i=1}^I$ is budget-balanced, so that the inequalities (18) are in fact equalities. Notice that these equalities can be interpreted as a marginal version of Akerlof (1970) pricing: each layer is priced at the expected cost of serving the types who trade it. As a result, the constraints (17) must all be binding, and the result follows.

Theorem 3 generalizes a similar but weaker entry-proofness result due to Glosten (1994, Proposition 7). His analysis of limit-order markets requires that the buyer's preferences be quasilinear, and that the entrant's tariff satisfy a property he dubs *single-crossing* and that captures a convexity requirement. By allowing for general preferences, Theorem 3 makes the result relevant for insurance markets, in which wealth effects may be significant.

Existence of an entry-proof convex market tariff obtains because each maximization problem in (ii) admits a solution under the Inada condition (1).[11] Hence budget-feasibility and entry-proofness are not conflicting requirements under nonexclusivity, in contrast with the pervasive nonexistence problems arising under exclusivity (Rothschild and Stiglitz (1976)). The difference is that, when competition is exclusive, the buyer's indirect utility functions no longer satisfy single-crossing: by offering a cream-skimming contract, the entrant can attract a type i without attracting types $j > i$, which allows him to target type i without worrying about adverse selection. The nonexistence of an entry-proof tariff is then arguably

[11] Or, when the buyer's preferences are linear, because of the imposition of a capacity constraint.

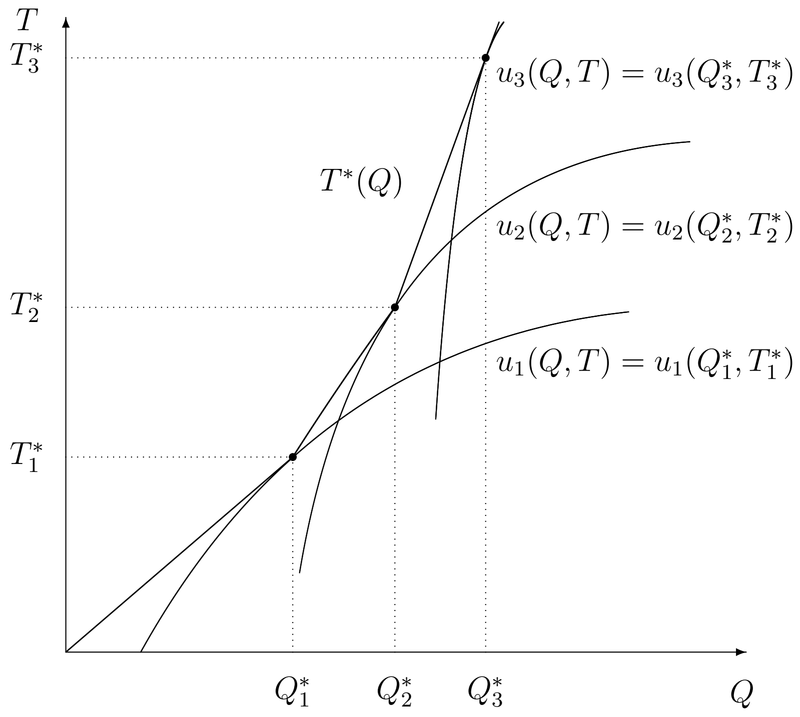

Figure 3: The JHG allocation and the JHG tariff for $I = 3$.

not due to private information or entry-proofness per se, but rather to this violation of single-crossing—or, to put it more provocatively, to the fact that the exclusive model does not capture the full extent of adverse selection.

Uniqueness of an entry-proof convex market tariff also follows if the solution to each maximization problem in (ii) is unique. This is the case if the buyer's preferences are strictly convex. If they are only weakly convex, multiple solutions may appear if the marginal rate of substitution of some type i is equal to $\overline{c}_i$ over a whole interval of quantities, but this is clearly a nongeneric phenomenon.

Theorem 3 thus characterizes an essentially unique allocation. Following AMS (2014, 2019a, 2021), we label this allocation, which was originally introduced in different contexts by Jaynes (1978), Hellwig (1988), and Glosten (1994), the *JHG allocation.*[12] Similarly, the *JHG tariff* T^* consists of a sequence of layers with unit prices $\overline{c}_i$, and features an upward kink at any quantity $Q_i^* \in (0, Q_I^*)$ such that $Q_{i+1}^* > Q_i^*$ and $\overline{c}_{i+1} > \overline{c}_i$. This sequence of layers can be interpreted as a family of limit orders with maximum quantities $Q_i^* - Q_{i-1}^*$ and unit prices $\overline{c}_i$. The JHG allocation and the JHG tariff are illustrated in Figure 3 in the case of three types with strictly convex preferences.

As an application, consider linear utility functions $u_i(Q,T) \equiv v_i Q - T$, subject to a capacity constraint $Q \in [0,1]$. Such linear preferences generalize those in Akerlof (1970) to the case of a divisible good; strict single-crossing requires that v_i be strictly increasing in

[12] The contributions of Jaynes (1978) and Hellwig (1988) are discussed in Section 5.3.

i. Each problem in (ii) admits a unique solution if $v_i \neq \overline{c}_i$ for all i, which we will assume for simplicity. To determine the JHG allocation and the JHG allocation, we apply (ii) in Theorem 3 recursively. By convention, $Q_0^* = 0$; then, beginning with type 1, Q_i^* remains zero as long as $v_i < \overline{c}_i$. If this inequality holds for all types, then the market is inactive; in that case, according to (iii), the essentially unique entry-proof convex market tariff is only defined at zero, with $T(0) = 0$. Otherwise, let i^* be the lowest type such that $v_i > \overline{c}_i$. Applying (ii) at i^* implies that type i^* trades up to capacity at unit price $\overline{c}_{i^*}$; moreover, types $i > i^*$ must also trade $Q_i^* = 1$, as the capacity constraint is binding in (ii). Finally, according to (iii), the unique entry-proof convex market tariff is linear, with $T(Q) = \overline{c}_{i^*}Q$ for all $Q \in [0, 1]$. The upshot from this discussion is that, when the buyer's preferences are linear, the JHG allocation generically features a single layer, and corresponds to the competitive-equilibrium allocation in Akerlof (1970) that maximizes the gains from trade.

The property that the indirect utility functions $(u_i^T)_{i=1}^I$ be ordered according to the weak single-crossing condition plays a key role in the above analysis. This property itself results from the two assumptions that the primitive utility functions $(u_i)_{i=1}^I$ be ordered according to the strict single-crossing condition, and that the market tariff be convex. Because this second assumption effectively constrains market outcomes, it is natural to ask whether it can be dispensed with. The answer is positive in the following three settings. In the two-type case, the proof of Theorem 2 follows from a direct argument that does not require that the market tariff be convex. When the buyer has linear preferences, as above, AMS (2011, p. 1888) also offers a direct proof. Finally, AMS (2021, Online Appendix F) show that the JHG allocation turns out to be the only budget-feasible allocation implemented by an entry-proof market tariff that is first convex and then concave. The general case raises however a difficult issue: in the absence of single-crossing, we do not know for sure whether a contract that attracts type i also attracts all types $j \geq i$, or only a subset of those with a more or less favorable expected cost; as a result, the entry-proofness constraint (17) need not hold. While entry-proofness per se selects a convex tariff in a large class of admissible tariffs allowing for quantity discounts, the general problem thus remains open.

3.4 Discussion

A noticeable feature of the JHG allocation is the relationship between demand and supply on each layer. On the first layer, the price is the expected cost of serving all types, and the quantity supplied is exactly the demand of the first type at this price. Indeed, supplying less would inefficiently ration demand, while supplying more would entail losses on the excess quantity. On the second layer, the first type is no longer active, and the same reasoning applies: the price is the expected cost of serving all types except the first, and the quantity supplied is exactly the residual demand of the second type at this price—and so on. Overall, the quantity supplied on each layer matches the residual demand of the marginal type, at a price equal to expected cost. On each layer but the last one, relatively low-cost types thus subsidize relatively high-cost ones, in contrast with the absence of cross-subsidies that characterizes candidate entry-proof allocations under exclusive competition.

It should also be noted that the JHG allocation typically allows for marginal rates of substitution to differ across types. For instance, in the two-type case under adverse selection, we have $\tau_1(Q_1^*, T_1^*) = \bar{c}_1 < c_2 = \tau_2(Q_2^*, T_2^*)$ when this allocation is interior and separating. This contrasts with private-value models where side trades take place on Walrasian markets, which calls for an equalization of marginal rates of substitution (Hammond (1979, 1987)). Yet, this difference does not create any opportunities for side trading, because the goods under consideration are not the same: for instance, in a Rothschild and Stiglitz (1976) insurance economy, insurance for a low-risk consumer is not the same good as insurance for a high-risk consumer. Indeed, supposing that consumers have access to the same constant-return-to-scale technology as firms, the opportunity cost for type 1 of selling additional coverage to type 2 is c_2, and at this price type 2 is not willing to buy. Similarly, the opportunity cost of selling coverage to type 1 is only c_1, but at this price all types would be attracted; hence the relevant unit cost is $\bar{c}_1$, and at this price type 1 is not willing to buy. The JHG allocation thus exhausts the incentive-compatible gains from trade. Together with entry-proofness, these features support the idea that the JHG allocation is a natural candidate for a competitive allocation.

Alternatively, AMS (2020) propose to reconsider this economy from the viewpoint of a social planner endowed with the same linear technology as the buyers and acting under asymmetric information. As in the classical setting of Harris and Townsend (1981), the planner is able to control all communication among the buyers, and in fact he optimally chooses to prohibit all forms of communication apart from a report each agent privately sends to him. Then the only constraints he faces are the incentive-compatibility constraints

$$\text{For all } i \text{ and } j,\ u_i(Q_i, T_i) \geq u_i(Q_j, T_j) \tag{19}$$

and the budget constraint (5). This leads to the classical definition of second-best allocations as Pareto-optima in the set of budget-feasible and incentive-compatible allocations. In general, such allocations form a non-degenerate continuum, according to the weight put on each type; moreover, under single-crossing, either the downward or the upward local incentive-compatibility constraints must be binding, apart from special cases.[13]

Let us now assume that the planner cannot monitor side trades between different buyers nor prevent the entry of a seller with the same technology. For simplicity, consider the two-type case, and suppose that preferences are strictly convex and satisfy strict single-crossing. Two consequences then follow for the set of allocations that the planner can implement. First, according to Theorem 2, this set collapses to a single allocation, namely, the JHG allocation. It is thus impossible for the planner to redistribute between types: both quantities and transfers are uniquely defined. Second, in general, this allocation is not second-best efficient, in the sense given above. Indeed, the incentive-compatibility constraints (19) are superseded by the entry-proofness constraints (17), which turn out to be necessary and sufficient to characterize the JHG allocation. When $Q_2^* > Q_1^*$, (17) implies that the local incentive-compatibility constraints do not bind, so that the JHG allocation is not second-best.

Overall, the uniqueness of the budget-balanced allocation robust to side trading contrasts with the multiplicity of second-best allocations, which form a nondegenerate frontier. The planner is thus severely constrained by his inability to monitor trades. As discussed in AMS (2020), this result has consequences for actual policies, because the possibility of side trading may undo their effects. For instance, the only possibility for public health insurance is to propose a single basic coverage, sold at average cost, and chosen so as to maximize the utility of low-risk consumers at that price. Private insurers can then compete to provide complementary coverage at price c_2. Another example is provided by bailout policies on financial markets. Under exclusivity, they aim at attracting only the least profitable borrowers, either through direct lending (Philippon and Skreta (2012)), or through the repurchasing of low-quality assets (Tirole (2012)). By contrast, when the borrower can complement a public program with private funds, the only possibility is to provide the same loan Q_1^* to all projects at expected cost, while the riskiest borrowers in addition turn to a competitive market for additional funding at price c_2.

4 Competitive Screening

Entry-proofness provides a parsimonious and tractable way of modeling perfect competition, which is relatively insensitive to the details of market interactions; for that reason, the

[13]See, for instance, Prescott and Townsend (1984) and Crocker and Snow (1985) for characterizations of second-best allocations in Rothschild and Stiglitz (1976) insurance economies.

JHG allocation characterized in Section 3 is arguably a natural and robust candidate for a competitive-equilibrium allocation of a nonexclusive market subject to adverse selection. Yet, by design, this approach does not shed light on how this allocation may be decentralized; a valuable complement to this approach would thus be to implement the JHG allocation as the unique equilibrium outcome of an extensive-form game in which strategic sellers compete to serve privately informed buyers.

To start with, and by way of comparison, we should observe that decentralization is easy to achieve in the standard case of exclusive competition. Indeed, in this context, the unique entry-proof allocation characterized in the insurance setting of Rothschild and Stiglitz (1976) or in the more general setting of Riley (1979) can be easily supported—as long as it exists—in a pure-strategy equilibrium of a competitive-screening game in which sellers first simultaneously post menus of contracts, from which the buyer then choose a single contract according to her type; specifically, there exists an equilibrium of this game in which two sellers offer a menu consisting of the trades comprised in this allocation. Therefore, the existence problem under exclusivity is not tied to decentralization per se, but to the fact that an entry-proof allocation may robustly fail to exist.[14]

By contrast, under nonexclusivity, we know that an entry-proof tariff exists, but its decentralization is much more delicate, because the buyer is now free to combine contracts issued by different sellers. As we shall now see, this generates novel strategic effects, which make it more difficult—indeed, in general, impossible—to implement the JHG allocation via competitive-screening games.

4.1 The Competitive-Screening Game

To clarify this issue, let us consider a general setting in which a finite number K of sellers simultaneously contract with a single buyer. We throughout assume that types are ordered according to the strict single-crossing condition and, when types are continuously distributed, that the mapping $(i, q, t) \mapsto u_i(q, t)$ is continuous. As discussed in Section 2.1, trade is nonanonymous and contracting is bilateral, in the sense that trade between a seller and a buyer can only be made contingent on the information reported by the buyer to this seller. In these situations, the menu theorems of Peters (2001), Martimort and Stole (2002), and Page and Monteiro (2003) allow us to restrict, with no loss of generality, to competition

[14] While the existence issue has been addressed by considering mixed strategy-equilibria (Rosenthal and Weiss (1984), Dasgupta and Maskin (1986), Farinha Luz (2017)) or by considering alternative extensive forms (Miyazaki (1977), Wilson (1977), Spence (1978), Riley (1979), Engers and Fernandez (1987), Netzer and Scheuer (2014), Mimra and Wambach (2019)), the corresponding equilibrium allocations typically do not coincide with the Rothschild and Stiglitz (1976) allocation; see Mimra and Wambach (2014) for a survey.

in menus or nonlinear tariffs. The corresponding extensive-form game, which we denote by G^{CS}, unfolds in two stages:

1. Each seller k offers a compact menu of contracts $C^k \subset \mathbb{R}_+ \times \mathbb{R}$ that contains at least the null trade $(0, 0)$.

2. After privately learning her type, the buyer selects a contract from each of the menus C^k offered by the sellers.

A pure strategy for type i is a function that maps every menu profile $(C^1, \ldots, C^K)$ into a contract profile $((q^1, t^1), \ldots, (q^K, t^K)) \in C^1 \times \ldots \times C^K$. The compactness of the sellers' menus ensures that every type i's utility-maximization problem

$$\max \left\{ u_i \left(\sum_k q^k, \sum_k t^k \right) : (q^k, t^k) \in C^k \text{ for each } k \right\}$$

always has a solution. The solution concept for G^{CS} is pure-strategy subgame-perfect Nash equilibrium. For future reference, we let

$$T(Q) \equiv \min \left\{ \sum_k t^k : (q^k, t^k) \in C^k \text{ for each } k \text{ and } \sum_k q^k = Q \right\}$$

be the market tariff associated to the equilibrium menus C^k, and we let, for each i,

$$U_i \equiv \max\{u_i(Q, T(Q)) : Q \geq 0\}$$

be the equilibrium utility of type i.

It should be noted that the set of strategies for the sellers in G^{CS} is the same as in a standard competitive-screening game under exclusivity. Yet the assumption that the buyer can simultaneously trade with several sellers has two implications for the set of potentially profitable contracts any seller may offer. On the one hand, it tends to expand this set, as this seller may choose to complement his competitors' offers by proposing additional trades to the buyer. On the other hand, it also gives his competitors more instruments to block his deviations, compared to when competition is exclusive; indeed, contracts that are not traded on the equilibrium path may become relevant in case a seller deviates, and in fact equilibria often require the presence of such latent contracts, as we shall now see.

4.2 Linear Preferences

Let us first assume that the buyer has linear preferences, subject to a capacity constraint. In this scenario, which extends Akerlof (1970) to the case of a divisible good, the JHG allocation features a single layer, and corresponds to the competitive-equilibrium allocation that maximizes the gains from trade. The following result, due to AMS (2011), shows that this allocation is uniquely supported in any equilibrium of G^{CS}:

Theorem 4 *Let $u_i(Q,T) \equiv v_iQ - T$ for $Q \in [0,1]$ and suppose that $v_i \geq c_i$ for all i. Then, generically, any equilibrium of G^{CS} implements the JHG allocation, and there exists a linear-pricing equilibrium with price $\overline{c}_{i^*}$, where i^* is the first type i such that $v_i > \overline{c}_i$.*

In equilibrium, all the buyer types with valuations $v_i > \overline{c}_{i^*}$ trade up to capacity, while all the buyer types with valuations $v_i < \overline{c}_{i^*}$ do not trade at all. Sellers thus earn zero expected profits, and none of them is indispensable to serve any buyer type. Finally, all trades take place at the same price in equilibrium, despite the fact that sellers can propose arbitrary nonlinear tariffs. Thus Theorem 4 provides a game-theoretic foundation for Akerlof's (1970) predictions in a setting where the traded good is divisible and, besides nonexclusivity, few restrictions on feasible trades or instruments are imposed. In particular, low-valuation types such that $c_i < v_i < \overline{c}_{i^*}$ are excluded from trade in equilibrium, unlike what would happen under exclusive competition. The existence and uniqueness of the equilibrium allocation described in Theorem 4 paves the way for many applications—notably in finance, where the divisibility assumption is natural.[15]

The driving intuition for these results is that the unobservability of the buyer's aggregate purchases limits the sellers' ability to screen types and thereby the effectiveness of cream-skimming deviations. Suppose, for instance, that the equilibrium price is high, so that low-valuation, and hence on average low-cost types are not served. A cream-skimming deviation targeted at these types must involve trading a relatively small quantity q at a relatively low price. However, this contract becomes also attractive to high-valuation, and hence on average high-cost types if, along with it, they can trade the additional quantity $1-q$ at the equilibrium price; this is exactly what the linear tariff allows for.

This reasoning illustrates the fact that deviations are blocked by latent contracts, that is, contracts that are not traded on the equilibrium path, but which the buyer may want to trade at the deviation stage.[16] In general, many such contracts are needed to support the equilibrium allocation. This is particularly striking when the distribution of types is discrete, because then only finitely many contracts are effectively traded, while infinitely many latent

[15]A case in point is the security-design model of DeMarzo and Duffie (1999), to which Theorem 4 directly applies; see Biais and Mariotti (2005) for an early result along these lines.

[16]The role of latent contracts has originally been stressed in moral-hazard environments; see, for instance, Hellwig (1983), Arnott and Stiglitz (1991), Bisin and Guaitoli (2004), and Attar and Chassagnon (2009).

contracts must be issued. In particular, no equilibrium can in this case be sustained through direct mechanisms, which provides a concrete example of a failure of the revelation principle in common-agency games (Peters (2001), Martimort and Stole (2002)).

4.3 Cournot-Convergence under Strictly Convex Preferences

Although Theorem 4 holds for general distributions of types,[17] it does not generally extend to the case of strictly convex preferences for the buyer. In a seminal article, Biais, Martimort, and Rochet (2000) consider a situation in which strategic market-makers (sellers) compete to serve a risk-averse insider (buyer) who has private but imperfect information about the value of an asset, and thus has both informational and hedging motives for trade. Assuming that the buyer has constant absolute risk-aversion α and faces residual Gaussian risk with variance σ^2, they show the following result:

Theorem 5 *Let $u_i(Q,T) \equiv v_i Q - \frac{\alpha\sigma^2}{2} Q^2 - T$ and let the buyer's type be continuously distributed. Then, under regularity conditions, G^{CS} admits a symmetric equilibrium in which sellers post the same strictly convex tariff and earn strictly positive expected profits. The equilibrium market tariff converges to the JHG tariff as the number K of sellers grows large.*

This equilibrium exhibits the Cournot-like feature that each seller is indispensable to serve any buyer type who trades a nonzero quantity in equilibrium. Specifically, the strict convexity and symmetry of the equilibrium tariffs implies that any such type has a unique best response that consists in evenly splitting her total purchases between the sellers. This contrasts with the equilibria that obtain in the linear case, in which no seller is indispensable and thus any buyer type who trades up to capacity has multiple best responses that involve trading with different sellers. Because sellers earn strictly positive expected profits in equilibrium, the aggregate equilibrium allocation does not coincide with the JHG allocation; yet, in analogy with classical Cournot-convergence theorems, it converges to the competitive JHG allocation as the number of sellers grows large.

The assumption in Theorem 5 that the buyer's type be continuously distributed is key to ensure that, despite being indispensable, a single seller cannot profitably raise his tariff. To illustrate this point, suppose that a seller deviates by replacing a portion of his strictly convex equilibrium tariff by the corresponding chord. This would increase his expected profit if the buyer's behavior remained the same. But such a change raises (lowers) the marginal price for relatively low-cost (high-cost) types who would choose trades in this portion of

[17] Specifically, it holds for any distribution for which $v_i \geq c_i$ at any atom.

the tariff. As a result, under adverse selection, trades change in an unfavorable way for the deviating market maker. This effect is reinforced by the fact that the buyer simultaneously trades with several sellers, as any increase in the quantity she purchases from a seller is compensated by a reduction in the quantity she purchases from his competitors.

4.4 Exclusion

We now argue that, in any game G^{CS}, exclusion is a robust feature of any equilibrium that shares two key properties of the Biais, Martimort, and Rochet (2000) equilibrium. To formulate these properties, we let, for each k,

$$T^{-k}(Q) \equiv \min\left\{\sum_{l\neq k} t^l : (q^l, t^l) \in C^l \text{ for each } l \neq k \text{ and } \sum_{l\neq k} q^l = Q\right\}$$

be the submarket tariff associated to the equilibrium menus C^l, $l \neq k$, and we let, for all i and k,

$$z_i^{-k}(q,t) \equiv \max\left\{u_i(q + Q^{-k}, t + T^{-k}(Q^{-k})) : Q^{-k} = \sum_{l\neq k} q^l \text{ for some } (q^l, t^l) \in C^l,\, l \neq k\right\}$$

be type i's indirect utility from trading (q, t) with seller k. The two properties we wish to emphasize can now be stated as follows.

P1 For each k, there exists i such that $U_i = z_i^{-k}(0, 0)$.

An equilibrium satisfies P1 if, for each seller, there exists at least one type for whom trading with this seller is not indispensable for her to obtain her equilibrium utility; this reflects the relatively weak requirement that, in equilibrium, the buyer's individual-rationality constraint in her dealings with each seller must bind for at least one type.

P2 For all k and $Q > 0$, $T(Q) < T^{-k}(Q)$.

An equilibrium satisfies P2 if trading with each seller is indispensable for each type who purchases a nonzero aggregate quantity.

The symmetric equilibrium characterized by Biais, Martimort, and Rochet (2000) satisfies both P1 and P2 because all sellers offer the same strictly convex tariff. Indeed, this implies that each seller is indispensable to minimize the cost of purchasing any strictly positive aggregate quantity, whence P2. This also implies that the indirect utility functions z_i^{-k} satisfy the strict single-crossing condition for all k, so that any seller k for whom the individual-rationality constraint were not binding could raise his tariff without affecting the buyer's incentives, whence P1.

The following theorem, a formal proof of which is provided in the appendix, abstracts from the parametric assumptions of Biais, Martimort, and Rochet (2000) to show that exclusion must more generally take place in any equilibrium of any game G^{CS} that satisfies P1–P2:

Theorem 6 *Consider an equilibrium of a game G^{CS} that satisfies P1–P2. Then there exists some type i_1 such that, in equilibrium, every type $i \leq i_1$ trades $q_i^k = 0$ with every seller k and obtains utility $U_i = u_i(0,0)$.*

In light of this result, it is worth noticing that Biais, Martimort, and Rochet (2000) assume that the continuous support of the buyer's type distribution includes an interior type i_0 such that $\tau_{i_0}(0,0) = c_{i_0}$ and thus for which there are no gains from trade.[18] This in turns ensures that there exists an interval of types at the bottom of the type distribution who are excluded from trade in equilibrium, as requested by Theorem 6. However, this assumption is fairly restrictive: it does not hold, for instance, in standard Rothschild and Stiglitz (1976) insurance economies, because a risk-averse consumer is always willing to purchase full coverage at the fair price, equal to her riskiness. As we shall now see, the existence of equilibria of G^{CS} games then becomes problematic.

4.5 The Existence Conundrum

As a starting point, let us consider the two-type case with $c_2 > c_1$ and strictly convex preferences for the buyer, and let us examine a candidate equilibrium of G^{CS} in which both types 1 and 2 purchase strictly positive quantities. AMS (2014) show that the aggregate equilibrium allocation then has the same structure as the JHG allocation. First, sellers earn zero expected profits. Second, the quantity Q_1 purchased by type 1 is priced at the average cost $\bar{c}_1$, while the additional quantity $Q_2 - Q_1$ purchased by type 2 is priced at the marginal cost c_2. Third, the quantity Q_1 purchased by type 1 is equal to her demand $D_1(\bar{c}_1)$.

This sounds promising but, because type 1 is not excluded from trade, Theorem 6 implies that P1 or P2 has to give way. It is intuitive that the weak requirement P1 should be maintained, and thus that P2 should be dropped. Specifically, it can be shown that P1 holds for type 1 and that $T(Q_1) = T^{-k}(Q_1) = \bar{c}_1 Q_1$ for all k, so that no seller is indispensable to

[18] Specifically, the insider can trade on both sides of the market, so that types $i > i_0$ ($i < i_0$) are willing to buy (sell) at price c_i. As in Back and Baruch (2013) and Biais, Martimort, and Rochet (2013), we focus on the ask side of the market.

provide type 1 with her equilibrium aggregate trade. Because any seller k who trades with both types 1 and 2 on the equilibrium path makes a profit with type 1 and a loss with type 2, this opens the way to a *lemon-dropping* deviation that essentially consists for seller k in convincing type 2 to trade the layer $(Q_1, \bar{c}_1 Q_1)$ with his competitors and the complementary layer $(Q_2 - Q_1, c_2(Q_2 - Q_1))$ with him, thus neutralizing his losses with type 2.

Specifically, the deviation involves two contracts. When $Q_2 > Q_1 > 0$, the first one is essentially that traded by type 1 with seller k on the equilibrium path, while the second one makes the quantity $Q_2 - Q_1$ available at a unit price slightly lower than c_2. Thus type 2 can strictly increase her utility by trading the second contract on top of the layer $(Q_1, \bar{c}_1 Q_1)$ made available by the sellers other than k; besides, seller k can break ties to make sure that type 2 strictly prefers this contract to the first one. As a result, seller k can make his loss with type 2 arbitrarily small while securing, as $\bar{c}_1 > c_1$, a strictly positive profit with type 1; hence the deviation is profitable. A similar but slightly more involved argument shows that seller k has a profitable deviation also when $Q_2 = Q_1 > 0$.

As a consequence, trade can take place in equilibrium only if type 1 is excluded from trade: in short, the possibility of cross-subsidizing between contracts at the deviation stage makes it impossible to support cross-subsidies between types on the equilibrium path. Specifically, the following result holds (AMS (2014, Theorems 1–2)):

Theorem 7 *Suppose there are two buyer types with strictly convex preferences, and that $c_2 > c_1$. Then any equilibrium of G^{CS} implements the JHG allocation, but an equilibrium exists if and only if $Q_1^* = 0$ in that allocation, that is, if and only if $\tau_1(0,0) \leq \bar{c}_1$. If an equilibrium exists, it can be sustained by each seller posting the JHG tariff, which consists of a single layer with unit price c_2.*

To allow for a finer comparison with the continuous-type model of Biais, Martimort, and Rochet (2000), let us now suppose as in Section 2.1 that there is an arbitrary but finite number I of types, each assumed to have strictly convex preferences. The same difficulty arises as for the characterization of entry-proof tariffs, however: when $I > 2$, the game G^{CS} with no restrictions on admissible menus is hardly tractable. As in Section 3.3, a convenient assumption is that sellers are restricted to post convex tariffs, which ensures that the indirect utility functions $(z_i^{-k})_{i=1}^I$ are quasiconcave and satisfy the weak single-crossing property for all k. The resulting convex-tariff game G^{CT} can be interpreted as a discriminatory auction in which sellers simultaneously bid quantities at each marginal price. In line with Back and Baruch (2013), this captures oligopolistic competition on a limit-order market where market-makers post collections of limit orders that are executed by an informed insider in order of price priority.

It should be noted that the set of deviations for the sellers is much smaller in G^{CT} than in G^{CS}; in particular, the deviation that led to Theorem 7 is no longer feasible. However, the following result, due to AMS (2019a, Theorems 2–3), shows that, in spite of this, the equilibrium-existence problem only becomes more acute when the number of types increases:

Theorem 8 *Suppose there are I buyer types with strictly convex quasilinear preferences, and that c_i is strictly increasing in i. Then any equilibrium of G^{CT} implements the JHG allocation, but an equilibrium exists if and only if $Q_i^* = 0$ for all $i < I$ in that allocation, that is, if $\tau_i(0,0) \leq \bar{c}_i$ for all $i < I$. If an equilibrium exists, it can be sustained by each seller posting the JHG tariff, which consists of a single layer with unit price c_I.*

The proof proceeds by showing that, in any candidate equilibrium, the market tariff T is piecewise linear and has a structure similar to that of the JHG tariff. That is, the quantity supplied on each layer but the last one matches the residual demand of the marginal type, at a price equal to the expected cost of serving the buyer types who trade along this layer. This implies that sellers earn zero expected profits, and also that the marginal type on any such layer exhausts the supply at the corresponding marginal price. As a result, each seller offering trades at this marginal price is indispensable for the marginal type and all higher types to reach their equilibrium utility. But one can hardly be indispensable and yet earn zero expected profit: hence any such seller could raise his tariff in a profitable way, a contradiction. This shows that the market tariff T must consists of a single layer and be such that no seller is indispensable to serve the buyer types who trade along it. However, each seller will then want to issue a limit order to hedge against the risk of large purchases emanating from the most costly types. This implies that all types except perhaps the last one must be excluded from trade.

The upshot from Theorems 5 and 7–8 is twofold. First, the structure of equilibria of discrete-type models, when they exist, is very different from that of the equilibria of continuous-type models that have been emphasized in the literature: namely, pricing is linear and only the last type can trade in equilibrium. Second, necessary and sufficient conditions for the existence of an equilibrium become increasingly stringent as the number of types increases: equilibria fail to exist when there are sufficiently many types with similar preferences, as when we approximate the continuous sets of types postulated by Biais, Martimort, and Rochet (2000) or Back and Baruch (2013). The pure-strategy-equilibrium correspondence thus fails to be lower hemicontinuous when we move from discrete-type models to continuous-type models. Overall, the predictions of competitive-screening models are very sensitive to fine modeling details, which makes them somewhat fragile.

4.6 Ways Out

A natural way to address the equilibrium-existence problem in discrete-type models is to weaken the equilibrium concept. This can be done in two ways.

First, we may consider mixed-strategy equilibria of G^{CS}, the existence of which follows from Carmona and Fajardo (2009). Preliminary investigations of the two-type case have led to a robust example of a mixed-strategy equilibrium that exists when the necessary and sufficient conditions for the existence of a pure-strategy equilibrium are not satisfied (Attar, Farinha Luz, Mariotti, and F. Salanié (2021)). The key point is that the strategic uncertainty faced by each seller regarding the tariffs offered by his competitors makes it now impossible for him to target specific types, unlike in the deviations used to derive Theorems 7–8. However, this equilibrium bears no obvious relationship with existing equilibrium candidates; the JHG allocation, in particular, does not emerge even when the number of sellers grows large. The systematic characterization of mixed-strategy equilibria nevertheless remains a fascinating—though hard—topic for future research.

Next, we may consider ε-equilibria of G^{CS}. AMS (2019a) show that, if every type i has quasilinear preferences $u_i(Q,T) \equiv v_i(Q) - T$, then, as the number K of sellers grows large, G^{CS} admits an ε-equilibrium, with ε of the order of $1/K^2$, that supports the JHG allocation. The intuition is that if $K-1$ sellers contribute to providing a fraction $1/K$ of the JHG tariff, the residual gains from trade for the remaining seller vanish when K grows large because the resulting market tariff is almost entry-proof. The reason why convergence takes place at rate $1/K^2$ is that these gains from trade are, for every type i, bounded above by

$$v_i(Q_i^*) - v_i\left(\frac{K-1}{K}\,Q_i^*\right) - \frac{1}{K}\,\bar{c}_i Q_i^*,$$

which is at most of the order of $1/K^2$ as $v_i'(Q_i^*) \leq \bar{c}_i$ at the JHG allocation. Thus we retrieve the convergence result of Theorem 5 for the competitive limit, albeit at the cost of relying on a notion of approximate equilibrium. Glosten (1994, Proposition 2) provides a similar result, assuming from the outset that there is an infinite number of sellers.

4.7 Regulation

An alternative route to decentralize the JHG allocation consists in explicitly introducing a market regulation. In this spirit, AMS (2019b) study a nonexclusive insurance market in which it is prohibited for sellers to cross-subsidize between contracts. The regulation thus bears on the total profit a seller earns on each contract, and is targeted at dumping practices; it can alternatively be interpreted as banning profits on basic-coverage contracts.[19] Specifically, let G^{CSR} be the regulated game that is obtained from G^{CS} by adding one final stage in which a seller's profit is confiscated whenever he makes a loss on any of the contracts he is trading. The following result then holds (AMS (2019b, Theorem 2)):

Theorem 9 *Suppose there are two buyer types with strictly convex preferences, and that* $c_2 > c_1$. *Then the JHG allocation is the unique candidate-equilibrium allocation of* G^{CSR}. *Moreover, under regularity conditions on the buyer's preferences,* G^{CSR} *has an equilibrium as long as there are sufficiently many sellers.*

There are two parts in this result. The necessity part states that the regulation has no anti-competitive implications. The intuition is that each seller aims at increasing his profit by complementing the aggregate coverage provided by its competitors, which gives rise to a form of Bertrand competition over each layer; as a result, the JHG allocation is the unique candidate-equilibrium outcome. The sufficiency part reflects that the regulation, by blocking the cross-subsidies between contracts that were at the root of Theorem 7, helps restore existence of an equilibrium even if both types trade in the JHG allocation.

A necessary feature of equilibrium is that sellers must issue latent contracts to discipline their competitors. As in AMS (2011), these contracts are not traded in equilibrium but are meant to block cream-skimming deviations. In the context of insurance, these contracts provide additional coverage that high-risk consumers are willing to combine with the coverage provided by any such deviation. This makes it impossible for a seller to profitably deviate by separating low-risk from high-risk consumers.

5 Alternative Extensive Forms

An important takeaway from the literature surveyed in Section 4 is that competitive-screening games generally fail to implement the JHG allocation. This failure can be traced back to a common source, namely, the paucity of instruments allowing to punish a deviating

[19] Several insurance markets are actually regulated along analogous lines. For instance, in health insurance, Germany and Switzerland rely on a central fund to redistribute costs among firms according to a risk-equalization scheme. These cost-sharing mechanisms, by pooling and redistributing costs among sellers of a standardized basic-coverage contract, prevent firms from earning abnormal profits on such coverage by dropping lemons on their competitors.

seller. Indeed, if sellers make their offers simultaneously, the only device available to block deviations consists in letting the buyer select latent contracts in the nondeviating sellers' menus. However, this device is effective only when the buyer has linear preferences, reflecting the very special property that, if latent contracts are issued at the equilibrium price, all the types who are willing to trade at this price have the same indirect utility function z_i^{-k}: they are willing to trade a contract issued by a deviating seller if and only if its unit price is less than the equilibrium price. This no longer holds when the buyer has strictly convex preferences and different types trade at different marginal prices.

The generic failure of latent contracts at sustaining equilibria in competitive-screening games—let alone at implementing the JHG allocation—suggests that we look for alternative extensive forms whereby the sellers sequentially receive information over the course of the game. Three kinds of extensive forms have been studied in the literature. The first one allows the buyer to signal her type by recontracting, which requires that all previously signed contracts be publicly observable. The second one lets the sellers bid through an ascending discriminatory auction, in which the offers made at previously quoted prices are publicly observable. The third one enables the sellers—and, possibly, the buyer as well—to voluntarily disclose information about the contracts selected by the buyer; which information eventually becomes available, and to whom, then depends on the agents' disclosure strategies.

5.1 Recontracting

Beaudry and Poitevin (1995) study a sequential game in which a risk-averse entrepreneur whose project can be of low or high riskiness can repeatedly solicit financing from successive cohorts of uninformed lenders before the realization of the project's return. An important feature of this game is that there is a potentially infinite number of recontracting rounds. Hence there is no last stage of the game in which the entrepreneur could commit to reject further offers: a lender can never be sure that she will not try to further diversify her risk by selling new claims on her project.[20] That is, nonexclusivity is distinctively linked to the absence of commitment and the sequential nature of contracting.

At each round of recontracting, the buyer can solicit further offers from a new cohort of lenders; if she does so, she has to provide a summary of all the contracts she has signed in previous rounds, though not of the contracts she has rejected. This observability assumption stands in contrast with the competitive-screening models surveyed in Section 4, in which, by design, no seller has information about existing contractual relationships. Another difference

[20]This feature is similar to Kahn and Mookherjee's (1998) moral-hazard model of nonexclusive contracting.

is that, at each round of recontracting, competition is exclusive, in the sense that the buyer can accept at most one offer from one lender. The assumption that lenders are short-lived and cannot observe previously rejected contracts is meant to limit the lenders' ability to sustain collusive outcomes.

In this setup, Beaudry and Poitevin (1995) show that, when the high-risk project has positive NPV and each project can be financed by riskless claims using the entrepreneur's collateralizable wealth, there exists a perfect Bayesian equilibrium that supports the JHG allocation: an entrepreneur with a low-risk project obtains the net claims corresponding to her preferred financial position among the contracts with nonnegative pooling profits, while an entrepreneur with a high-risk project manages to completely diversify her risk without pledging any of her wealth in the project.

This outcome can be supported without recontracting on the equilibrium path, with each seller posting in the first round a menu consisting of the trades comprised in the JHG allocation. If the entrepreneur solicits additional offers, then she is believed to have a high-risk project; for instance, she is offered $Q_2^* - Q_1^*$ at price c_2 if she initially accepted (Q_1^*, T_1^*), and she is offered $(0, 0)$ if she initially accepted (Q_2^*, T_2^*). No lender, therefore, can profit by deviating, for any offer that would be accepted by an entrepreneur with a low-risk project would also be accepted by an entrepreneur with a high-risk project in anticipation of future rounds of recontracting. Notice that it is essential for this reasoning that the buyer can only select a single contract at each round, that previously signed contracts be observable, and that there always be further opportunities of recontracting. It is fair to ask whether the first two assumptions are consistent with the intuitive notion of a nonexclusive market; in particular, a prediction of the model is that each entrepreneur trades with a single lender.

5.2 A Discriminatory Ascending Auction

An alternative approach consists in sticking more closely to competitive-screening games, while allowing punishments to be carried out by the sellers themselves. This requires, of course, that deviations be observable by the nondeviating sellers, in the spirit of the reactive-equilibrium literature cited in Footnote 14. In this respect, the recursive structure of the JHG allocation suggests that it be implemented sequentially, layer by layer, from the bottom up. To validate this intuition, AMS (2021) propose to model the strategic interactions between sellers as a discriminatory ascending auction.

In their model, the auctioneer quotes price sequentially, in increasing order, and according to a discrete price grid with a minimum tick size. Each time he quotes a new price, each seller publicly announces the maximum quantity he stands ready to trade with the buyer at this price; in other terms, he offers a limit order at the current price. Once this auctioning phase is completed, the buyer selects which quantities to purchase from which sellers at each price, according to her type. We denote by G^{DA} the corresponding extensive-form game. The solution concept for G^{DA} is pure-strategy subgame-perfect Nash equilibrium.

It should be noted that G^{DA} can be interpreted as a sequential version of the convex-tariff game G^{CT}. Indeed, as it is optimal for the buyer to take up the best price offers first, she in the end faces a collection of convex tariffs that aggregate into a convex market tariff T. From her perspective, the fact that T was built up sequentially is irrelevant.

For the sellers, by contrast, the fact that bids are made sequentially and publicly during the auctioning phase is crucial, as it allows them to react, at any price, to a deviation at a lower price; indeed, the key advantage of a sequential auction lies in its transparency, a point emphasized in other contexts by Milgrom (2000) and Ausubel (2004). Importantly, such reactions—which can take place almost immediately when the tick size is small—can only take the form of quantity increases or decreases at future prices, while the quantities supplied at lower prices cannot be withdrawn or augmented. This commitment assumption makes G^{DA} quite different from the Walrasian *tâtonnement* process.

From our implementation perspective, two questions immediately arise. First, does G^{DA} admit an equilibrium? Second, how do equilibrium allocations relate to the JHG allocation? The second question is especially pressing, because the dynamic nature of G^{DA} may perhaps allow to sustain equilibria with collusive outcomes.

The first result established by AMS (2021) is that G^{DA} admits a very simple Markov perfect equilibrium. The relevant states variables are the current price p and the aggregate quantity Q^- supplied at prices lower than p. Assuming for simplicity that the buyer has quasilinear preferences, with demand function D_i, the residual demand of type i in state (p, Q^-) is $\max\{D_i(p) - Q^-, 0\}$; observe that maximizing aggregate expected profits in any state (p, Q^-) asks for serving the residual demand of the type i such that $\overline{c}_i < p \leq \overline{c}_{i+1}$, which we shall call the *profitable residual demand* in state (p, Q^-). The following result then holds (AMS (2021, Theorem 3)):

Theorem 10 *Suppose there are I buyer types with strictly convex quasilinear preferences, and that $\overline{c}_i$ is strictly increasing in i. Then there exists a Markov perfect equilibrium of G^{DA} in which, in any state (p, Q^-),*

(i) if $p \leq \overline{c}_1$, each seller supplies a zero quantity;

(ii) if $\overline{c}_1 < p \leq \overline{c}_I$, each seller supplies a share $1/K$ of the profitable residual demand;

(iii) if $p > \overline{c}_I$, each seller supplies an infinite quantity.

Moreover, the resulting aggregate equilibrium allocation converges to the JHG allocation as the tick size goes to zero.

The mechanics of the equilibrium are very simple. First, at any price $p \leq \overline{c}_I$, no unilateral increase in supply is profitable if the profitable types at price p, that is, all the types j such that $p > \overline{c}_j$, rationally choose to ignore this deviation and carry on trading the same quantity with each seller; indeed, the deviation can then only lead to losses with unprofitable types at price p and reduce their residual demand at higher prices. Second, at any price $p \leq \overline{c}_I$, no unilateral decrease in supply is profitable, because the corresponding increase in the profitable residual demand at the next price will be shared with the other sellers in the continuation equilibrium. As a result, although each seller is indispensable to serve a strictly positive profitable residual demand, no seller has an incentive to wait for a higher price to be quoted. This stands in stark contrast with the convex-tariff game G^{CT}, in which a seller indispensable at price p can always secretly deviate by bidding a lower quantity at price p and a higher quantity at a slightly higher price.

Given a tick size Δ, the following equilibrium outcome obtains. As soon as the price reaches $\overline{c}_1 + \Delta$, the sellers serve the demand $D_1(\overline{c}_1 + \Delta)$ of type 1; this quantity will also be purchased by types $i > 1$. Then, as soon as the price reaches $\overline{c}_2 + \Delta$, the sellers serve the residual demand $\max\{D_2(\overline{c}_2 + \Delta) - D_1(\overline{c}_1 + \Delta), 0\}$ of type 2; this quantity will also be purchased by types $i > 2$—and so on, until the price reaches $\overline{c}_I + \Delta$, at which point the sellers flood the market. By construction, the resulting aggregate equilibrium allocation converges to the JHG allocation as Δ goes to zero.

The second result established by AMS (2021) is that every sequence of equilibria of G^{DA} satisfies this convergence property, modulo an intuitive refinement that can be described as follows. In any play of the game, every type i accepts all offers up to some price p_i. However, the sellers' aggregate supply at price p_i may well exceed type i's residual demand at this price; she can then break ties in many different ways, and her choice typically matters to the sellers. An equilibrium of G^{DA} is *robust to irrelevant offers* if every type i's trades at price p_i do not depend on offers made at prices $p > p_i$. Intuitively, the buyer never punishes a seller for deviating at a price at which she is not willing to trade.[21] The following result then holds (AMS (2021, Theorem 4)):

[21] The tie-breaking rules used in the proof of Theorem 10 are consistent with this refinement.

Theorem 11 *In any sequence of equilibria robust to irrelevant offers of G^{DA} associated to a sequence of tick sizes going to zero, the aggregate equilibrium allocations converge to the JHG allocation, and the equilibrium market tariffs converge to the JHG tariff.*

Leaving technical details aside, Theorem 11 results from a simple Bertrand undercutting argument. To see this, suppose, by way of contradiction, that, given the limit market tariff, strictly positive expected profits can be earned at some price p. Because the highest price at which trade takes place can be shown to be bounded along any sequence of equilibria when the tick size Δ goes to zero, we can focus on the highest such p. Continuation profits at higher prices must be zero: indeed, the robustness refinement ensures that, if they were strictly negative, then, for Δ small enough, some seller could profitably withdraw all his offers at such prices without affecting his expected profits up to price p. Now, the convergence of aggregate supply functions as Δ goes to zero implies that, for Δ small enough, aggregate supply in a left-neighborhood of p becomes negligible. Hence each seller can, almost without losing priority, undercut his competitors at a price arbitrarily close to p, and supply nothing afterwards; by doing so, he can appropriate almost all expected profits at price p, and the robustness refinement again ensures that his expected profits at lower prices remain the same. But then this deviation would be profitable for at least one seller, a contradiction. Overall, this argument shows that, given the limit tariff, strictly positive expected profits cannot be earned at any price p; given budget-balance, this exhaustion of gains from trade characterizes the JHG tariff, from which Theorem 11 follows.

Taken together, Theorems 10–11 provide an implementation of the JHG allocation as the essentially unique equilibrium outcome of competition when each seller can quickly react to his competitors' offers; an attractive feature of this implementation is that, in the spirit of Bertrand competition, it only requires that there be two competing sellers. From a market-design perspective, these positive results invite us to reconsider the role of continuous bidding for financial and insurance markets, and offer a useful complement to studies that advocate a transformation of continuous markets into batch auctions, so as to avoid possible inefficiencies linked to high-frequency trading (Budish, Cramton, and Shin (2015)).

5.3 Information Disclosure

A common feature of the recontracting game of Beaudry of Poitevin (1995) and of the ascending discriminatory auction of AMS (2021) is that the release of information to sellers about previously signed contracts or previously made offers is exogenous. By contrast, following the seminal contribution of Jaynes (1978), several authors have studied what happens when agents can voluntarily disclose information about the contracts they are engaged in, so that the information available to each seller is endogenously determined in equilibrium. Although the buyer can still in principle subscribe to multiple contracts issued by different sellers, each seller can then enforce exclusivity clauses contingent on the information disclosed by his competitors.[22]

Jaynes (1978) considers the following timing. First, each seller offers a menu of contracts, possibly specifying exclusivity clauses; besides, he commits to disclose to a subset of his competitors the contract selected by the buyer from his own menu. Next, the buyer selects a contract from each seller's menu, information is disclosed, and exclusivity clauses are enforced, which determines the contracts that are eventually executed. The JHG allocation turns out to be the only candidate-equilibrium allocation. In the two-type case, Jaynes' (1978) proposed equilibrium can be described as follows. First, two sellers offer a limit order at price $\bar{c}_1$ with maximum quantity Q_1^*; these sellers share their information and enforce exclusivity clauses, which ensures that they do not make losses by overselling to type 2. Second, two sellers offer a limit order at price c_2 with maximum quantity $Q_2^* - Q_1^*$; these sellers do not share their information.

As pointed out by Hellwig (1988), however, this candidate equilibrium is not robust to a cream-skimming deviation, whereby one of the sellers supposed to offer trades at price $\bar{c}_1$ secretly deviates by offering a contract at a price slightly less than $\bar{c}_1$ that, per se, attracts type 1, but would attract type 2 only in combination with contracts issued by the nondeviating sellers at price $\bar{c}_1$. Thanks to the information disclosed by the nondeviating sellers, the deviating seller is still able to enforce exclusivity on this contract; as a result, he is assured to only attract type 1, and the deviation is profitable. Intuitively, the idea is that sellers by themselves have no basis for treating deviating and nondeviating firms asymmetrically in their disclosure decisions.

In that respect, it should be noted that the above deviation is effective only if the nondeviating sellers are not aware of it. Otherwise, they could punish the deviating seller by concealing from him the contracts selected by the buyer from their menus. In that case, a seller attempting to cream-skim type 1 would no longer be able to enforce exclusivity; as a result, his contract would also become attractive for type 2, along with contracts issued by nondeviating sellers at price $\bar{c}_1$ and, possibly, c_2. Hence cream-skimming is impossible if the sellers' offers are public. This intuition is formalized by Hellwig (1988), who studies a multi-stage extensive-form game in which each seller can make his disclosure decisions

[22] In exclusive insurance markets, information sharing enables the construction of joint databases collecting information on each loss so as to ensure that the same loss is not indemnified twice.

contingent on his competitors' contract offers.

Jaynes (2011) and Stiglitz, Yun, and Kosenko (2020) have more recently argued that information sharing may allow to support the JHG allocation in equilibrium even if sellers cannot change their disclosure decisions in reaction to the offers of their competitors. The idea is that they can instead rely on information revealed by the buyer. In equilibrium, the firms' disclosure strategies induce the buyer to truthfully reveal her information to them, which in turn enables them to treat deviating and nondeviating firms asymmetrically.

6 Concluding Remarks: Empirical Perspectives

To conclude, we briefly examine the implications of the theoretical results surveyed in this article for empirical work. The discussion will focus on insurance markets, prominent examples of which—life-insurance, annuity, long-term-care, and, to some extent, health-insurance markets—are nonexclusive.

6.1 The Positive-Correlation Property

A standard way to test for the presence of adverse selection on insurance markets is to exploit the positive-correlation property, which states that, under adverse selection, the aggregate coverage purchased by a consumer and her riskiness should be positively correlated conditionally on observables (Chiappori and B. Salanié (2000)). This property is typically satisfied when a consumer's preferences over aggregate coverage-premia pairs (Q, T) only depend on her riskiness; indeed, it is then equivalent to the single-crossing condition, which precisely expresses the fact that riskier consumers are more eager to purchase more coverage. In our notation, this is the case if the riskiness c_i and the willingness-to-pay $\tau_i(Q, T)$ are both increasing in the consumer's type i, so that the demand Q_i for coverage is increasing in c_i.[23] Chiappori and B. Salanié (2000, 2003) have developed several econometric methods to test this prediction.

Under single-crossing, the positive-correlation property is a characteristic of consumer demand; as such, it is independent of whether competition on the market is exclusive or nonexclusive.[24] Yet the empiricist should care about the difference. Indeed, under

[23]This is the case, for instance, if consumer preferences have an expected-utility representation and face a binary loss (Rothschild and Stiglitz (1976)), as well as for many other specifications of consumer preferences and more than two loss levels (AMS (2021, Online Appendix C)).

[24]In a more general setting, Chiappori, Jullien, B. Salanié, and F. Salanié (2006) show that the positive-correlation property can alternatively be derived from a simple inequality on equilibrium profits, even if the single-crossing property is not satisfied.

exclusivity, it does not matter whether his data originate from firms or from consumers, as each consumer's aggregate coverage is provided by a single contract issued by a single firm. This explains the usual reliance on within-firm data, which are easier to obtain. However, under nonexclusivity, such data can be misleading, as low-coverage contracts may disproportionately attract high-risk consumers in combination with other contracts.[25] Thus the positive-correlation property may still hold at the consumer level, taking into account all sources of coverage; but the contracts sold by a firm may feature a negative correlation between the riskiness of its customers and the coverage it sells to them.

The validity of the positive-correlation property has been at the centre stage of empirical studies of nonexclusive insurance markets (Cawley and Philipson (1999), Finkelstein and Poterba (2004), Finkelstein and McGarry (2006)). However, because these studies typically take as a benchmark the exclusive-competition model, the above distinction between demand- and supply-side approaches is often overlooked. Rejecting adverse selection on these markets on the basis of the failure of the positive-correlation property is a decision that should, therefore, be taken with some care: in principle, we would need to collect comprehensive data at the consumer level about all sources of coverage. As pointed out by B. Salanié (2017), this is likely to be a demanding, though worthwhile task.

6.2 Exploiting Price and Cost Data

The analysis of entry-proof tariffs in Section 4 leads to a very sharp prediction for the competitive outcomes of nonexclusive insurance markets: each marginal unit of coverage available along the market tariff should be priced at the expected cost of serving the consumer types who choose to purchase it. This suggests an alternative empirical strategy exploiting price and cost data to compare the price of each layer of insurance to its average cost, as measured by the empirical loss frequency of the consumers who trade this layer.

To illustrate this approach, suppose that the loss is binary and that we have data on observationally equivalent consumers $n = 1, \ldots, N$, providing information about individual aggregate coverage-premia pairs (Q^n, T^n) and loss realizations $L^n \in \{0, 1\}$. Given this data, a natural two-step empirical procedure may run as follows.

The first step would be to construct an estimate of the market tariff T or, more precisely, of the marginal price schedule T'.[26] Although data on firms' offers are typically not available,

[25] This is actually a prediction of the regulated game G^{CSR} under free-entry (AMS (2019b, Theorem 3)): in equilibrium, basic contracts, traded by both low- and high-risk consumers, offer more coverage than complementary contracts, traded by high-risk consumers only. Thus, with data originating from a single firm, we could well observe a negative correlation between risk and coverage.

[26] This is in line with the classical problem of estimating a firm's production frontier.

we could, to this end, use the data on individual aggregate coverage and premia, assuming that each consumer strives to minimize the price she pays for her aggregate coverage. For instance, we could perform a nonparametric regression

$$T^n = T(Q^n) + \varepsilon^n,$$

with one-sided error terms ε^n capturing the idea that consumers may fail to combine the firms' offers optimally.

The second step would be to test whether the estimator $\hat{T}'$ of T' satisfies the property that each marginal quantity is priced at the expected cost of serving the consumers who purchase it. This would involve comparing, for each aggregate coverage level Q, the estimated marginal price $\hat{T}'(Q)$ with the empirical loss frequency

$$\hat{\bar{c}}\,(Q) \equiv \frac{\sum_n 1_{\{Q^n \geq Q, L^n = 1\}}}{\sum_n 1_{\{Q^n \geq Q\}}}$$

of the consumers whose aggregate coverage is at least Q.

Estimates of prices and costs play a crucial role in this procedure. This contrasts with tests of the positive-correlation property, which only rely on aggregate coverage amounts and loss realizations. The procedure is thus closer to that proposed by Einav, Finkelstein, and Cullen (2010) in a setting where consumers have a zero-one demand for coverage: evidence of adverse selection is obtained if the average cost of serving the consumers choosing to buy an additional layer of insurance is affected by the price of that layer. Our analysis suggests that the upper-tail conditional expectation function is the generalization of the firms' cost function in Einav, Finkelstein, and Cullen (2010) to richer environments where firms offer nonexclusive insurance contracts and consumers can choose different levels of coverage. An attractive feature of this approach is that it is fully nonparametric: there is no need to make assumptions about consumers' underlying utility functions nor about the distribution of their private information.

Appendix

Proof of Theorem 6. The proof consists of three steps.

Step 1 For each i in the support of the distribution of types, let $Q_i^{-k}(0,0)$ be a solution to the maximization problem that defines $z_i^{-k}(0,0)$. Then

$$z_i^{-k}(0,0) = u_i(Q_i^{-k}(0,0), T^{-k}(Q_i^{-k}(0,0))) \leq u_i(Q_i^{-k}(0,0), T(Q_i^{-k}(0,0))) \leq U_i, \tag{20}$$

where the first inequality follows from $T \leq T^{-k}$, and the second inequality follows from the fact that U_i is type i's equilibrium utility. Now, if $U_i = z_i^{-k}(0,0)$ for some type i and some seller k, all the inequalities in (20) are in fact equalities. This has two fundamental consequences. First, we have

$$u_i(Q_i^{-k}(0,0), T^{-k}(Q_i^{-k}(0,0))) = u_i(Q_i^{-k}(0,0), T(Q_i^{-k}(0,0))),$$

which implies

$$T^{-k}(Q_i^{-k}(0,0)) = T(Q_i^{-k}(0,0))$$

and, hence, by P2,

$$Q_i^{-k}(0,0) = 0. \tag{21}$$

Second, we have

$$U_i = u_i(Q_i^{-k}(0,0), T^{-k}(Q_i^{-k}(0,0)))$$

and, hence, by (21),

$$U_i = u_i(0,0). \tag{22}$$

Because $U_i \geq z_i^{-l}(0,0) \geq u_i(0,0)$ for all $l \neq k$, the upshot from this reasoning is that, if $U_i = z_i^{-k}(0,0)$ for some type i and some seller k, then, for this type i, $U_i = z_i^{-k}(0,0)$ for any seller k, and hence (21) must hold for all k. In other terms, the indispensability property P2 implies that, if, for some type, the individual-rationality constraint binds for some seller, then it must bind for all sellers.

Step 2 By P1 and Step 1, there exists some i such that $U_i = z_i^{-k}(0,0)$ for all k. Let (q_i^k, t_i^k) be the contract traded by such a type i with seller k in equilibrium, and let $Q_i^{-k} = \sum_{l \neq k} q_i^l$ be the quantity purchased by type i from the sellers other than k, so that

$$U_i = u_i(q_i^k + Q_i^{-k}, t_i^k + T^{-k}(Q_i^{-k})). \tag{23}$$

We claim that $Q_i^{-k} = 0$. Indeed, suppose, by way of contradiction, that $Q_i^{-k} > 0$. Because type i could abstain from trading with the sellers other than k, it must be that

$$u_i(q_i^k + Q_i^{-k}, t_i^k + T^{-k}(Q_i^{-k})) \geq u_i(q_i^k, t_i^k). \tag{24}$$

Similarly, consider the maximization problem that defines $z_i^{-k}(0,0)$, and whose unique solution, by (21), is $Q_i^{-k}(0,0) = 0$. As type i could instead purchase Q_i^{-k} from the sellers other than k, it must be that

$$u_i(0,0) > u_i(Q_i^{-k}, T^{-k}(Q_i^{-k})). \tag{25}$$

We know from (22)–(23) that the left-hand sides of (24)–(25) are both equal to U_i. Hence

$$u_i(0,0) \geq u_i(q_i^k, t_i^k). \tag{26}$$

Representing by $T = \phi(Q)$ the equilibrium indifference curve of type i, (25)–(26) amount to

$$T^{-k}(Q_i^{-k}) > \phi(Q_i^{-k}) \tag{27}$$

and

$$t_i^k \geq \phi(q_i^k). \tag{28}$$

Because ϕ is concave and $\phi(0) = 0$, ϕ is subadditive. Thus, by (27)–(28),

$$t_i^k + T^{-k}(Q_i^{-k}) > \phi(q_i^k + Q_i^{-k}),$$

which amounts to

$$u_i(0,0) > u_i(q_i^k + Q_i^{-k}, t_i^k + T^{-k}(Q_i^{-k})),$$

in contradiction to (22)–(23). Hence $Q_i^{-k} = 0$, as claimed. Because this is true for all k, we obtain that

$$\sum_k q_i^k = \frac{1}{K-1} \sum_k Q_i^{-k} = 0,$$

and thus that $q_i^k = 0$ for all k.

Step 3 We now verify that there exists some type i_1 such that $Q_i \equiv \sum_k q_i^k$ vanishes if and only if $i \leq i_1$, which concludes the proof. By P1 and Steps 1–2, we know that there exists at least one type i for which $Q_i = 0$. This implies in particular that

$$\text{For each } Q > 0,\ u_i(0,0) \geq u_i(Q, T(Q)),$$

and, hence, by strict single-crossing, that

$$\text{For all } j < i \text{ and } Q > 0,\ u_j(0,0) > u_j(Q, T(Q)),$$

so that $Q_j = 0$ for all $j < i$. Thus the set of types who are excluded from trade is an interval $\mathcal{I}^0$ at the bottom of the type distribution. We just need to check that $\mathcal{I}^0$, if it is not reduced to a single type, contains its least upper bound i_1. (This is obvious if the distribution of types is discrete.) By assumption, the mapping $i \mapsto u_i(0,0)$ is continuous, and so is the mapping $i \mapsto U_i$ by Berge's maximum theorem. Because $U_j = u_j(0,0)$ for all $j \in \mathcal{I}_0$, it follows that $U_{i_1} = u_{i_1}(0,0)$ as well, and hence that $U_{i_1} = z_{i_1}^{-k}(0,0)$ for all k. Proceeding as in Step 2 then shows that $Q_{i_1} = 0$, as desired. Hence the result. ■

References

[1] Akerlof, G.A. (1970): “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84(3), 488–500.

[2] Allen, F. (1985): “Repeated Principal-Agent Relationships with Lending and Borrowing,” *Economics Letters*, 17(1-2), 27–31.

[3] Arnott, R.J., and J.E. Stiglitz (1991): “Equilibrium in Competitive Insurance Markets with Moral Hazard,” NBER Working Paper No. 3588.

[4] Attar, A., and A. Chassagnon (2009): “On Moral Hazard and Nonexclusive Contracts,” *Journal of Mathematical Economics*, 45(9–10), 511–525.

[5] Attar, A., V. Farinha Luz, T. Mariotti, and F. Salanié (2021): “Mixed-Strategy Equilibria in Limit-Order Markets,” Mimeo, Toulouse School of Economics.

[6] Attar, A., T. Mariotti, and F. Salanié (2011): “Nonexclusive Competition in the Market for Lemons,” *Econometrica*, 79(6), 1869–1918.

[7] Attar, A., T. Mariotti, and F. Salanié (2014): “Nonexclusive Competition under Adverse Selection,” *Theoretical Economics*, 9(1), 1–40.

[8] Attar, A., T. Mariotti, and F. Salanié (2018): “Concurrence Non Exclusive et Sélection Adverse,” *Revue Economique*, 69(6), 1009–1023.

[9] Attar, A., T. Mariotti, and F. Salanié (2019a): “On Competitive Nonlinear Pricing,” *Theoretical Economics*, 14(1), 297–343.

[10] Attar, A., T. Mariotti, and F. Salanié (2019b): “Regulating Insurance Markets: Multiple Contracting and Adverse Selection,” Working Paper No. 19-1033, Toulouse School of Economics.

[11] Attar, A., T. Mariotti, and F. Salanié (2020): “The Social Costs of Side Trading,” *Economic Journal*, 130(630), 1608–1622.

[12] Attar, A., T. Mariotti, and F. Salanié (2021): “Entry-Proofness and Discriminatory Pricing under Adverse Selection,” Working Paper No. 17-788, Toulouse School of Economics.

[13] Ausubel, L.M. (2004): “An Efficient Ascending-Bid Auction for Multiple Objects,”

American Economic Review, 94(5), 1452–1475.

[14] Back, K., and S. Baruch (2013): "Strategic Liquidity Provision in Limit Order Markets," *Econometrica*, 81(1), 363–392.

[15] Baruch, S., and L.R. Glosten (2019): "Tail Expectation and Imperfect Competition in Limit Order Book Markets," *Journal of Economic Theory*, 183, 661–697.

[16] Beaudry, P., and M. Poitevin (1995): "Competitive Screening in Financial Markets when Borrowers Can Recontract," *Review of Economic Studies*, 62(3), 401–423.

[17] Biais, B., T. Mariotti (2005): "Strategic Liquidity Supply and Security Design," *Review of Economic Studies*, 72(3), 615–649.

[18] Biais, B., D. Martimort, and J.-C. Rochet (2000): "Competing Mechanisms in a Common Value Environment," *Econometrica*, 68(4), 799–837.

[19] Biais, B., D. Martimort, and J.-C. Rochet (2013): "Corrigendum to "Competing Mechanisms in a Common Value Environment"," *Econometrica*, 81(1), 393–406.

[20] Bisin, A., and P. Gottardi (2006): "Efficient Competitive Equilibria with Adverse Selection.," *Journal of Political Economy*, 114(3), 485–516.

[21] Bisin, A., and D. Guaitoli (2004): "Moral Hazard and Nonexclusive Contracts," *RAND Journal of Economics*, 35(2), 306–328.

[22] Budish, E., P. Cramton, and J. Shim (2015): "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130(4), 1547–1621.

[23] Carmona, G., and J. Fajardo (2009): "Existence of Equilibrium in Common Agency Games with Adverse Selection," *Games and Economic Behavior*, 66(2), 749–760.

[24] Cawley, J., and T. Philipson (1999): "An Empirical Examination of Information Barriers to Trade in Insurance," *American Economic Review*, 89(4), 827–846.

[25] Chiappori P.-A., B. Jullien, B. Salanié, and F. Salanié (2006): "Asymmetric Information in Insurance: Some Testable Implications," *RAND Journal of Economics*, 37(4), 783–798.

[26] Chiappori P.-A., and B. Salanié (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108(1), 56–78.

[27] Chiappori, P.-A., and B. Salanié (2003): "Testing Contract Theory: A Survey of Some Recent Work," in *Advances in Economics and Econometrics*: *Theory and Applications, Eighth World Congress*, *Volume 1*, ed. by M. Dewatripont, L.P. Hansen, and S.J. Turnovsky. Cambridge, UK: Cambridge University Press, 115–149.

[28] Cole, H.L., and N.R. Kocherlakota (2001): "Efficient Allocations with Hidden Income and Hidden Storage," *Review of Economic Studies*, 68(3), 523–542.

[29] Crocker, K.J., and A. Snow (1985): "The Efficiency of Competitive Equilibria in Insurance Markets with Asymmetric Information," *Journal of Public Economics*, 26(2), 207–219.

[30] Dasgupta, P., and E. Maskin (1986): "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications," *Review of Economic Studies*, 53(1), 27–41.

[31] DeMarzo, P., and D. Duffie (1999): "A Liquidity-Based Model of Security Design," *Econometrica*, 67(1), 65–99.

[32] Diamond, D.W., and P.H. Dybvig (1983): "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91(3), 401–419.

[33] Einav, L., A. Finkelstein, and M.R. Cullen (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 125(2), 877–921.

[34] Engers, M., and L. Fernandez (1987): "Market Equilibrium with Hidden Knowledge and Self-Selection," *Econometrica*, 55(2), 425–439.

[35] Farhi, E., M. Golosov, and A. Tsyvinski (2009): "A Theory of Liquidity and Regulation of Financial Intermediation," *Review of Economic Studies*, 76(3), 973–992.

[36] Farinha Luz, V. (2017): "Characterization and Uniqueness of Equilibrium in Competitive Insurance," *Theoretical Economics*, 12(3), 1349–1391.

[37] Finkelstein, A., and K. McGarry (2006): "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, 96(4), 938–958.

[38] Finkelstein, A., and J. Poterba (2004): "Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market," *Journal of Political Economy*, 112(1), 183–208.

[39] Glosten, L.R. (1989): "Insider Trading, Liquidity, and the Role of the Monopolist Specialist," *Journal of Business*, 62(2), 211–235.

[40] Glosten, L.R. (1994): "Is the Electronic Open Limit Order Book Inevitable?" *Journal of Finance*, 49(4), 1127–1161.

[41] Golosov, M., and A. Tsyvinski (2007): "Optimal Taxation with Endogenous Insurance Markets," *Quarterly Journal of Economics*, 122(2), 487–534.

[42] Hammond, P.J. (1979): "Straightforward Individual Incentive Compatibility in Large Economies," *Review of Economic Studies*, 46(2), 263–282.

[43] Hammond, P.J. (1987): "Markets as Constraints: Multilateral Incentive Compatibility in Continuum Economies," *Review of Economic Studies*, 54(3), 399–412.

[44] Harris, M., and R.M. Townsend (1981): "Resource Allocation under Asymmetric Information," *Econometrica*, 49(1), 33–64.

[45] Hellwig, M.F. (1983): "On Moral Hazard and Non-Price Equilibria in Competitive Insurance Markets," Discussion Paper No. 109, Sonderforschungsbereich 21, Universität Bonn.

[46] Hellwig, M.F. (1988): "A Note on the Specification of Interfirm Communication in Insurance Markets with Adverse Selection," *Journal of Economic Theory*, 46(1), 154–163.

[47] Hendren, N. (2013): "Private Information and Insurance Rejections," *Econometrica*, 81(5), 1713–1762.

[48] Hosseini, R. (2015): "Adverse Selection in the Annuity Market and the Role for Social Security," *Journal of Political Economy*, 123(4), 941–984.

[49] Jacklin, C.J. (1987): "Demand Deposits, Trading Restrictions, and Risk Sharing," in *Contractual Arrangements for Intertemporal Trade*, ed. by E.C. Prescott and N. Wallace. Minneapolis, MI: University of Minnesota Press, 26–47.

[50] Jaynes, G.D. (1978): "Equilibria in Monopolistically Competitive Insurance Markets," *Journal of Economic Theory*, 19(2), 394–422.

[51] Jaynes, G.D. (2011): "Equilibrium and Strategic Communication in the Adverse Selection Insurance Model," Economics Department Working Paper No. 91, Yale University.

[52] Kahn, C.M., and D. Mookherjee (1998): "Competition and Incentives with Nonexclusive Contracts," *RAND Journal of Economics*, 29(3), 443-465.

[53] Kehoe, T.J., D.K. Levine, and E.C. Prescott (2002): "Lotteries, Sunspots, and Incentive Constraints," *Journal of Economic Theory*, 107(1), 39–69.

[54] Leland, H.E., and D.H. Pyle (1977): "Informational Asymmetries, Financial Structure, and Financial Intermediation," *Journal of Finance*, 32(2), 371–387.

[55] Mailath, G., and G. Nöldeke (2008): "Does Competitive Pricing Cause Market Breakdown Under Extreme Adverse Selection?" *Journal of Economic Theory*, 140(1), 97–125.

[56] Malinvaud, E. (1972): "The Allocation of Individual Risks in Large Markets," *Journal of Economic Theory*, 4(2), 312–328.

[57] Martimort, D. (2006): "Multi-Contracting Mechanism Design," in *Advances in Economics and Econometrics*: *Theory and Applications*, *Ninth World Congress*, *Volume 1*, ed. by R. Blundell, W.K. Newey, and T. Persson. Cambridge, UK: Cambridge University Press, 57–101.

[58] Martimort, D., and L. Stole (2002): "The Revelation and Delegation Principles in Common Agency Games," *Econometrica*, 70(4), 1659–1673.

[59] Milgrom, P. (2000): "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108(2), 245–272.

[60] Milgrom, P., and C. Shannon (1994): "Monotone Comparative Statics," *Econometrica*, 62(1), 157–180.

[61] Mimra, W., and A. Wambach (2014): "New Developments in the Theory of Adverse Selection in Competitive Insurance," *Geneva Risk and Insurance Review*, 39(2), 136–152.

[62] Mimra, W., and A. Wambach (2019): "Contract Withdrawals and Equilibrium in Competitive Markets with Adverse Selection," *Economic Theory*, 67(4), 875–907.

[63] Mirrlees, J.A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38(2), 175–208.

[64] Miyazaki, H. (1977): "The Rat Race and Internal Labor Markets," *Bell Journal of Economics*, 8(2), 394–418.

[65] Myers, S.C., and N.S. Majluf (1984): "Corporate Financing and Investment Decisions when Firms Have Information that Investors Do not Have," *Journal of Financial Economics*, 13(2), 187–221.

[66] Myerson, R.B. (1983): "Mechanism Design by an Informed Principal," *Econometrica*, 51(6), 1767–1797.

[67] Netzer, N., and F. Scheuer (2014): "A Game Theoretic Foundation of Competitive Equilibria with Adverse Selection," *International Economic Review*, 55(2), 399–422.

[68] Page, F.H. Jr, and P.K. Monteiro (2003): "Three Principles of Competitive Nonlinear Pricing," *Journal of Mathematical Economics*, 39(1–2), 63–109.

[69] Pauly, M.V. (1974): "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection," *Quarterly Journal of Economics*, 88(1), 44–62.

[70] Peters, M. (2001): "Common Agency and the Revelation Principle," *Econometrica*, 69(5), 1349–1372.

[71] Philippon T., and V. Skreta (2012): "Optimal Interventions in Markets with Adverse Selection," *American Economic Review*, 102(1), 1–28.

[72] Prescott, E.C., and R.M. Townsend (1984): "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, 52(1), 21–45.

[73] Riley, J.G. (1979): "Informational Equilibrium," *Econometrica*, 47(2), 331–359.

[74] Riley, J.G. (2001): "Silver Signals: Twenty-Five Years of Screening and Signaling," *Journal of Economic Literature*, 39(2), 432–478.

[75] Rockafellar, R.T. (1970): *Convex Analysis*, Princeton, NJ: Princeton University Press.

[76] Rosenthal, R.W., and A. Weiss (1984): "Mixed-Strategy Equilibrium in a Market with Asymmetric Information," *Review of Economic Studies*, 51(2), 333–342.

[77] Rothschild, C. (2015): "Nonexclusivity, Linear Pricing, and Annuity Market Screening," *Journal of Risk and Insurance*, 82(1), 1–32.

[78] Rothschild, M., and J.E. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90(4), 629–649.

[79] Rustichini, A., P. Siconolfi, P. (2008): "General Equilibrium in Economies with Adverse Selection, *Economic Theory*, 37(1), 1–29.

[80] Salanié, B. (2017): "Equilibrium in Insurance Markets: An Empiricist's View," *Geneva Risk and Insurance Review*, 42(1), 1–14.

[81] Sheshinski, E. (2008): *The Economic Theory of Annuities*, Princeton, NJ: Princeton University Press.

[82] Spence, M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87(3), 355–374.

[83] Spence, M. (1978): "Product Differentiation and Performance in Insurance Markets," *Journal of Public Economics*, 10(3), 427–447.

[84] Stiglitz, J.E., J. Yun, and A. Kosenko (2020): "Bilateral Information Disclosure in Adverse Selection Markets with Nonexclusive Competition," NBER Working Paper No. 27041.

[85] Tirole, J. (2012): "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," *American Economic Review*, 102(1), 29–59.

[86] Wilson, C. (1977): "A Model of Insurance Markets with Incomplete Information," *Journal of Economic Theory*, 16(2), 167–207.

[87] Wilson, R.B. (1989): "Game-Theoretic Analyses of Trading Processes," in *Advances in Economic Theory*, *Fifth World Congress*, ed. by T.F. Bewley. Cambridge, UK: Cambridge University Press, 33–70.

[88] Wilson, R.B. (1993): *Nonlinear Pricing*, Oxford: Oxford University Press.

[89] Yamashita, T. (2010): "Mechanism Games With Multiple Principals and Three or More Agents," *Econometrica*, 78(2), 791–801.

Stress discounting*

Frédéric Cherbonnier†
Christian Gollier‡
Aude Pommeret§

July 25, 2022

Abstract

Standard evaluations of public policies involve discounting the flow of expected net benefits at a risk free discount rate. Consequently, they systematically ignore the insurance benefits of policies that hedge the aggregate risk, and the social cost of projects that raise the aggregate risk. Normative asset pricing theory recommends adjusting the discount rate to the project's risk, but few countries have attempted to implement this complex solution. We explore an equivalent approach based on the property that the value of a project under uncertainty equals the expected value of its state-contingent NPV, using the relevant state-contingent risk-free discount rate. Under this "stress discounting" approach, projects are evaluated under two polar risk-free economic scenarios, one business-as-usual scenario, and one low-probability catastrophic scenario, in the spirit of the now well-established banking regulation. Ramsey discounting should be performed in each scenario to estimate the corresponding scenario-contingent NPV, which is a simple and intuitive task. This approach automatically values the insurance benefits of projects whose net benefits are negatively correlated with economic growth. We extend this approach to value carbon mitigation projects, combining the two economic scenarios with two polar climatic scenario. We hope that this simpler and more intuitive method will induce more countries to better integrate the key value of risk when shaping optimal public policies, in particular those with long-lasting consequences.

Keywords: Discounting, carbon pricing, rare disasters, cost-benefit analysis, stochastic discount factor.

JEL codes: G12, H43, Q54.

*The research leading to this paper has received indirect funding from Ademe, Engie, EDF, Total, SCOR through TSE research contracts, and from the Chair "Sustainable Finance and Responsible Investments" at TSE. It has also been funded by the Agence Nationale de la Recherche under grants ANR-17-CE03-0010 - LONGTERMISM and ANR-17-EURE-0010 (Investissements d'Avenir program).

†Toulouse School of Economics, Sciences-Po Toulouse.

‡Toulouse School of Economics, University of Toulouse-Capitole.

§IREGE, University Savoie Mont Blanc.

Perfection is achieved, not when there is nothing more to add, but when there is nothing left to take away. Antoine de Saint-Exupéry

1 Introduction

Under the standard practice, a project is deemed socially desirable if the discounted value of its flow of expected net social benefits is positive, using a unique public discount rate. For example, the official discount rate is 3.5% in the U.K. (Treasury, 2020). But this procedure ignores the social cost of projects that increase the macroeconomic risk, and the social benefit of projects that reduce it. In other words, it ignores risk and risk aversion. Compare for example an investment in a railways infrastructure (which arguably is most useful in a growing economy) to another investment in a mass vaccination infrastructure (which arguably is most useful during a pandemic that kills the economy). The standard public valuation practice values both projects equally if their expected net benefits are the same. However, risk aversion necessarily implies that a marginal reallocation of capital from railways to mass vaccination infrastructures would reduce the aggregate risk at no cost, thereby increasing welfare.[1]

Modern asset pricing theory such as the consumption-based CAPM (CCAPM), interpreted in a normative way, strongly recommends solving this problem by adjusting the discount rate to the risk profile of each project, and potentially of each of its specific benefit (Bodie and Merton, 2000; Brealey et al., 2017). Under the CCAPM and its extensions, the sufficient statistic of the risk profile is the CCAPM, which is the income-elasticity of the benefit under consideration. In the case of a railway infrastructure for example, this procedure would require to estimate the beta of the demand for transportation, but also the social cost of carbon and of air pollution, or the value of time lost and lives saved. But this perfect solution has failed. The complexity of this procedure is likely to explain its low level of adoption in the public sector around the world.

In this paper, we propose a simpler and more intuitive solution based on the fundamental asset pricing principle (Lucas, 1978):[2] Any state-dependent benefit $B_t = B_{ts} \mid_{s \in S}$ materializing at date t has a Present Value at date 0 equaling

$$PV = E[e^{-r_t t} B_t], \tag{1}$$

where $r_t = r_{ts} \mid_{s \in S}$ is the stochastic discount rate associated to maturity t and E is the expectation operator with respect to the states of nature $s \in S$. The state-dependent discount rate r_{ts} is the rate at which benefits in t should be discounted conditional to state s, i.e., under certainty. The Ramsey rule can be used to determine it. Project analysts just need to perform a sequence of NPV estimations under certainty. This approach is much simpler than the CCAPM approach because it just requires to use the Ramsey discount rate contingent to each state considered, thereby eliminating the complexity of risk-adjusting the discount rate together with the task to estimate the betas. This alternative approach is also more intuitive because the Ramsey rule is by now well understood. Moreover, measuring the value creation of an asset through the expectation of its contingent NPV should also be easily understood

[1] This is only partly accounted for in the U.K. where the public discount rate for health projects is reduced to 1.5% (Treasury, 2020).

[2] This principle is a direct consequence of the double additivity of the Discounted Expected Utility model, with respect to the states and to time. It does not hold under recursive preferences.

by practitioners. Finally, this approach does not give up anything to the basic principles of asset pricing. In fact, the CCAPM is a special case of the Stochastic Discount Factor (SDF) approach when assuming a geometric Brownian motion to consumption growth. More generally, the fact that the contingent discount rate is smaller in bad states of nature, i.e., in low-consumption states, provides the right instrument to deliver a premium to assets that yield more benefits in those states. Finally, our approach provides information about the circumstances in which the project is most valuable, as it measures the conditional PV of the project in each state of nature. This may be useful for the transparency of the decision-making and the public debate.

The complexity of this alternative approach depends upon the number of states of nature (or macroeconomic contingencies) that evaluators will have to consider. To choose the number of scenarios to consider and their characteristics, let us start by recalling that the CCAPM fails to predict interest rates and risk premia in financial markets (Mehra and Prescott, 1985; Weil, 1989). In short, this means that the gaussian volatility of the growth rate of consumption is too small to explain why the interest rate has been so low in the XXth century, and why the aggregate risk premium has been so larger. Rietz (1988) and Barro (2006, 2009) and Weitzman (2007) showed that this failure of the CCAPM can be solved by recognizing that the lower tail of the distribution of economic growth is fatter than assumed in the standard CCAPM. In Barro's work, the equilibrium risk-free rates and risk premia are mainly determined by the probability of a macroeconomic catastrophe, and the gaussian volatility plays only a marginal role on these matters. In this paper, we push this idea to this logical end by proposing to perform the expected NPV valuation with only two states. In the Business-As-Usual (BAU) scenario, consumption grows at a constant positive rate that reflects the likely economic prosperity. But in the low-probability catastrophic scenario, consumption drops immediately by an amount similar to what is suggested by Barro, and then gradually recovers at a growth rate smaller than in the BAU scenario. Calculating the contingent discount rates from the Ramsey rule in this two scenarios is then trivial. The aggregate risk in this model has thus 4 easy-to-understand parameters: the probability of catastrophe, the drop in consumption in that state, and the constant growth rates in the two states. The calibration of these parameters should closely fit the observed asset prices in the economy.

Other sources of risk can easily be added into this framework. To illustrate, in the context of climate change, the debate on the Social Cost of Carbon has explored the role of uncertainty mostly from the point of view of an uncertain climate sensitivity (Dietz et al., 2018; Daniel et al., 2019). Conditional to each macroeconomic scenario, a specific emission path can be considered together with an uncertain climate sensitivity to calibrate climate damages. We show that it is appropriate in that case to consider 2 types of shocks (economic and climatic), and therefore four distinct scenarios, to correctly evaluate cliate mitigation projects. Our objective here is to illustrate the method and to prove its efficiency. We check that the method gives satisfactory results by benchmarking it with more elaborate approaches. This is what we do in section 3 and 4 of this document, before applying these methods to real-based cases.

Thus, in this article, we start from the failure of national governments and key public institutions (World Bank, European Investment Bank,...) to adopt an efficient approach to discounting. There are various reasons why economists failed to improve the public discounting systems in the western world. The obvious argument is that using a single discount rate makes life much easier for the evaluators who are well accustomed with this practice, as estimating CCAPM betas may be technically difficult (Cherbonnier and Gollier (2020)). The second argument is that there is much at stake for many lobbies, in particular in those sectors with large betas. They are the losers of risk-adjusting the standard discounting system. There is a clear agency problem associated to the existence of various asymmetric information problems in this context. More flexibility in choosing project-specific discount rates may favor strong lobbies equipped with a good understanding of how to estimate CCAPM betas. Third, many public economists working in the sphere of public policy evaluations are not experts in asset pricing theory. At the same time, most asset pricing experts in academia continue to ignore public sector finance. Finally, adjusting discount rates to risk is a common practice in financial markets, but is often considered as highly inefficient by public servants in charge of implementing cost-benefit analyses in the public sector. They may be right, in the sense that the risk-adjustments used by private institutions and their stakeholders are inefficient. But this does not help. It is important to disentangle the normative nature of modern asset pricing theory, which provides a strong argument for the risk adjustment, from how market participants do apply it in practice.

Fifty years of developments in asset pricing theory has clearly demonstrated the key role of covariance with aggregate consumption when measuring the welfare impact of an asset. These developments justifying adjusting the discount rate to the risk profile of the projects had a very limited impact on the public practice of discounting. We are aware of only three countries which have attempted to adapt their evaluation practices in that direction. In 2020, the Dutch government (Rijksoverheid (2020)) has implemented a discounting system that contains three risk-adjusted discount rates. The Norwegian government had also adopted in 1997 a discounting system with three discount rates allocated to three risk classes defined by the projects' contribution to the aggregate uncertainty, but this system has been abandoned in 2012 as choosing the suitable risk class for a project was deemed too arbitrary[3] Since then, Norway uses a single discount rate of 4%. Finally, France has introduced a CCAPM rule since 2013 with a risk-free discount rate of r_f=2.5% and an aggregate risk premium of Π=2% (Quinet (2013)). Evaluators must estimate the beta of their project to determine the rate $r_f + \beta\Pi$ at which the expected net benefit must be discounted. The experience shows that very few evaluators have tried to estimate the beta of their projects, even for those costing tens of billion euros in public funds. They rather used a default beta of one, yielding an implicit single discount rate of 4.5%. In Gollier (2021), the welfare cost of ignoring the risk-adjustment in the discount rate has been estimated to be large, equivalent to 15% of permanent consumption at least. It is time to propose an operational method of discounting that has two properties: (1) approximate the efficient solution in a robust way, and (2) have a better chance to be adopted by the public sector, i.e., be simple, transparent and intuitive.

The paper is organized as follows. We recall in section 2 the principles of the SDF approach for public investment valuation. The basic stress discounting method with two scenarios is

[3]See Hagen et al. (2012) page 77: "*Experience from previous practice with several risk classes suggests that many project analysts have been uncertain about the technical criteria for choosing the risk class, and that such choices may therefore at times seem somewhat arbitrary. These circumstances suggest that it may be preferable to recommend simple and transparent rules that capture the most important aspects of the matter, without being too complex to understand or to apply.*"

explained, calibrated and illustrated in section 3. This latter section also provides some benchmarking exercises by comparing the results and implications of this method with those of the cutting-edge analytical methods (that depart from the classical Gaussian CCAPM in order to explain the asset pricing puzzles). In section 4, we expand the approach to allow for four scenarios, which is especially relevant for projects with a climate dimension. As an illustration, we estimate the discounted social cost of carbon, i.e. the discounted expected value of avoided damage when emitting one ton of CO2 less at a given horizon. Benchmarking exercises are then provided using a numerical approach based on the DICE model. Section 5 presents an application on a French nuclear waste project.

2 Two equivalent valuation methods

In this section, we first summarize the SDF approach to asset pricing. We then recall the methodology of the standard approach to value public investment and policies, stressing its operational complexities.

2.1 Expected and contingent present values

There is a representative agent in the economy whose discrete flow of consumption is given by the stochastic process $(C_0, C_1, ..., C_t, ...)$. This agent extracts utility $u(C_t)$ from consuming C_t at date t. Social welfare at date 0 is measured by the discounted sum of temporal expected utility, using a rate of pure preference for the present δ. Let's consider an investment project that generates a flow of net benefits $(B_0, B_1, ..., B_t, ...)$ that are potentially correlated to consumption.[4]

The Present Value (PV) of the project is defined as the sure monetary benefit received today that has the same impact on social welfare as a marginal investment in that project. In other words, PV equals

$$PV = \frac{1}{u'(C_0)} \left.\frac{\partial}{\partial \varepsilon}\right|_{\varepsilon=0} \sum_{t=0}^{+\infty} e^{-\delta t} Eu(C_t + \varepsilon B_t). \tag{2}$$

By definition, it is socially desirable to invest in the project if and only if PV is positive. This condition can be rewritten as follows:

$$PV = E\left[\sum_{t=0}^{+\infty} b_t(C_t) e^{-r_t(C_t)t}\right], \tag{3}$$

where $b_t(C_t)$ is the expected benefit at date t conditional to C_t, and where the state-contingent discount rate $r_t(C_t)$ is defined as

$$r_t(C_t) = \delta - \frac{1}{t} \log\left(\frac{u'(C_t)}{u'(C_0)}\right). \tag{4}$$

This means that the value creation of an investment project is the expectation of the contingent present values $\sum b_t \exp(-r_t t)$, using a stochastic discount factor $\exp(-r_t t)$. We hereafter

[4] As is well-known, risks that are not correlated to aggregate consumption should not be priced. Therefore, B_t should be interpreted as the expected net benefit at date t conditional to C_t.

assume that the utility function u exhibits constant relative risk aversion γ. Let g_t denote the annualized growth rate of consumption:

$$g_t(C_t) = \frac{1}{t} \log \left(\frac{C_t}{C_0} \right) \tag{5}$$

Combining this definition with equation (4) yields the Ramsey rule:

$$r_t = \delta + \gamma g_t(C_t). \tag{6}$$

The pair of equations (3) and (6) fully describes an evaluation procedure in which the project analyst must perform three different tasks:

1. Characterize the flow of net expected benefits $(b_0, b_1, ...)$ conditional to each growth scenario;

2. Compute the contingent PV of this flow in each scenario, using the associated discount rate;

3. Compute the expectation of the contingent PVs to obtain the PV of the project.

Each of these tasks is intuitive and simple. The complexity of the procedure may emerge however if the number of scenarios to consider is large.

A special case is worthy examining in more details at this stage. Consider a risk-free project, or a project whose net benefits are independent of economic growth. In that case, equation (3) simplifies to

$$PV = \sum_{t=0}^{+\infty} B_t e^{-r_{ft} t}, \tag{7}$$

where the risk-free discount rate r_{ft} is defined as follows:

$$r_{ft} = -\frac{1}{t} \log \left(E e^{-r_t t} \right). \tag{8}$$

From this benchmark, one can observe that when projects are risky, their PV will be larger or smaller than the risk-neutral PV depending upon whether their net benefits are negatively or positively statistically linked to economic growth. More precisely, equation (3) directly implies that

$$\begin{aligned} PV &= \sum_{t=0}^{+\infty} E\left[B_t e^{-r_t t}\right] = \sum_{t=0}^{+\infty} \left(E\left[B_t\right] E\left[e^{-r_t t}\right] + cov(B_t, e^{-r_t t}) \right) \\ &\geq \sum_{t=0}^{+\infty} e^{-r_{ft} t} E\left[B_t\right] \end{aligned} \tag{9}$$

whenever the net benefit of the project and the stochastic discount factor covary positively. From the Ramsey rule (6), this is the case when the net benefit and consumption growth vary in opposite direction, i.e., are anti-comonotone. Inequality (9) states that the project has a negative risk premium in that case, i.e., its value creation is larger than if one would assume independence between its net benefit and aggregate consumption. The opposite result holds when they are comonotone.

2.2 Lessons from the standard approach

The tradition in the asset pricing literature is to value an asset as the discounted sum of its flow of expected benefits using a risk-adjusted discount rate:

$$PV = \sum_{t=0}^{+\infty} E\left[B_t\right] e^{-\rho_t t} \tag{10}$$

This approach is compatible with the SDF approach described in the previous section if and only the risk-adjusted discount rate ρ_t is defined as

$$\rho_t = -\frac{1}{t} \log\left(E\left[\frac{B_t}{E[B_t]} e^{-r_t t}\right]\right) \tag{11}$$

Further simplifications can be obtained by making two additional assumptions. First suppose that the net benefit of the project is linked to aggregate consumption through the following functional form:

$$\log\left(\frac{B_t}{B_0}\right) = a_t + \beta_t \log\left(\frac{C_t}{C_0}\right) + \epsilon_t, \tag{12}$$

where we assume exogeneity, so that $E\left[\epsilon_t \mid C_t\right] = 0$ for all C_t. Observe that we can interpret the project-specific β_t as the income-elasticity of its net benefit (at date t). Second, suppose that aggregate consumption follows a discrete version of a geometric brownian motion so that $\log(C_t/C_0)$ is normally distributed with mean μt and variance $\sigma^2 t$. In that case, it is well-known that the risk-adjusted discount rate equals

$$\rho_t = r_f + \beta_t \Pi, \tag{13}$$

with risk-free discount rate $r_f = \delta + \gamma\mu - 0.5\gamma^2\sigma^2$ and aggregate risk premium $\Pi = \gamma\sigma^2$. This is the standard CCAPM approach to discounting. Under that approach, the project analyst has an a priori simple task to perform. The analysis requires estimating the income-elasticity β_t of the project to determine the risk-adjusted rate to discount the flow of expected benefits.

In spite of its apparent simplicity, this standard approach faces serious operational difficulties. The most immediate one arises from the estimation of the beta. The flow of benefits may include several distinct components (for example a financial one and some externalities) each of them exhibiting an income-elasticity that requires estimation. The beta of the project is then the sum of the of the income-elasticity of these components. In addition, the income-elasticities need to be evaluated for each project. A difficulty that arises is due to the asset pricing puzzles that the CCAPM generates. With a growth process calibrated with growth rate $\mu = 2\%$ and volatility $\sigma = 3\%$ and with a CRRA $\gamma = 2$, we obtain a risk-free discount rate net of the rate of impatience of 3.82% and an aggregate risk premium 0.18%. In particular this risk premium is so low that it makes sense in practice to discount all project at the mean discount rate of 4%. Financial markets reveal much smaller interest rate on average (risk-free rate puzzle, Weil (1989), and much larger risk premia (equity premium puzzle, Mehra and Prescott (1985)). Over the last two decades, the asset pricing literature has solved these puzzles by considering growth stochastic processes that fatten the tails of the distribution of future consumption.[5] Bansal and Yaron (2004) has pioneered a new literature on "long run risks" in which trend of growth reverses to the mean and in which growth volatility is itself stochastic.[6] Rietz (1988) and Barro (2006, 2009) showed that the inclusion of rare disasters in the growth process can also solve these puzzles.

These strategies to solve the puzzles of the standard CCAPM make it much more complex to operationalize for project analysts. This is because the CCAPM formula (13) needs to be revised when exiting the Gaussian world. Indeed, as shown by Martin (2013), the efficient risk-adjusted discount rate becomes a polynomial function of the beta of the project, where the coefficient associated to the nth power of β is proportional to the $n+1$ cumulant of the annual change in log consumption. This raises some complexity to the analysis. Another source of complexity comes from the observation that the coherence of the calculation requires the project analyst to estimate the flow of expected benefits EB_t by using the complex growth stochastic processes that have been used to estimate the risk-free rate and the aggregate risk premium. In practice, this is infeasible without allowing some shortcuts that have not been provided in the literature up to now.

3 Basic project evaluation: stress discounting with two states of the economy

In the face of these complexities, should we go back to using a single discount rate? We believe not, because the cost of ignoring the social cost of risk in the economy is potentially large (Gollier (2021)). In this section, we develop an evaluation procedure based on the SDF methodology presented in Section 2.1. This alternative procedure aims at two objectives. First, we want it to be simple, intuitive and easy to operationalize. Second, we want it to generate valuations that approximate well the true value of assets and investment projects. Because these objectives goes in opposite directions, one should leave the precise procedure to implement to national circumstances, based on the willingness of project analysts and their principals to use a more complex procedure in exchange for more accurate results. The calibration of the procedure that we use in this paper should be interpreted as an illustration.

Inspired by Barro (2006, 2009), we propose a procedure based on the SDF approach using only two states or scenarios, a BAU scenario and a stressed scenario.

3.1 Projects maturing in one year

In this section, we limit our analysis to the evaluation of short projects, i.e., projects maturing within one period (here, one year). For a one-period horizon, our two-state uncertainty is characterized by three parameters: the growth rate g^b in the Business-As-Usual (BAU), the growth rate $g^s < 0$ in the catastrophic state, and the probability π of the catastrophe. We want to match three moments:

[5] Another path to solve the financial puzzle is to generalize the discounted expected utility framework into its Epstein-Zin-Weil extension. This approach raises two issues. First, Epstein et al. (2014) have shown that the calibration of EZW preferences necessary to solve the asset pricing puzzles generates a new puzzle, which is related to the implausibly large value of an early resolution of uncertainty. Second, under the veil of ignorance, risk aversion and the aversion to consumption fluctuations should be equivalent from a normative viewpoint.

[6] See Gollier (2018) for a discussion on the link between stochastic volatility and the fourth moment of the distribution of consumption.

δ	0	rate of pure preference for the present
γ	4	degree of relative risk aversion
r_f	1.0%	risk-free rate
Π	2.0%	aggregate risk premium
g^b	2.0%	growth rate of consumption in the BAU scenario
g^s	-33.3%	growth rate of consumption in the stress scenario
π	2.33%	probability of the stress scenario
G	1.3%	growth rate of expected consumption
r^b	8.0%	contingent discount rate in the BAU scenario
r^s	-133.4%	contingent discount rate in the stress scenario

Table 1: Benchmark calibration of the evaluation model.

- The growth rate g^b in the BAU. We assume a growth rate of $g^b = 2\%$ per year on the basis of the trend of growth in the western world over the last century.
- The risk-free rate r_f. We assume a risk-free rate of $r_f = 1\%$ corresponding to the average real rate of return of Bills in the western world over period 1880-2005 (Barro, 2009).
- The expected rate of return of a claim on aggregate consumption ($\beta = 1$). (Barro, 2009) documents an average real return on equity of 7.5%, whereas Bansal and Yaron (2004) assume a CCAPM beta of equity equaling $\beta = 3$. This is compatible with an aggregate risk premium Π around 2.17%. We calibrate our model to produce an aggregate risk premium of $\Pi = 2\%$. This means that the risk-adjusted discount rate for a share on aggregate consumption is 3% under our calibration.

We follow Barro (2009) by assuming a degree of relative risk aversion equaling $\gamma = 4$, which is an upper bound of what we recognize as a standard attitude toward risk. Finally, for moral reasons, we assume no discrimination across generations in the welfare function ($\delta = 0$). There is no ethical argument to penalize individuals on the basis of the generation to which they belong. It is in line with what was retained in Stern and Stern (2007) - a rate very close to zero, equal to 0.1% -, but the French and British authorities have for their part retained a rate equal to 1% and 1.5% respectively - cf. Quinet et al. (2013) and Treasury (2020).

Using the framework presented in the previous section, it is easy to verify that matching the above-mentioned moments univocally determines the remaining two parameters of our two-state model: The probability of catastrophe must be equal to $\pi = 2.33\%$, and the rate of growth in the catastrophic state must be equal to $g^s = -33.3\%$. Observe that the values of these two stress parameters are similar to those obtained by Barro (2009), as he estimated from international data a probability of catastrophe of 1.7% and an expected drop in GDP in that case of 29%. The efficient discounting system and its underlying economic context that supports it is summarized in Table 1. The expected consumption grow at rate 1.3%.

We now show that adding some gaussian noise to the growth process does not add much to the resolution of the asset pricing puzzle. Compared to our calibration for the 1-year

	risk free rate	aggregate risk premium
certainty	5.2%	0.0%
benchmark	1.0%	2.0%
benchmark with gaussian noise	0.7%	2.1%

Table 2: Effect of risk on asset prices. The "certainty case" corresponds to an economy growing at the sure rate $G = 1.3\%$. The benchmark case corresponds to the two-state growth model described in Table 1. In the "benchmark with gaussian noise", I add to the benchmark a gaussian noise in the BAU scenario of the benchmark model, with a standard deviation of 2%.

maturity, Barro (2009) has the additional ingredient of a gaussian noise around $g^b = 2\%$ in the BAU scenario. In his paper, the gaussian noise has a standard deviation of 2%. In Table 2, I compare the risk free rate and the aggregate risk premium under three economic growth process. The reference is an economy with no uncertainty at all, growing at the sure rate $G = 1.3\%$. In such an economy, the risk free rate equals 5.2%. We see that the two-state uncertainty examined in our benchmark dramatically reduces this interest rate to 1%. The addition of the gaussian noise with a volatility of 2% in the BAU scenario generates an additional reduction of the risk free rate to 0.7%. Concerning the aggregate risk premium, the two-state risk increases it to 2.0%, from 0 in the case of certainty. The addition of the gaussian noise has the marginal effect to increase it to 2.1%. This illustrates the fact that our two-state benchmark risk captures most of the asset pricing impact of the uncertainty of the Barro's model. In other words, removing the gaussian noise from the model has a marginal impact in the evaluation of projects.

Going back to our benchmark model, we know that there are two approach to the evaluation of projects. We support the simple SDF approach. Consider a 1-year project generating an expected benefit b_1^b conditional to the BAU scenario, and an expected benefit b_1^s conditional to the stress scenario. Using the exact Ramsey rule, we obtain the state-dependent discount rates $r^b = \gamma gs = 8.0\%$ and $r^s = \gamma\theta = -133.4\%$, Using the associated SDF e^{-r}, this yields the simple pricing formula

$$PV = (1-\pi)(0.92b_1^b) + \pi(3.80b_1^s) = 0.902b_1^b + 0.088b_1^s. \tag{14}$$

The strong asymmetry between the two state-dependent discount factors illustrates the valuation bonus for projects able to generate benefits in the stress scenario.

Let us compare this very simple valuation formula to what would be required in the CCAPM approach consisting in discounting the expected benefit at a risk-adjusted rate. The evaluator should first estimate the beta of the project. An estimation of this income-elasticity of the benefit is given by the following equation:

$$\beta = \frac{Log(b_1^b) - Log(b_1^s)}{g^b - g^s}. \tag{15}$$

This allows the evaluator to estimate in turn the risk-adjusted discount rate $(1 + 2\beta)\%$. Finally, the evaluator must discount the expected benefit $(1-\pi)b_1^b + \pi b_1^s$ at this rate. This is a long and obscure detour to produce a value for projects.

SDF Approach	
Present value	PV=$0.902 \times 2 + 0.088 \times 4 = 2.157$
Discount rate (eq. (11))	$\rho_1 = -\log(PV/Eb) = -5.25\%$
CCAPM approach	
beta	$\beta = (\log(b_1^b) - \log(b_1^s))/(g^b - g^s) = -1.96$
Discount rate (eq. (13))	$\rho_1 = (1 + 2\beta)\% = -2.92\%$
Present value	$=e^{-R}(0.9767 \times 2 + 0.0233 \times 4) = 2.107$

Table 3: Two approaches to the estimation of the value of a one-year project with state-dependent benefits ($b_1^b = 2, b_1^s = 4$) in the benchmark calibration of the two-state model of Table 1.

Let me illustrate the SDF and CCAPM approaches with a simple example under our benchmark calibration described in Table 1. We assume $b_1^b = 2$ and $b_1^s = 4$, which describes a negative-beta project. The expected benefit of the project equals 2.047. Following the SDF valuation rule (14), we immediately obtain a project value equaling 2.15. This is compatible with a risk-adjusted discount rate of -5.25%. Under the CCAPM approach, we first use equation (15) to obtain an estimation of the beta of the project. This gives $\beta = -1.96$, and thus a risk-adjusted discount rate of -2.92%. This yields a present value of 2.11. Beyond the unnecessary complexity of the CCAPM approach, notice the large discrepancy between the two estimations in terms of the (implicit) discount rate. this CCAPM approach does not generate an exact solution. Indeed, remember that the linear CCAPM discounting rule $r_f + \beta\Pi$ is exact only when the growth rate of consumption is gaussian, which is not the case as soon as one recognizes the existence of fat tails or catastrophic events.

We have already explained that outside the gaussian world, the efficient risk-adjusted discount rate is a polynomial function of the beta of the asset. Martin (2013) gives the following exact rule:

$$\rho_1 = \delta + \sum_{n=1}^{+\infty} \frac{\kappa_n}{n!} \left(\beta^n - (\beta - \gamma)^n \right), \tag{16}$$

where κ_n is the n-th cumulant of log consumption growth in the first period. The first cumulants are familiar: κ_1 is the mean, $\kappa_2 = \sigma^2$ the variance, $\kappa_3/Sigma^3$ the skewness, and κ_4/σ^4 is the excess kurtosis of $\log(C_1/C_0)$. In the gaussian case, only the first term in the summation operator of equation (16) is non-zero, which makes it equivalent to equation (13). In all other cases, higher-order cumulants are non-trivial and the risk-adjusted discount rate becomes a non-linear function of β. In our two-state benchmark, the log consumption growth is heavily negatively skewed, whereas $\beta^3 - (\beta - \gamma)^3$ is convex in β, so that this term introduces a concavity in the $(_1, \beta)$ relationship. In Figure 1, we represented the exact risk-adjusted discount rate and its CCAPM linear approximation. The special case described in Table 3 is illustrated in the left side of this figure.

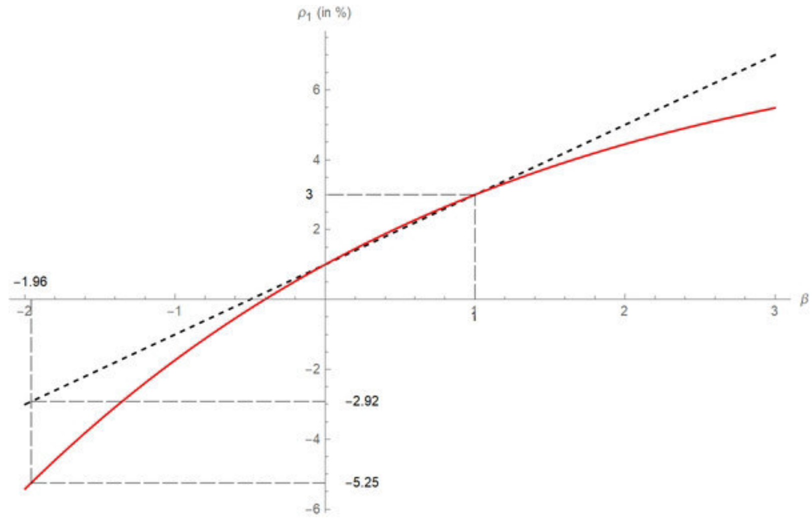

Figure 1: Risk-adjusted discount rate for a project with net benefit $b_1 = C_1^\beta$ as a function of β in the benchmark two-state model calibrated in Table 1. The red curve is the exact rate from equation (11), whereas the dashed line corresponds to its CCAPM approximation from equation (13).

3.2 Multiple-year projects

In this section, we generalize our model to time horizons larger than one year. In Barro's model, the growth process is i.i.d. over time, so that a catastrophe can occur every year. One advantage of this assumption is that the term structure of discount rates is flat, i.e., ρ_t is independent of t. The disadvantage is that evaluators face a myriad of scenarios to examine when considering long-lived projects. To keep our model as simple as possible, we continue to assume that there are only two possible scenarios, with a stress scenario occurring with probability $\pi = 2.33\%$, and a BAU scenario occurring with probability $1-\pi$. $\{C_t^s\} \mid_{t\in\mathbb{N}}$ and $\{C_t^b\} \mid_{t\in\mathbb{N}}$ describe the deterministic growth process of consumption in the stress and BAU scenarios respectively. The state-dependent annualized growth rate of consumption is denoted $g_t^i = (1/t)\log(C_t^i/C_0)$, $i \in \{s,b\}$. In this benchmark, we propose to select these two state-dependent growth processes in such a way that the term structures of risk-free rates and aggregate risk premia be flat, respectively at $r_f = 1\%$ and $\Pi = 2\%$ as in the one-year case. These two sets of conditions

$$r_f = \delta - \frac{1}{t}\log\left(\pi\exp(-\gamma g_t^s t) + (1-\pi)\exp(-\gamma g_t^b t)\right) \tag{17}$$

$$r_f + \Pi = \delta - \frac{1}{t}\log\left(\pi\exp((1-\gamma)g_t^s t) + (1-\pi)\exp((1-\gamma)g_t^b t)\right) + \frac{1}{t}\log\left(\pi\exp(g_t^s t) + (1-\pi)\exp(g_t^b t)\right), \tag{18}$$

for all $t \in \mathbb{N}/0$ univocally determine our two-state growth process characterized by $\{(g_t^s, g_t^b)\} \mid_{t\in\mathbb{N}/0}$. Of course, for $t = 1$, the solution is as described in the previous section, with $g_1^s = -33.3\%$ and $g_1^b = 2.0\%$. Consumption levels and annualized gowth rates in the two scenarios are described in Table 4 and in Figure 2. In the stress scenario, the annualized growth rate increases

t	g_t^s	g_t^b	C_t^s	C_t^b	r_t^s	r_t^b
1	-33.40	2.00	0.72	1.02	-133.00	-133.00
11	-6.07	1.27	0.51	1.15	-24.30	-24.30
21	-3.53	1.20	0.48	1.29	-14.10	-14.10
31	-2.49	1.18	0.46	1.44	-9.94	-9.94
41	-1.90	1.19	0.46	1.63	-7.59	-7.59
51	-1.52	1.20	0.46	1.85	-6.08	-6.08
61	-1.25	1.23	0.47	2.12	-5.01	-5.01
71	-1.05	1.26	0.47	2.45	-4.22	-4.22
81	-0.90	1.31	0.48	2.88	-3.60	-3.60
91	-0.78	1.35	0.49	3.43	-3.11	-3.11
101	-0.68	1.41	0.50	4.14	-2.71	-2.71
111	-0.60	1.46	0.52	5.05	-2.38	-2.38
121	-0.53	1.51	0.53	6.22	-2.11	-2.11
131	-0.47	1.56	0.54	7.72	-1.87	-1.87
141	-0.42	1.60	0.56	9.61	-1.67	-1.67
151	-0.37	1.65	0.57	12.00	-1.49	-1.49
161	-0.33	1.68	0.58	15.00	-1.34	-1.34
171	-0.30	1.71	0.60	18.80	-1.20	-1.20
181	-0.27	1.74	0.61	23.50	-1.08	-1.08
191	-0.24	1.77	0.63	29.40	-0.97	-0.97

Table 4: Growth process and state-dependent discount rates in our two-state benchmark model. The state-dependent annualized growth rates g_t^i and discount rates r_t^i are expressed in % per year. We assume $\gamma = 4$ and $\delta = 0$. This yields a constant risk-free interest rate of $r_f = 1\%$ and a constant aggregate risk premium of $\Pi = 2\%$, for all maturities.

over time. Consumption is at a minimum 30-50 years after the crash, a period during wich consumption is 54% smaller than today. Two centuries after the crash, consumption remains 37% smaller than today. In the BAU scenario, consumption has an exponential trajectory, but the annualized growth rate of consumption is hump-shaped, with a minimum annualized growth rate at 1.2% around a time horizon of 20-50 years.

The corresponding efficient state-contingent discount rates are immediately derived from the Ramsey rule:

$$r_t^i = \delta + \gamma g_t^i, \tag{19}$$

for all t and for $i \in \{s, b\}$. Table 4 and Figure 2 describe the term structures of these two state-dependent discount rates. The stress-specific discount rates has an increasing term structure, but it remains negative for all maturities under consideration (200 years). The term structure of the BAU discount rates is hump-shaped in parallel to the BAU annualized growth rates.

The discounting system described in Table 4 provides an easy workplace for the evaluator, who must estimate the flow of expected benefits $\{b_t^s, b_t^b\} \mid_{t\in\mathbb{N}/0}$ of the project under scrutiny

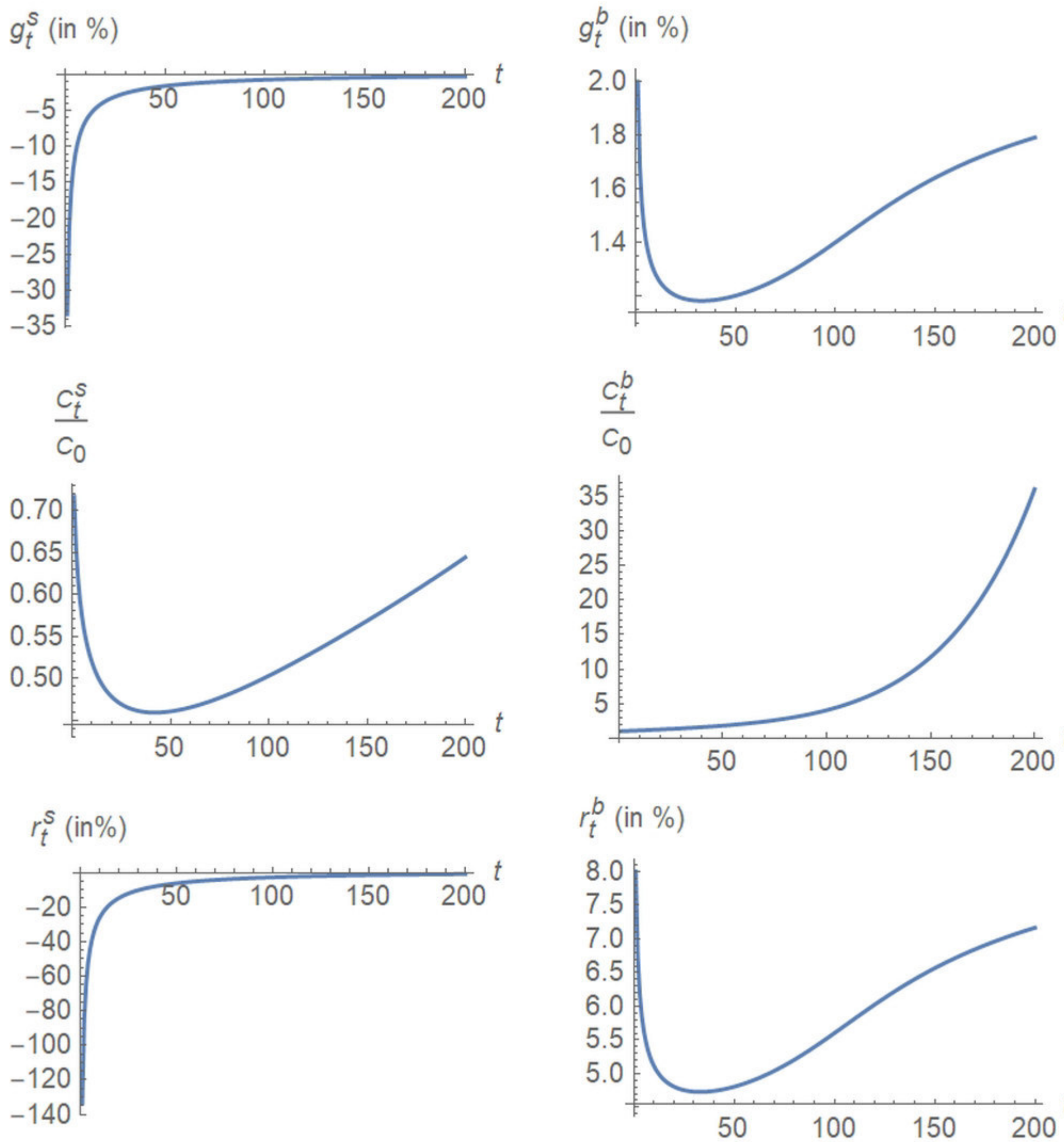

Figure 2: State-contingent annualized growth rates (line 1), consumption levels (line 2) and discount rates (line 3) in our benchmark model. This is a graphical representation of Table 4.

under the two scenarios. This allows the evaluator to compute the contingent PVs

$$PV^i = \sum_{t=0} b_t^i \exp(-r_t^i t), \tag{20}$$

and eventually the value of the project which is the expected PV $\pi PV^s + (1-\pi)PV^b$.

4 When climate matters: Stress discounting with four scenarios

In the context of climate change, future aggregate consumption will mostly depend upon two sources of uncertainty: the growth of Total Factor Productivity (TFP), and the climate sensitivity. A large climate sensitivity will raise climate damages, and will therefore adversely affect aggregate consumption. As discussed by Dietz et al. (2018), Cai and Lontzek (2019) and Lemoine (2021), these two sources of risk imply opposite correlations between consumption and the benefit of mitigation. If TFP uncertainty predominates, consumption and mitigation benefits will be positively correlated, as a larger growth yields more emissions and a larger marginal benefit of mitigation, assuming a convex damage function. If the uncertainty affecting the climate sensitivity dominates, consumption and mitigation benefits will be negatively correlated. Indeed, a larger climate sensitivity implies at the same time a larger mitigation benefit and a smaller future aggregate consumption. This means that the risk-adjusted discount rate to value mitigation efforts, i.e., to estimate the Social Cost of Carbon (SCC), is highly sensitive to the way these two sources of risk. According to Dietz et al. (2018), the so-called "climate beta", is smaller tan 1, and has a decreasing term structure. Indeed, the risk of climate damage is increasing over time while negatively correlated with economic growth. Consequently, the beta is decreasing over time, and could take low or even negative values on the long term. This means that climate mitigation projects might exhibit an insurance value over the long run.

In this section, we propose to re-examine these issue by considering two risks with two possible outcomes each, i.e., a four-state structure of risk.

4.1 Methodology

Applying the stress discounting procedure to value an investment project requires more than two scenarios to take into account both economic and climatic uncertainties. We follow Dietz et al. (2018) who analyzes 10 sources of uncertainty and shows that the two prevailing ones are shocks on TFP and climate sensitivity. Consequently, we consider two binary risks, that of an economic disaster as before, and that of a climatic catastrophe corresponding to a very high climate sensitivity. We combine these two sources of uncertainty, giving a total of four possible scenarios.

One of the key outputs is the expected discounted value of carbon benefit. In practice, the public authorities identify a tutelary value for carbon, that is a social cost of carbon path $\overline{SCC}_t$, equal to the consumption-equivalent at time t of the damage induced by emitting one more ton of CO_2 at the same period. What matters here is the present social value of a project avoiding carbon emissions. The evaluator can compute these values by discounting the social cost of carbon, using the official discount rate r defined by the public authorities.

This is what is shown in the left part of the following equation. In reality, the social cost of carbon SCC_t is not the same along the different states of nature (depending on growth g and climate sensitivity S), and the exact calculation should compute the expected weighted value, as shown in the right part of this equation:

$$e^{-rt}\overline{SCC}_t \simeq \frac{E[u'(C_t(S,g))SCC_t(S,g) \mid S,g]}{u'(C_0)}$$

As explained in the previous section, when the avoided carbon emissions are proportional to the consumption raised to the β power (as it would be the case for instance if they are linked to a volume of passenger transport), the expected value of environmental externality is then

$$e^{-rt}\overline{SCC}_t C_t^{\beta} \simeq \frac{E[u'(C_t(S,g))SCC_t(S,g)C_t^{\beta}(S,g) \mid S,g]}{u'(C_0)}$$

where the right-hand side corresponds to the approximate method currently used and the left-hand side gives the exact value when considering all the states of nature. The stress-discounting method amounts here to do this calculation by considering only four states of nature. It requires to provide the evaluator for each of the four scenarios with a social cost of carbon path and a basic assumption on economic growth (that translates into a constant discount rate). We will apply this method in section 5 but need first to calibrate then benchmark it. This is done in the two following sections.

4.2 Calibration

Before anything else, general assumptions are needed on risks and on climate damage. Regarding environment-economy modelling, we use the DICE 2016R2, with slightly modified parameters to allow easy comparisons with the analysis done in section 3.[7] We also assume that the environmental policy is set at a low level of ambition (which amount to stop all emissions in 2080), and is not modified to take into account any surprises on growth or climate. This amounts to saying that we do not become aware until too late if the sensitivity to the climate is much higher than expected (Pindyck (2021)) or if we enter in a period of secular stagnation with very low economic growth.[8]

As previously mentioned, we also follow Dietz et al. (2018) who show that two main sources of uncertainty matter, respectively on economic growth and on climate sensitivity, and use their calibration. More precisely, they assume that the growth rate of the total factor productivity (TFP) is driven by a fist-order autoregressive process with an uncertain trend growth rate g_0 that follows a normal law with mean 1.6% and standard deviation 0.9%. Regarding climate sensitivity, Dietz et al. (2018) consider as in Dietz and Stern (2015) that it follows a log-logistic function with a mean at $2.9C^{\circ}$, a standard deviation equal to $1.4C^{\circ}$, truncated from below at 0.75°.

As mentioned before, all these choices are only indicative - in particular, it is up to the public authorities to choose these assumptions, and possibly use a more recent Integrated

[7]We choose as previously a pure rate of preference equal to zero, instead of 1.5% as originally in DICE, and assume no deceleration of TFP, i.e. using the notations used in this model $dela = 0$ instead of 0.5%.

[8]In practice, this means that we maintain in these simulations the emission control rate (MIU parameter in DICE) at its optimal level as calibrated in the base scenario (we assumed here that emissions cannot become negative, following Anderson and Peters (2016)).

Assessment Model. Similarly, regarding the scenarios, we arbitrarily choose to retain disaster scenarios with a probability of occurring equal to 8%. Regarding climate, under the assumptions defined in the previous paragraph, there is a probability equal to 92% that the climate sensitivity is under 4.5° with a mean at 2.6°, and otherwise above this threshold with a mean at 5.7°. We analyze growth distribution using Monte Carlo simulations, and check that the average growth in the 8% worst case corresponds to a stressed scenario with a sudden drop of 43% and a subsequent yearly negative TFP growth equal to -0.2%. The average growth in the 92% remaining cases correspond to a TFP growth equal to 2.3%. We check that the consumption growth is relatively independent of the climate sensitivity, and equals respectively to 3.6% and -0.1% in the BAU and the stressed economic scenarios. We thus obtain the following four scenarios :

- Good climate and BAU (probability 84.64%) : $S = 2.6$ and $g = 3.6\%$;
- Bad climate and BAU (probability 7.36%) : $S = 5.7$ and $g = 3.6\%$;
- Good climate and Stress (probability 7.36%) : $S = 2.6$ and $g = -0.1\%$ after 43% drop of consumption;
- Bad climate and Stress (probability 0.64%) : $S = 5.7$ and $g = -0.1\%$ after 43% drop of consumption.

We can then estimate the paths of temperatures and the social cost of carbon for the four corresponding scenarios, as show in figure 3. As expected, in the scenario with an economic shock, the SCC is higher at the beginning (before the sudden drop in consumption) and lower later on ~~after~~ because it corresponds to the monetary equivalent of damage while future generations are poor in those stressed scenarios.

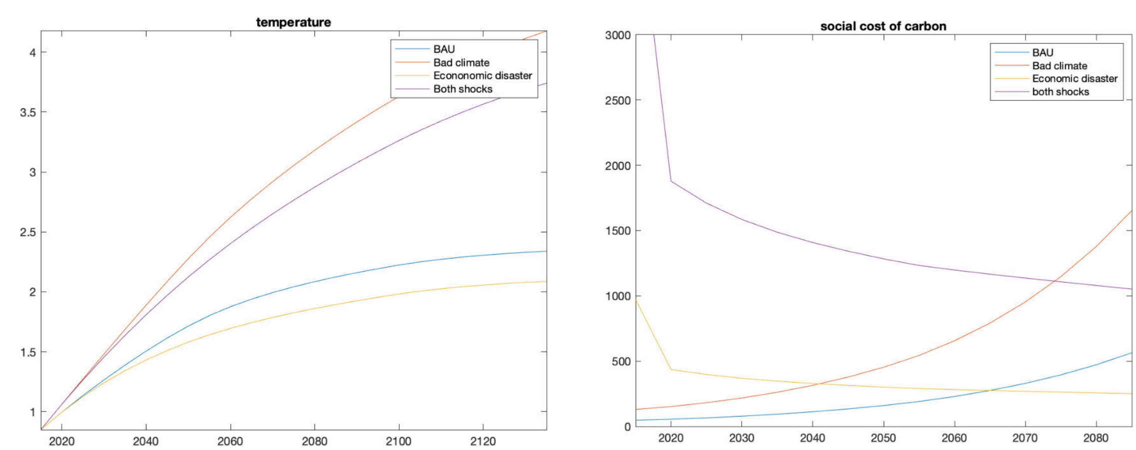

Figure 3: Temperature and social cost of carbon for the four scenarios.

4.3 Benchmark comparison

We can then compute the expected discounted value of carbon benefit, as explained in previous section, and benchmark it to the exact value obtained by running 5000 draws of a

Monte Carlo simulation. We also benchmark it to what would be obtained by using a single discount rate and a single social cost of carbon path, corresponding to the average scenario. The latter scenario is obtained by assuming a mean climate sensitivity equal to 2.9° and taking the mean TFP obtained from the Monte Carlo simulation. This is close to assuming that annual consumption growth is 2.7%, which corresponds to a non-risk adjusted social discount rate equal to 3.915%. The result is shown in the left part of the figure 4. We see that the value obtained using the stress-discounting method is very close to the exact value, whereas the method based on a single growth scenario and a single discount rate provides a value more than 50% too low. This shows that taking risk into account greatly raises the carbon value (this point is discussed in more detail on an example in the next section) and that limiting ourselves to four schematic scenarios is a good approximation. The graph on the right hand-side provides the risk-free rates according to those three approaches, and shows that the stress-discounting method provides, as seen in section 4, a declining term structure close to the real value.

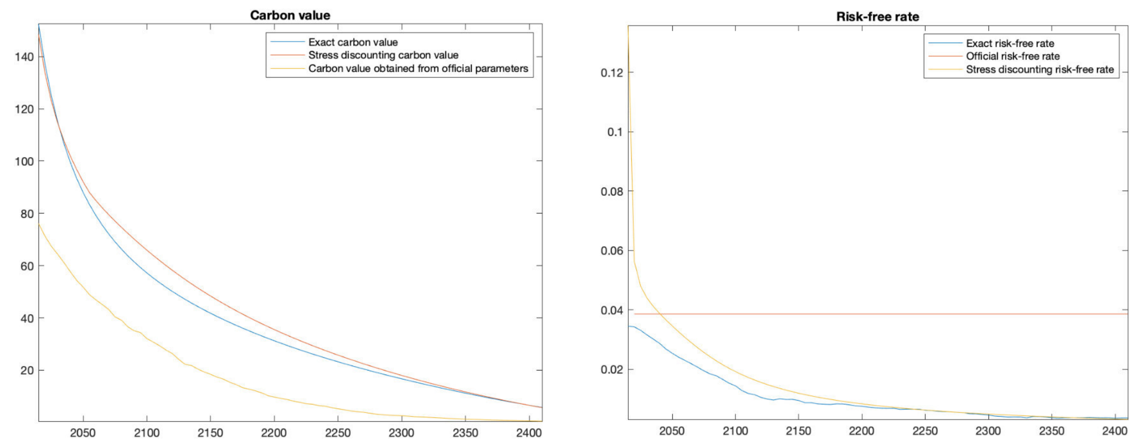

Figure 4: Discount rates and expected present carbon value obtained from a numerical estimation (with 5000 Monte Carlo simulations), from the stress-discounting method and from an approximation based on a single discount rate and a single carbon path (mimicking the method currently followed by public administration).

Another benchmark can be provided through the illustrative example of a project that avoid emitting 1 ton of CO2 each year during 50 years starting from now. The expected present value of such a project is presented in figure 5.

STRESS DISCOUNTING			
	probability	NPV	*contribution*
No negative shocks	*85%*	335	*283*
climate shock	*7%*	932	*69*
eco shock	*7%*	7322	*539*
both shocks	*1%*	31361	*201*
stress		1091	
BENCHMARK			
Exact value		1080	
base		600	

Figure 5: Valuation of a project avoiding 1tCO_2 per year

We observe that

- The stress discounting method gives a very good approximation of the exact value, whereas the official method is very wrong:
- The fourth scenario, despite his very low probability of occurrence, has a significant impact. When both shocks are present, the discounted carbon value is very high because of the high climate sensitivity and because future generations will be relatively poor.

5 Applications

5.1 Nuclear wastes

In France, the second generation of nuclear power plans (1970-2050) will produce a total of 83,000 m^3 of nuclear wastes of high activity or medium activity/long life. The current policy project is to build a geological repository at a depth of 500 meters in the French Ardennes. The site will take 10 years to be built, and the wastes will progressively be transferred in the repository over the next century, for a final irreversible closure around 2150. The flow of gross costs associated to this project is described in Figure 6, with a non-discounted sum of 25.5 billion euros. The code name of this project is Cigéo.[9] Its management has been delegated to ANDRA, the national agency in charge of nuclear waste management. There exist various alternative solutions to Cigéo. Many are either not technologically mature or prohibitively more expensive (Bouttes et al. (2021)). Let us examine the credible alternative of a Permanent Surface Storage (PSS). The PSS strategy consists in periodically repackaging the nuclear wastes to be stored on surface or subsurface, as currently practiced in all countries producing nuclear electricity. For France, the annual gross costs of PSS strategy are estimated at 100 million euros. Gross costs do not take account of the elasticity of these costs (of labour, cement, land use, capital,...) to changes in GDP. Following (Bouttes et al. (2021)), we assume an income-elasticity of these costs equal to 0.8. Notice that this implies that costs will be larger in the BAU scenario than in the stress scenario, and that Cigéo has no macro-hedging benefit if we limit the analysis to the income-elasticity of these costs.

[9] For "Centre Industriel de stockage GEOlogique" in French.

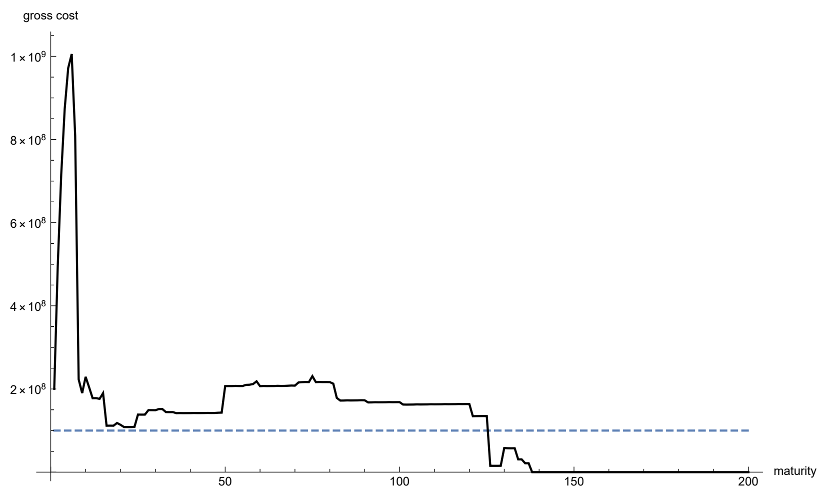

Figure 6: Flow (to be extrapolated to infinity) of the gross costs (in euros per year) of Cigéo (geological repository, plain curve) and PSS (surface storage, dashed curve). Source: ANDRA.

	contingent value		expected
	BAU	stress	value
PV Cigéo	10.61	47.00	12.43
PV PSS	4.68	163.69	12.63
PV PSS with damages	4.68	240.72	16.48

Table 5: Valuation (in billion euros) of Cigéo and PSS costs. In the last line, we add a permanent flow of health and environmental damages $x = 50$ million euros/year materializing in the stress scenario if the PSS option is implemented ex ante.

In Table 5, we summarize the stress discounting procedure to evaluate the competing options Cigéo and PSS. In the BAU scenario, the contingent PV of PSS costs is much smaller than the contingent PV of Cigéo costs. This is due to the observation that with a discount rate as high as $r^b = 3.6\%$, the PSS option is very attractive given the postponment of most expenditures. If one were sure that the BAU scenario would prevail, Cigéo should not be implemented. But the opposite conclusion should be made contingent to the stress scenario, because of the much smaller discount rate $r^s = 0.2\%$ combined with the large penalty $\Delta = 3.08$ to be used when evaluating costs in that scenario. When taking the expected value of these two pairs of contingent PV costs using the stress probability $\pi = 5\%$, we obtain basically the same PV of the costs, around 12.5 billion euros, with a small advantage for Cigéo.

An important piece of the story is missing in this comparison of the two options to manage nuclear wastes in the long run. The geological repository is used as a passive natural barrier to radionuclides. On the contrary, the surface storage of nuclear wastes requires an active maintenance to guarantee its safety. Cigéo has thus an important safety benefit compared to PSS, in particular in the stress scenario. Such a scenario is likely to be associated with a degradation of our democratic institutions and their ability to maintain the right level of supervision, protection and maintenance of the surface storage, as well as cope with the potential adverse consequences of a nuclear incident. To model this idea, let us assume that in the stress scenario, a permanent flow $\bar{B}_s$ of health and environmental damages is incurred by the local population if the PSS option had been selected ex ante. The contingent PV of these damages equals $\bar{B}_s\Delta/(1-\exp(-r^s))$. Multiplying this contingent PV by the probability π of the stress scenario yields the additional PV of costs of the PSS option. Suppose for example a flow $\bar{B}_s = 50$ million euros of annual damages. The contingent PV of this flow in the stress scenario is 77 billion euros. Consistent with equation (**??**), this adds 3.85 billion euros to the expected PV costs, thereby inducing a strong preference for Cigéo. The safety issue associated to the PSS option means that Cigéo is an insurance for future generations, and Table 5 makes that apparent. Although the risk of a chaotic evolution of our society is small, Cigéo should be implemented because of its insurance benefit against the large health and environmental risk of the alternative option.

Imagine how the project analyst would have proceeded if asked to use the standard CCAPM approach to evaluate Cigéo, using PSS as the default option. This analyst should first estimate the CCAPM-beta of the net benefit of Cigéo. It would include the positive income-elasticity of the costs and the negative beta associated to the safety issue examined above. The global CCAPM-beta of Cigéo has a term structure, that should be used to determine the maturity-specific Cigéo risk-adjusted discount rates using the term structures of the risk-free rate and aggregate risk premia examined in Section 3. Then, the analyst should estimate the flow of expected net benefits of Cigéo, using the information about the distribution of future incomes and the income-elasticity of the costs that include health and environmental damages. Finally, this flow should be discounted using the maturity-specific discount rates. It will generate a positive NPV for Cigéo. However, this standard approach is complex and does not provide the essence of the argument for why Cigéo should be implemented.

5.2 Climate mitigation project ?

transportation project with strong climate mitigation impact, such as grand paris express ?

6 Conclusion

[TBD]

References

Anderson, K. and G. Peters (2016): "The trouble with negative emissions," Science, 354, 182–183.

Bansal, R. and A. Yaron (2004): "Risks For the Long Run: A Potential Resolution of Asset Pricing Puzzles," Journal of Finance, 59, 1481–1509.

Barro, R. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," Journal of Finance, 121, 823–866.

——— (2009): "Rare disasters, asset prices, and welfare costs," Americcan Economic Review, 99, 243–264.

Bodie, Z. and R. Merton (2000): Finance, Pearson.

Bouttes, J.-P., C. Gollier, A.-L. M. Allemand, A. Pommeret, and E. Preud'homme (2021): "Contre-expertise de l'évaluation socioéconomique du projet de CIGEO," Tech. rep., Secrétariat Général pour l'Investissement.

Brealey, R., S. Myers, and F. Allen (2017): Principles of corporate finance, 12th edition, McGraw Hill Education.

Cai, Y. and T. S. Lontzek (2019): "The social cost of carbon with economic and climate risks," Journal of Political Economy, 127, 2684–2734.

Cherbonnier, F. and C. Gollier (2020): "Risk-adjusted social discount rates," TSE DP.

Daniel, K. W., R. B. Litterman, and G. Wagner (2019): "Declining CO_2 price paths," Proceedings of the National Academy of Science, 116, 20886–20891.

Dietz, S., C. Gollier, and L. Kessler (2018): "The climate beta," Journal of Environmental Economics and Management, 87, 258–274.

Dietz, S. and N. Stern (2015): "Endogenous growth, convexity of damage and climate risk: how Nordhaus' framework supports deep cuts in carbon emissions," The Economic Journal, 125, 574–620.

Epstein, L., E. Farhi, and T. Strzalecki (2014): "How Much Would You Pay to Resolve Long-Run Risk?" American Economic Review, 104, 2680–2697.

Gollier, C. (2018): "Stochastic volatility implies fourth-degree risk dominance: Applications to asset pricing," Journal of Economic Dynamics and Control, 95, 155–171.

——— (2021): "The welfare cost of ignoring the beta," TSE DP.

Hagen, K., S. Bernstein, B. Bye, L. Hltkrantz, K. Nyborg, K. Pedersen, M. Sandsmark, G. Volden, and G. Avistand (2012): "Cost-benefit analysis," Tech. rep., Official Norwegian Reports NOU 2012:16.

Lemoine, D. (2021): "The climate risk premium: how uncertainty affects the social cost of carbon," Journal of the Association of Environmental and Resource Economists, 8, 27–57.

Lucas, R. (1978): "Asset prices in an exchange economy," Econometrica, 46, 1429–46.

Martin, I. (2013): "Consumption-based asset pricing with higher cumulants," Review of Economic Studies, 80, 745–773.

Mehra, R. and E. C. Prescott (1985): "The Equity Premium: A Puzzle," Journal of Monetary Economics, 15, 145–161.

Pindyck, R. S. (2021): "What we know and don't know about climate change, and implications for policy," Environmental and Energy Policy and the Economy, 2, 4–43.

Quinet, E. (2013): "L'évaluation socioéconomique des investissements publics," Tech. rep., Commissariat Général à la Stratégie et à la Prospective.

Quinet, E. et al. (2013): "L'évaluation socioéconomique des investissements publics—Rapport final," Paris: Commissariat général à la stratégie et à la prospective (CGSP).

Rietz, T. (1988): "The equity risk premium: A solution," Journal of Monetary Economics, 22, 117–131.

Rijksoverheid (2020): "Rapport Werkgroep discontovoet 2020," Tech. rep., Dutch Ministerie van Financien.

Stern, N. and N. H. Stern (2007): The economics of climate change: the Stern review, cambridge University press.

Treasury (2020): "Green book: Central government guidance on appraisal and evaluation," Tech. rep., U.K. Treasury.

Weil, P. (1989): "The equity premium puzzle and the risk-free rate puzzle," Journal of Monetary Economics, 24, 401–421.

Weitzman, M. (2007): "Subjective expectations and asset-return puzzle," American Economic Review, 97, 1102–1130.

Consumer Protection in Retail Investments: Are Market Adjusted Damages Efficient?

Claude Fluet*
Université Laval

This version, September 2023

Abstract

Market adjusted damages (MAD) is the most common form of redress for retail investors awarded compensation for unsuitable financial advice or portfolio mismanagement. Damages are computed as the difference between realized returns and what would have been obtained with an ex ante suitable investment strategy given the investor's needs, adjusting for the actual performance of the market. I analyze the properties of this formula from three perspectives: (i) providing compensatory damages in the sense of making the investor 'whole' despite the unsuitable investment; (ii) as optimal insurance against erroneous advice; (iii) as efficient liability incentives for experts to deliver reliable advice and management. I show that each perspective yields a different variant of MAD. A common feature is that redress is costly – a cost ultimately born by investors – because of opportunistic behavior in seeking redress, given that ex post one is sometimes better off with an unsuitable portfolio.
Keywords: Financial advice, retail financial markets, investor redress, misselling, household finance.

JEL: D18, G24, G28, G52, L51.

*Email: claude.fluet@fsa.ulaval.ca

1 Introduction

It is well known that retail investors have limited understanding of financial products and often rely on the recommendations of experts when making investment decisions. The important losses suffered by small investors in the 2007-2008 financial crisis highlighted the deficiencies of the market for financial services, e.g., investment advisors, brokerage firms, financial planners, etc. A voluminous literature documents the mis-selling and mis-pricing of financial products, either because of conflicts of interest due to commissions and kickbacks from product designers or simply because of careless or incompetent financial advice.[1]

Accordingly, but already underway since the early 2000s, there has been a tightening of the regulatory framework for retail financial markets over the past fifteen years, e.g., the *Dodd–Frank Wall Street Reform and Consumer Protection Act* (2010) in the US or the MiFiD II (2018) regulations in the EU. The emphasis has been on business conduct rules for ensuring transparency, disclosure of appropriate information, and suitable recommendations through *Know Your Client* requirements. As noted by many[2], the analogy is with product safety regulations. To a lesser extent, redress mechanisms for wronged consumers have also been considered, with discussions on how to promote the right to seek compensation from investment advisors who recommended or sold unsuitable financial products, either through civil courts, industry arbitration panels or financial ombudsman authorities.[3] As in producer liability for safety defects, advisor liability would provide direct protection for investors and benefit them indirectly by strengthening the financial intermediaries' incentives to deliver reliable services, thus improving the attractiveness of retails investments.

Investor redress raises the issue of how to assess the harm suffered by wronged investors, given that investments are intrinsically risky. The most

[1] See, among others, Basel Committee on Banking Supervision (2008), Campbell (2006), Campbell et al (2011), Inderst and Ottaviani (2012a), Calcagno and Monticone (2015), Célérier and Vallée (2017), Gennaioli et al (2021), Egan (2019), Egan et al (2019), Linnainmaa et al (2021), Astous et al (2022), and the references therein.

[2] See for instance Warren (2008), Cherednychenko (2010), Moloney (2012).

[3] See CFA Institute (2014) for an international comparison of redress mechanisms and International Organization of Securities Commissions (2021).

common formula for determining the damages awarded, usually referred to as Market Adjusted Damages, is to consider what the investor would have received with a suitable portfolio or investment strategy, given the investor's time horizon and risk tolerance:

> This measure of damage allows the claimant to recover the difference between what the claimant's account made or lost versus what a well-managed account, given the investor's objectives, would have made during the same time period. (FINRA 2017, p. 67)

Specifically,

> These damages compensate an investor for losses caused by wrongful conduct in both a rising and falling market by adding or reducing return according to the actual performance of the market. (Aidikoff et al 2014, p. 135)

Similar formulations are used by many adjudicatory bodies and have been discussed by legal scholars.[4] The following formulation is particularly explicit:

> Where inappropriate financial advice has been provided, the purpose of compensation is to place the consumer in the financial position they would have been in if the financial adviser had provided appropriate financial advice...We need to consider what would have been a suitable alternative. We will look for an alternative portfolio of investments with the correct mix of defensive and growth assets. (Financial Ombudsman Service Australia 2014, p. 2 and 4)

Consensus over these damages formula is relatively recent as evidenced by the evolution of the notion in the legal literature. In the early 1970s, Cohen (1971, p. 1605, footnote 5) remarks that "there has been almost no discussion of the proper measure of damages in a suit for the loss caused by the

[4] See for instance Financial Regulatory Authority (2017), Vandendriesche (2015), Dolden and Newnham (2015), Stanton (2017).

recommendation of an unsuitably high risk" and then compares 'actual damage caused', interpreted as Market Adjusted Damages as defined above, to a rescission standard whereby the investor receives the purchase price of his investment plus risk-free interest since the date of purchase.[5] Easterbrook and Fischel (1985) discuss Market Adjusted Damages but express reservations about the concept. In the case where a client was recommended an excessively risky portfolio, they propose that the best measure of the harm suffered by the investor is the excess risk assessed ex ante, independently of ex post realized returns:

> The court could compute the extent to which the portfolio the broker put together was riskier than an appropriate target portfolio and award compensation that depends on how far a well-chosen portfolio would be expected to outperform the excessively risky one. Any client could obtain this compensation even if his portfolio later beat the market. The award should be based on excess risk viewed ex ante, not on how things turned out. (Easterbrook and Fischel 1985, p. 651).

Literally interpreted, Market Adjusted Damages compensate an investor for any ex post loss due to faulty advice, assuming that suitability is verifiable by the adjudicatory body and that alternative suitable investments (from an ex ante perspective) can be determined. The concept is attractive because it appears to offer an easy way to disentangle the risk of faulty advice from the intrinsic market risk of any risky investment. However, the notion is not without problems. First, an investor sold an unsuitable portfolio will file a claim only when the investment turns out to be unsuccessful. She will stay put when the unsuitable portfolio delivers returns greater than with an appropriate portfolio given her needs and risk tolerance. This will occur, for instance, when an unsuitably high risk portfolio with large expected returns was recommended and the market evolved favorably. With Market Adjusted Damages, the liability risk faced by the advice provider is

[5] See also FINRA (2017, p. 67) for the definition of rescission damages and Himes (1999) for the use of various measures in US court decisions.

therefore one-sided and materializes only in market downturns, with no compensation in market up-turns.[6] Secondly, and most importantly, the advice providers' liability costs will ultimately be born by investors, because they will be factored in the advice fees or the loads on the funds sold. This should be taken into account if the purpose of investors' right to claim damages is to improve their expected utility from investing in risky assets.

This paper analyzes damages formulas for investor redress from three perspectives. First, I characterize the formula that provides full compensatory expectation damages at least cost, in the sense of minimizing the liability cost incorporated in advice fees and given that advice providers will exert costly effort to deliver reliable recommendations. Secondly, I characterize the optimal insurance coverage against the risk of erroneous advice, given that the cost of coverage will be part of the advice fee and is therefore born by investors. Finally, I characterize the efficient liability scheme taking into account the dual function of liability, i.e., providing protection to investors against the risk of unsuitable advice and providing advisors with incentives to supply suitable advice. This third perspective connects with the standard model of producer liability, as developed by Spence (1977) and Shavell (1987, 2007) among others. I show that each of the three perspectives yields a different variant of the Market Adjusted Damages formula.

As in the analytical literature on the market for financial advice (Bolton et al 2007, Inderst and Ottaviani 2012a, 2012b, Carlin and Gervais 2012, Gennaioli et al 2015), I consider a setting where retail investors have difficulty in identifying their needs and have little knowledge of how to invest. The financial advisor's job is to identify the client's needs and to match clients with appropriate investment strategies and products. By contrast with the extant literature, I abstract from biased advice due to conflicts of interest and focus on the risk of mismatches due to the advisor's imperfect information about the clients' needs and the cost to the advisor of identifying correct matches. Another difference is that the extant analytical

[6] Easterbrook and Fischel (1985) make similar observations (see p. 649). The possibility of investor opportunism may explain the reluctance of many jurisdiction to allow redress for faulty advice, for fear of subjecting advice providers to excessive risk of liability; see Black (2010).

literature is too stylized to allow a meaningful study of damages formulas. I consider the recommendation of investment strategies in a setting with risky asset returns where damages, in case of unsuitable recommendations, may depend on the ex post realized returns of the investments. I characterize the appropriate damages formula in this setting.

The paper develops as follows. Section 2 describes the analytical framework. The sections 3 and 4 present two preliminary results, least cost expectation damages and optimal insurance, which are shown to differ. Section 5 characterizes the efficient liability scheme. Section 6 concludes.

2 Model

Consider an investment period, say from date 0 to date 1, in an economy with complete risk trading opportunities. Agents have an exponential utility function with respect to end of period wealth:

$$u_i(w_i) = -\frac{1}{\alpha_i} e^{-\alpha_i w_i} \tag{1}$$

where w_i is the date 1 wealth of agent i and α_i is the agent's absolute risk aversion coefficient. The date 0 market value of the prospect w_i is the expectation $E(mw_i)$ where m is the market stochastic discount factor. With exponential utility functions, it is well known that m is an exponential function of aggregate wealth.[7] Equivalently, in terms of wealth per capita,

$$m = Be^{-\alpha s} \tag{2}$$

where B is a positive constant, α is the harmonic mean of the α_i's, i.e., $1/\alpha$ is the average risk tolerance in the economy, and s is the random date 1 wealth per capita. The gross risk-free rate of return is R_f satisfying $E(m) = 1/R_f$, so that m can be rewritten as

$$m = \frac{e^{-\alpha s}}{R_f E(e^{-\alpha s})}. \tag{3}$$

[7] See for instance Bühlmann (1980), Wang (2003), and Johnston (2007).

Optimal portfolios. Consider now an agent whose date 1 wealth derives solely from the investment of an initial capital w^0. An optimal investment strategy for that agent maximizes $Eu_i(w_i)$ subject to the budget constraint $E(mw_i) \leq w^0$. The optimal strategy eliminates all idiosyncratic risks and the date 1 payoffs, written $w_i(s)$, satisfy the first-order condition

$$u_i'(w_i(s)) = \nu m(s), \; s \in S, \tag{4}$$

where u_i' denotes marginal utility, $\nu > 0$ is a Lagrange multiplier, and $m(s)$ is the stochastic discount factor for a dollar delivered in state s of the economy. Substituting from (1) and (3) into the first-order condition and then in the budget constraint yields

$$w_i^* = w^0 R_f + \left(\frac{\alpha}{\alpha_i}\right)(s - \overline{s}) \tag{5}$$

where

$$\overline{s} \equiv \frac{E(ms)}{E(m)}$$

is the 'risk-neutral' expected wealth per capita.[8]

The prospect w_i^* is the payoff of an optimal portfolio or investment strategy for an agent with absolute risk aversion α_i and wealth invested equal to w^0. Rather than a buy and hold investment, one can also view w_i^* as resulting from an optimal dynamic trading strategy over the investment period $[0, 1]$, as for instance in Palma and Prigent (2009).

My focus is a subset of unsophisticated agents, small or retail investors, who are "unable to fend for themselves" to use the language of the 1933 US Securities Act[9]. These investors have a vague understanding of their needs and are unable to differentiate between optimal and suboptimal portfolios, let alone design and manage complex trading strategies. They therefore seek the advice of experts, e.g., investment advisors, financial planners, brokers or other 'money doctors'. The term financial advisor will refer here to any intermediary with expertise whose job is to assess the investor's needs and

[8] The risk-neutral probability density function is $g^*(s) = [m(s)/E(m(s))]g(s)$, where $g(s)$ is the 'physical' density function.

[9] I borrow this from Carlin and Gervais (2012).

transform the initial capital w^0 into a prospect of date 1 payoffs.

To simplify, the retail investors all have the same initial capital and they belong to one of two categories in terms of risk tolerance: some are type l with risk aversion α_l, some are type h with risk aversion α_h where $\alpha_l > \alpha_h$, i.e., type h is the more risk tolerant. Accordingly, financial advisors design optimal portfolios for each risk tolerance category, for instance a 'conservative' versus an 'aggressive' strategy, and they match customers with the suitable returns profile.

If the fee for financial advice or load on funds is p, the amount effectively invested is $w^0 - p$ yielding the payoff

$$w_i = (w^0 - p)R_f + \left(\frac{\alpha}{\alpha_i}\right)(s - \bar{s}) = w_i^* - pR_f,\ i = l, h. \tag{6}$$

Figure 1 illustrates the net payoffs as a function of the state of the economy at date 1. The portfolio designed for the more risk tolerant has greater risk exposure, as captured by the steeper slope, but this is compensated by larger expected returns.[10]

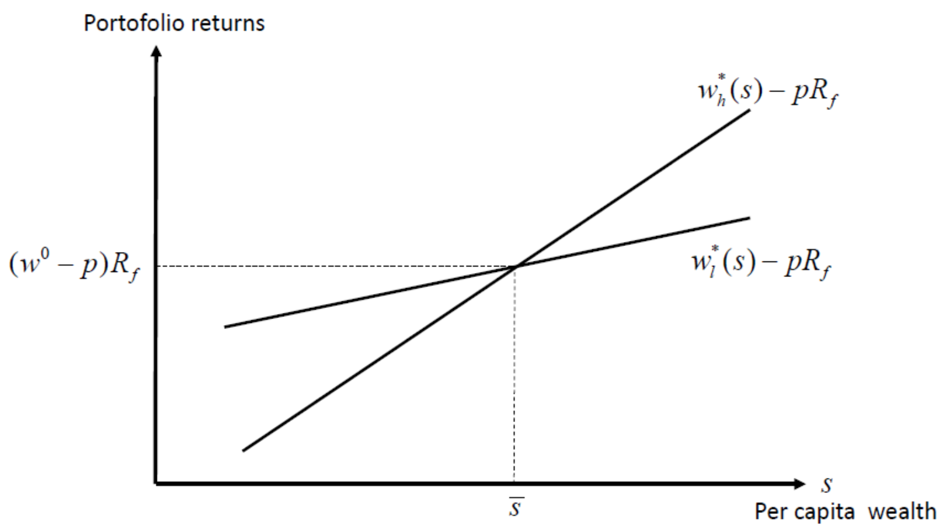

Fig. 1. Optimal net-of-fee returns for types l and h

[10] The expected payoff is $E(w_i) = R_f(w^0 - p) + (\alpha/\alpha_i)(E(s) - \bar{s})$ where $E(s) > \bar{s}$.

KYC ('Know Your Client'). However, the matching process is imperfect. With probability η_i, a type i investor is matched with her type-optimal portfolio; with probability $1-\eta_i$, this investor is matched with the portfolio designed for type j, $j \neq i$. The probability of correct matches depends on the quality of the information gathered about the customer's needs and on the advisor's matching strategy. As in the analysis of experts markets, the financial advisor exerts costly effort to diagnose the customer's needs, in the present case whether the customer's type is l or h, and then selects a 'treatment', here an investment strategy.[11]

The advisor's KYC effort is denoted by $e \geq 0$. The cost to the advisor is $c(e)$, an increasing and strictly convex function with $c'(0) = 0$. I interpret $e = 0$ as some exogenous minimal effort level when an adviser meets a client, with $c(0) \geq 0$. The information obtained by the advisor is summarized by a signal x with continuous densities $f_l(x, e)$ and $f_h(x, e)$ over the same support $X \subset \mathbb{R}$. I assume that $f_l \neq f_h$ for all $e \geq 0$. The advisor obtains some information even with minimal effort, but greater effort will improve the information as described below.

A matching strategy is a function $\varphi(x) \in [0, 1]$ representing the probability of matching the client with the type l optimal portfolio given the information x. Hence,

$$\eta_l = \int_X \varphi(x) f_l(x, e)\, dx \text{ and } \eta_h = \int_X [1 - \varphi(x)] f_h(x, e)\, dx,$$

Let

$$\beta(\eta_h, e) = \max_{\varphi} \int_X \varphi(x) f_l(x, e)\, dx$$

$$\text{subject to } \int_X [1 - \varphi(x)] f_h(x, e)\, dx = \eta_h,\ \eta_h \in [0, 1].$$

From well known results[12], $\beta(\eta_h, e)$ is decreasing and concave in η_h with $\beta(0, e) = 1$ and $\beta(1, e) = 0$.

The function $\eta_l = \beta(\eta_h, e)$ is the 'matching opportunity frontier' describ-

[11] See Dulleck and Kerschbamer (2006), Balafoutas and Kerschbamer (2020), Chen et al (2022) on general experts markets.

[12] See Lehmann and Romano (2005), chapter 3.

ing the trade-offs between the probabilities of correct matches, for a given level of KYC effort. $\beta(\eta_h, e)$ is strictly increasing in e for all $\eta_h \in (0,1)$, i.e., a larger KYC effort shifts the matching opportunity frontier upwards.[13] For tractability, $\beta(\eta_h, e)$ is strictly concave and twice differentiable.

Redress. Once date 1 returns are realized, an investor sold an unsuitable portfolio has the opportunity to file a claim in order to obtain redress or compensation for the unsuitable advice. I assume that investors can always ascertain ex post whether they were mismatched. This is verifiable by courts or specifically designed industry arbitration panels. Moreover, filing a claim is without cost. The redress for an investor sold an unsuitable portfolio may in general depend on the performance of other relevant portfolios, for instance the realizations of w_l^* and w_h^*. I denote by D_i the ex post compensation awarded to a mismatched investor of type $i = l, h$. The main issue in what follows is to characterize the appropriate redress formulas.

Welfare. Let p_i be the load on the portfolios designed for type $i = l, h$; that is, I allow for the possibilities of different loads. Given the quality of matches and the possibility of redress, the average ex-post utility of of a type l investor is

$$\overline{U}_l = \eta_l E u_l(w_l^* - p_l R_f) + (1-\eta_l) E u_l(w_h^* - p_h R_f + D_l). \tag{7}$$

Similarly, for a type h investor, it is

$$\overline{U}_h = \eta_h E u_h(w_h^* - p_h R_f) + (1-\eta_h) E u_h(w_l^* - p_l R_f + D_h) \tag{8}$$

An efficient arrangement is a Pareto-optimum with respect to $\overline{U}_l$ and $\overline{U}_h$ subject to the constraints

$$\eta_l \leq \beta(\eta_h, e),\ \eta_h \in [0,1], \tag{9}$$

$$c(e) + \sum_{i=l,h} \lambda_i (1-\eta_i) E(m D_i) \leq p_l[\lambda_l \eta_l + \lambda_h(1-\eta_h)] + p_h[\lambda_h \eta_h + \lambda_l(1-\eta_l)] \tag{10}$$

[13] That is, $e' > e$ yields a more informative signal in the sense of Blackwell (1951); see for instance Ganuza and Penalva (2010).

where λ_i is the proportion of type i in the population of retail investors considered.

The inequality (10) is the advisor's non negative profit constraint per customer: the loads on funds must cover the cost of KYC effort and the liability costs.[14] Given the matching strategy, the advisor faces a proportion $\lambda_l(1-\eta_l)$ of type l customers who will be mismatched and similarly a proportion $\lambda_h(1-\eta_h)$ of mismatched type h customers. The advisor hedges the risk by purchasing $\lambda_l(1-\eta_l)$ units of an asset (or portfolio) with date 1 payoffs equal to D_l and $\lambda_h(1-\eta_h)$ units of an asset with payoffs equal to D_h. The right-hand side of (10) is the income per customer, given the load on funds and the risk of mismatch.

Ex post investor opportunism. Consider the redress formulas $D_l = w_l^* - w_h^*$ and $D_h = w_h^* - w_l^*$. Then

$$E(mD_l) = E(mD_h) = 0$$

because w_l^* and w_h^* have the same date 0 market value of w^0. The following is therefore a Pareto-optimal arrangement: $p_l = p_h = p$ where $p = c(0)$, yielding the investor expected utility

$$\overline{U}_i = Eu_i(w_i^* - pR_f).$$

The quality of advice is irrelevant because errors in assigning portfolios can always be repaired ex post and this is ex ante without cost. Accordingly, the advisor exerts the minimal level of effort. The outcome is the same as with perfect matches.

However, the above is not feasible because investors are presumed to file a claim whenever a mismatch occurred, even though the amount awarded is negative. From Figure 1, this will arise for one type of investor or the other when $s \neq \overline{s}$. Welfare must therefore be maximized subject to the incentive compatibility or disclosure constraint that damages awarded are

[14] The liability cost with respect to type i can be written as $E(zmD_i)$. where z is an indicator variable with $z = 1$ when a mismatch is verified ex post, $z = 0$ otherwise. Because a mismatch is a purely idiosyncratic event, $E(m\,zD_i) = (1-\eta_i)E(m\,D_i)$.

non negative,

$$D_i \geq 0,\ i = l, h. \tag{11}$$

Given zero ex post litigation costs, investor redress is costly only because of the investors' ex post opportunism, i.e., investors with unsuitable portfolios will sometimes be better off ex post than with their type-optimal portfolio.

Miscellaneous. The following observations will be used repeatedly.

Observation 1. For any random x and y, the following statements are equivalent: $Eu(x) \geq Eu(y)$, $Eu'(x) \leq Eu'(y)$, and $Eu(x+k) \geq Eu(y+k)$ for any constant k.

Observation 2. Let $A \subset S$. Then, $u_i'(w(s)) = \nu m(s)$ for $s \in A$ and some constant ν if and only if $w(s) = w_i^*(s) + k$ for $s \in A$ and some constant k.

The first claim follows trivially from the specification of exponential utility functions, i.e., a constant payoff can be factored out. The second derives from the fact that the optimal payoffs for different amounts of initial wealth are parallel straight lines when expressed in terms of s; see (5).

3 Least-Cost Compensatory Damages

Before discussing arrangements, I consider two preliminary issues. The first is a liability rule that allows investors to claim full compensatory damages (from an ex ante perspective) for an unsuitable portfolio. The second, which turns out to yield a different specification, is the optimal insurance coverage that investors would want to subscribe against the risk of being assigned an unsuitable portfolio.

Expectation damages. For simplicity, let the load be the same across funds. Suppose that the law entitles mismatched investors to obtain damages satisfying

$$Eu_i(w_j^* - pR_f + D_i) = Eu_i(w_i^* - pR_f),\ i = l, h;\ j \neq i. \tag{12}$$

Per Observation 1, this is equivalent to

$$Eu_i(w_j^* + D_i) = Eu_i(w_i^*),\ i = l, h;\ j \neq i. \tag{13}$$

In expectation, investors are then in the same situation whether mismatches occur or not. Borrowing from contract law terminology, I will refer to D_i satisfying (13) as 'expectation damages'. D_i is restricted to be non negative, i.e., damages satisfy the disclosure constraint.

Facing the liability risk, advisors choose their KYC effort and matching strategy to minimize the per client cost

$$c(e) + \sum_{i=l,h} \lambda_i (1 - \eta_i) E(mD_i).$$

Competition between advisors will drive down the advice fee p to the resulting minimum cost per customer.

Suppose KYC effort can generate sufficiently precise information at reasonable cost, so that the advisor's cost minimization involves interior matching decisions, i.e., $\eta_h \in (0,1)$.[15] Substituting for $\eta_l = \beta(\eta_h, e)$, the advisor's effort and matching strategy satisfy the first-order conditions:

$$-\beta_{\eta_h}(\eta_h, e) = \frac{\lambda_h E(mD_h)}{\lambda_l E(mD_l)}, \tag{14}$$

$$\lambda_l \beta_e(\eta_h, e) E(mD_l) = c'(e), \tag{15}$$

where β_{η_h} and β_e denote partial derivatives.

The trade-off between the two possible types of errors depends on the relative costs of compensating type l and type h investors and KYC effort depends on the absolute level of these costs.[16] Increasing the probability

[15] Otherwise, only one type of portfolio would be sold and it would be optimal for the advisor to exert no effort.

[16] The left-hand side of (14) can be shown to define a critical value of the posterior odds (of type l versus type h) for classifying the investor as l rather than h. To interpret (15), note that, from the enveloppe theorem, (15) is equivalent to

$$d\left\{\lambda_l \beta(\eta_h, e) E(mD_l) + \lambda_h \eta_h E(mD_h)\right\}/de = c'(e).$$

where η_h is a function of e via (14).

of correct matches is always beneficial in terms of reducing liability costs. However, for a given quality of information, an increase in η_l must be traded off against a decrease in η_h. A larger KYC effort relaxes that trade-off, but there is then a trade-off between KYC effort and liability costs.

The investors' expected utility is then

$$\overline{U}_i = Eu_i(w_i^* - pR_f), \quad i = l, h,$$

where p is the advisors' minimized unit cost. Incentives to provide reliable advice are driven solely by the advisors' liability risk. The clients' inability to assess the quality of advice does not matter. They are indifferent because they suffer no loss from erroneous advice. They only search for the lowest price, which results from the trade-off between the cost of KYC effort and liability costs. This equilibrium replicates the simple model of producer strict liability for safety defects as developed for instance in Shavell (1987, 2007).

Least-cost damages. Although indifferent to the quality of advice, investors care about its price. So far, damages have been defined by the condition (13) but without further characterization. There is clearly an infinity of formulas satisfying that condition. I now look for the one with the smallest cost. At equilibrium, this will yield the lowest advice fee. The problem is then to choose D_i that minimizes

$$E(mD_i) \text{ subject to (13) and } D_i \geq 0.$$

Proposition 1 *The feasible least-cost compensatory damages for a mismatched type i investor are $D_i^C = \max\{w_i^* - w_j^* - \delta_i^C, 0\}$ where $\delta_i^C > 0$ solves*

$$Eu_i(\max\{w_i^* - \delta_i^C, w_j^*\}) = Eu_i(w_i^*). \tag{16}$$

The intuition for the form of the damages formula is that ex post compensation should be paid only when the benefit-cost ratio is highest. When $D_i^C > 0$, the mismatched type i investor gets $w_j^* + D_i^C = w_i^* - \delta_i^C$. Per Observation 2, the marginal utility-price ratio u_i'/m is then constant, where the constant depends on δ_i^C.

The resulting damages formula is 'Market Adjusted Damages' (MAD) but with a deductible. The investor is compensated for part of the ex post loss due to an unsuitable portfolio, provided the loss is above a threshold. As in the MAD formula, damages are computed as the difference between the realized returns under the unsuitable portfolio and the returns that would have been obtained under the appropriate portfolio, but minus a deductible.

Figure 2 provides an illustration of the post redress payoffs. The gray line depicts the payoffs to type l with the type-optimal portfolio; the broken dark line depicts the final payoffs to a mismatched type l investor. The damages $D_l(s)$ are expressed as a function of the state of the economy.

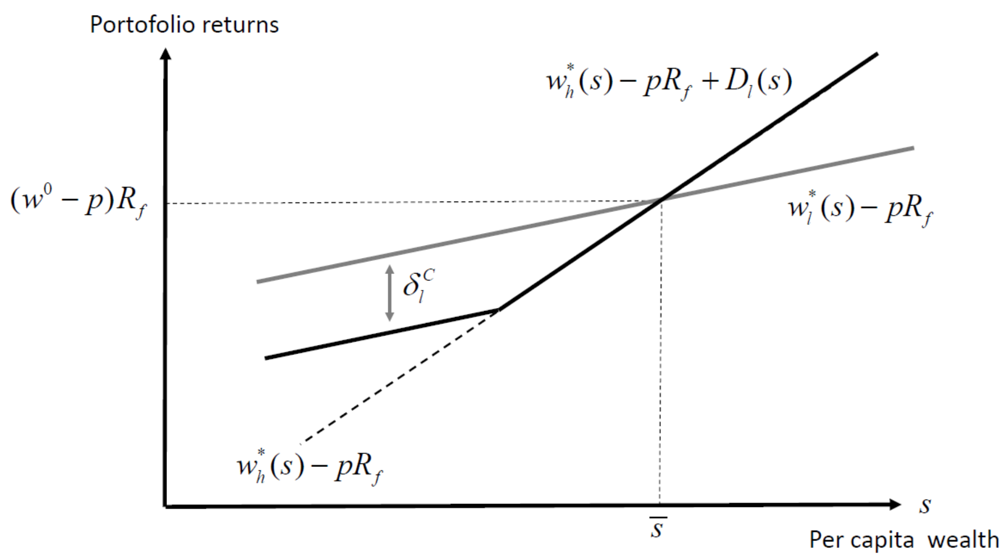

Fig. 2. Least-cost compensatory damages for a mismatched type l

Corollary 1 *If the law requires compensatory damages, the fee for financial advice is minimized by Market Adjusted Damages with the appropriate deductible.*

The result contrasts with other forms of compensation discussed in the literature. As noted in the introduction, Easterbrook and Fischel (1985,

p. 651) suggest that a wronged investor should be awarded compensation based on inadequate risk-taking viewed ex ante, irrespective of "how things turn out". Similarly, de Palma and Prigent (2009) use compensating variations based on certainty equivalents to quantify the losses from misaligned portfolios.

To illustrate, one can always find a constant compensation in case of mismatch, say $D_i \equiv d_i$, that solves (13). With exponential utility functions, the amount required is $d_i = w_{ii}^C - w_{ij}^C$ where w_{ii}^C is the certainty equivalent of the payoffs under the suitable portfolio and w_{ij}^C is the certainty equivalent under the wrong portfolio, i.e., $u_i(w_{ii}^C) = Eu_i(w_i^*)$ and $u_i(w_{ij}^C) = Eu_i(w_j^*)$. The above shows that, while differences in certainty equivalents is an appropriate measure of harm, it does not constitute the appropriate damages awarded ex post.

4 Mismatch Insurance

The preceding section did not discuss another feature of the simple model of producer liability, namely that damages equal to consumers' losses constitute efficient insurance coverage (Spence, 1977). Expectation damages obviously provide insurance against the risk of an unsuitable portfolio, but it does not follow that this is the optimal insurance coverage.

Let us consider the type i investors in isolation. Suppose they face the risk of a mismatch with exogenous probability $1 - \eta_i$, in which case they get the type j optimal portfolio. Without insurance, and assuming there is no advice fee, a type i investor has expected utility

$$U_i = \eta_i Eu_i(w_i^*) + (1 - \eta_i) Eu_i(w_j^*).$$

Because $Eu_i(w_j^*) < Eu_i(w_i^*)$, it follows that $Eu_i'(w_j^*) > Eu_i'(w_i^*)$ per Observation 1. Hence, the investors would want to transfer some wealth from the no-mismatch to the mismatch event, i.e., purchasing some coverage against the risk of mismatch is beneficial.

Abusing notation, I reinterpret p as the insurance premium paid up-front for the coverage D_i in case of a mismatch. The zero-profit insurance premium is then $p = (1 - \eta_i)E(mD_i)$. The investor's expected utility is

$$U_i = \eta_i E u_i(w_i^* - pR_f) + (1 - \eta_i)Eu_i(w_j^* - pR_f + D_i).$$

Feasible insurance policies must satisfy the disclosure constraint $D_i \geq 0$.

To gather intuition, consider the coverage scheme

$$D_i = \max\{w_i^* - w_j^* - \delta, 0\}. \tag{17}$$

This is again the MAD formula with a deductible. Least-cost expectation damages is the particular case with $\delta = \delta_i^C$. Let $p(\delta)$ be the insurance premium given the coverage (17). The investors' expected utility is then

$$\begin{aligned} U_i(\delta) \quad \equiv \quad & \eta_i E u_i(w_i^* - p(\delta)R_f) \\ & +(1 - \eta_i)Eu_i(\max\{w_i^* - \delta, w_j^*\} - p(\delta)) \end{aligned} \tag{18}$$

Because a larger deductible means less insurance coverage, $p(\delta)$ is a decreasing function. We have the following result.

Lemma 1 $U_i'(\delta) < 0$ *for* $\delta \geq 0$*.*

Investors would be willing to pay for an insurance coverage greater than the least-cost compensatory damages defined by the deductible $\delta = \delta_i^C$. Recall that the MAD formula literally interpreted entails that ex post losses due to unsuitable advice are compensated, which amounts to $\delta = 0$, equivalently $D_i = \max\{w_i^* - w_j^*, 0\}$. The above shows that this is better from the investors' point of view than least-cost expectation damages, even though they bear the cost of coverage. Note that expected utility is then greater with the unsuitable than with the type-optimal portfolio.

Optimal insurance. From the lemma, expected utility can be increased further by allowing coverage with a *negative* deductible. I show that a negative deductible is indeed the optimal policy and derive the result without exogenously imposing the form of coverage as in (17). The optimal mismatch insurance for type i investors solves

$$\max_{D_i,p} U_i = \eta_i E u_i(w_i^* - pR_f) + (1-\eta_i) E u_i(w_j^* - pR_f + D_i)$$

subject to $(1-\eta_i)E(m\,D_i) \le p$ and $D_i \ge 0$.

Proposition 2 *The optimal indemnity for a type i investor assigned the wrong portfolio is $D_i = \max\{w_i^* - w_j^* - \delta_i, 0\}$ for some $\delta_i < 0$.*

The end-of-period payoffs under the optimal coverage are depicted in Figure 3 for a type l investor. As in the MAD formula, the compensation for a mismatch depends on the difference in returns between the suitable and unsuitable portfolios. However, the indemnity is greater than the ex post loss due to the mismatch and an indemnity may be paid even though there is no ex post loss.

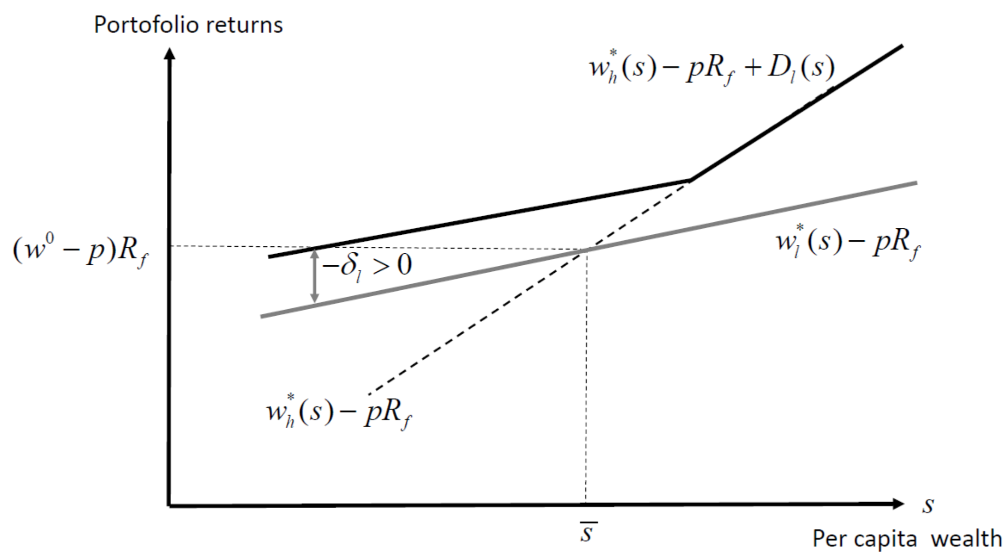

Fig. 3. Optimal mismatch insurance for type l

Two intuitions underlie the result. First, when compensation is paid, the indemnity should be adjusted so as to keep constant the marginal utility-price ratio u_i'/m. The rationale is the same as for the least-cost compensatory damages: the expected utility of mismatched investors is then maximized for a given insurance premium; conversely, the insurance premium is minimized for a given level of expected utility of mismatched investors. Per Observation 2, this requires a payoff equal to w_i^* up to a constant, thereby yielding the form of damages described in the proposition. Secondly, ex post overcompensation, i.e., $\delta_i < 0$, is desirable because it mitigates the inefficiencies imposed by the disclosure constraint. A negative deductible allows wealth to be transferred more often from the mismatch to the no-mismatch event, which is beneficial from an ex ante perspective. Although mismatched investors obtain greater expected utility than with the correct match, this nevertheless comes at a cost, i.e., the insurance premium. Investors would be better off ex ante if mistakes never occurred and they got w_i^* for sure.

5 Optimal Liability

I now turn to the Pareto-optimal arrangement. This is analyzed first without considering the advisor's incentives in making recommendation decisions. Next, incentives are taken into account in order to describe efficient liability schemes.

Optimal allocation. I characterize a Pareto-optimum with respect to $\overline{U}_l$ and $\overline{U}_h$ as defined as in (7) and (8). The constraints are the matching possibility set (9), the non negative profit constraint (10), and the disclosure constraints (11). The maximization is with respect to e, η_l, η_h, D_l, D_h, p_l, and p_h. As before, it is assumed that KYC effort generates enough information at reasonable cost for matching decisions to be interior. To shorten notation, I write

$$u_{ii} \equiv u_i(w_i^* - p_i R_f),\ u_{ij} \equiv u_i(w_j^* - p_j R_f + D_i),\ \overline{u}_i \equiv \eta_i u_{ii} + (1-\eta_i)u_{ij}, \quad (19)$$

u_{ii} is the utility of type i from a correct match and u_{ij} the utility (including redress) from a mismatch.

Proposition 3 *In a Pareto-optimal allocation:*
(i) redress for a type i investor sold the wrong portfolio is $D_i = \max\{(w_i^* -$

$w_j^* - \delta_i, 0\}$ *for some* $\delta_i < 0$, $i = l, h$;
(ii) matching decisions and advisor effort satisfy

$$-\beta_{\eta_h}(\eta_h, e) = \frac{\theta_h \dfrac{(Eu_{hh} - Eu_{hl})}{R_f E\overline{u}_h'} + \lambda_h(p_h - p_l + E(mD_h))}{\theta_l \dfrac{(Eu_{ll} - Eu_{lh})}{R_f E\overline{u}_l'} + \lambda_l(p_l - p_h + E(mD_l))}, \tag{20}$$

$$\beta_e(\eta_h, e)\left[\theta_l \frac{(Eu_{ll} - Eu_{lh})}{R_f E\overline{u}_l} + \lambda_l(p_l - p_h + E(mD_l))\right] = c'(e), \tag{21}$$

where $\theta_l = 1 - \theta_h$ *are weights attached to each type's expected utility.*

Figure 4 provides an illustration of the optimal redress. The broken black line is the final payoffs for a mismatched type l investor; the broken gray line, the final payoffs for a mismatched type h. For either type, the optimal redress is the MAD formula with a negative deductible.

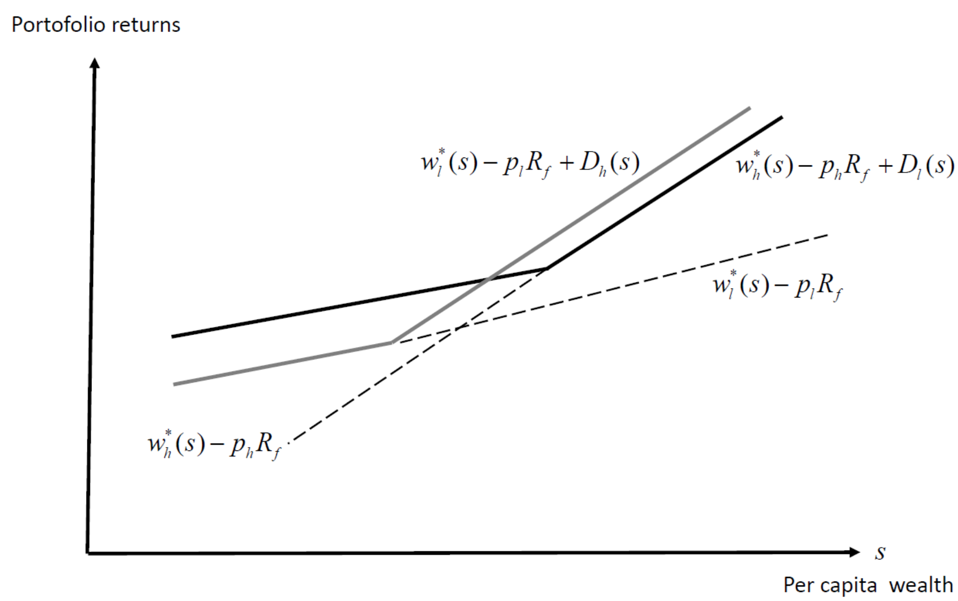

Fig. 4. Optimal redress for type l and h

The conditions for the advisor's matching decisions and KYC effort should be compared with (14) and (15) of Section 3. The difference is that the trade-off between η_l and η_h now does not depend solely on the relative costs of redress to advisor; similarly, KYC effort does not depend only on the absolute level of these costs. The right-hand side of (20) equals the relative social benefits of classifying the investor as one type or the other. The numerator is the benefit over type h clients of a marginal increase in η_h: with respect to the first term, $(Eu_{hh} - Eu_{hl})/(R_f E\overline{u}'_h)$ is the wealth equivalent, in date 0 dollars, of the difference in expected utility between a suitable and an unsuitable portfolio with redress; the rest of the numerator is the savings in redress costs to mismatched type h investors net of the difference in loads. The interpretation of the denominator is similar. Condition (21) states that the marginal cost of KYC effort equals the marginal social benefit. The expression inside the brackets must be positive because β_e and c' are positive. It follows that both the numerator and denominator on the right-hand side of (20) are positive.

Two-part liability schemes. Suppose advisors are liable for the optimal redress payments defined above. The advisor will then choose KYC effort and matching decisions to maximize

$$p_l[\lambda_l \eta_l + \lambda_h(1-\eta_h)] + p_h[\lambda_h \eta_h + \lambda_l(1-\eta_l)] - c(e) - \sum_{i=l,h} \lambda_i(1-\eta_i)E(mD_i)$$

Compared with (20) and (21), the advisor's decisions are then inefficient because effort and matching decisions would satisfy

$$-\beta_{\eta_h}(\eta_h, e) = \frac{\lambda_h(p_h - p_l + E(mD_h))}{\lambda_l(p_l - p_h + E(mD_l))}$$

$$\beta_e(\eta_h, e)\lambda_l[p_l - p_h + E(mD_l)] = c'(e)$$

Thus, the optimal redress payments differ from what provides appropriate incentives to advisors. From (21), because $Eu_{ll} < Eu_{lh}$ (and similarly $Eu_{hh} < Eu_{hl}$) making the advisor liable for the full amount of redress overstates the marginal benefit to clients of KYC effort.

As is well known, there may be a discrepancy between legal damages achieving optimal insurance and damages providing efficient incentives. This arises in particular when the harm suffered by consumers includes a non-pecuniary dimension (Spence, 1977; Shavell, 1987; Polinsky and Shavell, 2010). Liability equal to the consumers' optimal insurance coverage will then typically under-incentivize producers. Achieving efficient incentives requires additional instruments, e.g., fines imposed on producers contingent on the occurrence of harm. In the present case, by contrast, damages equal to the optimal insurance coverage will tend to over-incentivize advisors. The implication is that the liability costs faced by advisors should be decoupled from the insurance coverage provided to investors.

Consider, among other possibilities, a two-part liability scheme involving an insurance pool at the industry level. For instance, the industry arbitration panel handling complaints also operates an insurance pool. The pool is responsible for paying the optimal redress amounts D_l and D_h. It is financed by an ex ante per customer fee t imposed on advisors and by billing advisors ex post for part of the compensation paid out to investors who won a claim against them. Advisors are liable for the payments $\widehat{D}_i = \max\{(w_i^* - w_j^* - \widehat{\delta}_i, 0\}$, $i = l, h$, where the $\widehat{\delta}_i$'s are such that

$$\lambda_i E(m\widehat{D}_i) = \theta_i \frac{(Eu_{ii} - Eu_{ij})}{R_f E\overline{u}_i} + \lambda_i E(mD_l)) \tag{22}$$

where the right-hand side is set at the optimal values. Because $Eu_{ii} < Eu_{ij}$, (22) holds with $\widehat{\delta}_i > \delta_i$. Loosely speaking, advisors are then liable for the redress paid net of the cost of overcompensating mismatched investors, i.e., the advisors' liability cost is akin to expectation damages.

The fee levied by the pool satisfies

$$t = \sum_{i \in \{l,h\}} \lambda_i (1 - \eta_i) E[m(D_i - \widehat{D}_i)]$$

where the right-hand side is computed at the optimal values. The advisor then chooses KYC effort and matching decisions to maximize

$$p_l[\lambda_l \eta_l + \lambda_h(1-\eta_h)] + p_h[\lambda_h \eta_h + \lambda_l(1-\eta_l)] - t - c(e) - \sum_{i=l,h} \lambda_i (1-\eta_i) E(m\widehat{D}_i),$$

which yields the optimal decisions.

Proposition 4 *Optimal incentives for advisors are achievable by two-part liability with an industry compensation pool: the compensation paid to investors sold the wrong products is decoupled from the advisors' liability payments.*

The purpose of a two-part scheme with decoupling is to prevent advisors from facing too large a liability risk, which would unduly increase the cost of advice, while still allowing investors to be appropriately insured against the risk of wrong advice.

6 Concluding Remarks

Investors ultimately bear the cost of redress for unsuitable financial advice. Assuming litigation and verifiability costs are nil, redress is costly only because of investors' ex post opportunism, i.e., investors recommended an unsuitable portfolio will file a complaint only when the 'wrong' portfolio does worse ex post than the ex ante suitable portfolio, which will not always be the case. Literally interpreted, Market Adjusted Damages, i.e., allowing the claimant to recover the difference in returns between suitable and unsuitable portfolio, would overcompensate investors from an ex ante perspective. However, Market Adjusted Damages with the appropriate deductible is an efficient formula if the purpose of the law is to award expectation damages at least cost. By contrast, optimal insurance coverage against the risk of unsuitable recommendations will sometimes overcompensate ex post and will overcompensate from an ex ante perspective. The reason is that this mitigates the effects of ex post investor opportunism. When considering optimal liability for incentivizing financial advisors, however, making advisors liable for the full amount of insurance coverage creates too much incentives, resulting in too high advice costs. An optimal scheme decouples the advisor's liability cost from the insurance coverage provided to investors. Loosely speaking, the advisor should be liable for expectation damages and the excess insurance coverage funded at the industry level.

Appendix

Proof of Proposition 1: The Lagrangean is

$$\mathcal{L} = E(m(s)D_i(s)) + \gamma\{Eu_i(w_i^*(s)) - Eu_i(w_j^*(s) + D_i(s))\} - E(\mu(s)D_i(s))$$

where γ is the multiplier of (13) and $\mu(s) \geq 0$, $s \in S$, are the multipliers of the disclosure constraints. The Kuhn-Tucker conditions are $\mu(s)D_i(s) = 0$, $s \in S$, and

$$\gamma u_i'(w_j^*(s) + D_i(s)) = m(s) - \mu(s),\ s \in S. \tag{23}$$

When $D_i(s) > 0$, $\mu(s) = 0$ and therefore $u_i'(w_j^*(s) + D_i(s)) = m(s)/\gamma$. Per Observation 2, the latter implies

$$w_j^*(s) + D_i(s) = w_i^*(s) + k \tag{24}$$

for some k. The equality (24) cannot hold for all s. Suppose it does. Then (13) would be

$$Eu_i(w_i^*(s) + k) = Eu_i(w_i^*(s)),$$

implying that $k = 0$. But then (24) would imply $D_i(s) < 0$ for some s, contradicting the disclosure constraints.

From (24), damages are $D_i(s) = \max\{w_i^*(s) - w_j^*(s) + k, 0\}$ so that (13) becomes

$$Eu_i(w_j^*(s) + D_i(s)) = Eu_i(\max\{w_i^*(s) + k, w_j^*(s)\}) = Eu_i(w_i^*(s))$$

which can only be satisfied with $k < 0$, yielding the deductible $\delta_i^C = -k$ in the proposition.■

Proof of Lemma 1: We prove the claim for $i = l$; the logic is the same for $i = h$. Using (5),

$$w_l^*(s) - w_h^*(s) = \gamma(\overline{s} - s) \text{ where } \gamma = \alpha/\alpha_h - \alpha/\alpha_l > 0.$$

For a given δ, the premium is therefore

$$p(\delta) = (1-\eta_l)\int_{s_-}^{\widetilde{s}(\delta)} \left[\gamma(\overline{s}-s)-\delta\right] m(s)g(s)\,ds$$

where $\widetilde{s}(\delta) = \overline{s} - \delta/\gamma$. It follows that

$$p'(\delta) = -\left(1-\eta_l\right)\int_{s_-}^{\widetilde{s}(\delta)} m(s)g(s)\,ds. \tag{25}$$

Expected utility is

$$\begin{aligned}
U_l(\delta) \;=\; & \eta_l \int_{s_-}^{s_+} u_l(w_l^*(s) - p(\delta)R_f)g(s)\,ds \\
& +(1-\eta_l)\int_{s_-}^{\widetilde{s}(\delta)} u_l(w_l^*(s) - p(\delta)R_f - \delta)g(s)\,ds \\
& +(1-\eta_l)\int_{\widetilde{s}(\delta)}^{s_+} u_l(w_h^*(s) - p(\delta)R_f)g(s)\,ds.
\end{aligned}$$

Therefore,

$$\begin{aligned}
\frac{U_l'(\delta)}{p'(\delta)R_f} \;=\; & -\eta_l \int_{s_-}^{s_+} u_l'(w_l^*(s) - p(\delta)R_f)g(s)\,ds \\
& -(1-\eta_l)\int_{s_-}^{\widetilde{s}(\delta)} u_l'(w_l^*(s) - p(\delta)R_f - \delta)g(s)\,ds \\
& -(1-\eta_l)\int_{\widetilde{s}(\delta)}^{s_+} u_l'(w_h^*(s) - p(\delta)R_f)g(s)\,ds \\
& -\frac{(1-\eta_l)\int_{s_-}^{\widetilde{s}(\delta)} u_l'(w_l^*(s) - p(\delta)R_f - \delta)g(s)\,ds}{p'(\delta)R_f}.
\end{aligned} \tag{26}$$

With the type-optimal portfolio, $u_l'(w_l^*(s) - p(\delta)R_f) = \widehat{\nu}m(s)$ for all s and some $\widehat{\nu}$. In the case of a mismatch and for $s \in [s_-, \widetilde{s}(\delta)]$, we have $u_l'(w_i^*(s) - p(\delta)R_f - \delta) = \nu m(s)$ for some $\nu \geq \widehat{\nu}$, where the inequality follows from $\delta \geq 0$. Substituting in (26), using (25), and recalling that

$E(m) = 1/R_f$,

$$
\begin{aligned}
\frac{U_l'(\delta)}{p_l'(\delta)R_f} &= -\eta_l\widehat{\nu}/R_f - (1-\eta_l)\nu\int_{s_-}^{\widetilde{s}(\delta)} m(s)g(s)\,ds \\
&\quad -(1-\eta_l)\int_{\widetilde{s}(\delta)}^{s_+} u_l'(w_h^*(s) - p(\delta)R_f)g(s)\,ds + \nu/R_f \\
&= \eta_l(\nu-\widehat{\nu})/R_f + (1-\eta_l)\int_{\widetilde{s}(\delta)}^{s_+}[\nu m(s) - u_l'(w_h^*(s) - p(\delta)R_f)]g(s)\,ds \\
&> 0.
\end{aligned}
$$

The sign follows from $\nu \geq \widehat{\nu}$ and $\nu m(s) > u_l'(w_h^*(s) - p(\delta)R_f)$ for $s > \widetilde{s}(\delta)$. Because $p'(\delta) < 0$, we get $U_l'(\delta) < 0$.■

Proof of Proposition 2: Let $\nu R_f \geq 0$ and $\mu(s) \geq 0$ for all s be the multipliers associated with $(1-\eta_i)E(m\,D_i) \leq p$ and $D_i \geq 0$ respectively. The Lagrangian is

$$\mathcal{L} = U_i + \nu R_f\,[p - (1-\eta_i)E(m(s)\,D_i(s))] + E(\mu(s)D_i(s)). \tag{27}$$

The Kuhn-Tucker conditions are

$$\eta_i E u_i'(w_i^*(s) - pR_f) + (1-\eta_i)E u_i'(w_j^*(s) - pR_f + D_i(s)) = \nu, \tag{28}$$

$$u_i'(w_j^*(s) - p_iR_f + D_i(s)) = \nu R_f m(s) - \mu(s)/(1-\eta_i) \text{ for all } s, \tag{29}$$

$$\mu(s)D_i(s) = 0 \text{ for all } s.$$

First, we show that $\mu(s) = 0$ for all s is not possible. Suppose the contrary. Taking the expectation of (29) then implies

$$E u_i'(w_j^*(s) - pR_f + D_i(s))] = \nu = E u_i'(w_i^*(s) - pR_f). \tag{30}$$

where the second equality is obtained by substituting the first equality in (28). Now (29) and Observation 2 applied to the right-hand side of (30) imply

$$\frac{u_i'(w_j^*(s) - pR_f + D_i(s))}{m(s)} = \nu R_f = \frac{u_i'(w_i^*(s) - pR_f)}{m(s)},\ s \in S.$$

Hence, $D_i(s) = w_i^*(s) - w_j^*(s)$ for all s, contradicting (**??**). Thus, $\mu(s) > 0$ over a set with positive measure. From (29) it then follows that

$$Eu_i'(w_j^*(s) - pR_f + D_i(s)) < \nu$$

which from (28) implies

$$Eu_i'(w_i^*(s) - pR_f) = \widehat{\nu} \text{ for some } \widehat{\nu} > \nu.$$

But then, using (29) and Observation 2 again, when $D_i(s) > 0$,

$$\frac{u_i'(w_j^*(s) - pR_f + D_i(s))}{m(s)} = \nu R_f < \widehat{\nu} R_f = \frac{u_i'(w_i^*(s) - pR_f)}{m(s)}. \tag{31}$$

Applied to the left-hand side of (31), Observation 2 implies that $w_j^*(s) - pR_f + D_i(s) = w_i^*(s) + k$ for some k. From the inequality in (31), $k > -pR_f$. Equivalently, $k = -pR_f - \delta$ where $\delta < 0$.■

Proof of Proposition 3: A Pareto-optimum allocation maximizes $V \equiv \gamma_l \lambda_l \overline{U}_l + \gamma_h \lambda_h \overline{U}_h$ for some weights γ_l and γ_h attached to the expected utility of types l and h. The rest of the argument is then similar to that of Proposition 2 and is therefore omitted. The weights θ_l and θ_h in the proposition satisfy

$$\theta_l = \frac{\gamma_l \lambda_l E\overline{u}_l'}{\gamma_h \lambda_h E\overline{u}_h' + \gamma_l \lambda_l E\overline{u}_l'}, \quad \theta_h = 1 - \theta_l.■$$

References

[1] Aidikoff, P., Uhl, R., Bakhtiari, R., Boice, K. (2014). Market adjusted damages in the FINRA forum. Piaba Law Journal, 21, 135-150.

[2] d'Astous, Ph., Gemmo, I., Michaud P.-C. (2022). The Quality of Financial Advice: What Influences Client Recommendations? Working Paper 30205, NBER.

[3] Balafoutas, L., Kerschbamer, R. (2020). Credence goods in the literature: What the past fifteen years have taught us about fraud, incentives, and the role of institutions. Journal of Behavioral and Experimental Finance 26, 100285.

[4] Basel Committee on Banking Supervision (2008). Customer suitability in the retail sale of financial products and services. Bank for International Settlements.

[5] Black, B. (2010). How to Improve Retail Investor Protection after the Dodd-Frank Wall Street Reform and Consumer Protection Act. Paper 184, University of Cincinnati College of Law.

[6] Blackwell, D. (1951). Comparison of Experiments. In J. Neyman (ed.), BerkeleyProceedings of the Second Symposium on Mathematical Statistics and Probability, Berkeley, University of California Press, 93-102.

[7] Bolton, P., Freixas, X., Shapiro, J. (2007). Conflicts of interest, information provision, and competition in the financial services industry. Journal of Financial Economics 85, 297–330.

[8] Bülhmann, H. (1980). An Economic Premium Principle. Astin Bulletin 11, 52-60.

[9] Calcagno, R., Monticone, C. (2015). Financial literacy and the demand for financial advice. Journal of Banking & Finance, 50, 363-380.

[10] Campbell, J. Y. (2006). Household Finance. Journal of Finance, 61,1553–1604.

[11] Campbell, J. Y., Jackson, H., Madrian, B., Tufano, P. (2011), Consumer Financial Protection. Journal of Economic Perspectives, 115, 707-747.

[12] Carlin, B.I., Gervais, S. (2012). Legal Protection in Retail Financial Markets. Review of Corporate Finance Studies, 1, 68-108.

[13] Célérier, C., Vallée, B. (2017). Catering to Investors through Security Design: Headline Rate and Complexity. Quarteerly Journal of Economics 132, 1469-1508.

[14] CFA Institute (2014). Redress in Retail Investment Markets. International Perspectives and Best Practices.

[15] Chen, Y., Li, J., Zhang, J. (2022). Efficient Liability in Expert Markets. International Economic Review, 63, 1717-1744.

[16] Cherednychenko, O. (2010). The Regulation of Retail Investment Services in the EU: Towards the Improvement of Investor Rights? Journal of Consumer Policy 33, 403-424.

[17] Cohen, S.B. (1971). The Suitability Rule and Economic Theory, Yale Law Journal, 804, 1604-1634.

[18] De Palma, A., Prigent, J.-L. (2009). Standardized versus customized portfolio: a compensating variation approach. Annals of Operations Research, 165, 161-185.

[19] Dolden, E., Newnham, T. (2015). Liability for Financial Advisors in Canada, Dolden, Wallace & Follick.

[20] Dulleck, U., Kerschbamer, R. (2006). On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods. Journal of Economic Literature, 44, 5-42.

[21] Egan, M. (2019). Brokers vs. retail investors: Conflicting Interests and Dominated Products. Journal of Finance 74, 1217–1260.

[22] Egan, M., Matvos, G., Seru, A. (2019). The Market for Financial Adviser Misconduct. Journal of Political Economy 127, 233–295.

[23] Easterbrook, F.H., Fischel, D.R. (1985). Optimal Damages in Securities Cases. The University of Chicago Law Review , 52, 611-652.

[24] Financial Industry Regulatory Authority (2017), Office of Dispute Resolution. *Arbitrator's Guide.*

[25] Financial Ombudsman Service Australia (2014). The FOS Approach to Calculating Loss in Financial Advice Disputes.

[26] Ganuza, J.J., Penalva, J.S. (2010). Signal Ordering Based on Dispersion and the Supply of Private Information in Auctions. Econometrica 78, 1007-1030.

[27] Gennaioli, N., Shleifer, A., Vishny, R. (2015). Money doctors. Journal of Finance 70, 91–114.

[28] Himes, S.F. (1999). Measuring Damages for Fraud-Based Mismanagement of a Securities Portfolio. Securities Regulation Journal, 47, 74-127.

[29] Inderst, R., Ottaviani, M. (2012a). Financial Advice. Journal of Economic Literature, 50, 494-512.

[30] Inderst, R., Ottaviani, M. (2012b). How (Not) to Pay for Advice: A Framework for Consumer Financial Protection. Journal of Financial Economics, 105, 393-411.

[31] International Organization of Securities Commissions (2021). Complaint Handling and Redress System for Retail Investors.

[32] Johnston, M. (2007). Extension of the Capital Asset Prcing Model to Non Normal Dependence Structures. Astin Bulletin, 37, 35-52.

[33] Lehmann E.L., Romano J.P. (2005) Testing Statistical Hypotheses, 3rd ed., Springer.

[34] Linnainmaa, J.T., Melzer, B.T., Previtero, A. (2021). The Misguided Beliefs of Financial Advisors. Journal of Finance, 76, 587-621.

[35] Moloney N. (2012). The investor model underlying the EU's investor protection regime: Consumers or investors? European Business Organization Law Review, 13, 169-193.

[36] Polinsky, A.M., Shavell, S. (2010). The Uneasy Case for Product Liability. Harvard Law Review, 123, 1438-1492.

[37] Shavell, S. (1987). *Economic Analysis of Accident Law.* Harvard University Press, Cambridge, MA.

[38] Shavell, S. (2007). Liability for Accidents. In A.M Polinsky and S. Shavell (eds.), Handbook of Law and Economics, Vol. 1, pp. 139-182, Elsevier.

[39] Spence, M. (1977). Consumer Misperceptions, Product Failure, and Producer Liability. Review of Economic Studies, 44, 561-572.

[40] Stanton, K. (2017). Investment advice: the statutory remedy. Tottel's Journal of Professional Negligence, 33(2), 153-174.

[41] Vandendriescche E. (2015). *Investor Losses. A Comparative Legal Analysis of Causation and Assessment of Damages in Investor Litigation*, Intersentia.

[42] Wang, S.S. (2003). Equilibrium Pricing Transforms: New Results Using Bühlmann's 1980 Economic Model. Astin Bulletin, 33, 57-73.

[43] Warren E. (2008). Product safety regulation as a model for financial services regulation. Journal of Consumer Affairs, 42, 452-46.

New Safety Technologies and Vehicle Safety

Michael Hoy
University of Guelph, Guelph, Canada

Rachel Huang
National Central University,
Taoyuan City, Taiwan

Mattias K. Polborn
Vanderbilt University

Kili Wang
Tamkang University, Taiwan

September 8, 2023

Abstract

There has been a significant amount of research, especially empirical, on the effects that vehicle safety technologies such as seatbelts and airbags have on driver behaviour and resulting accindent rates. In this paper we investigate the impact of vehicle safety technologies which either reduce the probability of an accident or the size of loss associated with an accident should one occur. We do this in an environment with heterogeneous individuals who differ either by their (subjective) cost of taking effort to avoid accidents or by their (subjective) size of loss should an accident occur. We investignate both selection effects (i.e., who will value more highly the technology and so purchase it) and the effect on driving behaviour. The latter is the so-called offsetting or risk compensation effect. Using a data set that combines information from two sources: one about equipment levels of vehicles and the other from insurance experience (i.e., accidents, changes in bonus malus) we investigate the effects of selection (adverse versus advanatageous recruitment) and offsetting behaviour for varying quality airbags and braking systems.

Keywords: Value of research, externalities.

1 Introduction

There has been substantial empirical research on the effect that improved safety devices, such as seatbelts and airbags, have on driving behaviour and resulting accident rates (e.g., see Peltzman, 1975 and Harless and Hoffer, 2003). If the adoption of a safety device reduces the size of loss of an accident to individuals, then the marginal value of exerting effort to avoid accidents falls and one expects some reduction in safe driving behaviour which in turn increases the risk of an accident to others. This type of reaction has been termed the offsetting or risk compensation effect and should be taken into account when valuing improved safety devices or measures (see Gossner and Picard, 2005). In the case of voluntary purchase of such devices, the presence of any externality due to the offsetting effect is also a relevant policy concern (see Hoy and Polborn, 2015). The optimal intervention depends on the extent to which adoption of the safety device reduces the level of care that individuals take as well as the strength of the resulting externality effect. In a setting with heterogeneous preferences, one cannot draw clear conclusions about the strength of any offsetting behaviour created by voluntary adoption of an improved safety technology by simply comparing accident rates or driving records of adopters to non-adopters. The reason is that one needs to decompose this difference into a selection (or recruitment) effect and a behavioural effect. It is this feature of safety technologies that we investigate here.

In a classic paper, Peltzman (1975) identified an important offsetting effect due to mandatory seat belt legislation. Many empirical papers since have investigated the existence and strength of offsetting behaviour across a wide range of technologies and environments. The finding that intended improvements to safety from such regulation may be reduced, entirely eliminated, or even reversed due to offsetting behaviour is an important policy consideration. A plethora of recent developments of safety technologies,[1] which are available for voluntary purchase with select automobiles, makes further study of their effects on driving behaviour important. An important policy concern is determining whether individuals should be allowed to make their own decisions about which vehicle safety features to adopt and, if so, what role can taxes an subsidies play to improve welfare.

We develop a theoretical model with two key features which allows us to organize how to investigate the empirical effects of improved safety technologies with an eye towards providing input for policy. Firstly, assuming no change in driving behaviour, we classify technologies based on whether they have an effect on the size of loss should an accident

[1] The estimated fitment rate for recently developed safety features for 2017 global passenger vehicle production includes 14% with automated emergency braking, 8% with lane keeping assist, 11% with blind spot monitoring, and 7% with adaptive cruise control (IHS Markit - quoted in CARANDDRIVER, Nov, 2017, p. 82).

occur or, alternatively, an effect on the likelihood of incurring an accident. Safety devices such as airbags affect the size of loss due to an accident while others, such as lane departure warning systems or improved braking systems, presumably reduce the probability of an accident. We refer to these types of safety technologies as loss mitigation (LMT) and probability reduction (PRT) technologies, respectively.[2] Secondly, it is important to be able to understand the forces (individual preferences) that motivate some vehicle owners to adopt improved safety techologies and how adoption affects their driving behaviour. To this end, we assume individuals may differ either by their (subjective) size of loss should an accident occur or by their preceived cost of own effort (diligence in safe driving habits). There are, of course, many other potential behavioural traits that may influence such choices. We discuss some of these later in the paper.

One important consideration is that, depending on the extent of any possible offsetting effect, a PRT may provide for either a positive or negative externality. As long as any reduced level of attentiveness to safe driving does not completely neutralize or reverse the inherent effect of the reduced probability of a vehicle with improved PRT causing an accident, the technology provides a positive externality and so a subsidy is in order. If the adopter reduces his own efforts at safe driving so much that he becomes more likely to cause an accident, then the adoption of the PRT leads to a negative externality and should be taxed. As in the classic case of a mandated LMT (e.g., safety belts), adoption of any LMT in our model also leads to a reduction of safe driving efforts. Thus, adopters of improved LMTs generate a negative externality due to the expected offsetting effect and so a tax is in order.[3]

Compared to measuring the effects of safty innovations when they are mandatory (e.g., seat belt laws) or publicly provided (e.g., improved road barriers), there are more complications in an environment of voluntary purchase/adoption. Many questions are raised which require a careful analysis of data in any empirical exercise. What type of individuals will purchase these devices? A priori to adoption decisions, will those at higher risk of accident or lower risk of accident adopt improved safety devices? Conditional on no offsetting effects, would adopters (ex post) display higher or lower accident rates; that is, will there be adverse or advantageous recruitment in addition to possible offsetting effects? How will adoption affect driving habits in regards to safety given different reasons for choosing a particular type of safety device? Given the relevant externalities associated with offsetting behaviour, valuing such technologies requires separating the selection (recruitment) effects from behavioural effects from adoption. For example, if there is a positive correla-

[2]Some technologies may affect both the size and probability of loss, although we do not model such a mixed possibility here.

[3]See Hoy and Polborn (2015) for analysis of optimal taxation of safety technologies for both LMTs and PRTs in a setting with homogeneous individuals.

tion between adoption and accident rates but this is due entirely to selection effects and not behavioural effects, then one will draw a different conclusion about the value of such devices than if the positive correlation is due in part to some offsetting behaviour. These complications are absent when all drivers adopt the safety measure either actively through mandates or passively through public provision.

In order to generate useful intution on these matters, we consider an increasingly complex environment of safety technology adoption: Firstly, we consider mandatory adoption of technologies; second, we consider voluntary adoption of a single type of technology (e.g., LMT) while holding the level of the other technology (e.g., PRT) fixed; third, we consider the voluntary and simultaneous adoption of both types of safety technology. Each of these three settings relate well to different policy environments as discussed later in the paper. Although we do not propose explicit policy recommendations, our results point to appropriate tax/subsidy policies that would enhance welfare.

These three scenarios represent alternative polic environments regarding regulations about vehicle safety. There are many instances of specific safety features being made mandatory, such as seat belts, minimal quality airbag requirements, rear view cameras, etc.. Of particular note are laws passed by the EU requiring from 6 July 2022 that all new pasenger vehicles be fitted with a suite of features including reversing detection with camera or sensors, attention warining in case of driver drowsiness, lane keeping assist and also, between 2024 and 2025, a plan to include advanced driver distraction warning (see https://ec.europa.eu/docsroom/documents/50774). The first scenario of exogenous improvements applies to these instances. Given existing types of echnologies which are mandated, it is useful to know how those affect the value of other newly developed technologies that individuals may choose to adopt voluntarily. This is the second scenario in which we analyze introducing one or the other new PRT or LMT technology. Finally, our third scenario considers the choice problem for both types of technology offered simultaneously.

Although we do not develop specific policy conclusions, our work points in some useful directions by identifying various possible externalities from voluntary (or mandatory) adoption. We address some of these in the discussion section of the paper.

We provide an empirical application using a data set acquired from the Taiwan Insurance Institute (TII). This data set provides detailed information on insureds' claims and driving records. This data is supplemented with information from vehicle records regarding various vehicle characteristics including two safety technologies: quality of airbag systems and quality of braking systems. We designate as a high quality airbag system any vehicle with airbags equipped for both front and back seats, while we designate as a high quality braking system any vehicle which is equipped not only with an anti-lock brake system, which is standard for virtually all cars in our sample, but is also equipped with a traction control system, vehicle stability control system, acceleration slip regulation as well as down-hill assist control, and hill-start assist control. We treat the high quality airbag system as a LMT and the high quality braking system as a PRT.[4] We also consider choice of a SUV as an enhanced LMT since, in comparison to a sedan, a SUV is larger and heavier and so provides better protection to its occupants.

The data covers two years (2011 and 2012) and contains 2,371,730 observations. It is an unbalanced panel. We perform two empirical exercises. Firstly, we treat the data in a cross-sectional manner to estimate the relationship between claims arising from third party losses and various vehicle and driver characteristics including quality of braking and airbag systems. These results should be treated as descriptive of the relationship between safety technologies and accident claims since recruitment effects are not separated from behavioural effects. Second, we extract observations from the data set for those individuals who are present in both years and have an identifiable change in automobile. For these individuals we can determine if they have purchased a new (different) vehicle with higher, lower, or same quality of both airbag system (LMT) and braking system (PRT). This allows us to estimate behavioural effects of the adopted technologies without the confounding implications of possible recruitment effects.

Our propositions lead to implications on whether positive or negative correlations between adopters of a safety technology are consistent with the presence of heterogeneous cost or loss size types in the population (i.e., the presence of adverse or advantageous recruitment). There are, however, substantial challenges in drawing conclusions about actual choices of vehicle safety technologies based on such preferences. Vehicle choice is not solely driven by safety technology present in the chosen vehicle but other features of the vehicle as well which may be bundled together with safety technologies. In the case of purchasing a SUV, it seems appropriate to view the choice to be based on both safety considerations (bigger is safer) and other features (bigger means more storage space).

The paper is structured as follows. A brief literature review follows. Section 3 of this paper provides the basic theoretical model and propositions relating choice of technology to individuals based on each type of heterogeneity (i.e., differing costs of precaution and differing subjective size of loss). Section 4 describes the data and our empirical analysis. In the final section we provide a discussion of our findings.

[4] A high quality braking system presumably also has some characteristics of a LMT since, conditional on being in a potential accident scenario, a better braking system may not allow one to avoid the accident but would reduce the speed of the impact and hence reduce the size of loss.

2 Literature Review

Much of the literature about the phenomenon of offsetting behaviour has been directed at determining empirically its size in a wide variety of economic settings. We focus here on those papers relating to traffic safety.[5] Papers most closely related to ours include Harless and Hoffer (2003), who investigate the recruitment and offsetting effects of voluntary adoption of airbags and Winston, et al. (2006) who consider both adoption of airbags and antilock braking systems as do we.

In comparison to the vast range and depth of empirical research on the offsetting hypothesis, there is relatively little theoretical analysis of the phenomenon. Our model should be thought of as futher developing this stream of research. Of particular relevance to our work is the paper by Blomquist (1986). He develops a general model of driver safety behaviour and demonstrates the result that "under plausible conditions a change in exogenous safety, which is beyond driver control, causes a compensatory change in driver effort in the opposite direction", (Blomquist, 1986, p. 371). His model has both dimensions of safety as does ours (i.e., safety technologies and endogenous driver safey choice) and provides a useful comparative static result describing conditions under which the choice of exogenous safety may reduce the driver's own effort to avoid bad outcomes. However, he does not explicity model the two types of technology that we do and he also does not address welfare implications.

Neill (1993) also develops a model to determine conditions under which the probability of an accident increases or decreases as a result of an increase in the level of an imposed safety technology or regulation. As in our model, this depends on how the increase in the imposed safety technology affects the marginal benefit of individuals' own levels of precaution. His paper investigates how this relationship between the safety technology and the individual's effort to avoid accidents impact on the choice of self-insurance (LMT in our terminology and safety devices in his). However, he does not address the normative implications of imposed safety technologies and restricts his attention to $LMTs$.

Hause (2006) also develops a general model of the offsetting phenomenon. He points out (pp. 689-690) that "Despite accumulating evidence on the empirical relevance of OB (offsetting behaviour), none of the theoretical literature has provided a model determining formal conditions under which dominant or partial OB occurs, much less the magnitude of the OB effect on expected accident loss". By a dominant effect Hause means that the OB effect (change in own effort of accident avoidance) results in no net change in the expected accident loss. By a partial OB effect is meant that the net effect of the safety regulation or technology is a reduction in the net expected accident loss, but less than the

[5] For example, workplace safety (e.g., Lanoie (1992)), sports (e.g., Potter (2011) on formula 1 racing and McCannon (2011) on basketball), food safety (e.g., Miljkovic (2011), et al.).

direct effect.

Another paper that has some of the same properties and objectives as our paper is that of Gossner and Picard (2005). Their goal is to investigate how to value the benefit of an improvement in road safety in the presence of an offsetting effect. The loss in their model is financial and the source of externalities is through the insurance market They consider a similar problem as in our paper by taking into account how changes in road safety affect precautionary effort levels of individuals. Due ot the fact that losses are financial, they, also investigate the implications of drivers'risk aversion on the value of improvements to road safety. In our model, our "uninsured losses" are meant to cover uncompensated pain and suffering as well as uninsured financial losses.

The most important advantage of our model is that we combine the elements of an explicit treatment of (1) optimal choice of safety features including consideration of whether specific features (for a PRT) are strategic complements or substitutes, (2) how the safety technology affects the marginal value of precaution, and (3) whether the adopted safety technology is an LMT (mitigates loss) or a PRT (reduces probability of loss). Importantly, we allow for heterogeneity of preferences in our model in one of two dimensions (cost of driving more safely and size of loss due to an accident). These features allow us to consider most carefully the interplay between adoption decisons (recruitment) and offsetting behaviour.

3 Models

In this section we first develop the individual's objective function based on the level of each type of technology (PRT and LMT) and two possible types of preference heterogeneity. We allow for individuals to differ either by size of loss amount due to an accident as well as differeing cost of precaution. As noted earlier in the paper, we develop our model in the context of three regulatory environments. The first involves describing the effect of an exogenous increase to one or the other type of technology while in the second we treat the case where the individual chooses a level of each type of technology while holding the level of the other technology fixed. Finally, we allow for simultaneous choice of levels of PRT and LMT.

We analyze separately each scenario for individuals who differ due to heterogeneous cost of precaution and due to heterogeneous perception of size of loss. In regards to generating the possibilities of advantageous versus adverse recruitment, the source of heterogeneity is crucial. The implications for analyzing the relationship between levels of these safety technologies and driving behaviour in the data are, of course, complicated by the effects of offsetting behaviour.

We assume each individual chooses (or required to adopt) a level of PRT, θ, and LMT, λ. For a given level of own care, p, a higher level of θ reduces the probability of an accident while a higher level of λ reduces the size of loss should an accident occur. Although some safety technologies no doubt have both effects, we do not model such a possibility.[6]

The probability of an accident (claim) is $D(p, \theta) \in (0, 1)$ with partial derivatives D_p, $D_\theta < 0$, D_{pp}, $D_{\theta\theta} > 0$. Given that a lower value of D_p (resp. D_θ) means p (resp. θ) is at the margin more productive in reducing the probability of an accident, it follows that $D_{p\theta} > 0$ implies that a higher value of θ reduces the marginal productivity of p or, equivalently, a lower value of θ increases the marginal productivity of p (and vice versa). In this case we say that own care and the PRT are substitutes. It seems plausible that a technology like lane departure warning would be a substitute for own care as it could give confidence to drive while more tired and/or pay less attention to one's location on the road. Therefore, if one person has a higher cost of own care then we might expect such a person to acquire a higher level of PRT when it is a substitute for own care (i.e., when $D_{p\theta} > 0$). It seems at least possible that choosing a higher quality ABS system, which is one of the variables of interest in our data set, may improve effectiveness of own care since more dangerous situations can be avoided if one is both more alert and has better brakes - an example of complementarity (i.e., $D_{p\theta} < 0$). On the other hand, it is also possible that better brakes reduces the benefit of driving at modest speeds since the braking distance to a stationary (or slower) vehicle is less and so collisions can be avoided at higher speeds. The sign of this cross-partial not surprisingly is important and so we investigate both possibilities. It seems intuitively appealing that if $D_{p\theta} < 0$, then purchasing a higher level of PRT may actually have a reverse offsetting effect (i.e., lead to an increase in own care and so a reinforcement of the reduction in loss probability). We also assume $D(p, \theta)$ is a strictly convex function.

We also acknowledge here, but do not explicitly model, that the level of care of other drivers will have an effect on an individual's probability of an accident and also may well affect the marginal benefit of both the individual's level of precaution (D_p) and PRT (D_θ). This is explicitly taken into account for a much simpler model with homogeneous individuals and only one type of safety technology in Hoy and Polborn (2015). In that paper, the equilibrium level of choice variables is the same as each individual's optimal value. With heterogeneous preferences, each individual generally chooses a different level for all variables and so equilibrium analysis and formal comparative statics analysis becomes unmanageable. We do, however, return to this issue when addressing our empirical

[6] An improved braking system seems a good candidate for possessing both effects. Being able to brake in a shorter distance (and in a more controled manner) should reduce the probability of being involved in an accident and, conditional on being involved in an accident, may well reduce the consequences.

strategy.

The size of the loss is $L(\lambda) \geq 0$ and depends on the level of LMT (λ) with $L_\lambda < 0$. We assume $L_{\lambda\lambda} < 0$. The loss is not meant to be a financial loss but is measured in monetary equivalent utility terms.[7] Similarly, we let the cost of own care be measured in monetary equivalent terms and represented by $c(p)$ with $c', c'' > 0$. The financial cost of PRT and LMT levels are represented by $k_R(\theta)$ and $k_M(\lambda)$ with both being increasing and strictly convex functions. Each individual chooses $\{p, \lambda, \theta\}$ to minimize

$$\Omega(p, \lambda, \theta) = D(p, \theta)L(\lambda) + c(p) + k_R(\theta) + k_M(\lambda) \tag{1}$$

In the scenario in which differential costs of precaution is the source of individual heterogeneity, we replace $c(p)$ with $(1 + \tau)c(p)$, $\tau \geq 0$ where higher τ represents an individual having higher cost of precaution. In the scenario in which individual heterogeneity is the result of differential size of loss, we replace $L(\lambda)$ with $(1 + \nu)L(\lambda)$, $\upsilon \geq 0$ where higher ν represents an individual having a higher loss from an accident.[8]

We do not include in our objective function characteristics of risk preferences beyond the subjective parameter which can reflect either differences in (subjective) size of loss or a weighting parameter on the probability of loss. In the context of our data, the set of insurance contracts available to consumers is tightly regulated by the Taiwan Insurance Institute and so we believe risk preferences over financial losses resulting from accidents is not an important factor to model. For other settings/countries, this would not be a reasonable assumption and alterations to the objective function which account for alternative risk preferences over financial outcomes would be important to include.

In each model, our reference to expected losses includes whatever costs are pertinent to the individual's choice problem (i.e., the cost of precaution for all models and also the cost of the PRT and LMT when those are purchased voluntarily. We first present the models for exogenous changes in levels θ, λ.

3.1 Exogenous changes to PRT and LMT

Governments often introduce mandatory use of safety equipment (e.g., helmets, safety belts, airbags, rear cameras, winter tires, etc.) or make safety improvements to roads (e.g., rumble strips, crash barriers, illuminated lines, lighting, etc.). To represent the effects of such policies, we write θ and λ as exogenously set at values $\overline{\theta}$ and $\overline{\lambda}$, respectively, and note that individuals incur no (direct) cost to these changes. We first consider case in which the source of heterogeneity is due to a differential cost of precaution. Therefore,

[7] See Hoy and Polborn (2015) for discussion.

[8] We assume these are uninsurable losses of accident victims. A good discussion of what these may can be found in Gossner and Picard (2005). We discuss a potential role for insurance later in this paper.

we assume an individual chooses precaution, p, to minimize the expected loss Ω where:

$$\Omega(p, \overline{\lambda}, \overline{\theta}) = D(p, \overline{\theta})L(\overline{\lambda}) + (1+\tau)c(p) \tag{2}$$

and refer to this scenario as Model A1.

The first order condition is

$$F_p(p, \overline{\lambda}, \overline{\theta}) = D_p L + (1+\tau)c' = 0 \tag{3}$$

Upon totally differentiating with respect to p and $\overline{\theta}$, we obtain:

$$\frac{dp}{d\overline{\theta}} = -\frac{D_{p\theta}L}{[D_{pp}L + (1+\tau)c'']} \tag{4}$$

Given that the denominator is positive (i.e., both D_{pp} and c'' are positive), it follows that the sign of $\frac{dp}{d\overline{\theta}}$ is the same as the sign of $D_{p\theta}$. This is intuitively pleasing since if own care and the level of the PRT are substitutes ($D_{p\theta} > 0$) then one would expect an increase in θ would lead to a reduction in p and vice versa if they are complements.

We can follow the same procedure to determine the effect of an exogenous change in the level of LMT ($\overline{\lambda}$) on own care. We obtain:

$$\frac{dp}{d\overline{\lambda}} = -\frac{D_p L_\lambda}{[D_{pp}L + (1+\tau)c'']} < 0 \tag{5}$$

The above represents a classic offsetting effect (e.g., Peltzman, 1975); that is, a reduction in the size of loss due to an exogenous policy intervention like mandatory seatbelts leads to a reduction in own care.

It is also interesting to see how, given exogenous levels of PRT and LRT, individual care varies according to the level of an individual's extra cost of care, τ. It is straightforward to show upon totally differentiating (3) with respect to p and τ, one obtains

$$\frac{dp}{d\tau} = -\frac{c'}{[D_{pp}L + (1+\tau)c'']} < 0 \tag{6}$$

This is intuitively pleasing since one would expect those with higher cost of precaution would engage in less precaution.

Proposition 1. *Suppose individuals differ according to cost of precaution. At given levels of PRT and LMT, individuals with higher cost of precaution choose a lower level of precaution. An exogenous increase in the level of PRT technology leads to a reduction (increase) in precaution if the PRT is a substitute (complement) to precaution. An exogenous increase in the level of LMT will lead to a reduction in the chosen level of precaution.*

Consider the case in which the source of heterogeneity is due to a differential size of loss, should an accident occur. The individual chooses precaution, p, to minimize the expected loss Ω where:

$$\Omega(p, \overline{\lambda}, \overline{\theta}) = D(p, \overline{\theta})[(1+\nu)L(\overline{\lambda})] + c(p) \tag{7}$$

and refer to this scenario as Model A2.

The first order condition is

$$F_p(p, \overline{\lambda}, \overline{\theta}) = D_p(1+\nu)L + c' = 0 \tag{8}$$

Upon totally differentiating with respect to p and $\overline{\theta}$, we obtain:

$$\frac{dp}{d\overline{\theta}} = -\frac{D_{p\theta}(1+\nu)L}{[D_{pp}(1+\upsilon)L + c'']} \tag{9}$$

As in the case for heterogeneity due to differential cost of precaution, the sign of $\frac{dp}{d\overline{\theta}}$ is the same as the sign of $D_{p\theta}$.

We can follow the same procedure to determine the effect of an exogenous change in the level of LMT ($\overline{\lambda}$) on own care. We obtain:

$$\frac{dp}{d\overline{\lambda}} = -\frac{D_p(1+\lambda)L_{\lambda}}{[D_{pp}(1+\upsilon)L + c'']} < 0 \tag{10}$$

which represents a classic offsetting effect (e.g., Peltzman, 1975).

It is also interesting to see how, given exogenous levels of PRT and LMT, individual care varies according to the level of an individual's extra loss from an accident, ν. It is straightforward to show upon totally differentiating (8) with respect to p and ν, one obtains

$$\frac{dp}{d\nu} = -\frac{D_p L(\overline{\lambda})}{[D_{pp}L + (1+\tau)c'']} > 0 \tag{11}$$

and so, as one would expect, those with higher loss from an accident engage in more precaution. We summarize these results in the following proposition.

Proposition 2. *Suppose individuals differ according to the size of loss due to an accident. At given levels of PRT and LMT, individuals with higher loss choose a higher level of precaution. An exogenous increase in the level of PRT technology leads to a reduction (increase) in precaution if the PRT is a substitute (complement) to precaution. An exogenous increase in the level of LMT will lead to a reduction in the chosen level of precaution.*

We see that if levels of PRT and LMT are determined exogenously, whether heterogeneity is due to differential cost of precaution or differential size of loss due to accident, the effect of an increase in PRT is to lead to an reduction (increase) in precaution if the PRT and precaution are substitutes (complements). The effects of heterogeneity on levels of precaution for any given (exogenously fixed) levels of PRT and LMT are as expected: higher cost individuals choose lower levels of precaution while higher loss individuals choose higher levels of precaution. Comparing levels of precaution across heterogeneous types is less straightforward when all variables are endogenously determined (i.e., chosen at a financial cost by individuals). Moreover, comparing the resulting loss probabilities (accident rates) across heterogeneous individuals is also straight forward: higher cost individuals experience higher accident rates while higher loss individuals experience lower accident rates. As we show below, such comparisons are not so straightforward when individuals choose levels of PRT and LMT. Selection effects combined with behavioural (offsetting) effects lead to a more complicated determination of such comparisons.

3.2 Endogenous Choice of one of PRT or LMT

We now return to the main concern of this paper which is the endogenous choice of safety technologies. In order to better develop intuition, it is helpful to begin with the restriction that each individual chooses a level of safety technology of only one type with the other type set at a fixed level. First we explore the scenario in which individuals differ according to cost of precaution and are faced with a fixed value of PRT ($\theta = \overline{\theta}$) and choose $\{p, \lambda\}$ to maximize:

$$\Omega(p, \lambda, \overline{\theta}) = D(p, \overline{\theta})L(\lambda) + (1+\tau)c(p) + k_R(\overline{\theta}) + k_M(\lambda) \tag{12}$$

Note the change in order of variables with $p \longrightarrow 1$, $\lambda \rightarrow 2$, and $\tau \longrightarrow 3$. This leads to first-order conditions:

$$F_1(p, \lambda, \overline{\theta}) = D_p(p, \overline{\theta})L(\lambda) + (1+\tau)c'(p) = 0 \tag{13}$$

$$F_2(p, \lambda, \overline{\theta}) = D(p, \overline{\theta})L_\lambda(\lambda) + k'_M(\lambda) = 0 \tag{14}$$

Total differentiation gives:

$$dF_1 = F_{11}dp + F_{12}d\lambda + F_{13}d\tau = 0 \tag{15}$$

$$dF_2 = F_{21}dp + F_{22}d\lambda + F_{23}d\tau = 0 \tag{16}$$

where $F_{11} = D_{pp}L + (1+\tau)c'' > 0$, $F_{12} = F_{21} = D_pL_\lambda > 0$, $F_{22} = DL_{\lambda\lambda} + k''_M > 0$, $F_{13} = c'$, $F_{23} = 0$. Thus, with $\mid F \mid > 0$, we have

$$\begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} \begin{bmatrix} dp/d\tau \\ d\lambda/d\tau \end{bmatrix} = \begin{bmatrix} -F_{13} \\ -F_{23} \end{bmatrix} = \begin{bmatrix} -c'(p) \\ 0 \end{bmatrix} \tag{17}$$

which gives (reasons stated below):

$$\frac{dp}{d\tau} = \frac{\begin{vmatrix} -c'(p) & F_{12} \\ 0 & F_{22} \end{vmatrix}}{\mid F \mid} = \frac{-c'(p)F_{22}}{\mid F \mid} < 0 \tag{18}$$

$$\frac{d\lambda}{d\tau} = \frac{\begin{vmatrix} F_{11} & -c'(p) \\ F_{21} & 0 \end{vmatrix}}{\mid F \mid} = \frac{c'(p)F_{21}}{\mid F \mid} > 0 \tag{19}$$

The higher cost types choose a lower level of precaution which, since that effect increases the marginal productivity of the LMT (λ), leads to them choosing a higher level of λ.

Consider the same scenario except for the situation in which individuals differ due to size of loss parameter υ. Individuals choose $\{p, \lambda\}$ to maximize:

$$\Omega(p, \lambda, \overline{\theta}) = D(p, \overline{\theta})(1 + \upsilon)L(\lambda) + c(p) + k_R(\overline{\theta}) + k_M(\lambda) \tag{20}$$

This leads to first-order conditions:

$$F_1(p, \lambda; \upsilon) = D_p(p, \overline{\theta})(1 + \upsilon)L(\lambda) + c'(p) = 0 \tag{21}$$

$$F_2(p, \lambda; \upsilon) = D(p, \overline{\theta})(1 + \upsilon)L_\lambda(\lambda) + k'_M(\lambda) = 0 \tag{22}$$

Total differentiation gives:

$$dF_1 = F_{11}dp + F_{12}d\lambda + F_{13}d\upsilon = 0 \tag{23}$$

$$dF_2 = F_{21}dp + F_{22}d\lambda + F_{23}d\upsilon = 0 \tag{24}$$

where $F_{11} = D_{pp}(1+\upsilon)L+c'' > 0$, $F_{12} = F_{21} = D_p(1+\upsilon)L_\lambda > 0$, $F_{22} = D(1+\upsilon)L_{\lambda\lambda}+k''_M > 0$, $F_{13} = D_pL < 0$, $F_{23} = DL_\lambda < 0$.

Thus, with $\mid F \mid > 0$, we have

$$\begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} \begin{bmatrix} dp/d\upsilon \\ d\lambda/d\nu \end{bmatrix} = \begin{bmatrix} -F_{13} \\ -F_{23} \end{bmatrix} = \begin{bmatrix} -D_pL \\ -DL_\lambda \end{bmatrix} \tag{25}$$

which gives (reasons stated below):

$$\frac{dp}{d\nu} = \frac{\begin{vmatrix} -D_pL & F_{12} \\ -DL_\lambda & F_{22} \end{vmatrix}}{\mid F \mid} = \frac{-D_pLF_{22} + DL_\lambda F_{12}}{\mid F \mid} \tag{26}$$

$$\frac{d\lambda}{d\upsilon} = \frac{\begin{vmatrix} F_{11} & -D_pL \\ F_{21} & -DL_\lambda \end{vmatrix}}{\mid F \mid} = \frac{-DLF_{11} + D_pLF_{12}}{\mid F \mid} \tag{27}$$

In both equations above, the first term of the numerator is positive but the second is negative. Therefore, neither has a definitive sign. A higher value of υ leads to an increase in the marginal productivity of both the LMT (λ) and precaution (p). It may then be the case, for example, that an individual with higher υ will find it worthwhile to choose a sufficiently higher value of λ so as to lower the size of loss of life enough that the marginal productivity of precaution falls.

We refer to these as Models B1.i and B2.i and we summarize the results in the following two pairs of equations with associated propositions.

Model B1.i: Individuals differ according to the cost of precaution and face a fixed level of PRT. Higher cost individuals choose a lower level of precaution and a higher level of LMT and, on average, experience both a higher accident rate and worse driving record.

$$\frac{dp}{d\tau} = \frac{-c'(p)F_{22}}{\mid F \mid} < 0, \quad \frac{d\lambda}{d\tau} = \frac{c'(p)F_{21}}{\mid F \mid} > 0 \tag{28}$$

Proposition 3. *Suppose individuals differ according to cost of precaution, the level of available PRT is fixed, and individuals choose level of precaution and LMT to minimize expected loss. Individuals with higher cost of precaution choose a lower level of precaution and a higher level of LMT and, on average, experience both a higher accident rate and worse driving record.*

Model B2.i: Individuals differ according to the size of loss and face a fixed level of PRT. The net effects on their levels of precaution and LMT are ambiguous. Although marginal productivity of each choice variable is higher for higher loss individuals, responding with a higher choice of one of the variables may lead to a reduction in the optimal choice of the other variable. Therefore, we cannot predict whether higher loss individuals will, on average, experience a lower or higher accident rate or a better or worse driving record.

$$\frac{dp}{d\nu} = \frac{-D_pLF_{22} + DL_\lambda F_{12}}{\mid F \mid}, \quad \frac{d\lambda}{d\upsilon} = \frac{-DLF_{11} + D_pLF_{12}}{\mid F \mid} \tag{29}$$

Proposition 4. *Suppose individuals differ according to size of loss, the level of available PRT is fixed, and individuals choose level of precaution and LMT to minimize expected loss. The relationships between the size of loss and any of the other variables of interest (i.e., level of precaution, level of LMT, average driving record or accident rate) are indeterminate.*

We now explore the scenario in which individuals differ according to cost of precaution and are faced with a fixed value of LMT ($\lambda = \overline{\lambda}$) and choose $\{p, \theta\}$ to maximize:

$$\Omega(p, \theta, \overline{\lambda}) = D(p, \theta)L(\overline{\lambda}) + (1 + \tau)c(p) + k_R(\theta) + k_M(\overline{\lambda}) \tag{30}$$

Note the change in order of variables with $p \longrightarrow 1$, $\theta \rightarrow 2$, and $\tau \longrightarrow 3$. The first-order conditions are:

$$F_1(p, \theta, \overline{\lambda}) = D_p(p, \theta)L(\overline{\lambda}) + (1 + \tau)c'(p) = 0 \tag{31}$$

$$F_2(p, \theta, \overline{\lambda}) = D_\theta(p, \theta)L(\overline{\lambda}) + k'_R(\theta) = 0 \tag{32}$$

Total differentiation gives:

$$dF_1 = F_{11}dp + F_{12}d\theta + F_{13}d\tau = 0 \tag{33}$$

$$dF_2 = F_{21}dp + F_{22}d\theta + F_{23}d\tau = 0 \tag{34}$$

where $F_{11} = D_{pp}L + (1+\tau)c'' > 0$, $F_{12} = F_{21} = D_{p\theta}L$ (sign of $D_{p\theta}$), $F_{22} = D_{\theta\theta}L + k_R'' > 0$, $F_{13} = c'$, $F_{23} = 0$. Thus, with $\mid F \mid> 0$, we have

$$\begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} \begin{bmatrix} dp/d\tau \\ d\theta/d\tau \end{bmatrix} = \begin{bmatrix} -F_{13} \\ -F_{23} \end{bmatrix} = \begin{bmatrix} -c'(p) \\ 0 \end{bmatrix} \tag{35}$$

which gives (reasons stated below):

$$\frac{dp}{d\tau} = \frac{\begin{vmatrix} -c'(p) & F_{12} \\ 0 & F_{22} \end{vmatrix}}{\mid F \mid} = \frac{-c'(p)F_{22}}{\mid F \mid} < 0 \tag{36}$$

$$\frac{d\theta}{d\tau} = \frac{\begin{vmatrix} F_{11} & -c'(p) \\ F_{21} & 0 \end{vmatrix}}{\mid F \mid} = \frac{c'(p)D_{p\theta}L}{\mid F \mid} \begin{Bmatrix} > 0 \text{ if } D_{p\theta} > 0 \\ < 0 \text{ if } D_{p\theta} < 0 \end{Bmatrix} \tag{37}$$

The higher cost types choose a lower level of precaution and a higher (lower) level of the PRT (θ) if the PRT and precaution are substitutes (complements).

Consider the same scenario except for the situation in which individuals differ due to size of loss parameter υ. Individuals choose $\{p, \theta\}$ to maximize:

$$\Omega(p, \theta, \overline{\lambda}) = D(p, \theta)(1+\upsilon)L(\overline{\lambda}) + c(p) + k_R(\theta) + k_M(\overline{\lambda}) \tag{38}$$

Note the change in order of variables with $p \longrightarrow 1$, $\theta \rightarrow 2$, and $\upsilon \longrightarrow 3$. This leads to first-order conditions:

$$F_1(p, \theta, \overline{\lambda}) = D_p(p, \theta)(1+\upsilon)L(\overline{\lambda}) + c'(p) = 0 \tag{39}$$

$$F_2(p, \theta, \overline{\lambda}) = D_\theta(p, \overline{\theta})(1+\upsilon)L(\lambda) + k_R'(\theta) = 0 \tag{40}$$

Total differentiation gives:

$$dF_1 = F_{11}dp + F_{12}d\theta + F_{13}d\upsilon = 0 \tag{41}$$

$$dF_2 = F_{21}dp + F_{22}d\theta + F_{23}d\upsilon = 0 \tag{42}$$

where $F_{11} = D_{pp}(1+\upsilon)L + c'' > 0$, $F_{12} = F_{21} = D_{p\theta}(1+\upsilon)L$ (same sign as $D_{p\theta}$), $F_{22} = D_{\theta\theta}(1+\upsilon)L + k_R'' > 0$, $F_{13} = D_pL < 0$, $F_{23} = D_\theta L < 0$.

Thus, with $\mid F \mid> 0$, we have

$$\begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} \begin{bmatrix} dp/d\upsilon \\ d\theta/d\nu \end{bmatrix} = \begin{bmatrix} -F_{13} \\ -F_{23} \end{bmatrix} = \begin{bmatrix} -D_pL \\ -D_\theta L \end{bmatrix} \tag{43}$$

which gives (reasons stated below):

$$\frac{dp}{d\nu} = \frac{\begin{vmatrix} -D_p L & F_{12} \\ -D_\theta L & F_{22} \end{vmatrix}}{\mid F \mid} = \frac{-D_p L F_{22} + D_\theta L D_{p\theta}(1+\upsilon)L}{\mid F \mid} \tag{44}$$

$$\frac{d\theta}{d\upsilon} = \frac{\begin{vmatrix} F_{11} & -D_p L \\ F_{21} & -D_\theta L \end{vmatrix}}{\mid F \mid} = \frac{-D_\theta L F_{11} + D_p L D_{p\theta}(1+\upsilon)L}{\mid F \mid} \tag{45}$$

In both equations above, the first term of the numerator is positive. The second term is positive (negative) if precaution and the PRT are complements (substitutes). Therefore, if precaution and PRT are complements, then individuals with a higher loss choose both a higher level of precaution and the PRT. In this case, those with higher levels of the safety technology (PRT) will be observed to have a lower accident rate and an improved driving record. If the PRT and precaution are substitutes, then it is possible that the negative second term will dominate the positive first term and one of the two partial derivatives will be negative. Thus, it is possible that a higher loss type might choose a lower level of precaution and end up with a worse driving record. By simple observation of the first order condition, it is not possible that a higher loss type would choose both a lower level of precaution AND lower level of PRT since that would imply optimal choices leading to higher marginal benefit than marginal cost for both choice variables.

We refer to these two scenarios analyzed above (i.e., heterogeneous cost types and heterogeneous loss types) as Models B1.ii and B2.ii. We summarize the results with the following propositions.

Proposition 5. *Suppose individuals differ according to cost of precaution, the level of available LMT is fixed, and individuals choose level of precaution and PRT to minimize expected loss. Individuals with higher cost of precaution choose a lower level of precaution and a higher (lower) level of the PRT if precaution and the PRT are complements (substitutes). As a result, one cannot infer whether those individuals holding a higher level of the safety technology (PRT) have a higher or lower accident rate or better or worse driving record.*

Proposition 6. *Suppose individuals differ according to size of loss, the level of available LMT is fixed, and individuals choose level of precaution and PRT to minimize expected loss. Individuals with higher loss from an accident will choose both higher levels of precaution and the PRT if these are complements. In that case we would observe lower accident levels and better driving records for those who choose higher levels of the safety technology (PRT). However, if the precaution and PRT are substitutes, it is possible that individuals with higher loss from an accident will choose either a lower level of precaution or a lower* *level of PRT. Therefore, it is possible that we would observe worse driving records, but better accident records, for those who hold higher levels of the safety technology (PRT).*

The scenario in which individuals choose simultaneously their levels of precaution, PRT, and LMT, comparative static results mirroring those in the above propositions are more complex mathematically and do not generate any definitive derivative signs of interest. Although the above analysis helps in understanding the intuition for the outcomes when choice variables are made simultaneously, one must rely on empirical analysis to draw any conclusions about changes in offsetting effects due to improved technologies for PRT and LMT becoming available. Therefore, we relegate the exercise of comparative statics determination to the appendix and move on to the empirical analysis in the following section.

4 Empirical Application

In this section, we examine our theoretical predictions by adopting a detailed individual level data of almost all passenger automobile liability insurance contracts sold in Taiwan during 2011 and 2012. To maintain the homogeneity in the incentive on the demand in safety equipment and the purpose of using the vehicles as much as is possible, only private passenger vehicles are included. Our data comprises safety information of the insured vehicles, the characteristics of the policyholders and the insurance experiences, including claims experiences and bonus-malus adjustment. With this unique and complete data, we can empirically investigate the relationship between the adoption in safety technology and the accident rate.

It is important to note that the source of the negative externality in our problem differs from the adverse selection problem in insurance. In that model, the externality arises due to the insurer not being able to identify the risk level of consumers and so high risk types can mimic low risk types which generates a negative externality which typically involves low risk types receiving too little insurance coverage. In our problem, the externality from offsetting behaviour associated with the voluntary adoption of safety technologies arises whether the perpretrators' identities are known or not. Of course, if behaviour were observable, some agent (e.g., the government) could intervene in an effective way to improve welfare.[9]

[9] Our paper follows closely the methodology in Hoy and Polborn (2015) which is essentially an application of the phenomenon of moral hazard in teams. See Holmstrom (1982) for a general characterization of this problem and Cooper and Ross (1985), Lanoie (1991), Pedersen (2003), and Risa (1992, 1995) for useful applications.

4.1 Data

Our data is acquired from Taiwan Insurance Institute (TII), which is a data and research platform of the insurance industry governed by the Financial Supervisory Commission in Taiwan. The TII data includes the type of insurance contracts, claims made, the characteristics of the policyholders and the insured vehicles but the safety information of the vehicles is absent. Thus, by matching the vehicle type, brand, model, and year of manufacture recorded in the TII data, we hand collect the information regarding the safety technology for each insured private passenger automobile via auto magazines, auto manufacture reports and the web sites of all possible resources. To reduce the burden of data collection, we focus on the top four vehicle brands in Taiwan, which accounts for more than 80% of the market share. In total, we have 2,371,730 observations during the data period 2011 to 2012. We are able to link some individuals present in the two years as described below.

For the first part of the empirical analysis we use the complete set of observations. For the second part of our empirical analysis, we extract only those observations based on individuals present in the data set for both years and with relevant information available. There are 1,786,490 observations from 893,245 individuals who are in both years of the data. Thus, we lose 2,371,730 − 1,786,490 = 585,240 observations. Some of the excluded individuals were insured in only one of the years 2011 and 2012. Of the 893,245 individuals who we can track and identify whether they replace their vehicle, there are 17,002 owners who purchased a new (or rather different) vehicle in 2012 and chose the same insurance company as in 2011. For this group, we can determine whether they switched to a vehicle with higher or lower or same quality of airbag and brake system.

Taken at face value, these numbers imply an unreasonably low fraction of individuals purchasing a new vehicle in a given year of 1.9% (i.e. 17,002/893,245 = 0.019). On average, about 8% of vehicle owners purchase a new car each year. The reason for this discrepancy is that an over-represented set of individuals who make up the 585,240 excluded observations purchased new cars at a disproportionately higher rate but have been excluded because we cannot track the specific change in vehicle characteristics. Consistency implies that approximately 20% (118,386) out of the 585,240 excluded observations purchased a new car but the information is lost as they switched insurers; i.e., (17,002 + 59,193)/(893,245 + 59,193) = 0.08.[10]

We use the claim on compulsory liability insurance as a proxy for traffic accidents .

[10] In Taiwan, when people purchase a new car, the car dealer will "recommend" an insurance company to the new car owner. Many car owners take up this recommendation. As a result, we lose the information on these new car owners if they purchased the new car and accepted the recommendation to the target insurance company in year 2012.

Compulsory liability insurance is designed to provide basic coverage for the third party's life as well as bodily injury caused by the usage of vehicles.[11] Every vehicle must be insured with this type of insurance and so our data set is comprehensive. In addition, we only included accidents involving a third party. Therefore, problems involving unclaimed accidents, a common feature of this data, are irrelevant to our study. Occupants of cars with airbags will presumably less frequently suffer bodily injury or death when involved with an accident. Including them would create a bias (see Harless and Hoffer (2003) for a discussion of this issue which contaminates their data set). It is an advantage that we can include only those accidents that involve a third party since the added protection of a LMT on any car that causes the accident would create a bias. However, a vehicle which triggers a third party claim creates no problem of bias based on whether it has a higher quality airbag. To compensate in the event of insufficient coverage under compulsory insurance, individuals can further purchase voluntary third party bodily injury or property damage liability insurance, which respectively covers bodily injury or property damage sustained by the third party.

Individuals with a higher degree of risk aversion may demand more insurance. Therefore, we divide our sample into two subsamples: one includes observations covered by voluntary third party bodily injury liability insurance (about 57%), and the other includes those without this type of additional insurance coverage (about 43% possess only compulsory third party insurance). Note that the risk covered in voluntary third party property damage liability insurance is different from the risk covered by compulsory liability insurance. Thus, we do not divide our sample according to the decision on the choice of voluntary third party property damage liability insurance.

As for many other countries, Taiwan has a bonus-malus system to provide an incentive for careful driving. The bonus malus coefficients in compulsory insurance could be 0.7, 0.74, 0.82, 1, 1.1, 1.2, 1.3, 1.4, 1.5, and 1.6. New drivers start at 1. If they remain claim free in the current policy year, then their bonus-malus coefficient become 0.82 in the following year, which means that they will enjoy an 18% price discount. If they remain claim free, then their coefficient will fall to 0.74 in the following year. The lowest coefficient possible is 0.7, which implies a 30% price discount. If a new driver has at least one claim in the current policy year, then the coefficient becomes 1.1 in the following year. Since this variable is determined by past driving records, we treat bm as a proxy for the risk type of an individual.

Two types of safety equipment are considered: airbag and braking system. Since

[11] Under the compulsory liablility insurance, the coverage for life is NT\$1,600,000. For bodily injury, the coverage depends on the degree of incapacity, ranging between NT\$40,000 and NT\$1,600,000. It also covers medical expenses, including the costs of first aid and treatment with an upper-limit of NT\$200,000.

almost every vehicle has at least one airbag and every vehicle has at least a standard anti-lock braking system, we examine the effect of the demand for high quality airbag and braking systems on claims. The high quality airbag system means the airbag is equipped in both front and back seats, while a high quality braking system means that the vehicle is equipped not only with an anti-lock brake system, but is also equipped with traction control system, vehicle stability control system, acceleration slip regulation as well as down-hill assist control, and hill-start assist control. We view the high quality airbag system as an improved LMT, and the high quality breaking system as an improved PRT.

We also consider the choice of purchasing a SUV rather than a car. We consider two aspects of SUVs on safety. On one hand, SUVs have a size advantage and thus afford protection to the driver and passengers of the vehicle. On the other hand, SUVs are heavy vehicles, which usually have a high impact in a traffic accident. If another vehicle or a pedestrian is hit by a SUV, there is a high chance that the passengers in the hit vehicle or the pedestrian would be seriously injured. This has a positive effect on probability of a claim involving third party bodily injury or death. Thus, SUVs could be viewed as having properties of both a higher level of LMT (lower loss to its occupants should an accident occur) and a lower level of PRT in that any accident is more likely to trigger a third party claim for bodily injury or death. Our data includes information on the policyholders, such as gender, marital status, and age. Other information on the insured vehicles are also included, such as vehicle age, and the vehicle registration location. Table 1 shows the definition of all variables used in our study.

As noted above, our sample is a two-year unbalanced panel data set which covers years 2011 and 2012. Panels A, B, and C of Table 2 respectively show the basic statistics of our variables for the whole sample, the subsample that is covered by voluntary third party bodily injury insurance, and the subsample that is not covered by voluntary third party bodily injury insurance. Panel A shows that about 47% of our research sample is in year 2011. The average claim rate (claim = 1) is 1%. About 80% of the observations have been rewarded by the bonus malus system and get a 30% discount (Dbm = 0). We classify the rest 20% of the observations (Dbm = 1) as high risk type according to past driving records. For safety equipment, about 10% of the observations have a high quality airbag system (airbag_high = 1), 39% have a high quality braking systems (brake_high = 1), and 8% of are SUVs (veh_suv = 1). Panel A further shows that fewer than 1% of the vehicles are equipped with both a high quality airbag and a high quality braking systems, whereas there are 51% of the vehicles are equipped with both standard quality airbag and braking systems. Females (female = 1) account for 60% of the registered car owners. About 75% of the insured are married. Age is highly concentrated in the 30 to 60 years old group (age3060 = 1). New cars (carage0 = 1) account for about 7% of the sample, and about 61% of the cars are more than 4 years old. From Panels B and C, we see that the subsample with voluntary third party bodily injury insurance have a 28% higher claim frequency, have a (slightly) better bonus-malus score and are less likely to have high quality airbag or breaking systems. These observations are suggestive of individuals with higher insurance levels being of lower risk (i.e., advantageous selection in the insurance context). However, much more attention to this issue is required to draw any strong conclusions.

Table 3 reports the correlations between the proxies of accident risk (third party claims), driver risk type (bonus-malus score) and the safety technologies for LMT and PRT. We see that claim and Dbm are significantly positively correlated, as one would expect, implying that individuals with worse accident history (higher bonus-malus coefficient) have a higher chance to file a claim in a given year. The correlation coefficient between *claim* and *airbag_high* and between Dbm and *airbag_high* is -0:005 and -0:004, respectively. Both of the coefficients are significant. Since the high quality airbag systems could be viewed as a LMT, this finding provides preliminary evidence for advantageous recruitment with any offsetting effect not strong enough to counter the recruitment effect. Table 3 also shows that *brake_high* is significantly negatively correlated with both Dbm and claim. This supports the view that high quality braking systems are purchased by more cautious drivers, which is also consistent with advantageous recruitment. Moreover, the combined effect of adopting this high quality PRT (i.e., recruitment effect net of any possible offsetting effect) is a reduction of the probability of a claim. Interestingly, the correlation between *veh_suv* and *claim* is insignificant, but *veh_suv* is significantly positively correlated with Dbm. One possible reason for the relationship could be that individuals with a high bonus-malus coefficient (high risk drivers) purchase SUVs to protect themselves as well as any other passengers in the vehicle. However, the safety advantage of SUVs is offset by the increased risk effect from drivers of SUVs. Of course, there are many other possible explanations for all of these tentative conclusions. For example, people who purchase SUVs may be from an age group which includes individuals with different driving abilities. The following two subsections investigate these matters more thoroughly through the use of probit regression equations.

In the first of these (subsection 4.2 below), we preform a probit regression with dependent variable claims[12] for the entire data set treated as a cross-section. In the second exercise (subsection 4.3 below) we investigate the impact of changes in the quality level of airbags and braking systems on drivers' claim experience for those who purchase a new car in 2012 in order to generate a more convincing test for offsetting behaviour.

[12]Recall that claims are from accidents involving bodily injuries and deaths to third parties as covered by compulsory liability insurance.

4.2 Statistical Evidence: Part 1 - Crossectional Analysis

In this subsection we investigate the relationship between accident rates and the choice of safety technologies based on our entire sample (treated as a cross-section). Doing so provides a better understanding of the relationships between accident risk and vehicle safety technologies than simple correlations. However, the results are still descriptive in that we cannot separate recruitment and behavioural (offsetting) effects of improved LMT and PRT technologies. We may tentatively identify how certain observables (e.g., age, marital status, gender) relate to the demand for improved safety technologies but unobservable preference heterogeneity is not revealed in this exercise. The analysis is still of interest since observing the combined effects of these forces on accidents (claims) in conjunction with the analysis of the following two subsections help us to better understand the various issues raised in this paper. We first employ the following Probit model:

$$\begin{aligned} \Pr(claim_{it} &= 1 \mid airbag\_high_{it}, brake\_high_{it}, beh\_suv_i, bm_{it}, X_{it}) \qquad (46) \\ &= F(airbag\_high_{it}\beta_1 + brake\_high_{it}\beta_2 + veh\_suv_{it}\beta_3 + bm_{it}\beta_4 + X_{it}\beta_5) \end{aligned}$$

In Equation (46), $claim_{it} = 1$ when the insured i has filed a claim based on compulsory automobile liability insurance during the policy year t, otherwise $claim_{it} = 0$. F denotes the cumulative distribution function of the Probit regression, and is assumed to be normally distributed. The variables $airbag\_high_{it}$, $brake\_high_{it}$, and veh_suv_{it} are the safety technologies. bm_{it} is the bonus-malus value of the insured i at time t. The vector X_{it} denote the explanatory variables, including gender, marital status, age of the policyholder, vehicle age, the vehicle registration location, and a year dummy ($year2011$) to control the time effect. β's are the corresponding coefficients.

Table 4 shows that for the whole sample, as well as in each subsample, the coefficient on $airbag\_high$ is significantly negative. This finding differs from some previous research (e.g., Peterson, Hoffer, and Millner, 1995; Harless and Hoffer, 2003) which finds that drivers of vehicles equipped with airbags are more likely to be at fault in accidents. However, we find that the coefficient of $airbag\_high$ is significantly negative; i.e., drivers in a vehicle equipped with high quality airbag systems are less likely to cause accidents. In other words, our findings suggest that high quality airbag systems are associated with advantageous recruitment and any offsetting effect that may exist is not sufficiently strong to counteract the recruitment effect.

Table 4 also shows that the coefficient of $brake\_high$ is significantly negative in all groups of samples, which is at least not inconsistent with advantageous recruitment. Treating the high quality braking system as a higher level of PRT means that adoption per se of this technology should lead to a reduction in the probability of an accident. If there is an offsetting effect, this is not strong enough to reverse the accident mitigation effect of the higher quality PRT. There could be an adverse recruitment effect in this case if it is not strong enough to reverse the net effect of the two forces described above. However, we cannot conclude one way or the other about the recruitment effect from these results.

The coefficient for veh_suv is not significant. In other words, we do not find evidence that choice of an SUV implies a change in the driver's probability of an accident.

As expected, individuals with a higher bonus-malus coefficient have a higher probability of sustaining an accident. On average, females have a higher accident rate than males. Married individuals have a lower probability than singles. For different age groups, we find that the young policyholders (younger than 25 years old) have the highest probability of sustaining an accident. New cars and cars with age younger than 4 years old have a higher probability of sustaining an accident than do cars older than 4 years old in the whole sample and the subsample without voluntary third party bodily injury insurance. The differences among different car age groups are not significant in the subsample with voluntary third party bodily injury insurance.

4.3 Statistical Evidence: Part 2 - Panel Data Estimates

In this section we use the data only for those vehicle owners who were present in both years of the sample period and purchased a new (different) vehicle in the second year. By tracking whether they retained the same quality safety technologies (LMT and PRT levels) or upgraded or downgraded we can estimate the impact of these decisions on claims (accidents). Assuming that preferences do not change over the two years, any observed change in claim experience resulting from a change in equipment is independent of recruitment effects. If an individual purchases a new vehicle with an upgraded braking system AND does not adjust behaviour, then we should observe a reduction in claim probability. If there is offsetting behaviour (i.e., the individual drives less carefully), then this will mitigate the intrinsic safety effect of the improved braking system. As long as the mitigation is only partial, there will still be a negative relationship between the adoption of the higher quality brake system and accidents caused and so a subsidy is in order. The size of the subsidy, however, should be reduced according to the extent of any mitigation effect. If the mitigation effect is more than 100%, then we would observe a positive relationship between purchasing a vehicle with an improved braking system and the probability of a claim. In this case a tax on cars with upgraded brakes would be in order.

If an individual purchases a new vehicle with an improved airbag system, there is no intrinsic effect on safety and so any change in claim experience may be considered attributable to behavioural change. From the theoretical perspective, we do expect at least a small increase in the probability of claim. The size of the offsetting effect would determine the size of the appropriate tax to deal with the negative externality.

As noted in section 4 (Data), there are 17,002 individuals who switched vehicles in 2012. The relevant data on these vehicles, as well as other variables used in this section, is summarized in Tables 5 and 7 with definitions of "new" variables given in Table 6. By "new" cars, we mean new purchases including individuals who purchase a used vehicle in 2012. From Table 7, we see that the age of newly purchased vehicles is, on average, 2.25 years less (newer) than the vehicles previously owned. A larger fraction of new vehicles had reduced quality braking systems (13.4% with increased quality and 17.8% reduced quality) while more vehicles had increased quality airbags (9.2% with increased quality and 5.5% decreased quality). 5.48% of new purchases represent a change from a car to a SUV while 2.25% involved the reverse change.

We estimate the following logistic regression equation:

$$\log\left(\frac{1-p}{p}\right) = inc_{brk}\beta_1 + dec_{brk}\beta_2 + inc_{arbg}\beta_3 + dec_{arbg}\beta_4 + inc_s\beta_5 + dec_s\beta_6 \quad (47)$$

$$+ \; < \; < int\_safe\_size >> + delta\_carage\_\beta_7 + X\beta_X + \Delta Externality \quad (48)$$

and estimate separately the set of observations including an increase in claims in 2012 compared to 2011 (*riskier*) and those observations including a decrease in claims in 2012 compared to 2011 (*lessrisky*). The definition of "becoming *riskier*" includes: (1) *riskier_ clm*: no claim in first year and at least one claim in second year, (2) *riskier_ clmtimes*: claim times increase from first year to second year, (3) *riskier_ clmamt*: increase in dollar amount of claims. A similar set of definitions is used for the *lessrisky* variables.

The X vector includes variables already used (e.g., female, married, age2530,age3060, ageabv60, carage0, carage1to4, city, north, south, east. We also include regional differences in the safety variables which pose changes in the driving environment and so are described as $\Delta Externality$, which is a vector that includes the change in the mean value of changes in safety equipment of vehicles, the mean value of the number of tickets issued in the registration county, and so forth as described in Table 6.

In order to use the Logistic Regression, we compare separately those individuals displaying a higher experience of claims to the those with no change in claims. So, for example, we define the dummy variable *riskier_ clm* which equals 1 for any individual who experienced a claim in 2012 but did not in 2011 and assign a value of 0 otherwise. Alternatively, we compare those individuals displaying a reduced claims' experience to those with no change in claims. In this case we define the dummy variable lessrisky_clm which equals 1 for any individuals who experienced a claim in 2011 but not in 2012 and assign a value of 0 otherwise. We then regress these dummy variables against the various independent variables to determine whether adoption of either a higher or lower quality brake system (or airbag system) is statistically related to an increased or decreased risk of making claims (i.e., being the cause of an accident resulting in a third party claim for bodily injury). For the sake of robustness, we also use as dependent variable a change in the amount of claims created between the two years and the number of times a claim is made)[13]. We find no important differences between the results.

In our regression that investigates possible reasons for increased claims (riskier driving - Table 8), we find a statistically significant negative relationship between vehicles with decreased quality airbag systems. This is consistent with classic offsetting behaviour as people who become more exposed to risk of injury in their vehicles through purchase of a vehicle with lower quality airbag system are less likely to create an accident claim. The reverse, however, does not hold; i.e., there is no statistically significant relationship between the variable inc_{arbg} and claims. There is no statistically significant relationship between our measures of riskier driving and either the purchase of a vehicle with increased or decreased quality braking systems. Of course, this does not mean there isn't a change in people's driving behaviour. For example, a person who purchases a new vehicle with improved brake system may drive less carefully because of this feature while the intrinsic improvement in safety from the improved brakes effectively mitigates the reduced caution in driving behaviour and so offsetting behaviour may be present.

In our regression that investigates possible reasons for a reduction in claims (less risky driving - Table 9), there is a statistically significant relationship between the reduction in claims and purchase of a vehicle with an improved brake system. As noted at the beginning of this section, such a result is consistent with some offsetting behaviour if the extent of the offsetting behaviour is not so strong as to reverse the beneficial effect on safety from the improved braking technology. In any case, if the net effect is an improvement in safety (reduction in claims), then a subsidy on vehicles with improved brake systems is warranted. Perhaps surprisingly, there is no statistically significant relationship between the purchase of a vehicle with a lower quality braking system and any of the measures of improved (reduced) claims experience. This may be explained by individuals who purchase vehicles with lower quality braking driving with even greater care to offset the intrinsic reduction in safety from the change in brake system.

There are a few results of interest from other variables which are significantly statistically related to the dependent variables. The change in car age (i.e., the bigger the difference - typically negative - in the age of the newly purchased car and the car owned in 2011) is positively related to our measures of reduced level of claims (e.g., $lessrisky\_clm$) and negatively related to our measures of increased level of claims (e.g., $riskier\_clm$), although not statistically significantly in the latter case. This could be due to drivers becoming more careful in how they drive their "newer" vehicles and/or due to overall increase

[13] There are only a few instances, for example, of a person having one claim in 2011 and two claims in 2012 that would trigger a value of 1 for the dummy variable $riskier\_clmtimes$.

safety (e.g., unmeasured characteristics such as better overall handling characteristics).

5 Discussion

A large number of empirical studies have investigated the effect of improvements in safety technologies over a wide range of phenomenon. We have developed a model using a classification of such technologies based on whether adopting the technology leads, ceteris paribus, to a reduction in the probability of an accident or, conditional on an accident occurring, reduces the extent of the consequences or size of loss due to the accident. We refer to the former as a probability reduction technology (PRT) and the latter as a loss mitigation technology (LMT). Our model also considers two possible sources of heterogeneity among potential adopters of improved technologies. In one case we consider that individuals differ by their perceived loss due to an accident while in the other, some individuals display a higher cost of taking precautions to avoid accidents (i.e., effort to drive more safely). In order to understand the relative safely levels of drivers who end up adopting improved technologies compared to those who do not, both before and after adoption, one must understand the reason for adoption (i.e., the source of heterogeneity in preferences). We investigated these issues theoretically and discuss in what follows how to draw policy conclusions based on observations driven by the various possibilities. We also examine the challenges in interpreting data linking accident rates to adoption decisions both from a theoretical perspective and through our empirical application.

Vehicle owners (drivers) are assumed to differ according either to their perceived loss or concern with being involved in an accident or to their personal cost of taking preventive actions (i.e., the extent of safe driving habits). These two dimensions of the model help in unraveling the relationship between riskiness of adopters of the different types of safety technologies both before and after adoption of improved safety technologies. For example, individuals who perceive higher losses due to accidents will, at least ex ante to adoption, drive more carefully and so be less likely to be involved in accidents. Such drivers will more likely adopt either an LMT or PRT. If the extent of any offsetting behaviour is not too large, then adopters will continue to display lower accident risk levels ex post to adoption. However, individuals who possess a higher cost to safe driving behaviour will also be more likely to adopt either type of technology but have higher accident rates ex ante. These individuals who adopt an improved LMT will have further incentive to reduce their safe driving habits and so have an even higher accident rate ex post to adoption. The effect of adopting a higher quality PRT for either type of driver depends on whether the PRT is a substitute or complement to safe driving behaviour (i.e., the offsetting effect may be the typical one of reducing safe driving or have the opposite effect of increasing safe driving behaviour). Interpreting the relationship between accident rates and adoption of safety technologies - the so-called recruitment effect - requires careful analysis of the relationship between both ex ante and ex post accident rates of adopters versus non-adopters. It is important to understand these relationships in order to draw appropriate conclusions about the extent of offsetting effects from empirical analysis.

The classic interpretation of a negative externality arising from the offsetting effect due to adoption of an LMT (such as seatbelts) arises from the reasonable presumption that the inherent reduction in the negative consequences due to accidents reduces the marginal value to exerting safe driving behaviour. More care must be made when considering adoption of $PRTs$. If, for example, the source of a positive relationship between accident rates and adoption of a PRT is due to a preference by those with a higher cost of careful driving wanting to balance their higher risk of accident by use of the improved technology, then it does not necessarily follow that there will be a negative externality effect resulting from the improved technology despite the observed higher accident rate of adopters compared to nonadopters. If the PRT is complementary to individuals' own safe driving efforts, then there will be an even greater impact on overall safety even if the accident rates of adopters is observed to be higher than nonadopters. This will happen if the combined effect of the PRT and enhanced safe driving efforts does not make adopters accident rates fall below that of nonadopters which, given the assumption that adopters have a higher cost of precaution, is possible. Even if the PRT and safe driving efforts are substitutes, the net effect on adopters' accident rates may still lead to a reduction in the probability of them causing an accident. Therefore, despite a perception of an offsetting effecting, there may exist an overall positive externality created by such technologies and so, from a welfare perspective, such a technology should be subsidized in such cases. Of course, if the overall effect of adoption of a PRT and resultant change in driving behaviour leads to an increase in adopters' probability of causing an accident, then a tax on the technology is in order.

Suppose, on the other hand, that the reason for those who adopt a PRT is that they perceive a higher loss due to any accident that may occur. Such individuals would, ex ante to adoption, display lower accident rates than nonadopters. If the PRT is a complement to safe driving habits, then a reverse offsetting effect would occur and the accident rate of adopters would be even lower ex post to adoption. If the PRT and safe driving habits are substitutes, then the usual offsetting effect can be expected. However, the ex post accident rate for adopters may remain below that for nonadopters. This could be observed even if the offsetting effect leads to an increase in the accident rate of adopters provided the offsetting effect was not so strong as to lead to adopters to have a higher (ex post) accident rate than the (ex ante and ex post) accident rate of nonadopters. In this scenario, the

PRT creates a negative externality and so should be taxed even though adopters display a lower accident rate than that of nonadopters.

As is evident from the above discussion, as well as the formal propositions in this paper, one must take care in drawing conclusions from observations of accident rates and vehicle (safety) characteristics in regards to the presence and extent of offsetting behaviour in conjunction with recruitment effects as well as the type of heterogeneity of preferences that exists in the population of vehicle owners. This is crucial information in determining appropriate policy considerations in regards to the appropriate tax (or subsidy) to apply to safety technologies as well as deciding which technologies to make mandatory. As has been noted before (e.g., Harless and Hoffer, 2003), data which allows one to follow individuals' driving records and accident histories over time can be very useful in this regard. We have analyzed an unbalanced panel data set to illustrate how our theoretical analyses can guide one to understand better these important issues. Ignoring the panel nature of the data, simple correlations indicate a negative relationship between accident rates and both the adoption of higher quality airbags and higher quality brake systems. According to the classic offsetting hypothesis, adopting a higher quality LMT is expected to lead to a reduced level of safe driving care and so, ignoring possible recruitment effects, a positive relationship between the quality of the technology and accident rates. The observed negative correlation points towards advantageous recruitment (i.e., safer or less risky drivers choose the better technology).

Drawing conclusions about $PRTs$ is more complicated. Adoption of a higher quality PRT by its nature leads to a reduction in the probability of an accident provided there is no overwhelming offsetting effect. If the PRT is a complement to safe driving behaviour, then one expects a reverse offsetting effect which strengthens the negative relationship between the level of PRT and accident rate. However, if the reason for individuals purchasing a higher quality PRT is due to having a higher cost of (own) precaution, the recruitment effect may look different depending on whether contemporaneous or historical accident rates are being observed. From our empirical analysis, we find a negative correlation using either current or past measures of accident rates (riskiness of drivers) and so again the negative correlations point toward advantageous recruitment.

There are many challenges to any study about the effects of safety vehicles for vehicles that also apply to our work. As mentioned earlier in the paper, some safety features may display both LMT and PRT effects. This is likely for higher quality brake systems. Also, although our example of a LMT (high quality airbag) presumably decreases the harm to occupants in any substantial impact, minor accidents may involve higher financial losses for such vehicles as more complex airbag systems, if triggered unnecessarily, may be more expensive to reset. Thus, the implication of the classic offsetting hypothesis that such a safety device would incentivize drivers to be less cautious may in fact be incorrect.

Another important challenge is to consider the role of insurance and traffic enforcement. Safer vehicles may be less expensive to insure and this provides an incentive to purchase such vehicles and so at least in part behaves as an appropriate subsidy, albeit not necessarily in a complete manner. Also, insurers may use experience rating in a way that lessens moral hazard, for example, by people who purchase vehicles with enhance $LMTs$. To be fully effective in welfare terms, however, such experience rating would need to be designed according to vehicle and driver types. Also, traffic enforcement is not likely to involve policies, including fines, which differ according to safety features of vehicles.

More generally, when we refer to the individual's level of precaution we mean things such as attentiveness to road hazards while driving, maintaining alertness, driving at safe speeds, and so forth. These are assumed unobservable to the social planner (or government). Our analysis is designed to consider how such choices create externalities for others under various scenarios of available PRT and LMT technologies and for individuals with two possible sources of heterogeneous preferences which lead them to value such technologies differently. Although certainly worthy of future research, we do not consider the may direct and indirect measures used for imperfectly observing (and controlling) individual choices of level of care or precaution. These include police enforcement of traffic regulations (fines for speeding, following too closely, etc.) and measures such as experience rating by insurers, that others have studied (e.g., Boyer and Dionne, 1987). We leave aside these sorts of issues, although they are all well worth exploring in future work.

Although our model advances the literature by allowing for two dimensions of preferences (cost of own effort towards safe driving and size of loss due to an accident), there are many more possible dimensions that one could explore. Some of these may be approximated reasonably well by our chosen dimensions, but others deserve greater attention. For example, our objective function implies risk neutrality. However, allowing for individuals to vary in their perceived size of loss due to an accident may approximate a difference in risk aversion with more risk averse individuals holding a higher degree of loss. Admittedly, though, the implications of risk aversion on choice of self-protection or level of safety technologies is a complicated matter. The difficulty of determining the effect of varying the degree of risk aversion on the optimal level of precaution is well known. It would also be difficult to determine the effect of risk aversion on choice of a PRT. Since increasing the degree of LMT reduces the size of loss and so increases income in the loss state of the world, this may pose less problematic. Allowing for differences in income levels would also create complications in our model.[14] We also do not explore the possibility of

[14] Our data set does not have information on income levels.

innate differences in driving ability that may be reflected in the probability of an accident occurring (i.e., the $D(p, \theta)$ function).

There are many other alternative assumptions one could make about preferences. We implicitly assume risk neutral expected utility preference. Many alternative behavioural models could of course also be explored. Our model does, however, allow for weighted probabilities. The factor $(1 + v)$ in the expected loss term, $D(p, \theta)(1 + v)L(\lambda)$, which reflects a multiplicative term on the size of the loss could also be treated as a weighting factor on the probability of the loss state. Other more heuristic models have also been suggested.

Another important consideration is whether individuals are well informed about the relative safety features of different vehicles. There are many such features to understand and trade off between models. Examples include visibility, handling, crash worthiness, relative effectiveness of the myriad safety features (including so-called nanny devices) that can be purchased between models of a given brand of vehicle and between brands. Moreover, people may consider other features of a vehicle important that may have to be traded off with safety features, such as storage compartments, comfort of seats, quality of sound system, etc..

Finally, it can be difficult to assess the extent to which a feature is advantageous in preventing high loss accidents. A good example is the decision to purchase a SUV. Although its size is an advantage in reducing the extent of harm to occupants should an accident occur, the size may be a disadvantage in avoiding an accident in the first place. Being both larger and having a higher centre of gravity implies a higher rollover risk as well as a longer stopping distance.

6 Conclusions

We have developed a model of decision making by owners/drivers of vehicles that allows for two sources of heterogeneity in preferences as potential reasons why people purchase vehicles with differing quality safety features. We also explicitly introduce two types of such safety features. One type, such as high quality airbag systems, offer greater protection against harms to individuals should an accident occur while the other, such as high quality brake systems, offer intrinsic reduction in the probability of being involved in an accident. We refer to these, respectively, as loss mitigation technologies (LMT) and probability reduction technologies (PRT). We show that the demand for these two types of technologies and the implications on the relationship between their adoption and accident probabilities both ex ante to adoption (recruitment effects) and ex post (recruitment plus offsetting behaviour effects) differs in interesting ways. We believe our model could, with extensions, be useful for studying the myriad of newly developed safety technologies for vehicles as well as in other domains involving changes to safety protocols and technology.

Using data from the Taiwan Insurance Institute (TII), supplemented with detailed information on insureds' claims and driving records, we illustrate our model with an empirical application involving these two types of safety features; i.e., quality of airbag systems (a LMT) and quality of braking systems (a PRT). Both simple correlations and cross-sectional regressions generated a negative statistical relationship between accident claims caused by drivers of vehicles and high-quality airbags or high-quality brake systems of those vehicles.[15] Consider first the case of airbags. Although causality cannot be inferred from these results, they are at least consistent with advantageous recruitment (i.e., less risky drivers are more likely to obtain vehicles with higher quality airbags). Given the classical offsetting effect due to adoption of an LMT, which would generate a positive relationship between the safety feature of high-quality airbag and level of safe driving, this result in principle makes advantageous selection a plausible conclusion. Our regressions based on the subset of owners present in both periods reveal a negative statistical relationship between drivers becoming riskier and purchase of a new vehicle with lower quality airbags which is consistent with the classic offsetting hypothesis (for a LMT). There is, however, no complementary effect for drivers who have purchased new vehicles that have upgraded quality of airbags (i.e., there is no statistically significant relationship between inc_arbg and increased claims experience).

Our results point in the direction of advantageous recruitment for both high-quality airbags and high-quality brake systems, the LMT and PRT investigated here, and that any offsetting effect from the adoption of high-quality brake systems is not strong enough to reverse the inherent improvement in the accident rate due to the nature of the PRT. On the basis of this finding, one can make the case that a subsidy on this PRT would improve welfare.

7 Appendix

We now consider the scenario in which each individual chooses simultaneously his level of precaution, PRT and LMT. Given what we have learned for the case of being able to choose only one of PRT and LMT (i.e., singly), it is not surprising that performing comparative statics leads in many cases to ambiguous results. For the case of heterogenous cost of precaution (Model C1), we have that each individual chooses $\{p, \lambda, \theta\}$ to minimize

$$\Omega(p, \lambda, \theta) = D(p, \theta)L(\lambda) + (1 + \tau)c(p) + k_R(\theta) + k_M(\lambda) \tag{49}$$

[15] This is the case both for contemporaneous claims and historical claims as measured by the bonus malus measure.

For convenience, we assign variable numbers 1, 2, 3, to p, λ, θ, respectively, and so write the first-order conditions for the optimization problem as follows.

$$F_1(p, \lambda, \theta) = D_p L + (1+\tau)c' = 0 \tag{50}$$

$$F_2(p, \lambda, \theta) = DL_\lambda + k'_M = 0 \tag{51}$$

$$F_3(p, \lambda, \theta) = D_\theta L + k'_R = 0 \tag{52}$$

Upon totally differentiating the above system we get, using standard notation,

$$\begin{bmatrix} F_{11} & F_{12} & F_{13} \\ F_{21} & F_{22} & F_{23} \\ F_{31} & F_{32} & F_{33} \end{bmatrix} \begin{bmatrix} dp/d\tau \\ d\lambda/d\tau \\ d\theta/d\tau \end{bmatrix} = \begin{bmatrix} -c'(p) \\ 0 \\ 0 \end{bmatrix} \tag{53}$$

where

$$F_{11} = D_{pp}L + (1+\tau)c'' > 0, \ F_{12} = D_p L_\lambda > 0, \ F_{13} = D_{p\theta}L? \tag{54}$$

$$F_{21} = D_p L_\lambda > 0, \ F_{22} = DL_{\lambda\lambda} + k''_M > 0, \ F_{23} = D_\theta L_\lambda > 0 \tag{55}$$

$$F_{31} = D_{p\theta}L?, \ F_{32} = D_\theta L_\lambda > 0, \ F_{33} = D_{\theta\theta}L + k''_R > 0 \tag{56}$$

Note that the sign of $F_{13}(F_{31})$, indicated by ?, is the same as the sign of $D_{p\theta}$ and so depends on whether precaution and the PRT are substitutes or complements. From the above, we have the following comparative statics results. Again, $|F| > 0$ and so signs are the same as the signs of the numerators.

$$\frac{d\lambda}{d\tau} = c'[(D_p L_\lambda)(D_{\theta\theta}L + k''_R) - (D_{p\theta}L)(D_\theta L_\lambda)]/\,|F| \tag{57}$$

The term $(D_p L_\lambda)(D_{\theta\theta}L + k''_R) > 0$ contributes to a positive relationship between λ and τ. It follows that $\frac{d\lambda}{d\tau} > 0$ if $D_{p\theta} \leq 0$ (i.e., if own care and the PRT are complements). This follows since a higher cost of p reduces the incentive to provide own care which in turn reduces the effectiveness of θ (the PRT) when they are complements. Lowering both p and θ leads to an increase in the probability of loss (D) which in turn increases the marginal value of the LMT and so any reduction in θ (in addition to a reduction in p) reinforces the incentive to increase λ. However, if precaution and the PRT are substitutes, then a lower choice of p due to a higher cost would lead to a higher productivity of θ which would lead to a reduction in the loss probability. This is demonstrated by the following result and explanation.

$$\frac{d\theta}{d\tau} = -c'[(D_p L_\lambda)(D_\theta L_\lambda) - (D_{p\theta}L)(DL_{\lambda\lambda} + k''_M)]/\,|F| > 0 \tag{58}$$

The second term in brackets represents a positive effect of an increase in τ on θ when p and θ are substitutes ($D_{p\theta} > 0$). This accords with intuition since in this case any reduction in (more costly) p makes θ (PRT) more effective in reducing the probability of loss. If p and θ are complements ($D_{p\theta} < 0$), then any reduction in p reduces the effectiveness of θ and so in that case the second term represents a negative effect of an increase in τ on θ. Note that any increase in the probability of loss due to either a decrease in p or θ would increase the marginal value of the LMT. Given these instruments (λ and θ) are chosen simultaneously, an induced increase in λ would reduce the size of the loss and so have a a negative effect on the marginal productivity of θ. The first term in square brackets, $(D_p L_\lambda)(D_\theta L_\lambda)$, is positive and so captures this negative effect of an increase of τ on θ. Therefore, the net effect of an increase in τ on θ depends on the relative strength of all of these effects. Notice that this second unambiguously positive effect is stronger the higher is the effect of increasing λ on the size of loss (i.e. on the magnitude of $|L_\lambda|$) and in fact disappears as $L_\lambda \to 0$ which corresponds to the results when θ is the only choice variable.

$$\frac{dp}{d\tau} = -c'[(DL_{\lambda\lambda} + k''_M)(D_{\theta\theta}L + k''_R) - (D_\theta L_\lambda)^2]/\,|F| > 0 \tag{59}$$

The part $(DL_{\lambda\lambda} + k''_M)(D_{\theta\theta}L + k''_R)$ (in square brackets) is positive and contributes to a negative relationship between τ and p, the effect one would expect from simply having the cost of own care increasing in τ. However, the term $-(D_\theta L_\lambda)^2$ reduces this effect and, if strong enough, may even lead to a positive relationship between τ and p.

Proposition 7. *Suppose individuals differ according to cost of precaution and choose (simultaneously) levels of PRT (θ) and LMT (λ) along with their level of precaution (p) to minimize expected loss. Individuals who face higher cost of precaution increase their level of LMT if the PRT is a complement to precaution (i.e., $D_{p\theta} < 0$). The effect is indeterminate if precaution and the PRT are substitutes ($D_{p\theta} > 0$). The relationship between cost of precaution and the other variables of interest (precaution, p, and the PRT,θ) are indeterminate.*

We now develop Model 3B in which the heterogeneity is due to differential size of loss should an accident occur. Recall that the objective function is

$$\Omega(p, \lambda, \theta) = D(p, \theta)(1 + \upsilon)L(\lambda) + c(p) + k_R(\theta) + k_M(\lambda) \tag{60}$$

and so the first-order conditions for the optimization problem are as follows.

$$F_1(p, \lambda, \theta) = D_p(1 + \upsilon)L + c' = 0 \tag{61}$$

$$F_2(p, \lambda, \theta) = D(1 + \upsilon)L_\lambda + k'_M = 0 \tag{62}$$

$$F_3(p, \lambda, \theta) = D_\theta(1 + \upsilon)L + k'_R = 0 \tag{63}$$

Upon totally differentiating the above system we get, using standard notation,

$$\begin{bmatrix} F_{11} & F_{12} & F_{13} \\ F_{21} & F_{22} & F_{23} \\ F_{31} & F_{32} & F_{33} \end{bmatrix} \begin{bmatrix} dp/d\upsilon \\ d\lambda/d\upsilon \\ d\theta/d\upsilon \end{bmatrix} = \begin{bmatrix} -D_p L \\ -DL_\lambda \\ -D_\theta L \end{bmatrix} \tag{64}$$

where

$$F_{11} = D_{pp}(1+\upsilon)L + c'' > 0,\ F_{12} = D_p(1+\upsilon)L_\lambda > 0,\ F_{13} = D_{p\theta}(1+\upsilon)L? \tag{65}$$

$$F_{21} = D_p(1+\upsilon)L_\lambda > 0,\ F_{22} = D(1+\upsilon)L_{\lambda\lambda} + k''_M > 0,\ F_{23} = D_\theta(1+\upsilon)L_\lambda > 0 \tag{66}$$

$$F_{31} = D_{p\theta}(1+\upsilon)L?,\ F_{32} = D_\theta(1+\upsilon)L_\lambda > 0,\ F_{33} = D_{\theta\theta}(1+\upsilon)L + k''_R > 0 \tag{67}$$

From the above, we have the following comparative statics results (noting that $|F| > 0$).

$$\frac{dp}{d\upsilon} = \left\{-D_p L[F_{22}F_{33} - (F_{23})^2] + DL_\lambda[F_{12}F_{33} - F_{13}F_{32}] - D_\theta L[F_{12}F_{23} - F_{22}F_{13}]\right\} / |F| \tag{68}$$

$$\frac{d\lambda}{d\upsilon} = \left\{D_p L[F_{21}F_{33} - F_{23}F_{31}] - DL_\lambda[F_{11}F_{33} - (F_{13})^2] + D_\theta L[F_{11}F_{23} - F_{21}F_{13}]\right\} / |F| \tag{69}$$

$$\frac{d\theta}{d\upsilon} = \left\{-D_p L[F_{21}F_{33} - F_{22}F_{31}] + DL_\lambda[F_{11}F_{32} - F_{12}F_{31}] - D_\theta L[F_{11}F_{22} - (F_{12})^2]\right\} / |F| \tag{70}$$

An increase in parameter ν leads to increased productivity of each choice variable. However, an increase in λ would reduce the marginal productivity of each of the other choice variables, (p, θ). Similarly, an increase in either p or θ would reduced the marginal productivity of λ. Finally, an increase in θ would decrease or increase the marginal productivity of p (or vice versa) depending on whether the two variables are substitutes or complements (i.e., whether $D_{p\theta}$ is positive or negative, respectively). Given these relationships, none of the comparative statics results can be signed definitively even if we make an assumption about the sign of $D_{p\theta}$.

Proposition 8. *Suppose individuals differ according to size of loss and choose (simultaneously) levels of PRT (θ) and LMT (λ) along with their level of precaution (p) to minimize expected loss. An increase in the loss size parameter (υ) increases the marginal productivity of each choice variable. However, any increase in λ reduces the marginal productivity of both precaution and the PRT. Moreover, any increase in one of precaution or the PRT increases or decreases the marginal value of the other depending on whether they are complements or substitutes. As a result, none of the signs of the comparative static relationships are determinate.*

8 References

Arnott, A. and J. Stiglitz (1986): "Moral Hazard and Optimal Commodity Taxation", *Journal of Public Economics*, vol. 29, pp. 1-24.

Boyer, M. and G. Dionne (1987): "The Economics of Road Safety", *Transportation Research B*, vol. 21B, pp. 413-31.

Blomquist, G. (1986): "A utility Maximization Model of Driver Traffic Safety Behavior", *Accident Analysis and Prevention*, vol. 18(5), pp. 371-375.

Cooper, R. and T. W. Ross (1985): "Product Warranties and Double Moral Hazrd", *Rand Journal of Economics*, pp. 103-113.

Diamond, P. (1974): "Single Activity Accidents", *Journal of Legal Studies*, vol. 3, pp. 107-164.

Ehrlich, I., and G. Becker (1972): "Market Insurance, Self-insurance, adn Self-protection", *Journal of Political Economy*, vol. 80, no. 4, pp. 623-48.

Gossner, O. and P. Picard (2005): "On the Consequences of Behavioral Adaptations in the Cost-Benefit Analysis of Road Safety Measures", *Journal of Risk and Insurance*, vol. 72, no. 4, pp. 557-599.

Harless, D. W., and G. E. Hoffer (2003): "Testing for Offsetting Behavior and Adverse Recruitment Among Drivers of Airbag-Equipped Vehicles", *Journal of Risk and Insurance*, vol. 70, no. 4, pp. 629-650.

Hause, J. C. (2006): "Offsetting Behavior and the Benefits of Safety Regulations", *Economic Inquiry*, vol. 44(4), pp. 689-698.

Holmstrom, B. (1982): "Moral Hazard in Teams", *Bell Journal of Economics*, vol. 13, pp. 324-340.

Lanoie, P. (1991): "Occupational Safety and Health: a Problem of Double or Single Moral Hazard", *Journal of Risk and Insurance*, vol. 58, no. 1, pp. 80-100.

Neill, J. R. (1993): "A Theoretical Reappraisal of the Offsetting Behavior Hypothesis", J*ournal of Regulatory Economics*, vol. 5, pp. 435-440.

P. A. Pedersen (2003): "Moral Hazard in Traffic Games", *Journal of Transport Economics and Policy*, vol. 37, pt.1, pp. 47-68.

S. Peltzman (1975): "The Effects of Automobile Safety Regulation", *Journal of Political Economy*, vol. 83, no. 4, pp. 677-725.

Risa, A. (1992): "Public Regulation of Private Accident Risk: The Moral hazard of Technological Improvements", *Journal of Regulatory Economics*, vol. 4, no. 4, pp. 335-346.

A. E. Risa (1995): "The Welfare State as Provider of Accident Insurance in the Workplace: Efficiency and Distribution in Equilibrium", *The Economic Journal*, vol. 105, no. 428, pp. 129-44.

S. Shavell (1982): "On Liability and Insurance", *Bell Journal of Economics*, vol. 13, no.

1, pp. 120-132.

Winston, C., Maheshri, V. and F. Mannering (2006), “An exploration of the offset hypothesis using disaggregate data: The case of airbags and antilock brakes,” *Journal of Risk and Uncertainty*, vol 32, pp. 83–99.

Acknowledgements

We would like to thank seminar participants at the American Risk and Insurance Meetings (20XX) and at the University of Montreal (20XX). In particular, we thank Marcel Boyer, Martin Boyer, and Georges Dionne for insightful comments.

Table 1 Variable definitions: Set 1

Variables	Definition
claim	A dummy variable, it equals 1 when the insured have ever filed the claim in compulsory liability insurance within one policy year; otherwise it equals 0.
bm	The value of bonus malus coefficient
Dbm	A dummy variable, it equals 1 when the bonus malus coefficient of the insured is larger than 0.7; otherwise it equals 0.
brake_high	A dummy variable, it equals 1 when the insured vehicle is equipped with high quality brake system; otherwise it equals 0. The high quality brake system means the vehicle is equipped not only with anti-lock brake system, but also equipped with the traction control system\vehicle stability control system\acceleration slip regulation as well as down-hill assist control \the hill-start assist control.
airbag_high	A dummy variable, it equals 1 when the insured vehicle is equipped with high quality airbag system; otherwise it equals 0. The high quality airbag system means there are airbags equipped for both front seats and equipped for the back seats.
high_brk_high_abg	A dummy variable, it equals 1 when the insured vehicle is equipped with high quality brake system as well as high quality airbag system; otherwise it equals 0.
high_brk_low_abg	A dummy variable, it equals 1 when the insured vehicle is equipped with high quality brake system, but equipped with low standard airbag system; otherwise it equals 0.
low_brk_high_abg	A dummy variable, it equals 1 when the insured vehicle is equipped with low standard brake system, but equipped with high quality airbag system; otherwise it equals 0.
low_brk_low_abg	A dummy variable, it equals 1 when the insured vehicle is equipped with low standard brake system as well as low standard airbag system; otherwise it equals 0.
veh_suv	A dummy variable, it equals 1 when the insured vehicle is a sport utility vehicle (SUV); otherwise it equals 0.

Table 1 Variable definitions: Set 1 (continued)

Variables	Definition
female	A dummy variable, it equals 1 when the insured is female; otherwise it equals 0.
married	A dummy variable, it equals 1 when the insured is in marriage status; otherwise it equals 0.
age2530	A dummy variable, it equals 1 when the insured equals or older than 25 years old and younger than 30 years old; otherwise it equals 0.
age3060	A dummy variable, it equals 1 when the insured equals or older than 30 years old and younger than 60 years old; otherwise it equals 0.
ageabv60	A dummy variable, it equals 1 when the insured equals or older than 60 years old; otherwise it equals 0.
carage0	A dummy variable, it equals 1 when the insured vehicle is brand new; otherwise it equals 0.
carage1to4	A dummy variable, it equals 1 when the insured vehicle is more than 1 year and not over 4 years; otherwise it equals 0.
city	A dummy variable, it equals 1 when the insured vehicle is registered in city area; otherwise it equals 0.
north	A dummy variable, it equals 1 when the insured vehicle is registered in northern part of Taiwan; otherwise it equals 0.
south	A dummy variable, it equals 1 when the insured vehicle is registered in southern part of Taiwan; otherwise it equals 0.
east	A dummy variable, it equals 1 when the insured vehicle is registered in eastern part of Taiwan; otherwise it equals 0.

Table 2 Summary statistics

	Mean	Std	Obs
Panel A Whole sample			
claim	0.0100	0.0995	2255157
Dbm	0.2004	0.4003	2255157
airbag_high	0.0952	0.2935	2255157
brake_high	0.3939	0.4886	2255157
veh_suv	0.0771	0.2667	2255157
high_brk_high_abg	0.0012	0.0345	2255157
high_brk_low_abg	0.3927	0.4883	2255157
low_brk_high_abg	0.0940	0.2918	2255157
low_brk_low_abg	0.5122	0.4999	2255157
female	0.6008	0.4897	2255157
married	0.7618	0.4260	2255157
age2530	0.0440	0.2050	2255157
age3060	0.8306	0.3751	2255157
ageabv60	0.1144	0.3183	2255157
carage0	0.0668	0.2497	2255157
carage1to4	0.3227	0.4675	2255157
city	0.6850	0.4645	2255157
north	0.4405	0.4964	2255157
south	0.3040	0.4600	2255157
east	0.0421	0.2009	2255157
year2011	0.4724	0.4992	2255157

Table 2 Summary statistics (continued)

	Mean	Std	Obs
Panel B Sample with voluntary third party bodily injury insurance			
claim	0.0110	0.1045	1286309
Dbm	0.1975	0.3981	1286309
airbag_high	0.0909	0.2874	1286309
brake_high	0.3791	0.4852	1286309
veh_suv	0.0957	0.2941	1286309
high_brk_high_abg	0.0014	0.0374	1286309
high_brk_low_abg	0.3777	0.4848	1286309
low_brk_high_abg	0.0895	0.2854	1286309
low_brk_low_abg	0.5315	0.4990	1286309
female	0.6583	0.4743	1286309
married	0.7533	0.4311	1286309
age2530	0.0386	0.1926	1286309
age3060	0.8485	0.3586	1286309
ageabv60	0.1055	0.3072	1286309
carage0	0.0786	0.2691	1286309
carage1to4	0.3675	0.4821	1286309
city	0.6872	0.4636	1286309
north	0.4396	0.4963	1286309
south	0.3042	0.4601	1286309
east	0.0449	0.2070	1286309
year2011	0.4600	0.4984	1286309

Table 2 **Summary statistics (continued)**

	Mean	**Std**	**Obs**
Panel C Sample without voluntary third party bodily injury insurance			
claim	0.0086	0.0926	968848
Dbm	0.2042	0.4031	968848
airbag_high	0.1009	0.3012	968848
brake_high	0.4135	0.4925	968848
veh_suv	0.0524	0.2229	968848
high_brk_high_abg	0.0009	0.0302	968848
high_brk_low_abg	0.4126	0.4923	968848
low_brk_high_abg	0.1000	0.3000	968848
low_brk_low_abg	0.4865	0.4998	968848
female	0.5246	0.4994	968848
married	0.7731	0.4188	968848
age2530	0.0511	0.2202	968848
age3060	0.8069	0.3947	968848
ageabv60	0.1262	0.3320	968848
carage0	0.0511	0.2202	968848
carage1to4	0.2633	0.4404	968848
city	0.6821	0.4657	968848
north	0.4416	0.4966	968848
south	0.3036	0.4598	968848
east	0.0386	0.1925	968848
year2011	0.4887	0.4999	968848

Table 3 Correlation coefficients

	claim	Dbm	airbag_high	brake_high	veh_suv
claim	1.000				
Dbm	0.012***	1.000			
airbag_high	-0.005***	-0.004***	1.000		
brake_high	-0.004***	-0.072***	-0.253***	1.000	
veh_suv	0.0004	0.019***	0.036***	-0.144***	1.000

Table 4 Pooled Probit regression of compulsory liability claim

	Whole sample		With voluntary third party bodily injury insurance		Without voluntary third party bodily injury insurance	
	Est.	P value	Est.	P value	Est.	P value
Intercept	-4.7884	<.0001	-4.4624	<.0001	-5.1590	<.0001
airbag_high	-0.1805	<.0001	-0.1693	<.0001	-0.1643	<.0001
brake_high	-0.0540	0.0003	-0.0341	0.0681	-0.0866	0.0004
veh_suv	-0.0073	0.7748	-0.0329	0.2709	-0.0163	0.7348
bm	0.8938	<.0001	0.8982	<.0001	0.9788	<.0001
female	0.1328	<.0001	0.0563	0.0020	0.1750	<.0001
married	-0.1242	<.0001	-0.1592	<.0001	-0.0608	0.0230
age2530	-0.2767	<.0001	-0.3356	<.0001	-0.2875	0.0004
age3060	-0.3357	<.0001	-0.4301	<.0001	-0.3382	<.0001
ageabv60	-0.3321	<.0001	-0.3906	<.0001	-0.3768	<.0001
carage0	0.1226	<.0001	-0.0092	0.7863	0.2900	<.0001
carage1to4	0.1631	<.0001	0.0288	0.1319	0.3370	<.0001
city	0.0660	<.0001	0.0698	0.0004	0.0558	0.0253
north	-0.4957	<.0001	-0.5021	<.0001	-0.4619	<.0001
south	-0.0498	0.0039	-0.0728	0.0009	-0.0125	0.6587
east	-0.1529	<.0001	-0.1826	<.0001	-0.1057	0.0823
year2011	0.0152	0.2568	-0.0114	0.5006	0.0902	<.0001
-2LogL	250843.91		155185.99		95213.045	
Observations	2255157		1286309		968848	

Table 5 Overview of safety technology decisions (New vehicle purchases)#

	suv	suv+ABS	small	small+ABS
increase airbag	17	0	566	1
decrease airbag	4	0	414	0
	suv	suv+airbag	small	small+airbag
increase brake	13	0	1612	1
decrease brake	3	0	1912	0
no change	9754			

#There are 2,706 observations not accounted for in this table. They are distributed into many descriptive cells, too numerous to include here.

Table 6 Variable definitions: Set 2

Variables	**Definition**
delta_clm	equals *clm_2* minus *clm_1* (*clm_1* and *clm_2* are the dummy variables which represent whether there is claim filed in first or second year)
delta_clmtimes	equals *clmtimes_2* minus *clmtimes_1* (*clmtimes_1* and *clmtimes_2* represents the claim times in first year or in second year)
delta_clmamt	equals *clmamt_2* minus *clmamt_1* (*clmamt_1* and *clmamt_2* represents the claim amount in first year or in second year)
riskier_clm	A dummy variable which equals 1 if *delta_clm*>0, otherwise 0.
riskier_clmtimes	A dummy variable which equals 1 if *delta_clmtimes*>0, otherwise 0.
riskier_clmamt	A dummy variable which equals 1 if *delta_clmamt*>0, otherwise 0.
lessrisky_X	A dummy variable which equals 1 if *delta_X*<0, otherwise zero – for *X* = *clm, clmtimes, clmamt*
inc_brk	A dummy variable, it equals 1 when the car owner switched vehicle from a low quality brake system to a high quality brake system; otherwise it equals 0.
dec_brk	A dummy variable, it equals 1 when the car owner switched vehicle from a high quality brake system to a low quality brake system; otherwise it equals 0.
inc_arbg	A dummy variable, it equals 1 when the car owner switched vehicle from a low quality airbag system to a high quality airbag system; otherwise it equals 0.
dec_arbg	A dummy variable, it equals 1 when the car owner switched vehicle from a high quality airbag system to a low quality airbag system; otherwise it equals 0.
inc_s	A dummy variable, it equals 1 when the car owner switched vehicle to a sport utility vehicle (SUV); otherwise it equals 0.
dec_s	A dummy variable, it equals 1 when the car owner switched vehicle from a sport utility vehicle (SUV) to other type of vehicle; otherwise it equals 0.

*inc_brk*inc_s*	Interaction term = 1 if new vehicle has *inc_brk* and *inc_s*, otherwise = 0.
*inc/dec_X*inc/dec_s*	Completes the set of interaction terms as described above depending on increase or decrease either brk or arbg along with increase or decrease s
delta_carage	A variable which equals the age of the new vehicle (in year 2012) minus the age of the old vehicle (in year 2011).
deltam_brk	A variable which equals the mean value of high quality brake system vehicles in the registration administrative area corresponding to each vehicle in year 2012 minus the mean value of high quality brake system vehicles in the registration administrative area corresponding to each vehicle in year 2011.
deltam_arbg	A variable which equals the mean value of high quality airbag system vehicles in the registration administrative area corresponding to each vehicle in year 2012 minus the mean value of high quality airbag system vehicles in the registration administrative area corresponding to each vehicle in year 2011.
deltam_s	A variable which equals the mean value of sport utility vehicles (SUV) in the registration county corresponding to each vehicle in year 2012 minus the mean value of sport utility vehicles (SUV) in the registration county corresponding to each vehicle in year 2011.
deltam_tkt	A variable which equals the mean value of the number of tickets in the registration county corresponding to each vehicle in year 2012 minus the mean value of the number of tickets in the registration county corresponding to each vehicle in year 2011.
deltam_clm	A variable which equals the mean value of claim probability in the registration county corresponding to each vehicle in year 2012 minus the mean value of claim probability in the registration county corresponding to each vehicle in year 2011.

Table 7 Detailed Summary Statistics on Panel Data Set (New vehicle purchases)

	Mean	StD
delta_clm	0.0036	0.1401
delta_clmtimes	0.0034	0.1432
delta_clmamt	86.3870	61184.3400
riskier_clm	0.0116	0.1073
riskier_clmtimes	0.0116	0.1073
riskier_clmamt	0.0118	0.1081
increase_brake	0.1343	0.3410
decrease_brake	0.1781	0.3826
increase_airbag	0.0918	0.2887
decrease_airbag	0.0548	0.2275
increase_size	0.0565	0.2308
decrease_size	0.0225	0.1482
delta_carage	-2.2557	6.1281
female	0.5722	0.4948
married	0.7355	0.4411
age2530	0.0503	0.2187
age3060	0.8348	0.3713
ageabv60	0.1010	0.3013
carage0	0.1800	0.3842
carage1to4	0.3824	0.4860
city	0.6825	0.4655
north	0.4505	0.4976
south	0.2987	0.4577
east	0.0383	0.1919
delta_mean_ABS	0.0114	0.0105
delta_mean_airbag	0.0068	0.0035
delta_mean_suv	0.0046	0.0141
delta_mean_ticket	-0.0002	0.0462
delta_mean_clm	-0.0001	0.0014

Table 8 Logistic Regression (*riskier_clm/clmtimes/clmamt*)

	Model 1		Model 2	
	coef	P value	coef	P value
Intercept	-4.4502	<.0001	-114.6000	0.4636
inc_brk	0.0020	0.9935	0.0084	0.9731
dec_brk	-0.0357	0.8741	-0.0390	0.8628
inc_arbg	0.1114	0.8098	0.1154	0.8031
dec_arbg	-1.0528	0.0419	-1.0640	0.0399
inc_s	-0.0179	0.9763	-0.0204	0.9731
dec_s	0.5123	0.3447	0.5137	0.3435
inc_brk*inc_s	-0.1177	0.9165	-0.1180	0.9163
inc_brk*dec_s	-0.3477	0.7014	-0.3563	0.6944
dec_brk*inc_s	-0.7644	0.3740	-0.7590	0.3775
dec_brk*dec_s	0.9637	0.4112	0.9622	0.4120
inc_arbg*inc_s	-0.9362	0.4043	-0.9435	0.4006
inc_arbg*dec_s	-12.2443	0.9792	-12.2417	0.9792
dec_arbg*inc_s	1.3452	0.2689	1.3459	0.2686
dec_arbg*dec_s	1.5811	0.1947	1.6003	0.1895
delta_carage	-0.0160	0.2463	-0.0159	0.2498
female	0.0345	0.8153	0.0359	0.8083
married	-0.0721	0.6602	-0.0776	0.6370
age2530	0.4819	0.5307	0.4754	0.5364
age3060	0.3032	0.6736	0.3050	0.6718
ageabv60	0.7031	0.3428	0.7059	0.3409
carage0	0.1855	0.4266	0.1922	0.4111
carage1to4	0.0961	0.5972	0.0980	0.5905
city	-0.0379	0.8168	0.0128	0.9466
north	-0.5056	0.0525	4.8039	0.5231
south	0.1094	0.6906	-4.6704	0.4875
east	-1.3178	0.0409	-2.1767	0.1228
deltam_brk	-1.3960	0.9111	1.0591	0.9374
deltam_arbg	-24.9420	0.5106	-15.5020	0.6947
deltam_s	1.7572	0.8601	3.0193	0.7748
deltam_tkt	1.6056	0.5267	1.6441	0.5387
deltam_clm	21.1059	0.7637	14.1497	0.8588
airbag_high_2	-0.1762	0.6416	-0.1784	0.6376

brake_high_2	0.1008	0.5996	0.0997	0.6038
veh_suv2	0.0920	0.8186	0.1011	0.8011
mean_brk 2			279.4000	0.4790
mean_arbg2			0.0000	.
mean_suv 2			-4.4810	0.5986
mean_tkt2			0.4901	0.8459
mean_clm2			9.4466	0.9117

Table 9 Logistic Regression (*lessrisky_clm/clmtimes/clmamt*)

	lessrisky_clm		lessrisky _clmtimes		lessrisky _clmamt	
	coef	P value	coef	P value	coef	P value
Intercept	-3.9595	<.0001	-3.9595	<.0001	-3.9614	<.0001
inc_brk	0.4300	0.0786	0.4300	0.0786	0.4851	0.0440
dec_brk	-0.0178	0.9462	-0.0178	0.9462	-0.0188	0.9432
inc_arbg	0.4494	0.1306	0.4494	0.1306	0.4586	0.1227
dec_arbg	-0.4896	0.2994	-0.4896	0.2994	-0.5199	0.2701
inc_s	-1.3866	0.1976	-1.3866	0.1976	-0.9595	0.2663
dec_s	-0.4842	0.6326	-0.4842	0.6326	0.2327	0.7482
inc_brk*inc_s	-12.7069	0.9848	-12.7069	0.9848	-13.0072	0.9849
inc_brk*dec_s	-0.0675	0.9627	-0.0675	0.9627	-0.8231	0.5123
dec_brk*inc_s	1.0315	0.4117	1.0315	0.4117	0.2878	0.7817
dec_brk*dec_s	-11.9444	0.9930	-11.9444	0.9930	-12.6209	0.9923
inc_arbg*inc_s	-0.1694	0.8941	-0.1694	0.8941	0.5351	0.6124
inc_arbg*dec_s	-12.8803	0.9896	-12.8803	0.9896	-13.5697	0.9889
dec_arbg*inc_s	2.9505	0.0567	2.9505	0.0567	2.5654	0.0685
dec_arbg*dec_s	-12.1934	0.9917	-12.1934	0.9917	-12.5094	0.9913
delta_carage	0.0323	0.0660	0.0323	0.0660	0.0311	0.0730
female	0.1840	0.3177	0.1840	0.3177	0.2202	0.2287
married	-0.1141	0.5676	-0.1141	0.5676	-0.1545	0.4297
age2530	-1.4930	0.0150	-1.4930	0.0150	-1.4943	0.0148
age3060	-1.0563	0.0157	-1.0563	0.0157	-1.0394	0.0174
ageabv60	-1.4403	0.0080	-1.4403	0.0080	-1.3237	0.0130
carage0	0.9162	0.0006	0.9162	0.0006	0.8640	0.0011
carage1to4	0.1536	0.4855	0.1536	0.4855	0.1541	0.4791
city	-0.0788	0.7023	-0.0788	0.7023	-0.0527	0.7968
north	-0.0221	0.9386	-0.0221	0.9386	-0.0350	0.9021
south	0.3672	0.2075	0.3672	0.2075	0.3439	0.2333
east	-0.0244	0.9587	-0.0244	0.9587	-0.0450	0.9238
deltam_brk	-27.1920	0.0619	-27.1920	0.0619	-26.9912	0.0604
deltam_arbg	44.5352	0.2072	44.5352	0.2072	42.5135	0.2275
deltam_s	-30.1583	0.0050	-30.1583	0.0050	-29.1234	0.0064
deltam_tkt	3.5103	0.2190	3.5103	0.2190	3.7310	0.1857
deltam_clm	158.8000	0.0296	158.8000	0.0296	166.4000	0.0191

A Simple Approach for Measuring Higher-Order Risk Attitudes

Cary Deck* Rachel J. Huang†

Larry Y. Tzeng‡ Lin Zhao§

April 16, 2023

Abstract

While the importance of the Arrow-Pratt coefficients of second order risk aversion is well established in economics, the importance of higher-order risk attitudes has only recently begun to be recognized. In this paper, we provide the first approach to measure the higher-order Arrow-Pratt coefficients with choices between compound lotteries. Specifically, we provide a theoretical basis for using risk apportionment to reveal the intensity of higher-order risk attitudes, and then draw upon our theoretical results to develop a simple, systematic, and generalizable procedure for eliciting Arrow-Pratt coefficients of prudence, temperance, and other higher-order risk attitudes. We demonstrate our approach in a laboratory experiment and find that the modal subject exhibits mild prudence and mild temperancein addition to mild risk aversion. Further, we find that degrees of risk aversion are positively correlated across orders. Finally, while our approach is non-parametric, we note that behavior is broadly consistent with subjects having utility described by the exponential-power function.

Keywords: higher-order Arrow-Pratt risk aversion, risk apportionment, comparative risk attitude, measuring intensity of prudence and temperance, laboratory experiment.

*Corresponding Author, Department of Economics, University of Alabama, Tuscaloosa, AL 35487, United States. & Economic Science Institute, Chapman University. Email: cdeck@cba.ua.edu. phone: +1 205 348 8972. fax: +1 205 348 8983.

†Department of Finance, National Central University, Taoyuan 32001, Taiwan. Email: rachel@ncu.edu.tw.

‡Department of Finance, National Taiwan University, Taipei 10617, Taiwan. Email: tzeng@ntu.edu.tw.

§School of Business Administration, Faculty of Business Administration, Southwestern University of Finance and Economics, Chengdu 611130, China. Email: zhaoliniss@swufe.edu.cn.

1 Introduction

Risk attitudes of an economic agent have long been a fundamental issue in economics. Under expected utility, risk aversion equates to a negative second derivative of the utility function. The attitudes associated with higher order derivatives, referred to as higher-order risk attitudes, are now understood to be essential to economic decisions as well. For example, prudence (equated with a positive third derivative under expected utility) entails aversion to greater downside risks (Menezes, Geiss and Tressler 1980), and a stronger saving motive when future wealth becomes riskier (Kimball 1990). Temperance (equated with a negative fourth derivative under expected utility) implies less willingness to take on risk in the presence of greater background risk (Kimball 1993), and higher risk premium when the volatility of consumption growth increases (Gollier 2018).

While the direction of risk attitudes (equated with the sign of the appropriate derivative) serves as a primitive determinant for economic decisions, the intensity of risk attitudes is indispensable for quantitative analyses. For second-order risk attitude, Arrow (1971) and Pratt (1964) introduced the coefficients of absolute and relative risk aversion. For third order, Kimball (1990) introduced the coefficient of absolute prudence and linked it to the strength of the precautionary saving motives. For fourth order, Kimball (1992) introduced the coefficient of absolute temperance, which is useful for analyzing economic decisions involving two or more independent risks (Kimball 1993; Eeckhoudt, Gollier and Schlesinger 1996; Gollier and Pratt 1996). The above manner of quantifying the intensity of risk attitudes has been extended to all higher orders (Caballé and Pomansky 1996; Denuit and Eeckhoudt 2010). The resulting sequence of coefficients, which can be universally named as the coefficients of Arrow-Pratt absolute or relative risk aversion for a given order, play an important role in a wide range of economic applications including investment (Guiso, Jappelli and Terlizzese 1996), saving (Eeckhoudt and Schlesinger 2008), asset pricing (Gollier 2001), bargaining (White 2008), auctions (Esö and White 2004) and so on.

Despite the importance of higher-order risk attitudes, there have been relatively few empirical attempts to measure higher-order Arrow-Pratt coefficients of risk aversion. For third order, the coefficient of relative prudence has been estimated assuming a life-cycle consumption model and using savings data (Dynan 1993; Gourinchas and Parker 2002; Carroll and Kimball 2008). For the fourth order, we are not aware of any empirical study using naturally occurring data to estimate coefficients of absolute or relative temperance. In the laboratory, Ebert and Wiesen (2014) jointly measured the intensity of risk aversion, prudence, and temperance based on risk compensations while Noussair, Trautmann and van de Kuilen (2014) rely on the number of prudent choices as a measure of the degree of prudence, but neither of these papers are able to provide measures of the Arrow-Pratt coefficients.

In this paper, we provide the first approach to measure higher-order Arrow-Pratt risk aversion coefficients using choices between compound lotteries. In a seminal paper, Eeckhoudt and Schlesinger (2006) show apportioning zero mean risks, choosing between disaggregating and aggregating such risks in compound lotteries, can be used to identify the sign of the n^{th} derivative of utility.[1] We show that apportionment of non-zero-mean risks provides non-zero boundaries on the Arrow-Pratt coefficients of risk aversion. That is, while Eeckhoudt and Schlesinger (2006) provide theoretical justification for a method to determine the direction of higher-order risk attitudes, we provide a theoretical justification for a method to measure the Arrow-Pratt coefficients of risk aversion.[2]

Our paper is related to the literature on comparative risk aversion. Under a choice-based framework, Chiu (2005) and Denuit and Eeckhoudt (2010) construct lotteries such that the higher-order intensity of risk aversion between two individuals can be compared through their lottery choices.[3] The major difference between our paper and previous work is that we adopt the technique of risk apportionment. Furthermore, the approaches of both Chiu (2005) and Denuit and Eeckhoudt (2010) require mixed risk aversion, whereas our approach can be applied to decision makers who are risk seeking or imprudent or intemperate.

Our framework for identifying the intensity of higher-order risk preferences has several desirable properties. First, our method does not require any parametric assumption about the specific functional form of the decision makers' utility. While some empirical papers have tried to esti-

[1]This approach has served as the main method for studying higher-order risk preferences. Eeckhoudt, Rey and Schlesinger (2007) rely on the approach when considering bivariate utility functions; Eeckhoudt, Schlesinger and Tsetlin (2009) generalize risk apportionment to a broader class of lotteries; and Crainich, Eeckhoudt, and Trannoy (2013) apply the method to mixed risk lovers. Also, the approach of Eeckhoudt and Schlesinger (2006) has become a mainstream tool for experimental studies of higher-order risk preferences (e.g. Deck and Schlesinger 2010, 2014; Haering, Heinrich and Mayrhofer 2020).

[2]Recently, Jindapon, Liu and Neilson (2021) and Schneider and Sutter (2021) promote alternative measures to evaluate the strength of n^{th} degree risk apportionment. However, the measures proposed in those papers differ from the Arrow-Pratt measure of absolute and relative higher-order risk aversion. By contrast, our paper connects risk apportionment with the Arrow-Pratt coefficients.

[3]There have been other approaches to comparing the intensity of higher-order risk attitudes across individuals as well. For example, Jindapon and Neilson (2007) propose a comparative statics approach for n^{th}-degree risk aversion in an optimal effort decision. The approach of Liu and Meyer (2013) involves a comparison of matching probabilities.

mate the intensity of higher-order risk attitudes (e.g. Dynan 1993 and Noussair, Trautmann and van de Kuilen 2014), thus far such efforts have had to rely upon additional assumptions regarding preferences, which arbitrarily confines the degrees of freedom to describe an individual's behavior. Second, our approach is simple, systematic and generalizable. It is simple in that it only involves comparisons between two lotteries that are themselves composed of combinations of certain losses and fifty-fifty lotteries. It is systematic in that it involves a series of incremental comparisons, similar to the multiple price list approach popularized by Holt and Laury (2002) for measuring second-order risk attitudes. It is generalizable in that it can be readily adapted to risk preferences at any order, which stands in contrast to approaches such as Cohen and Einav (2007) whose approach is context-dependent and difficult to adapt to arbitrarily high orders of risk attitude. Finally, our framework provides bounds on intensity rather than a mean based estimation as developed by Noussair, Trautmann and van de Kuilen (2014).

As a demonstration of our approach, we implement it in a controlled laboratory experiment and measure the intensity of the higher-order risk attitudes of our subjects. In the experimental task, we ask subjects to apportion a sequence of nine risks including negative-mean, zero-mean, and positive-mean risks. The observed patterns for the direction of second-, third-, and fourth-order risk attitudes as identified by the apportionment of zero-mean risks are consistent with those reported previously by Deck and Schlesinger (2014) and Noussair, Trautmann, and van de Kuilen (2014). The modal behavior among our subjects is a mild degree of risk aversion consistent with a large experimental literature, a mild degree of prudence consistent with the relatively small literature attempting to measure the intensity of third order risk preferences (e.g. Noussair, Trautmann and van de Kuilen 2014 and Dynan 1993), and a mild degree of temperance which is a novel finding in the literature. However, we also observe a fraction of subjects who exhibit more extreme prudence and temperance as well as sizable fractions of subjects who exhibit moderate to extreme levels of imprudence or intemperance. Additionally, we extend Jindapon and Neilson (2007) to cases when decision makers are risk-loving or imprudent to demonstrate the implications of our laboratory findings.

Our approach also allows us to consider how degree of risk aversion are related across orders. Among our subjects, their degrees of second and third order risk aversion are positive and significant as are their degrees of third and fourth order risk aversion, but the greatest correlation is between second and fourth order degrees of risk aversion. Finally, we conduct a calibration exercise to determine how well common utility functions match observed behavior. Ultimately, we find that the behavior of most subjects can described with a exponential-power utility function.

2 Intensity of Higher-Order Risk Preferences

How to characterize the intensity of risk preferences within the expected utility framework is a fundamental question that has been extensively investigated in the economics literature. To briefly review the theoretical progress, let $u(x)$ be a von Neumann-Morgenstern utility function of wealth x that is defined on $(0, \infty)$ and continuously differentiable up to the desired order. For $n = 1, 2, ...$, denote by $u^{(n)}(x)$ the n^{th} derivative of $u(x)$.

Arrow (1971) and Pratt (1964) introduce the coefficient of absolute risk aversion $-\frac{u^{(2)}(x)}{u^{(1)}(x)}$ as a measure of second-order risk attitude, which is well used in risk-taking decisions, e.g., investment and insurance choices. To examine the precautionary savings motive, Kimball (1990) introduces the coefficient of absolute prudence $-\frac{u^{(3)}(x)}{u^{(2)}(x)}$. The higher the coefficient of absolute prudence, the higher the strength of the precautionary saving motives. Kimball (1992) introduces the coefficient of absolute temperance $-\frac{u^{(4)}(x)}{u^{(3)}(x)}$ and shows that it is related to how strongly an individual is inclined to avoid binding one risk with another unavoidable independent risk. This pattern is extended to higher-order risk attitude by Caballé and Pomansky (1996) who define $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)}$ as the coefficient of the n^{th}-order absolute risk aversion. Following Caballé and Pomansky's terminology, $-\frac{u^{(2)}(x)}{u^{(1)}(x)}$, $-\frac{u^{(3)}(x)}{u^{(2)}(x)}$ and $-\frac{u^{(4)}(x)}{u^{(3)}(x)}$ can be relabeled as absolute risk aversion of second order, third order and fourth order, respectively.[4] As shown by Jindapon and Neilson (2007), given a non-monetary cost of effort, the strength of the willingness to invest in effort to reduce risk depends on the coefficient of n^{th}-order absolute risk aversion.

In parallel, Pratt (1964) introduced the coefficients of relative risk aversion $-x\frac{u^{(2)}(x)}{u^{(1)}(x)}$, while relative prudence $-x\frac{u^{(3)}(x)}{u^{(2)}(x)}$ is proposed by Kimball (1990). Eeckhoudt and Schlesinger (2008) extend these coefficients to relative temperance $-x\frac{u^{(4)}(x)}{u^{(3)}(x)}$, as well as higher orders via defining $-x\frac{u^{(n)}(x)}{u^{(n-1)}(x)}$ as the coefficient of the n^{th}-order relative risk aversion. They show that the condition $-x\frac{u^{(n)}(x)}{u^{(n-1)}(x)} \geq n-1$ is crucial to guarantee an increase in precautionary saving when there is an increase in risk

[4]Gollier (2018) labels absolute risk aversion of fifth order as absolute edginess when investigating aversion to risks on the variance of consumption.

in the return on saving.

In a seminal paper, Eeckhoudt and Schlesinger (2006) show that the sign of $u^{(n)}(x)$ can be revealed with choices between simple 50-50 lotteries composed of pure losses and zero-mean risks. However, their work remains silent on the intensity of n^{th}-order risk preferences. In this section we extend the approach of Eeckhoudt and Schlesinger (2006) for comparative higher-order risk aversion with choices between simple lotteries.

2.1 Intensity of Second-Order Risk Aversion

Let $w > 0$ denote an initial wealth level and $\tilde{\varepsilon}$ be a zero-mean risk. An individual is risk averse on a pre-specified interval $[a, b] \subset (0, \infty)$, if and only if for all lottery pairs supported on $[a, b]$ taking the form of $w + \tilde{\varepsilon}$ and w, w is always preferred to $w + \tilde{\varepsilon}$. Within the expected utility framework, risk aversion on $[a, b]$ is equivalent to $u^{(2)} \leq 0$ on $[a, b]$ (Rothschild and Stiglitz 1970).

Replacing the zero-mean risk $\tilde{\varepsilon}$ with a general non-zero-mean risk $\tilde{\delta}$, we can elicit a bound for $-\frac{u^{(2)}(x)}{u^{(1)}(x)}$ from a choice between

$$A_2 = w + \tilde{\delta} \quad \text{and} \quad B_2 = w. \tag{1}$$

Proposition 1. *Let A_2 and B_2 take the form of (1). For u and v that are twice continuously differentiable with $u^{(1)} > 0$ and $v^{(1)} > 0$, the following statements are equivalent:*

(i) For all $x \in [a, b]$, $-\frac{u^{(2)}(x)}{u^{(1)}(x)} \geq -\frac{v^{(2)}(x)}{v^{(1)}(x)}$;

(ii) For all A_2 and B_2 supported on $[a, b]$, $\mathbb{E}v(A_2) = \mathbb{E}v(B_2)$ always implies $\mathbb{E}u(A_2) \leq \mathbb{E}u(B_2)$.

All proofs are relegated to Appendix 7. Intuitively, Proposition 1 can be obtained from Pratt (1964). To see this, one can rewrite $\tilde{\delta} = \mathbb{E}\tilde{\delta} + (\tilde{\delta} - \mathbb{E}\tilde{\delta})$ where the first term is the mean and the second term is a zero-mean risk. If under v one is indifferent between A_2 and B_2, it means that the mean of the risk $\tilde{\delta}$ is exactly the compensating premium necessary to bear the zero-mean risk $\tilde{\delta} - \mathbb{E}\tilde{\delta}$. Since more risk aversion requires a greater compensating premium, B_2 would be more preferable than A_2 under u.

Proposition 1 demonstrates that comparative risk aversion can be revealed with simply lottery choices. In particular, Proposition 1 shows u is more risk averse than v, if and only if u always favors the risky lottery less than v. An analogous characterization of $-\frac{u^{(2)}(x)}{u^{(1)}(x)} \leq -\frac{v^{(2)}(x)}{v^{(1)}(x)}$ is available through reversing the inequality in statement (ii). In the special case where v is a linear (risk neutral) utility function, the equation $\mathbb{E}v(A_2) = \mathbb{E}v(B_2)$ amounts to requiring $\tilde{\delta}$ to have a zero mean, reproducing the equivalence of $u^{(2)} \leq 0$ with w always being preferred to $w + \tilde{\varepsilon}$.

The idea of comparing coefficients of risk aversion based on choice behavior was previously explored by Jewitt (1989) and Chiu (2005). However, in both of those papers the second derivative of the utility function is required to be negative. In contrast, our analysis imposes no condition on the second derivative. Thus, relative to Jewitt (1989) and Chiu (2005), we extend the comparison of risk attitudes to utility functions exhibiting risk seeking behavior while employing simpler lottery pairs.

2.2 Intensity of Third-Order Risk Aversion

Let $k > 0$ be a constant and recall that $\tilde{\varepsilon}$ denotes a zero-mean risk. Denote by $[x; y]$ a lottery with a 50-50 chance of receiving either outcome x or outcome y, where x and y themselves may be lotteries. According to Eeckhoudt and Schlesinger (2006), an individual is called prudent on $[a, b]$, if and only if for all lottery pairs taking the form of $[w; w - k + \tilde{\varepsilon}]$ and $[w + \tilde{\varepsilon}; w - k]$ supported on $[a, b]$, the latter is always preferred to the former. That is, a prudent individual prefers putting a zero-mean risk at the higher wealth level than at the lower wealth level. Prudence captures aversion to aggregating a loss with a zero mean risk. Within the expected utility framework, prudence on $[a, b]$ is equivalent to $u^{(3)} \geq 0$ on $[a, b]$.

As we show in Proposition 2, replacing the zero-mean risk $\tilde{\varepsilon}$ with a general non-zero-mean risk $\tilde{\delta}$, we can elicit a bound for $-\frac{u^{(3)}(x)}{u^{(2)}(x)}$ from choices between

$$A_3 = \left[w; w - k + \tilde{\delta}\right] \quad \text{and} \quad B_3 = \left[w + \tilde{\delta}; w - k\right]. \tag{2}$$

To do this, we need to distinguish between cases in which $u^{(2)} > 0$ and $u^{(2)} < 0$.

Proposition 2. *Let A_3 and B_3 take the form of (2). For u and v that are continuously differentiable up to the third order with $u^{(2)} \neq 0$ and $v^{(2)} \neq 0$, consider the following statements:*

(i) For all $x \in [a, b]$, $-\frac{u^{(3)}(x)}{u^{(2)}(x)} \geq -\frac{v^{(3)}(x)}{v^{(2)}(x)}$;

(ii) For all A_3 and B_3 supported on $[a,b]$, $\mathbb{E}v(A_3) = \mathbb{E}v(B_3)$ always implies $\mathbb{E}u(A_3) \leq \mathbb{E}u(B_3)$;

(iii) For all A_3 and B_3 supported on $[a,b]$, $\mathbb{E}v(A_3) = \mathbb{E}v(B_3)$ always implies $\mathbb{E}u(A_3) \geq \mathbb{E}u(B_3)$.

When $u^{(2)} < 0$ and $v^{(2)} < 0$, (i) and (ii) are equivalent; when $u^{(2)} > 0$ and $v^{(2)} > 0$, (i) and (iii) are equivalent.

Proposition 2 is analogous to Proposition 1, but for third-order risk attitude. Assuming risk aversion, u is more prudent than v, if and only if u exhibits a stronger propensity to disaggregate the loss and the risky lottery than v. An analogous characterization of $-\frac{u^{(3)}(x)}{u^{(2)}(x)} \leq -\frac{v^{(3)}(x)}{v^{(2)}(x)}$ is available through reversing the inequalities in statements (ii) and (iii). In the special case where v is a quadratic (prudence neutral) utility function, the equation $\mathbb{E}v(A_3) = \mathbb{E}v(B_3)$ amounts to requiring $\tilde{\delta}$ to have a zero mean.[5] This reproduces the equivalence of $u^{(3)} \geq 0$ with preferring $[w + \tilde{\varepsilon}; w - k]$ over $[w; w - k + \tilde{\varepsilon}]$ and the equivalence of $u^{(3)} \leq 0$ with preferring $[w; w - k + \tilde{\varepsilon}]$ over $[w + \tilde{\varepsilon}; w - k]$ as characterized by Eeckhoudt and Schlesinger (2006).

Chiu (2005) also uses choice behavior to compare coefficients of prudence. However, in Chiu's analysis the second and third derivatives of the utility function are required to be negative and positive, respectively. Our result shows that it is possible to identify the intensity of prudence without presupposing the signs of the second and third derivatives of the utility function although our lotteries are special cases of those in Chiu (2005). That is, our approach can accommodate utility functions exhibiting risk averse or seeking and prudent or imprudent behavior in the same manner. Moreover, the choices can be presented as simple lotteries with equiprobable outcomes.

2.3 Intensity of Higher-Order Risk Aversion

Comparative risk aversion can be identified for any order with choices between simple lotteries following the idea of risk apportionment. To do this, we first recall the lotteries introduced by Eeckhoudt and Schlesinger (2006), whose purpose is identifying the direction of preferences. Let $\{\tilde{\varepsilon}_i\}$ denote an indexed set of zero-mean risks that are mutually independent. For $w > 0$ and $k < 0$,

[5] To see this, we write $\mathbb{E}v(A_3) - \mathbb{E}v(B_3) = \frac{1}{2}[\mathbb{E}v_1(w + \tilde{\delta}) - v_1(w)]$, where $v_1(x) = v(x-k) - v(x)$ is linear in x when $v(x)$ is quadratic.

Eeckhoudt and Schlesinger define

$$\hat{A}_1 = w - k, \quad \hat{B}_1 = w,$$
$$\hat{A}_2 = w + \tilde{\varepsilon}_1, \quad \hat{B}_2 = w,$$

and

$$\hat{A}_n = \left[\hat{A}_{n-2} + \tilde{\varepsilon}_{\mathrm{Int}(n/2)}; \hat{B}_{n-2}\right], \quad \hat{B}_n = \left[\hat{A}_{n-2}; \hat{B}_{n-2} + \tilde{\varepsilon}_{\mathrm{Int}(n/2)}\right]$$

for $n \geq 3$ where $\mathrm{Int}(n/2)$ denotes the greatest integer not exceeding $n/2$. Based on this definition, we have

$$\hat{A}_3 = \left[\hat{A}_1 + \tilde{\varepsilon}_1; \hat{B}_1\right] = \left[w - k + \tilde{\varepsilon}_1; w\right], \quad \hat{B}_3 = \left[\hat{A}_1; \hat{B}_1 + \tilde{\varepsilon}_1\right] = \left[w - k; w + \tilde{\varepsilon}_1\right],$$
$$\hat{A}_4 = \left[\hat{A}_2 + \tilde{\varepsilon}_2; \hat{B}_2\right] = \left[w + \tilde{\varepsilon}_1 + \tilde{\varepsilon}_2; w\right], \quad \hat{B}_4 = \left[\hat{A}_2; \hat{B}_2 + \tilde{\varepsilon}_2\right] = \left[w + \tilde{\varepsilon}_1; w + \tilde{\varepsilon}_2\right].$$

In the above, $\hat{B}_1$ and $\hat{B}_2$ represent a fixed state, while $\hat{A}_1$ represents a certain loss relative to $\hat{B}_1$ and $\hat{A}_2$ represents a risky state relative to $\hat{B}_2$. For $n \geq 3$, $\hat{B}_n$ and $\hat{A}_n$ are lotteries involving $\hat{B}_{n-2}$ and $\hat{A}_{n-2}$ where $\hat{B}_n$ attaches an independent zero-mean risk to $\hat{B}_{n-2}$ and $\hat{A}_n$ attaches that zero-mean risk to $\hat{A}_{n-2}$. Eeckhoudt and Schlesinger prove that $(-1)^{n+1}u^{(n)} \geq 0$ if and only if for any zero-mean risks $\hat{B}_n$ is always preferred to $\hat{A}_n$. Such a preference is termed as "risk apportionment" of order n. Risk apportionment captures the aversion to combining "bad"—the additional pure risk $\tilde{\varepsilon}_{\mathrm{Int}(n/2)}$—with "bad"—the more risky state $\hat{A}_{n-2}$.

Our lotteries aimed at identifying the intensity of preferences are related to Eeckhoudt and Schlesinger's lotteries by replacing the zero-mean risk $\tilde{\varepsilon}_{\mathrm{Int}(n/2)}$ with a general non-zero-mean risk $\tilde{\delta}$. Assuming independence between $\tilde{\delta}$ and $\{\tilde{\varepsilon}_i\}$, we define

$$\begin{aligned} A_1 &= w - k, \quad B_1 = w, \\ A_2 &= w + \tilde{\delta}, \quad B_2 = w, \end{aligned} \tag{3}$$

and

$$A_n = \left[\hat{A}_{n-2} + \tilde{\delta}; \hat{B}_{n-2}\right], \quad B_n = \left[\hat{A}_{n-2}; \hat{B}_{n-2} + \tilde{\delta}\right], \tag{4}$$

for $n \geq 3$. Based on this definition, we have

$$A_3 = \left[\hat{A}_1 + \tilde{\delta}; \hat{B}_1\right] = \left[w - k + \tilde{\delta}; w\right], \quad B_3 = \left[\hat{A}_1; \hat{B}_1 + \tilde{\delta}\right] = \left[w - k; w + \tilde{\delta}\right],$$
$$A_4 = \left[\hat{A}_2 + \tilde{\delta}; \hat{B}_2\right] = \left[w + \tilde{\varepsilon}_1 + \tilde{\delta}; w\right], \quad B_4 = \left[\hat{A}_2; \hat{B}_2 + \tilde{\delta}\right] = \left[w + \tilde{\varepsilon}_1; w + \tilde{\delta}\right].$$

Allowing $\tilde{\delta}$ to have a non-zero mean and letting

$$\tilde{\varepsilon}_i = [-k_i; k_i] \text{ where } k_i > 0, \tag{5}$$

we can explore the intensity of the n^{th}-order risk aversion based on the choice between A_n and B_n.

Theorem 1. *For $n \geq 2$, let A_n and B_n be defined as in (3) and (4), with $\tilde{\varepsilon}_i$ specified in (5). For u and v that are continuously differentiable up to order n with $u^{(n-1)} \neq 0$ and $v^{(n-1)} \neq 0$, consider the following statements:*

(i) For all $x \in [a, b]$, $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)} \geq -\frac{v^{(n)}(x)}{v^{(n-1)}(x)}$;

(ii) For all A_n and B_n supported on $[a, b]$, $\mathbb{E}v(A_n) = \mathbb{E}v(B_n)$ always implies $\mathbb{E}u(A_n) \leq \mathbb{E}u(B_n)$;

(iii) For all A_n and B_n supported on $[a, b]$, $\mathbb{E}v(A_n) = \mathbb{E}v(B_n)$ always implies $\mathbb{E}u(A_n) \geq \mathbb{E}u(B_n)$.

When $(-1)^n u^{(n-1)} > 0$ and $(-1)^n v^{(n-1)} > 0$, (i) and (ii) are equivalent; when $(-1)^n u^{(n-1)} < 0$ and $(-1)^n v^{(n-1)} < 0$, (i) and (iii) are equivalent.

Theorem 1 generalizes Propositions 1 to 2 to higher orders. For example, the bound for temperance measured by $-\frac{u^{(4)}(x)}{u^{(3)}(x)}$ can be elicited from the choice between A_4 and B_4. An analogous characterization of $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)} \leq -\frac{v^{(n)}(x)}{v^{(n-1)}(x)}$ is available through reversing the inequalities in statements (ii) and (iii). In the special case where $v^{(n-1)} \equiv 0$, we can prove by induction that the equation $\mathbb{E}v(A_n) = \mathbb{E}v(B_n)$ amounts to requiring $\tilde{\delta}$ to have a zero mean, reproducing the characterization of $(-1)^{n+1}u^{(n)} \geq (\leq) 0$ that $\hat{B}_n$ is always more (less) preferable than $\hat{A}_n$.[6]

To our knowledge, Denuit and Eeckhoudt (2010) were the first to extend the equivalence between comparative risk aversion and a binary choice behavior to higher orders. Their lottery pairs,

[6]For example, on fourth order when v is a temperance neutral (cubic) utility function, we write $\mathbb{E}v(A_4) - \mathbb{E}v(B_4) = \frac{1}{2}[\mathbb{E}v_2(w + \tilde{\delta}) - v_2(w)]$, where $v_2(x) = \frac{1}{2}[v(x - k) + v(x + k)] - v(x)$ is linear in x when $v(x)$ is cubic. Thus, the equation $\mathbb{E}v(A_4) = \mathbb{E}v(B_4)$ amounts to requiring $\tilde{\delta}$ to have a zero mean, reproducing the characterization by Eeckhoudt and Schlesinger (2006) that $u^{(4)} \leq 0$ iff $[w + \tilde{\varepsilon}_1; w + \tilde{\varepsilon}_2]$ is always preferred to $[w; w + \tilde{\varepsilon}_1 + \tilde{\varepsilon}_2]$.

however, are designed only for utility functions that have positive odd numbered derivatives and negative even numbered derivatives up to the relevant order, which is a typical feature of the so-called "mixed risk averse" utility functions (Caballé and Pomansky 1996). In contrast, we construct lottery pairs by iteration of simple 50-50 lotteries, and hence impose no condition on the sign of derivatives of up to $n-1$. The iterative approach also allows the sign of the n^{th} derivative of utility functions to be either positive or negative, and simplifies the choices to involve only 50-50 lotteries.

3 An Implementable Procedure

Theorem 1 provides a means to compare the n^{th}-order absolute and relative risk aversion between two utility functions. We can bound the n^{th}-order absolute risk aversion of u by comparing u with v_1 and bound the n^{th}-order relative risk aversion of u by comparing it with v_2, where v_1 and v_2 satisfy

$$-\frac{v_1^{(n)}(x)}{v_1^{(n-1)}(x)} = \theta_1 \quad \text{and} \quad -x\frac{v_2^{(n)}(x)}{v_2^{(n-1)}(x)} = \theta_2, \quad \theta_1, \theta_2 \in \mathbb{R},$$

respectively. While Theorem 1 is stated in terms of comparing the intensities of absolute risk aversion, it works equally well for relative risk aversion when $x > 0$. Indeed, when comparing u with v_2 by Theorem 1, we get $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)} \geq$ or $\leq -\frac{v_2^{(n)}(x)}{v_2^{(n-1)}(x)} = \frac{\theta_2}{x}$, which is equivalent to $-x\frac{u^{(n)}(x)}{u^{(n-1)}(x)} \geq$ or $\leq \theta_2$. The structural assumptions on v_1 and v_2 serve as bases for bounding a subject's risk attitude, which depends on u. Thus, we do not assume that one's utility function u exhibits constant n^{th}-order absolute or relative risk aversion and instead compare one's utility function to functions with those specific forms at a given level of wealth. In fact, our procedure proposes no assumption on the form of individual's utility function.

For a given order n, a subject faces a task that involves a series of risk apportionment choices between A_n and B_n that systematically varies $\tilde{\delta}$ holding other parameters fixed. Formally, for choice j in a Task of Order n, we construct A_n and B_n following (3) and (4) with task-specific values of $w > 0$, $k > 0$, $\tilde{\varepsilon}_i$ as in (5), and

$$\tilde{\delta}_j = [-h; h] + l_j = [-h + l_j; h + l_j], \tag{6}$$

where $h > 0$, $-h = l_1 < l_2 < ... < l_{J-1} < l_J = h$ and $J \geq 3$. As j increases from 1 to J, $\tilde{\delta}_j$ moves from $[-2h; 0]$ to $[0; 2h]$. We use $A_n(j)$ and $B_n(j)$ to indicate explicitly the dependence of A_n and B_n constructed in this way on j. Then, a Task of Order n is formulated as

$$\text{Task of Order } n = \{(A_n(j), B_n(j)) : j = 1, ..., J\}, \tag{7}$$

in which we present subjects a sequence of lottery pairs and ask them to select their preferred option in each pair.

To take a numerical third-order example, consider the series of 9 lottery pairs: $(A_3(j), B_3(j))$, $j = 1, 2, ..., 9$ with $A_1 = 13$, $B_1 = 23$, $h = 4$ and $l_{j+1} - l_j = 1$. Thus, for $j = 1$, we have $\delta_1 = [-8; 0]$, and Choice 1 is between the lottery pair

$$A_3(1) = [13 + [-8; 0]; 23] \ \text{ and } \ B_3(1) = [13; 23 + [-8; 0]].$$

For $j = 2$, we have $\delta_2 = [-7; 1]$. Choice 2 is between the lottery pair

$$A_3(2) = [13 + [-7; 1]; 23] \ \text{ and } \ B_3(2) = [13; 23 + [-7; 1]].$$

For $j = 9$, we have $\delta_9 = [0; 8]$ and Choice 9 is between the lottery pair

$$A_3(9) = [13 + [0; 8]; 23] \ \text{ and } \ B_3(9) = [13; 23 + [0; 8]].$$

The lottery comparison of the form used by Eeckhoudt and Schlesinger (2006) occurs when $j = 5$, yielding $\delta_5 = [-4; 4]$ and a choice over the lottery pair

$$A_3(5) = [13 + [-4; 4]; 23] \ \text{ and } \ B_3(5) = [13; 23 + [-4; 4]].$$

For any sequence of lotteries constructed through the process described above, for $j = 2, ..., J-1$, there exist unique n^{th}-order constant absolute and relative risk aversion coefficients, denoted by

$\Theta_1(n,j)$ and $\Theta_2(n,j)$ respectively, such that

$$\mathbb{E}v_1(A_n(j)) = \mathbb{E}v_1(B_n(j)) \text{ under } \theta_1 = \Theta_1(n,j), \text{ and}$$
$$\mathbb{E}v_2(A_n(j)) = \mathbb{E}v_2(B_n(j)) \text{ under } \theta_2 = \Theta_2(n,j).$$

These coefficients make individuals with utility functions v_1 or v_2 indifferent between $A_n(j)$ and $B_n(j)$.[7] It is shown in Lemma A3 in Appendix A that both $\Theta_1(n,j)$ and $\Theta_2(n,j)$ are strictly increasing in j. As a convention, we set $\Theta_1(n,1) = \Theta_2(n,1) = -\infty$ and $\Theta_1(n,J) = \Theta_2(n,J) = \infty$.

We can identify both upper and lower bounds on risk attitude at some wealth level from a single task. This approach to providing bounds on risk attitude is justified by the following corollary.

Corollary 1. *For $n \geq 2$, consider a Task of Order n as in (7) that is supported on $[a,b] \subset (0,\infty)$. Let u be a utility function that is continuously differentiable up to order n. When $(-1)^n u^{(n-1)} > 0$, there exists a unique $j^* \leq J-1$ such that the individual prefers $B_n(j)$ to $A_n(j)$ for $j \leq j^*$, but $A_n(j)$ to $B_n(j)$ for $j \geq j^*+1$. When $(-1)^n u^{(n-1)} < 0$, similar behavior holds with the individual preferring $A_n(j)$ to $B_n(j)$ for $j \leq j^*$, but $B_n(j)$ to $A_n(j)$ for $j \geq j^*+1$. For both cases, there exist $x_1, x_2 \in [a,b]$ such that*

$$\Theta_1(n,j^*) \leq -\frac{u^{(n)}(x_1)}{u^{(n-1)}(x_1)} \leq \Theta_1(n,j^*+1),$$
$$\Theta_2(n,j^*) \leq -x_2\frac{u^{(n)}(x_2)}{u^{(n-1)}(x_2)} \leq \Theta_2(n,j^*+1).$$

Corollary 1 demonstrates what we can infer from a single task. Let us recall the previous example to illustrate Corollary 1. Suppose that an individual prefers $B_3(j)$ to $A_3(j)$ for $j=1$ and $j=2$ and $A_3(j)$ to $B_3(j)$ for $j=3$ to $j=9$. The choices of $B_3(1)$ and $A_3(9)$ reveal the individual is risk-averse. An individual with a degree of absolute prudence equal to -0.69 and relative prudence equal to -14.26 would be indifferent between $A_3(2)$ and $B_3(2)$. That is $\Theta_1(3,2) = -0.69$ and $\Theta_2(3,2) = -14.26$. Furthermore, an individual with a degree of absolute prudence equal to -0.31 and relative prudence equal to -5.82 would be indifferent between $A_3(3)$ and $B_3(3)$. That is $\Theta_1(3,3) = -0.31$ and $\Theta_2(3,3) = -5.82$. Thus, from Corollary 1, we know that the individual who

[7] Here, v_1 and v_2 are unique up to scaling and addition of polynomials up to order $n-2$. Since the moments of A_n and B_n are the same up to order $n-2$, the scaling factor and the polynomial terms do not affect the subsequent analysis.

prefers $B_3(j)$ to $A_3(j)$ only for $j \leq 2$ exhibits

$$-0.69 \leq -\frac{u^{(3)}(x_1)}{u^{(2)}(x_1)} \leq -0.31 \text{ and } -14.26 \leq -x_2\frac{u^{(3)}(x_2)}{u^{(3)}(x_2)} \leq -5.82$$

for some x_1 and x_2 within $[5, 31]$.

Corollary 1 does not impose any assumption about the functional form of u, so that risk attitude elicited in a wealth range does not imply anything on risk attitude on other wealth ranges. To elicit the bound of absolute or relative risk aversion of order n, there is no need to rely on separate information about any risk attitude of a lower order.

A premise for this corollary to work is that $u^{(n-1)}$ does not vary in sign over the relevant domain. Since risk aversion coefficients, $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)}$ or $-x\frac{u^{(n)}(x)}{u^{(n-1)}(x)}$, are definable only for segments with $u^{(n-1)} \neq 0$, this premise is not excessively demanding. Under this premise, the difference between the expected utility of $A_n(j)$ and that of $B_n(j)$ changes monotonically with j, yielding a single switch point from preferring $B_n(j)$ to preferring $A_n(j)$ under $(-1)^n u^{(n-1)} > 0$, or from preferring $A_n(j)$ to preferring $B_n(j)$ under $(-1)^n u^{(n-1)} < 0$. As with other techniques for measuring risk aversion, such as the multiple price list approach of Holt and Laury (2002), individuals whose preferences are captured by standard functional forms for utility, should exhibit a single switch point when going through the J choices of the task.

4 Experimental Design

To demonstrate the implementation of our procedure, we conducted a controlled laboratory experiment. For illustrative purposes, our experimental investigation will focus on the second, third and fourth order. For each order, we present subjects a task as described in (7), with each task consisting of 9 choices involving $J = 9$ pairs of lotteries. In total, subjects were asked to make 27 choices (9 choices / order $\times$ 3 orders). At the end of the experiment, one of these 27 choices was randomly selected and used to determine the subject's earnings.

Before continuing, we note that having 9 choices for a task means subjects can be placed into 10 continuous bins that only overlap at their endpoints.[8] The width of each bin depends on the specific lotteries that are used in the task and one could construct a finer or coarser set of bounds

[8] The bins overlap at the endpoints as a subject could be indifferent for a given choice.

Table 1: All Decision Tasks

Task of	Option $A_n(j)$	Option $B_n(j)$	$\tilde{\delta}_1$	$\tilde{\delta}_9$	$l_{j+1} - l_j$	Average Payoff
Order 2	$18 + \tilde{\delta}_j$	18	$[-8; 0]$	$[0; 8]$	1	18
Order 3	$[13 + \tilde{\delta}_j; 23]$	$[13; 23 + \tilde{\delta}_j]$	$[-8; 0]$	$[0; 8]$	1	18
Order 4	$[18 + [-5; 5] + \tilde{\delta}_j; 18]$	$[18 + [-5; 5]; 18 + \tilde{\delta}_j]$	$[-8; 0]$	$[0; 8]$	1	18

Note: This table reports the numerical payoffs used to construct tasks as formulated in (7). Recall that $[x; y]$ denotes a lottery where there is a 50% chance of receiving x and a 50% chance of receiving y. In $A_n(j)$ and $B_n(j)$, $\tilde{\delta}_j = [-h; h] + l_j$, with $\tilde{\delta}_1 = [-2h; 0]$, $\tilde{\delta}_9 = [0; 2h]$, and $l_{j+1} - l_j$ given in the second to the last column. For example, in our Task of Order 3, $\tilde{\delta}_1$ is $[-8; 0]$, $\tilde{\delta}_2$ is $[-7; 1]$, and $\tilde{\delta}_3$ is $[-6; 2]$, and so on.

by using more or fewer lotteries, respectively. This is also true for the multiple price list approach for measuring second-order risk aversion popularized by Holt and Laury (2002).[9]

4.1 Construction of Tasks

For all tasks, $\tilde{\delta}_1 = [-2h; 0]$ in Choice 1 only involves a loss, $\tilde{\delta}_5 = [-h; h]$ in Choice 5 involves 50-50 of equal sized gain and loss, and $\tilde{\delta}_9 = [0; 2h]$ in Choice 9 only involves a gain. Thus, for a Task of Order n Choice 1 and Choice 9 reveal the direction of the $(n-1)^{th}$-order risk attitude, while Choice 5 is consistent with Eeckhoudt and Schlesinger (2006) and reveals the direction of the n^{th}-order risk attitude.

Table 1 provides all 3 tasks used in the experiment, where payoffs are in US dollars. Our Task of Order 3 is the basis for the numerical example used in the previous section. The numerical payoffs in the tasks are designed such that risk attitudes associated with indifference between choices are in the neighborhood of the risk attitudes that have been reported previously in the literature (e.g, Holt and Laury 2002; Bliss and Panigirtzoglou 2004; Noussair, Trautmann and van de Kuilen 2014). Thus, our specific tasks are not calibrated to identify particularly extreme levels of risk attitude, although one could design tasks to partition more extreme risk attitudes using the same technique.

4.2 Steps for Making Choices

Our subjects face multiple choices in each task. To help subjects understand the relationship between the choices on a task, all nine lotteries for a given order are displayed on the screen at the

[9] We also note that one could attempt to identify the $\tilde{\delta}$ that makes a respondent indifferent between A_n versus B_n, using a method similar to that of Becker, Degroot, and Marschak (1964). However, our approach only requires subjects to make binary comparisons.

Table 2: The Degree of Risk Aversion Making $A_n(j)$ and $B_n(j)$ Indifferent

Task of	Choice j						
	$j=2$	$j=3$	$j=4$	$j=5$	$j=6$	$j=7$	$j=8$
	Panel A: Constant Absolute Degrees $\Theta_1(n,j)$						
Order 2	-0.69	-0.31	-0.14	0.00	0.14	0.31	0.69
Order 3	-0.69	-0.31	-0.14	0.00	0.14	0.31	0.69
Order 4	-0.69	-0.31	-0.14	0.00	0.14	0.31	0.69
	Panel B: Constant Relative Degrees $\Theta_2(n,j)$						
Order 2	-11.81	-5.19	-2.27	0.00	2.40	5.72	12.96
Order 3	-14.26	-5.82	-2.36	0.00	2.21	4.97	10.48
Order 4	-13.44	-5.51	-2.31	0.00	2.30	5.28	11.21

Note: In this table, Panel A and Panel B report the constant degrees of n^{th}-order absolute and relative risk aversion that make $A_n(j)$ and $B_n(j)$ in Choice j within the Task of Order n indifferent. Here, $n = 2, 3, 4$ and $j = 2, ..., 8$.

same time. The subject then makes choices about apportioning the lotteries in order from $\tilde{\delta}_1$ to $\tilde{\delta}_9$. Figure 1 provides an example of a subject facing Choice 5 in our Task of Order 3 where they are asked to apportion $\tilde{\delta}_5$. Once the subject has made all 9 choices for a task, a button appears on the screen enabling the subject to submit all 9 responses simultaneously. A subject is free to change their decision regarding the apportionment of a lottery at any point prior to pressing the submit button.[10]

The presentation of each choice follows Deck and Schlesinger (2014) and is meant to facilitate a choice as deciding to combine "good" with "good" or to combine "good" with "bad." For example, in the fifth lottery shown in Figure 1 the decision is if one wants to combine a 50-50 lottery that pays either -\$4 or \$4 with the bad \$13 outcome or the good \$23 outcome of an independent 50-50 lottery. A risk averse person would view the -\$4 or \$4 lottery as a bad while a risk-loving person would view it as a good. A prudent person would opt to combine the -\$4 or \$4 lottery with the \$23 outcome. For a person who is risk averse this is combining a good and a bad, whereas for a risk-loving person this is combining a good with a good. The top lottery shown in Figure 1 also demonstrates how choice 1 an n^{th}-order task each identify the subject's $(n-1)^{th}$-order preferences. The example task measures third-order risk attitude, but for the first choice the preferred option

[10]If a subject changes their decision for Choice i with $i \leq 8$ the software proceeds to Choice $i+1$ and continues sequentially from that point. Additionally, the software imposes that a subject make a single switch on each task; however, the subjects are not informed of this in the instructions. Rather, those that attempt to provide responses that do not conform to a single switch rule are notified of this requirement. This allows us to identify how many subjects naturally follow a single switch rule. 30% of subjects never exhibited behavior inconsistent with a single switch. 47% and 8% of the subjects exhibited single switch only for the second and only for the second and third order task, respectively. The other 15% of subjects did not exhibit single switch on the second order task.

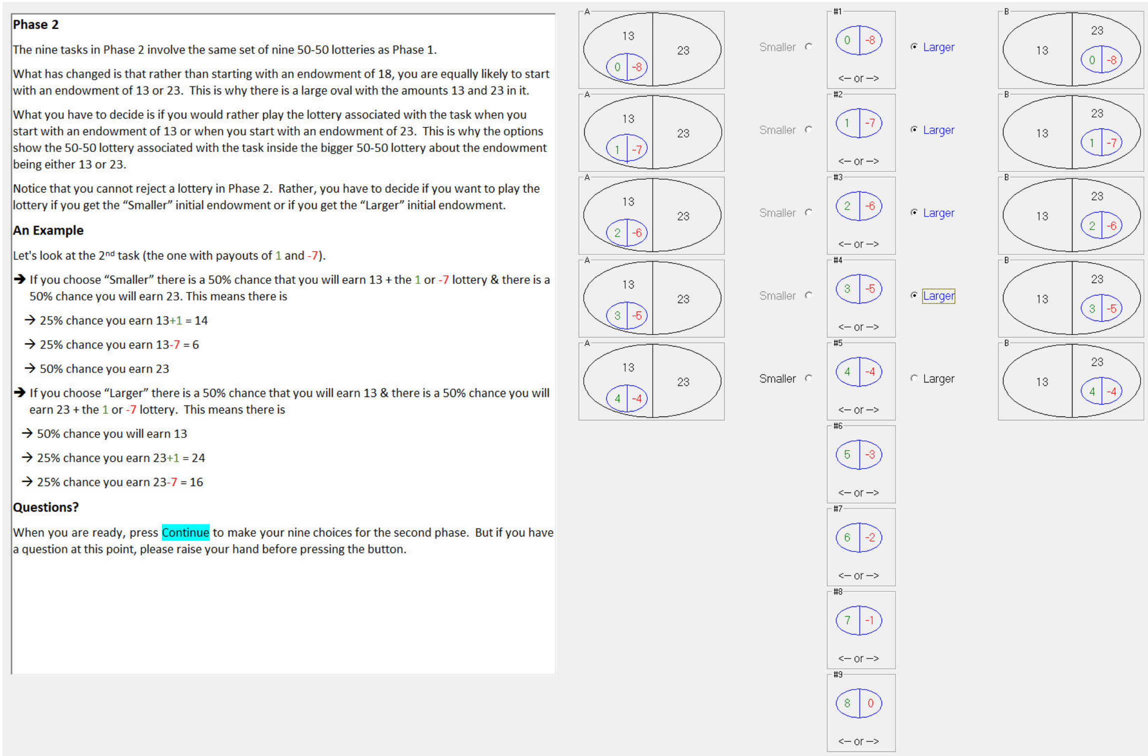

Figure 1: Subject Interface.

only depends on one's second-order risk attitude since the 50-50 lottery with -\$8 and \$0 is a pure loss. Similarly, the ninth choice, which is not visible in Figure 1 apportions the 50-50 lottery with \$0 and \$8, which is a pure gain and thus depends only one's second-order risk attitude.

4.3 Procedures

The study was conducted at The University of Alabama's TIDE Lab. One hundred subjects were recruited from the lab's standing pool of volunteers.[11] While many of the subjects had participated in other unrelated studies, none had participated in a study about risk. The average salient earnings were \$18.74 (with a minimum of \$7 and a maximum of \$31). Subjects also received a \$5 payment for participating in the study.

Data were collected during fourteen sessions with subjects being recruited for 30 minutes. Each session involved between 4 and 14 subjects; however, subjects did not interact with each other during a session. At the start of a session, subjects were seated at individual computer stations

[11]Because the experiment is meant to demonstrate the implementation of the procedure laid out in the previous section and to provide a general sense of the observed degrees of relative and absolute prudence and temperance rather than testing specific hypotheses, the sample size is arbitrary and not based on statistical power.

separated by privacy dividers. Subjects read general computerized instructions.[12] The paid portion of the study in which they completed the 3 tasks shown in Table 1 was self-paced. To facilitate subject understanding, the tasks were presented in order. Task specific instructions were presented just prior to the subjects making their decisions for that task. These instructions remained visible on the left portion of the screen throughout the time the subject was making decisions.

One the subject completed all 27 choices over the 3 tasks, the computer randomly selected one choice from one task to be used in determining the subject's payment. Any 50-50 lottery required to determine the outcome of the selected option was resolved through the use of a physical spinner as had been explained previously to the subjects.[13] Each subject complete a survey that consisted of a single question about gender and was then paid in private and dismissed from the study.

5 Experimental Results

The main results of the experiment are captured in Figure 2, which presents histograms for each order of the choice at which subjects changed their apportionment decisions. The relevant risk attitudes associated with a particular switch point can be found in Table 2. But as preliminary point, we note that the behavior we observe is consistent with the directional results of Deck and Schlesinger (2014) and Noussair, Trautmann, and van de Kuilen (2014). Effectively, those experiments only considered choices with zero mean lotteries (i.e. apportionment decisions of the form in Choice 5 of each of our tasks) and we find that a majority of our subjects indicate a preferences for $B_n(5)$ over $A_n(5)$ for all n = 2, 3, and 4.

Overall, 60% of our subjects made choices indicating some degree of 2^{nd} order risk aversion. For the second order task, the most common switching point was at Choice 6 indicating that 31% of the subjects are slightly risk averse exhibiting 2^{nd} order absolute risk aversion between 0.00 and 0.14 and 2^{nd} order relative risk aversion between 0.00 and 2.40. Further, 75% of the subjects exhibiting 2^{nd} order absolute risk aversion between -0.14 and 0.31 and 2^{nd} order relative risk aversion between -2.27 and 5.72 indicating that few subjects have extreme second order risk attitudes. As an aside, we also report that 2^{nd} order behavior did no differ by gender (p-value for χ^2 test = 0.598).

[12]The instructions are available in the supplementary materials.

[13]The subjects were invited to inspect the spinner before beginning the paid portion of the experiment and again before their final payoff was determined.

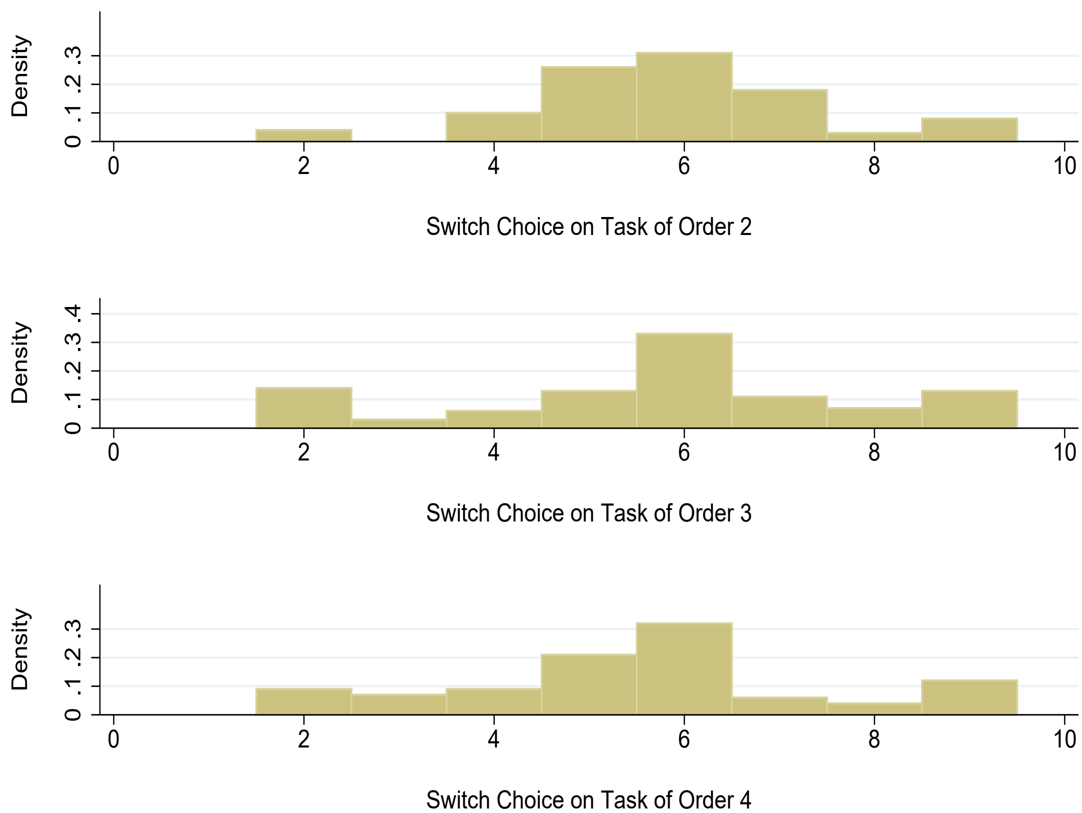

Figure 2: Histogram of Switching in Task of Order n

An approximately two-thirds majority of the subjects exhibited some degree of prudence. As with the second order task, for the third order task the most common switching point was at Choice 6 indicating that 33% of the subjects are slightly prudent exhibiting 3^{rd} order absolute risk aversion between 0.00 and 0.14 and 3^{rd} order relative risk aversion between 0.00 and 2.21. Unlike what was observed for second order risk, on the third order tasks a sizeable fraction of the subjects exhibit extreme attitudes. Only 57% of the subjects exhibiting 3^{rd} order absolute risk aversion between -0.14 and 0.31 and 3^{rd} order relative risk aversion between -2.36 and 4.97, while 14% (13%) of the subjects exhibit 3^{rd} order absolute risk aversion below -0.69 (above 0.69) and 3^{rd} order relative risk aversion below -14.26 (above 10.48). We also note that there is no gender difference in third order behavior (p-value for χ^2 test = 0.266).

For the fourth order task, only a slight majority of 54% exhibited temperance. The modal response was again to switch at choice 6 indicating that this 32% of the subjects exhibited 4^{th} order absolute risk aversion between 0.00 and 0.14 and 4^{th} order relative risk aversion between 0.00 and 2.30. As with the prudence, the degree of temperance exhibited by the subjects is more extreme than the degrees of second order risk aversion that were exhibited. Only 59% of the subjects exhibited 4^{th} order absolute risk aversion between -0.14 and 0.31 and 4^{th} order relative risk aversion between -2.31 and 5.28. Finally, as with second and third order behavior, there is no difference in 4^{th} order behavior by gender (p-value for χ^2 test = 0.582).

As detailed in Section 2, the degree of n^{th} order risk aversion depends on the sign of the $(n-1)^{th}$ derivative of the utility function. The structure of the tasks we implement is such that Choice 1 (or Choice 9) of a Task of Order n is sufficient to determine the sign of the $(n-1)^{th}$. Hence, in Figure 3 we separate the behavior of the subjects based on their selection of $A_n(1)$ or $B_n(1)$. Because every subject selected $B_2(1)$ rather than $A_2(1)$, indicating that all of our subjects exhibit behavior consistent with a monotonically increasing utility function, the figure only shows behavior for the Tasks of Orders 3 and 4. The top-right portion of Figure 3 shows that risk-averse subjects most commonly switch at Choice 6. It also indicates that far more subjects switch at Choices 6 through 9 than at Choices 2 through 5. The top-left portion of Figure 3 indicates that risk-loving subjects most commonly switch at Choice 6 as well. However, by contrast with what is observed for risk-averse subjects, the numbers of risk-loving subjects who switch at Choices 6 through 9 is similar to the number who switch at Choices 2 through 5.

What is clear from the top portion of the figure and supported statistically is that those who are risk averse are more likely to have greater degrees of absolute and relative prudence than are those who are risk seeking (p-value for χ^2 test $= 0.027$). However, as suggested by the lower portion of the figure prudent and imprudent people do not exhibit substantially different degrees of absolute and relative temperance (p-value for χ^2 test $= 0.596$).

One implication of our results draws upon Jindapon and Neilson (2007) who examine comparative risk aversion in a model where decision makers can exert effort to shift an initial wealth distribution to a preferred distribution. Specifically, Jindapon and Neilson (2007) showed that if the initial distribution differs from the preferred distribution by a simple increase in 3^{rd} degree risk, then a risk-averse agent with preferences captured by the utility function u would invest more effort than another risk-averse agent whose preferences are captured by the utility function v if and only if u has a higher degree of absolute prudence than v.[14] Appendix 7 extends Jindapon and Neilson (2007) by assuming that agents are risk-loving instead of risk averse and showing that when the initial distribution differs from the preferred distribution by a simple increase in 3^{rd} degree risk, then a risk-loving agent with utility function u would invest more effort than another risk-loving agent with utility v if and only if u has a lower degree of absolute prudence than v. Note that subjects who switch at Choice $j + 1$ has a higher degree of absolute prudence than subjects who switch at Choice j regardless of second order risk preferences. Thus, in the top-right portion of Figure 3 subjects who switch at Choice $j+1$ would exert more efforts than the subjects who switch at Choice j or earlier whereas in the top-left portion of Figure 3 subjects who switch at Choice $j+1$ would exert less efforts than subjects who switch at Choice j or earlier.

Jindapon and Neilson (2007) and our Appendix 7 also provide a basis for understanding the implications of observed 4^{th} order behavior. Jindapon and Neilson (2007) showed that if the initial distribution differs from the preferred distribution by a simple increase in 4^{th} degree risk, then a prudent agent with higher degree of absolute temperance is willing to exert more effort to pursue the preferred distribution than is a prudent agent with a lower degree of absolute temperance. Thus, in the bottom-right portion of Figure 3 subjects who switch at Choice $j + 1$ would exert more effort than than subjects who switch at Choice j or earlier. Our Appendix 7 shows that an imprudent agent with a lower degree of absolute temperance is willing to exert more effort to pursue

[14] See Appendix 7 for the definition of a simple increase (decrease) in 3^{rd} degree risk.

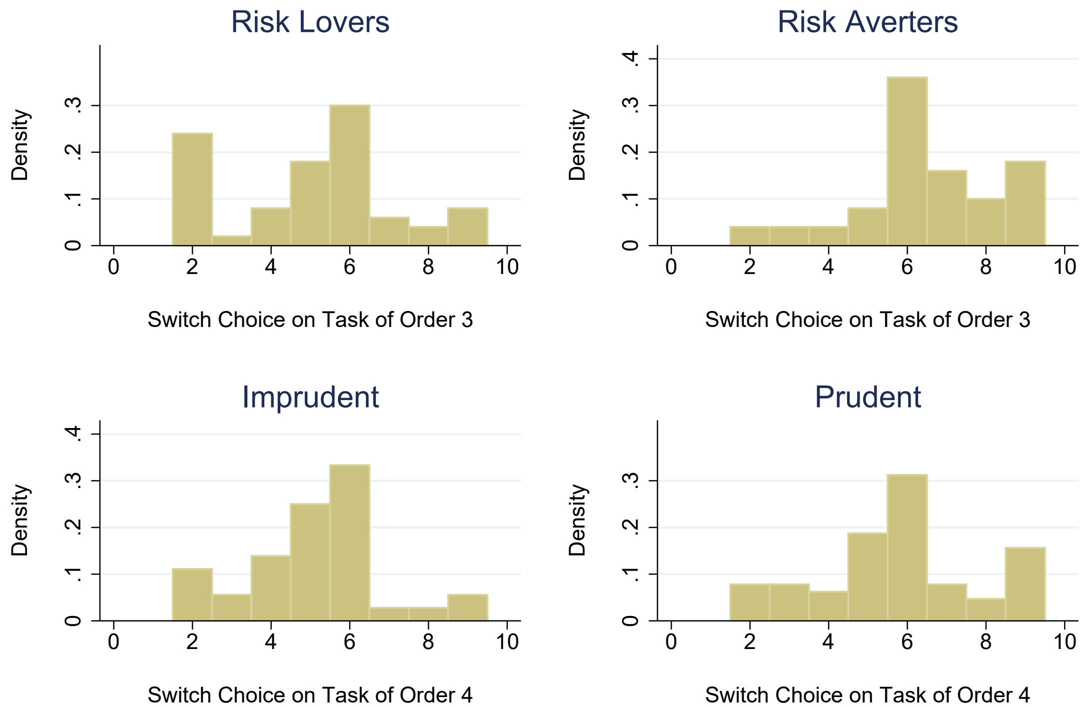

Figure 3: Histogram of Switching in Task of Order n Given Sign of $(n-1)^{th}$ Order Risk Attitude

the preferred distribution than is an imprudent agent with a higher degree of absolute temperance. Thus, in bottom-left portion of Figure 3 subjects who switch at Choice $j+1$ would exert less efforts than subjects who switch at Choice j or earlier.

While not the main focus of our study, the within-subject nature of our data allows us to examine how degrees of risk aversion are related across orders. Specifically, the correlation between switching points on the Tasks of Orders 2 and 3 is 0.209 (p-value $= 0.037$) indicating that a greater degree of absolute (relative) second order risk aversion is associated with a greater degree of absolute (relative) prudence. Similarly, the degrees of absolute (relative) prudence and absolute (relative) temperance are positively correlated among our subjects; the correlation between switching points on the Tasks of Orders 3 and 4 is 0.258 (p-value $= 0.010$). But the strongest relationship that we observer is between the degree of second order risk aversion and the degree of temperance; the correlation between switching points on the Tasks of Orders 2 and 4 is 0.422 (p-value < 0.001). That the connection between even order risk attitudes is stronger than the relationship between even and odd orders is consistent with the notion of people being mixed risk averters and mixed risk seekers.

Finally, while our approach is non-parametric, we report a calibration exercise to determine how well different utility functions describe observed behavior. For each subject for each utility function, we identify the parameter values that best fit the observed behavior of the subject across all three orders.[15] Table **??** reports the mean and standard deviation of the subject specific model parameters. For example, when considering the exponential utility function, the average value of γ across the subjects is 0.15.[16] The table also reports the mean and standard deviation of the accuracy rate for a function when parameter values are subject specific. That is, allowing for each subject to have a unique value of γ, on average 80% of a subject's choices are consistent with the exponential utility function. Among the utility functions listed in Table **??**, the exponential-power utility function has the greatest average accuracy at 91%.[17]

For behavior to be consistent with a well behaved utility function, then a subject should prefer $A_1(5)$ to $B_1(5)$ iff they prefer $A_2(1)$ to $B_2(1)$ as both choices depend only on the sign of $u^{(1)}$. Similarly, the choices between $A_2(5)$ and $B_2(5)$ and $A_3(1)$ and $B_3(1)$ both identify the sign of $u^{(3)}$. If attention is restricted to those subjects who make consistent decisions over these two pairs of choices, then exponential-power utility function's average accuracy rate increases to 98% although this performance is not substantially better than that of the power utility function that has an accuracy rate of 95% among these subjects as shown in the top portion of Table **??**. Interestingly, as shown in the lower portion of Table **??**, the subjects who are not consistent on either pair of choices behave as if they are close to risk neutral on average, which imply they were indifferent for all third and fourth order choices.

6 Conclusions

This paper introduces a simple and systematic procedure for identifying the intensity of risk attitudes using the notion of risk apportionment. Our process is systematic in that it involves a series of

[15]Matlab code for the identification is available in the supplementary materials.

[16]This does not imply that 0.15 is the value of γ that best fits all of the data for the exponential utility function.

[17]Because options in the third order task have equal means and options in the fourth order task have equal means and variances, we do not consider linear or quadratic utility functions as these wold imply indifference over all third and fourth order choices and all fourth order choices, respectively.

binary comparisons where each comparisons differs from the others in the same incremental manner. The process is simple in that it only involves comparisons between two lotteries that are themselves composed of combinations of certain losses and fifty-fifty lotteries. Further, our approach can be used to identify both relative and absolute risk aversion of any arbitrary degree without relying upon assumptions regarding the respondent's underlying preference structure.

We also demonstrate the implementation of our approach in a laboratory setting. Consistent with previous work, we find that a majority of our subjects are non-satiated, risk averse, prudent, and temperate. Our approach allows us to go further and identify that the typical behavior of our subjects is modest relative and absolute risk aversion, modest relative and absolute prudence, and modest relative and absolute temperance. Further, we find that higher order degrees of risk aversion are positively correlated, although the strongest relationship that we observe is between second order risk aversion and prudence. Finally, our calibration exercise suggests that behavior is generally consistent with the exponential-power utility function.

7 Acknowledgments

The authors are grateful to Manel Baucells, Henry Chiu, David Crainich, Paan Jindapon, Liqun Liu, Diego Nocetti, Richard Peter and Ilia Tsetlin for helpful comments. All authors share first authorship, but are listed alphabetically by convention.

References

Arrow, K.J., 1971. *Essays in the theory of risk-bearing.* Markham Publishing, Chicago.

Becker, G.M., DeGroot, M.H. and Marschak, J., 1964. Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3), 226-232.

Bliss, R.R. and Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. *Journal of Finance*, 59(1), 407-446.

Caballé, J. and Pomansky, A., 1996. Mixed risk aversion, *Journal of Economic Theory* 71(2), 485-513.

Carroll, C.D. and Kimball, M.S., 2008. Precautionary Saving and Precautionary Wealth. in Durlauf, N.S. and Blume, L.E. (eds), *The New Palgrave Dictionary of Economics*, 2nd edn. (London: MacMillan).

Chiu, W.H., 2005. Skewness preference, risk aversion, and the precedence relations on stochastic changes. *Management Science*, 51(12), 1816-1828.

Cohen, A. and Einav, L., 2007. Estimating risk preferences from deductible choice. *American Economic Review*, 97(3), 745-788.

Crainich, D., Eeckhoudt, L. and Trannoy, A., 2013. Even (mixed) risk lovers are prudent. *American Economic Review*, 103(4), 1529-1535.

Deck, C. and Schlesinger, H., 2010. Exploring higher order risk effects. *Review of Economic Studies*, 77(4), 1403-1420.

Deck, C. and Schlesinger, H., 2014. Consistency of higher order risk preferences. *Econometrica*, 82(5), 1913-1943.

Denuit, M.M. and Eeckhoudt L., 2010. A general index of absolute risk attitude. *Management Science*, 56, 712-715.

Dynan, K.E., 1993. How prudent are consumers?. *Journal of Political Economy*, 101(6), 1104-1113.

Ebert, S. and Wiesen, D., 2014. Joint measurement of risk aversion, prudence, and temperance. *Journal of Risk and Uncertainty*, 48(3), 231-252.

Eeckhoudt, L., Gollier, C. and Schlesinger, H., 1996. Changes in background risk and risk taking behavior. *Econometrica*, 64(3), 683-689.

Eeckhoudt, L., Rey, B. and Schlesinger, H., 2007. A good sign for multivariate risk taking. *Management Science*, 53(1), 117-124.

Eeckhoudt, L. and Schlesinger, H., 2006. Putting risk in its proper place. *American Economic Review*, 96, 280-289.

Eeckhoudt, L. and Schlesinger, H., 2008. Changes in risk and the demand for saving. *Journal of Monetary Economics*, 55(7), 1329-1336.

Eeckhoudt, L., Schlesinger H. and Tsetlin I., 2009. Apportioning of risks via stochastic dominance. *Journal of Economic Theory*, 144, 994-1003.

Esö, P. and White, L., 2004. Precautionary bidding in auctions. *Econometrica*, 72(1), 77-92.

Gollier, C., 2001. Wealth inequality and asset pricing. *The Review of Economic Studies*, 68(1), 181-203.

Gollier, C., 2018. Stochastic volatility implies fourth-degree risk dominance: Applications to asset pricing. *Journal of Economic Dynamics and Control*, 95, 155-171.

Gollier, C. and Pratt, J.W., 1996. Risk vulnerability and the tempering effect of background risk. *Econometrica*, 64(5), 1109-1123.

Gourinchas, P.O. and Parker, J., 2002. Consumption over the life cycle. *Econometrica*, 70(1), 47-89.

Guiso, L., Jappelli, T. and Terlizzese, D., 1996. Income risk, borrowing constraints, and portfolio choice. *American Economic Review*, 86(1), 158-172.

Haering, A., Heinrich, T. and Mayrhofer, T., 2020. Exploring the consistency of higher order risk preferences. *International Economic Review*, 61(1), 283-320.

Holt, C.A. and Laury, S.K., 2002. Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644-1655.

Jewitt, I., 1989. Choosing between risky prospects: The characterization of comparative statics results, and location independent risk. *Management Science*, 35(1), 60-70.

Jindapon, P. and Neilson, W.S., 2007. Higher-order generalizations of Arrow-Pratt and Ross risk aversion: A comparative statics approach. *Journal of Economic Theory*, 136(1), 719-728.

Jindapon, P., Liu, L. and Neilson, W. S., 2021. Comparative risk apportionment. *Economic Theory Bulletin*, 9(1), 91-112.

Kimball, M.S., 1990. Precautionary saving in the small and in the large. *Econometrica*, 58(1), 53-73.

Kimball, M.S., 1992. Precautionary motives for holding assets. In P. Newman, M. Milgate and J. Eatwell (eds.) *The New Palgrave Dictionary of Money and Finance*. Vol. 3, London: The Macmillan Press Limited, pp. 158-161.

Kimball, M.S., 1993. Standard risk aversion. *Econometrica*, 61(3), 589-611.

Liu, L. and Meyer, J., 2013. Substituting one risk increase for another: A method for measuring risk aversion. *Journal of Economic Theory*, 148(6), 2706-2718.

Menezes, C., Geiss, C. and Tressler, J., 1980. Increasing downside risk. *American Economic Review*, 70(5), 921-932.

Noussair, C.N., Trautmann, S.T. and van de Kuilen, G., 2014. Higher order risk attitudes, demographics, and financial decisions. *Review of Economic Studies*, 81(1), 325-355.

Pratt, J.W., 1964. Risk aversion in the small and in the large. *Econometrica*, 32(1), 122-136

Rothschild, M. and Stiglitz, J.E., 1970. Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3), 225-243.

Schneider, S. O. and Sutter, M., 2021. Higher order risk preferences: New measures, determinants and field behavior. Working Paper, MPI Collective Goods Discussion Paper, (2020/22).

White L., 2008. Prudence in bargaining: The effect of uncertainty on bargaining outcomes. *Games and Economic Behavior*, 62(1), 211-231.

Appendix A Mathematical Proofs

To prove our formal results, we first establish three technical lemmas.

Lemma A1. *For u and v that are twice continuously differentiable with $u^{(1)}$ and $v^{(1)}$ having the same sign, the following statements are equivalent:*

(i) $-\frac{u^{(2)}(x)}{u^{(1)}(x)} \geq -\frac{v^{(2)}(x)}{v^{(1)}(x)}$ *for all* $x \in [a, b]$*;*

(ii) $-\frac{u^{(1)}(x)-u^{(1)}(x-k)}{u(x)-u(x-k)} \geq -\frac{v^{(1)}(x)-v^{(1)}(x-k)}{v(x)-v(x-k)}$ *for all* $x, x-k \in [a, b]$ *with* $k > 0$.

Proof. Assume first that $u^{(1)} > 0$ and $v^{(1)} > 0$. Since (i) is a direct consequence of (ii) after letting $k \to 0$, we concentrate on the proof that (i) implies (ii). Statement (i) is equivalent to $u(x) = \varphi(v(x))$, where φ is twice differentiable with $\varphi^{(1)} > 0$ and $\varphi^{(2)} \leq 0$ (Pratt 1964). By the mean value theorem, $u(x) - u(x-k) = \varphi(v(x)) - \varphi(v(x-k)) = \varphi^{(1)}(v(x-\theta k))[v(x) - v(x-k)]$ with some $\theta \in (0, 1)$, and accordingly,

$$
\begin{aligned}
-\frac{u^{(1)}(x) - u^{(1)}(x-k)}{u(x) - u(x-k)} &= -\frac{\varphi^{(1)}(v(x))v^{(1)}(x) - \varphi^{(1)}(v(x-k))v^{(1)}(x-k)}{\varphi^{(1)}(v(x-\theta k))[v(x) - v(x-k)]} \\
&\geq -\frac{v^{(1)}(x) - v^{(1)}(x-k)}{v(x) - v(x-k)},
\end{aligned}
$$

which follows from $\varphi^{(1)}(v(x)) \leq \varphi^{(1)}(v(x-\theta k)) \leq \varphi^{(1)}(v(x-k))$.

To address the alternative case with $u^{(1)} < 0$ and $v^{(1)} < 0$, we introduce $\hat{u} = -u$ and $\hat{v} = -v$, which satisfy $\hat{u}^{(1)} > 0$ and $\hat{v}^{(1)} > 0$. Because $-\frac{\hat{u}^{(2)}(x)}{\hat{u}^{(1)}(x)} = -\frac{u^{(2)}(x)}{u^{(1)}(x)}$, $-\frac{\hat{u}^{(1)}(x)-\hat{u}^{(1)}(x-k)}{\hat{u}(x)-\hat{u}(x-k)} = -\frac{u^{(1)}(x)-u^{(1)}(x-k)}{u(x)-u(x-k)}$ and similar equations also hold for $\hat{v}$ and v, the result follows straightforwardly by adapting the former analysis to $\hat{u}$ and $\hat{v}$. Q.E.D.

Lemma A2 below extends Lemma A1 to higher orders.

Lemma A2. *For $n \geq 3$, let $\hat{A}_n$ and $\hat{B}_n$ be the lotteries introduced by Eeckhoudt and Schlesinger (2006) with $w = 0$ and $\tilde{\varepsilon}_i$ as in (5). For u and v that are continuously differentiable up to order n with $u^{(n-1)}$ and $v^{(n-1)}$ having the same sign, define*

$$
\begin{aligned}
u_{n-2}(x) &\equiv \mathbb{E}u(x + \hat{A}_{n-2}) - \mathbb{E}u(x + \hat{B}_{n-2}), \\
v_{n-2}(x) &\equiv \mathbb{E}v(x + \hat{A}_{n-2}) - \mathbb{E}v(x + \hat{B}_{n-2}).
\end{aligned}
$$

The following statements are equivalent:

(i) $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)} \geq -\frac{v^{(n)}(x)}{v^{(n-1)}(x)}$ *for all* $x \in [a,b]$*;*

(ii) $-\frac{u^{(2)}_{n-2}(x)}{u^{(1)}_{n-2}(x)} \geq -\frac{v^{(2)}_{n-2}(x)}{v^{(1)}_{n-2}(x)}$ *for all* $x + \hat{A}_{n-2}, x + \hat{B}_{n-2} \in [a,b]$.

Proof. For $n = 3$, $u_1(x) = u(x-k) - u(x)$ and $v_1(x) = v(x-k) - v(x)$. We apply Lemma A1 with $u^{(1)}$ and $v^{(1)}$ to get the equivalence between (i) and (ii).

For $n = 4$, $u_2(x) = \frac{1}{2}[u(x-k_1) + u(x+k_1)] - u(x)$ and $v_2(x) = \frac{1}{2}[v(x-k_1) + v(x+k_1)] - v(x)$. Accordingly, (ii) implies (i) after letting $k_1 \to 0$. To prove that (i) implies (ii), we apply Lemma A1 to $u^{(2)}$ and $v^{(2)}$, and obtain that (i) implies

$$-\frac{u^{(3)}(x) - u^{(3)}(x-k_1)}{u^{(2)}(x) - u^{(2)}(x-k_1)} \geq -\frac{v^{(3)}(x) - v^{(3)}(x-k_1)}{v^{(2)}(x) - v^{(2)}(x-k_1)}.$$

Similarly, applying Lemma A1 to $\hat{u}(x) \equiv u^{(1)}(x) - u^{(1)}(x-k_1)$ and $\hat{v}(x) \equiv v^{(1)}(x) - v^{(1)}(x-k_1)$ leads to

$$-\frac{\hat{u}^{(1)}(x+k_1) - \hat{u}^{(1)}(x)}{\hat{u}(x+k_1) - \hat{u}(x)} \geq -\frac{\hat{v}^{(1)}(x+k_1) - \hat{v}^{(1)}(x)}{\hat{v}(x+k_1) - \hat{v}(x)},$$

which yields (ii) as $\hat{u}(x+k_1) - \hat{u}(x) = 2u_2^{(1)}(x)$ and $\hat{v}(x+k_1) - \hat{v}(x) = 2v_2^{(1)}(x)$.

For $n \geq 5$ the proof is by induction. Suppose that the equivalence between (i) and (ii) holds true for all orders up to $n-1$. For order n, recall that $\tilde{\varepsilon}_{\text{Int}(n/2)-1} = [-k_{\text{Int}(n/2)-1}; k_{\text{Int}(n/2)-1}]$ and

$$\begin{aligned}
u_{n-2}(x) &= \frac{1}{4}\left[u_{n-4}\left(x - k_{\text{Int}(n/2)-1}\right) + u_{n-4}\left(x + k_{\text{Int}(n/2)-1}\right)\right] - \frac{1}{2}u_{n-4}(x), \\
v_{n-2}(x) &= \frac{1}{4}\left[v_{n-4}\left(x - k_{\text{Int}(n/2)-1}\right) + v_{n-4}\left(x + k_{\text{Int}(n/2)-1}\right)\right] - \frac{1}{2}v_{n-4}(x).
\end{aligned}$$

Applying the equivalence for the fourth order to u_{n-4} and v_{n-4}, we obtain that (ii) is equivalent to

$$-\frac{u^{(4)}_{n-4}(x)}{u^{(3)}_{n-4}(x)} \geq -\frac{v^{(4)}_{n-4}(x)}{v^{(3)}_{n-4}(x)} \quad \text{for all } x + \hat{A}_{n-4}, x + \hat{B}_{n-4} \in [a,b].$$

We further apply the equivalence for the $(n-4)^{th}$ order to $u^{(2)}$ and $v^{(2)}$ to obtain that (ii) is equivalent to (i) for order n. Q.E.D.

Lemma A3. *For $n \geq 2$, consider a Task of Order n as specified in (7) that is supported on $[a, b] \subset (0, \infty)$. For each $j = 2, ..., J-1$, there exists a unique constant absolute risk aversion coefficient $\Theta_1(n, j)$ such that $\mathbb{E}v_1(A_n(j)) = \mathbb{E}v_1(B_n(j))$ under $\theta_1 = \Theta_1(n, j)$. Moreover, $\Theta_1(n, j)$ increases strictly in j. The same statement holds true for $\Theta_2(n, j)$, the constant relative risk aversion coefficient such that $\mathbb{E}v_2(A_n(j)) = \mathbb{E}v_2(B_n(j))$.*

Proof. We proceed under the assumption that $(-1)^n v_1^{(n-1)} > 0$ as the alternative case $(-1)^n v_1^{(n-1)} < 0$ can be addressed by adapting the analysis to $-v_1$. By induction, we have $v_{1(n-2)}^{(1)} > 0$ and $\mathbb{E}v_1(B_n(j)) - \mathbb{E}v_1(A_n(j)) = v_{1(n-2)}(w) - \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_j)$.

For $j = 2, ..., J-1$, the existence and uniqueness of $\Theta_1(n, j)$ follow from the monotonicity of $v_{1(n-2)}^{-1}(\mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_j))$ with respect to θ_1 (Pratt 1964), together with the facts that $\lim_{\theta_1 \to -\infty} v_{1(n-2)}^{-1}(\mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_j)) = w + \text{ess sup}(\tilde{\delta}_j)$ and $\lim_{\theta_1 \to \infty} v_{1(n-2)}^{-1}(\mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_j)) = w + \text{ess inf}(\tilde{\delta}_j)$.

As j increases, $\mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_j)$ increases strictly, but $\mathbb{E}v_{1(n-2)}(w)$ does not change. Thus, for any $j_1 < j_2$, $\mathbb{E}v_{1(n-2)}(w) = \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_{j_1})$ under $\theta_1 = \Theta_1(n, j_1)$ implies $\mathbb{E}v_{1(n-2)}(w) < \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_{j_2})$ under the same θ_1, which further implies $\mathbb{E}v_{1(n-2)}(w) < \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_{j_2})$ under all $\theta_1 \leq \Theta_1(n, j_1)$ by virtue of Proposition 1. Accordingly, to achieve $\mathbb{E}v_{1(n-2)}(w) = \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_{j_2})$ under $\theta_1 = \Theta_1(n, j_2)$, it must hold that $\Theta_1(n, j_2) > \Theta_1(n, j_1)$, which proves the monotonicity of $\Theta_1(n, j)$ with respect to j.

When $j = 1$, $\mathbb{E}v_{1(n-2)}(w) > \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_1)$ for any finite $\Theta_1 \in \mathbb{R}$ yielding the convention $\Theta_1(n, 1) = -\infty$. When $j = J$, $\mathbb{E}v_{1(n-2)}(w) < \mathbb{E}v_{1(n-2)}(w + \tilde{\delta}_J)$ for any finite $\Theta_1 \in \mathbb{R}$ yielding the convention $\Theta_1(n, j) = \infty$. Q.E.D.

In what follows, we prove our main results.

Proof of Proposition 1. To prove that (i) implies (ii), let F and G be the cumulative distribution functions of B_2 and A_2, respectively. Then, integration-by-parts yields $\mathbb{E}v(B_2) - \mathbb{E}v(A_2) =$

$\int_a^b v^{(1)}(x)[G(x) - F(x)]dx$ and

$$\begin{aligned}
&\mathbb{E}u(B_2) - \mathbb{E}u(A_2) \\
&= \int_a^b \frac{u^{(1)}(x)}{v^{(1)}(x)} \left(-\frac{u^{(2)}(x)}{u^{(1)}(x)} + \frac{v^{(2)}(x)}{v^{(1)}(x)} \right) \int_a^x v^{(1)}(y)[G(y) - F(y)]dydx \\
&\quad + \frac{u^{(1)}(b)}{v^{(1)}(b)} \int_a^b v^{(1)}(x)[G(x) - F(x)]dx.
\end{aligned}$$

Since F intersects G from below once, $\mathbb{E}v(B_2) = \mathbb{E}v(A_2)$ implies $\int_a^b v^{(1)}(y)[G(y) - F(y)]dy = 0$ but $\int_a^x v^{(1)}(y)[G(y) - F(y)]dy \geq 0$ for all $x \in [a, b]$, which in turn implies $\mathbb{E}u(B_2) \geq \mathbb{E}u(A_2)$.

To show that (ii) implies (i), we argue by contradiction. If $-\frac{u^{(2)}(x_0)}{u^{(1)}(x_0)} < -\frac{v^{(2)}(x_0)}{v^{(1)}(x_0)}$ for some $x_0 \in [a, b]$, then continuity implies $-\frac{u^{(2)}(x)}{u^{(1)}(x)} < -\frac{v^{(2)}(x)}{v^{(1)}(x)}$ on $(x_0 - \varepsilon, x_0 + \varepsilon) \cap [a, b]$ for some $\varepsilon > 0$. For $\tilde{\delta}$ supported on $(x_0 - \varepsilon, x_0 + \varepsilon) \cap [a, b]$, we can follow the proof for (i) implying (ii) to obtain that $\mathbb{E}v(B_2) = \mathbb{E}v(A_2)$ always implies $\mathbb{E}u(B_2) \leq \mathbb{E}u(A_2)$ and a strict inequality $\mathbb{E}u(B_2) < \mathbb{E}u(A_2)$ holds when $\tilde{\delta}$ is not constant. This a contradiction to (ii). Q.E.D.

Proof of Proposition 2. For $k > 0$, recall $u_1(x) = u(x-k) - u(x)$ and $v_1(x) = v(x-k) - v(x)$. For individuals with utility function u, the choice between A_3 and B_3 is based on comparing $\mathbb{E}u_1(A_2)$ and $\mathbb{E}u_1(B_2)$ with a similar result holding for v. We apply Proposition 1 to obtain that (ii) is equivalent to $-\frac{u_1^{(2)}(x)}{u_1^{(1)}(x)} \geq -\frac{v_1^{(2)}(x)}{v_1^{(1)}(x)}$ for all $x, x - k \in [a, b]$, and further apply Lemma A2 at the third order to obtain that (ii) is equivalent to (i). Q.E.D.

Proof of Theorem 1. For $n \geq 3$, let u_{n-2} and v_{n-2} be defined as in Lemma A2. For an individual with utility function u, the choice between A_n and B_n is based on comparing $\mathbb{E}u_{n-2}(A_2)$ and $\mathbb{E}u_{n-2}(B_2)$. A similar statement holds for v. We apply Proposition 1 to obtain that (ii) is equivalent to $-\frac{u_{n-2}^{(2)}(x)}{u_{n-2}^{(1)}(x)} \geq -\frac{v_{n-2}^{(2)}(x)}{v_{n-2}^{(1)}(x)}$ for all $x + A_{n-2}, x + B_{n-2} \in [a, b]$, and further apply Lemma A2 to obtain that (ii) is equivalent to (i). Q.E.D.

Proof of Corollary 1. We assume $(-1)^n u^{(n-1)} > 0$ as the case $(-1)^n u^{(n-1)} < 0$ can be addressed by adapting the above analysis to $-u$. Recall that for an individual with utility function u, the choice between $A_n(j)$ and $B_n(j)$ is based on comparing $\mathbb{E}u_{n-2}(w + \tilde{\delta}_j)$ and $\mathbb{E}u_{n-2}(w)$. We prove by induction that $u_{n-2}^{(1)} > 0$. When $j = 1$, $\tilde{\delta}_1 = [-2h; 0]$ is a first-degree deterioration relative to 0 and thus w is preferred to $w + \tilde{\delta}_1$ by u_{n-2}; when $j = J$, $\tilde{\delta}_J = [0; 2h]$ is a first-degree improvement relative to 0 and thus $w + \tilde{\delta}_J$ is preferred to w by u_{n-2}. As j increases from 1 to J, $\mathbb{E}u_{n-2}(w + \tilde{\delta}_j)$

increases but $\mathbb{E}u_{n-2}(w)$ does not change, yielding a single point at which the individual switches from preferring w to preferring $w+\tilde{\delta}_j$ under u_{n-2}, or equivalently, from preferring $B_n(j)$ to preferring $A_n(j)$ under u.

When the individual prefers $B_n(j)$ to $A_n(j)$ for $j \leq j^*$ and prefers $A_n(j)$ to $B_n(j)$ for $j \geq j^* + 1$, we have $\mathbb{E}u(B_n(j^*)) \geq \mathbb{E}u(A_n(j^*))$ and $\mathbb{E}u(B_n(j^*+1)) \leq \mathbb{E}u(A_n(j^*+1))$. If otherwise $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)} < \Theta_1(j^*)$ or $-\frac{u^{(n)}(x)}{u^{(n-1)}(x)} > \Theta_1(j^*+1)$ for all $x \in [a,b]$, then Theorem 1 would imply $\mathbb{E}u(A_n(j^*)) > \mathbb{E}u(B_n(j^*))$ or $\mathbb{E}u(A_n(j^*+1)) < \mathbb{E}u(B_n(j^*+1))$, which is a contradiction. Q.E.D.

Appendix B Effort-Making Problem

Jindapon and Neilson (2007) examined comparative risk aversion in a model where decision makers can exert effort to shift an initial wealth distribution (G) to a preferred distribution (F). Given that u and v are utility functions that exhibit risk-aversion and/or prudence, their paper provides conditions on u and v for an individual whose preference is captured by u to exert more effort than an individual whose preference is captured by v. In this appendix, we extend their model to consider individuals who are risk loving and/or imprudent.

Assume that an individual with utility function u can invest in an effort $e \in [0,1]$ with a non-monetary cost of effort $c(e)$ where $c'(e) > 0$ and $c''(e) > 0$ such that her wealth distribution will become $eF(x) + (1-e)G(x)$, where $x \in [a,b]$. Thus, the objective of the individual is as follows:

$$\max_{e\in[0,1]} \int_a^b u(x)[edF(x) + (1-e)dG(x)] - c(e).$$

The first-order condition is

$$\int_a^b u(x)d[F(x) - G(x)] - c'(e) = 0,$$

and the second-order condition holds automatically due to $-c''(e) < 0$. Since the second-order condition does not depend on the sign of $u^{(2)}$, there always exists a unique interior solution if $\int_a^b u(x)d[F(x) - G(x)] > 0$, regardless of whether or not the individual is risk averse or risk loving.

Define $F^{(2)}(x) = \int_a^x F(t)dt$ and $G^{(2)}(x) = \int_a^x G(t)dt$. Following Jindapon and Neilson (2007), we assume that F differs from G by a simple decrease in third order risk, i.e., $F^{(2)}(x)$ crosses $G^{(2)}(x)$

only once from below and $E_F(x) = E_G(x)$. Taking integration by parts, the first-order condition becomes

$$\int_a^b [-u^{(2)}(x)][G^{(2)}(x) - F^{(2)}(x)]dx - c'(e) = 0.$$

To compare two individuals with utility functions u and v, denote the optimal effort levels corresponding to u and v by e_u^* and e_v^*, respectively.

Proposition B1. *Regarding the effort-making problem, we have:*

(i) Under $u^{(2)} < 0$ and $v^{(2)} < 0$, $e_u^* \geq e_v^*$ for any F and G such that F differs from G by a simple decrease in third degree risk, if and only if $-\frac{u^{(3)}(x)}{u^{(2)}(x)} \geq -\frac{v^{(3)}(x)}{v^{(2)}(x)}$ for all $x \in [a, b]$;

(ii) Under $u^{(2)} > 0$ and $v^{(2)} > 0$, $e_u^* \geq e_v^*$ for any F and G such that F differs from G by a simple decrease in third degree risk, if and only if $-\frac{u^{(3)}(x)}{u^{(2)}(x)} \leq -\frac{v^{(3)}(x)}{v^{(2)}(x)}$ for all $x \in [a, b]$.

Statement (i) in the above proposition is owing to Jindapon and Neilson (2007). Here, we examine the "if" part of statement (ii). Assume $F^{(2)}(x) \leq G^{(2)}(x)$ for $x \leq x_0$ and $F^{(2)}(x) \geq G^{(2)}(x)$ for $x \geq x_0$, and scale u and v so that $u^{(2)}(x_0) = v^{(2)}(x_0)$. In light of the first-order condition, $e_u^* \geq e_v^*$ if and only if

$$\int_a^b \left[-\frac{u^{(2)}(x)}{u^{(2)}(x_0)} + \frac{v^{(2)}(x)}{v^{(2)}(x_0)} \right] [G^{(2)}(x) - F^{(2)}(x)]dx \geq 0.$$

If $-\frac{u^{(3)}(x)}{u^{(2)}(x)} \leq -\frac{v^{(3)}(x)}{v^{(2)}(x)}$ for all $x \in [a, b]$, then it holds $\frac{u^{(2)}(y)}{u^{(2)}(z)} \leq \frac{v^{(2)}(y)}{v^{(2)}(z)}$ for all $z \geq y$. Accordingly, we have $F^{(2)}(x) \leq G^{(2)}(x)$ and $\frac{u^{(2)}(x)}{u^{(2)}(x_0)} \leq \frac{v^{(2)}(x)}{v^{(2)}(x_0)}$ for $x \leq x_0$ and $F^{(2)}(x) \geq G^{(2)}(x)$ and $\frac{u^{(2)}(x)}{u^{(2)}(x_0)} \geq \frac{v^{(2)}(x)}{v^{(2)}(x_0)}$ for $x \geq x_0$, which in turn implies the desired inequality. The "only if" portion of the statement can be proved using the same approach as in the proof of Theorem 3 in Jindapon and Neilson (2007).

By the same token, we can extend the above analysis to the fourth order. We say F differs from G by a simple decrease in fourth degree risk, if $F^{(3)}(x) = \int_a^x F^{(2)}(t)dt$ crosses $G^{(3)}(x) = \int_a^x G^{(2)}(t)dt$ only once from below, and moreover, $E_F(x) = E_G(x)$ and $E_F(x^2) = E_G(x^2)$ hold. The result is formally stated as follows, and the proof is omitted.

Proposition B2. *Regarding the effort-making problem, we have:*

(i) *Under* $u^{(3)} > 0$ *and* $v^{(3)} > 0$, $e_u^* \geq e_v^*$ *for any* F *and* G *such that* F *differs from* G *by a simple decrease in fourth degree risk, if and only if* $-\frac{u^{(4)}(x)}{u^{(3)}(x)} \geq -\frac{v^{(4)}(x)}{v^{(3)}(x)}$ *for all* $x \in [a, b]$;

(ii) *Under* $u^{(3)} < 0$ *and* $v^{(3)} < 0$, $e_u^* \geq e_v^*$ *for any* F *and* G *such that* F *differs from* G *by a simple decrease in fourth degree risk, if and only if* $-\frac{u^{(4)}(x)}{u^{(3)}(x)} \leq -\frac{v^{(4)}(x)}{v^{(3)}(x)}$ *for all* $x \in [a, b]$.

A Complete Characterization of Downside Risk Preference

Donald C. Keenan
University of Cergy-Pontoise
dkeenan@uga.edu

Arthur Snow
University of Georgia
snow@uga.edu

Abstract

We characterize third-order risk preference in expected utility theory by utility transformations and by rankings of risk-preference measures. At the second order, a risk-averse transformation is exactly opposite to a risk-loving transformation, and is replicated by the ranking of Arrow-Pratt measures r. However, at the third order, transformations that introduce aversion correspond to rankings by utility measures that are not opposites, as u being more averse than risk-neutral utility i is equivalent to Kimball's prudence measure p being positive, but i being less averse than u requires that p exceed three times r. We resolve this paradox and shed light on previously reported comparative statics predictions based these extremes.

1 Introduction

In expected utility theory, risk preferences are dictated by the derivatives of the von Neumann-Morgenstern utility function *u* defined on income $y > 0$. As any utility function *u* is a transformation φ of risk-neutral utility $i(y) = y$, these derivatives are exactly those of the transformation generating *u*. Assuming that marginal utilities are always positive, reflecting non-satiation, *direction* of *n*th-order risk preference is indicated by the sign of the *n*th utility derivative divided by the first. Thus, at the second order, aversion to bearing risk for $u = \varphi(i)$ is identified with a positive Arrow-Pratt index $r_u = -u'' / u'$, because decision makers with concave utility $(u'' < 0)$ always dislike any increase in risk [Rothschild & Stiglitz (1970)].[1] At the third order, direction is indicated by the sign of the measure $d_u = u''' / u'$, introduced by Crainich & Eeckhoudt (2008), because these decision makers always dislike any increase in downside risk [Menezes et al. (1980)].

As emphasized by Eeckhoudt (2012), risk preference refers to both direction and *intensity*. At the second order, a transformation of utility $\varphi(v)$ increases the intensity of risk aversion if and only if φ is itself risk averse, that is, $r_\varphi > 0$ [Pratt (1964)], and successive risk-averse transformations produce a strict partial *ordering* of utilities by greater risk aversion. At the third order, however, successive downside risk-averse

[1] To accommodate the definitions of increased risk and increased downside risk, we assume that income *y* belongs to a nonempty, compact positive interval. Throughout, inequalities are assumed to hold for all incomes in the interval, and we use primes to denote derivatives.

transformations, satisfying $d_\varphi > 0$, do not necessarily yield strict partial orderings.[2] This deficiency is remedied if the transformations are required to be risk averse as well as downside risk averse. With $r_\varphi > 0$ and $d_\varphi > 0$, successive transformations yield an ordering of utilities by greater risk-averse and downside risk-averse preference [Keenan & Snow (2016)].

When the reference for comparison v is risk neutral, we find that utility $u = \varphi(i)$ is risk averse and downside risk averse if and only if $r_u = r_\varphi > 0$ and $d_u = d_\varphi > 0$. These conditions are equivalent to $r_u > 0$ and $p_u > 0$, where $p_u = -u'''/u''$ is the index of prudence introduced by Kimball (1990), since $d_u = p_u r_u$. However, we also find that risk-neutral utility is less risk averse and less downside risk averse than $u = \varphi(i)$ if and only if $r_u > 0$ and $p_u - 3r_u > 0$. Thus, the measure conditions required for less aversion are stronger than those required for greater aversion.

In this paper, we resolve this paradox between greater and less aversion, and in so doing develop a complete characterization of downside risk preference that encompasses greater and less aversion. In the next section, we characterize direction and intensity for third-order risk preference, and conditions necessary and sufficient for strict partial orderings by greater intensity in terms of restrictions on the risk preferences of utility transformations. It is rarely possible to determine the transformation that converts the

---

[2] A strict partial ordering is asymmetric and transitive. Lacking these properties, a ranking of utilities cannot consistently yield unambiguous comparative statics predictions. A case in which a ranking by $d_\varphi > 0$ is symmetric rather than asymmetric is provided by utility $u = \varphi(i) = -1/y$ and utility $i = \psi(u) = -1/u$, which are downside risk-averse transformations of each other. Liu & Meyer (2012) provide examples illustrating both intransitivity and symmetry.

risk preferences of one utility function into those of another. Hence, a characterization of risk preference defined in terms of utility transformations is lent tractability when it has a parallel representation in terms of utility *measures* as illustrated by the preceding discussion. At the second order, this role is served by the Arrow-Pratt index r_u, which indicates direction and intensity, and yields partial orderings of utility functions by greater risk aversion. Characterizing conditions for downside risk preference defined in terms of utility measure are identified in section 3.

In section 4, we investigate reversibility for third-order risk preference, an intrinsic although seldom recognized property of greater risk aversion. Specifically, reversing the risk preference embodied in a transformation from averse to loving, reverses the resulting utility ranking. For greater downside risk aversion, *reversibility* ensures that predictions for less aversion reverse those for greater aversion. Hence, by ensuring reversibility, we resolve the conflict between the conditions for greater and less aversion outlined above.

Finally, a complete characterization identifies a *comparative statics* thought experiment that identifies greater aversion toward bearing risk. At the second order, this role is served by the risk premium, the decision maker's willingness to pay to avoid risk, which is always greater after a risk-averse transformation of utility. At the third order, it is common to associate downside risk aversion with prudence and greater prudence with greater downside risk aversion, as observed by Crainich & Eeckhoudt (2008). In section 5, we contrast direction and comparative statics predictions for greater prudence and greater downside risk aversion. Conclusions are offered in section 6.

2 Direction, Intensity, and Ordering for Downside Risk Preference

We begin with a transformation of utility *v*, $u = \varphi(v)$, and the relationship between the attitude toward risk-bearing embodied in the transformation φ and the risk aversion measures for *v* and *u*,

$$r_u = r_\varphi v' + r_v, \tag{1}$$

obtained by dividing the second derivative of *u* by its first. Unless otherwise specified, the risk preferences assumed for the transformation φ are independent of those for either the final utility *u* or the reference utility *v*. Equation (1) shows that, in contrast, second-order risk preference for *u* is conditional on the second-order preference of both the transformation and the reference utility.

However, whether *v* or *u* is chosen as the reference has no bearing on either the direction of risk preference for *u* or its preference intensity relative to *v*, both as indicated by their Arrow-Pratt measures. Thus, with the inverse transformation denoted by $\psi = \varphi^{-1}$, we have $v = \psi(u)$, and

$$r_v = r_\psi u' + r_u. \tag{2}$$

Together, equations (1) and (2) imply

$$r_\varphi = -r_\psi / \psi', \tag{3}$$

indicating that φ is risk averse if and only if its inverse ψ is risk loving. For the special case in which the reference *v* is risk neutral, utility $u = \varphi(i)$ takes on the risk preferences of φ.

Extending the link between risk-averse transformations $(r_\varphi > 0)$ and greater risk aversion from the second to the third order ties downside risk-averse transformations

$(d_\varphi > 0)$ to greater downside risk aversion [Keenan & Snow (2002), (2009)]. However, as noted, a ranking of utility functions by successive transformations satisfying $d_\varphi > 0$ is not necessarily either asymmetric or transitive, and therefore cannot produce reliable comparative statics predictions. By requiring that the transformations are risk averse as well as downside risk averse, that is,

$$r_\varphi > 0 \text{ and } d_\varphi > 0, \tag{4}$$

rankings generated by these transformations are asymmetric and transitive, and therefore constitute strict partial orderings [Keenan & Snow (2016)].

When the reference utility is risk neutral, the inequality conditions stated at (4) imply that utility $u = \varphi(i)$ exhibits risk averse and downside risk averse preferences, $r_u = r_\varphi > 0$ and $d_u = d_\varphi > 0$. Accordingly, we associate conditions (4) with *direction* of downside risk-averse preference. When the reference utility *v* satisfies these direction conditions, we associate transformations satisfying conditions (4) with greater *intensity* of downside risk-averse preference and with *orderings* by greater downside risk aversion.

3 Measures for Downside Risk Preference

To derive measure representations of the transformation conditions $r_\varphi > 0$ and $d_\varphi > 0$, we obtain

$$d_u = d_\varphi v'^2 + 3r_\varphi r_v v' + d_v \tag{5}$$

from the first and third derivatives of $u = \varphi(v)$, then substitute for r_φ from equation (1), add and subtract $3r_u^2$, and rearrange terms to arrive at

$$d_\varphi = [(d_u - 3r_u^2) - (d_v - 3r_v^2) + 3r_u(r_u - r_v)] / v'^2. \tag{6}$$

To consolidate notation, we introduce the measure

$$D_v = d_v - 3r_v^2. \tag{7}$$

The following is now an immediate consequence of equations (1) and (6), equating inequality restrictions (4) for greater downside risk aversion with restrictions on utility measures.

Proposition 1 [Keenan & Snow (2022)] Given $u = \varphi(v)$, we have

$$r_\varphi > 0 \text{ and } d_\varphi > 0 \text{ if and only if } r_u > r_v \text{ and } (D_u - D_v) + 3r_u(r_u - r_v) > 0.$$

Thus, a ranking of utility functions defined by transformations that are risk averse and downside risk averse is equally represented by restrictions on the changes in the utility *measures* r and D. The Proposition identifies restrictions sufficient for $r_\varphi > 0$ and $d_\varphi > 0$ as, with $r_v > 0$, if both r_v and D_v increase after $u = \varphi(v)$ replaces v, then $r_\varphi > 0$ and $d_\varphi > 0$, implying that the transformation increases downside risk aversion.

4 Reversibility for Downside Risk Preference

A transformation of utility that increases risk aversion is reversible since the transformation must be risk averse and its inverse, a risk-loving transformation, yields the reverse ranking by less risk aversion. As a complement to Proposition 1, the following is a further consequence of equations (1) and (6) establishing *reversibility* for an ordering by greater downside risk-averse preference.

Proposition 2 [Keenan & Snow (2022)] Given $u = \varphi(v)$, we have

(a) $r_\varphi > 0$ and $d_\varphi > 0$ if and only if $r_u > r_v$ and $D_u - D_v + 3r_u(r_u - r_v) > 0$;

(b) $r_\varphi < 0$ and $d_\varphi < 0$ if and only if $r_u < r_v$ and $D_u - D_v + 3r_u(r_u - r_v) < 0$.

Part (a) restates Proposition 1 characterizing greater downside side risk aversion, while part (b) reverses the inequality conditions (4) and characterizes less downside risk aversion. The two parts provides sufficient conditions in terms of the measures r and D for greater and less aversion, respectively. By exploiting the relation $d_v = p_v r_v$, we can rewrite equation (7) as $D_v = r_v(p_v - 3r_v)$. Then the final inequality in part (a) is satisfied if $r_u(p_u - 3r_u) > r_v(p_v - 3r_v)$, while the reverse inequality is sufficient for the final inequality in part (b).

Corollary 1 Given $u = \varphi(v)$, we have

(a) if $r_u > r_v > 0$ and $p_u - 3r_u > p_v - 3r_v > 0$, then $r_\varphi > 0$ and $d_\varphi > 0$;

(b) if $0 < r_u < r_v$ and $0 < p_u - 3r_u < p_v - 3r_v$, then $r_\varphi < 0$ and $d_\varphi < 0$.

Thus, conditional on $r_v > 0$ and $p_v - 3r_v > 0$, if these measures both increase when $u = \varphi(v)$ replaces v, then φ increases downside risk aversion, and φ reduces downside risk aversion if $r_u > 0$ and $p_u - 3r_u > 0$, and both increase when the reference utility v replaces the final utility u.

Reversibility for transformation φ at the third order, however, does not imply reversibility for its inverse $\psi = \varphi^{-1}$, as it does at the second order where reversibility follows from equation (3). For the inverse transformation, d_{ψ} is given by

$$d_{\psi} = [(D_v - D_u) + 3r_v(r_v - r_u)] / u'^2, \tag{8}$$

in parallel with equation (6). The next result characterizes reversibility at the third order for the inverse transformation ψ.

Proposition 3 Given $v = \psi(u)$, we have

(a) $r_{\psi} < 0$ and $d_{\psi} < 0$ if and only if $r_u > r_v$ and $D_u - D_v + 3r_v(r_u - r_v) > 0$;

(b) $r_{\psi} > 0$ and $d_{\psi} > 0$ if and only if $r_u < r_v$ and $D_u - D_v + 3r_v(r_u - r_v) < 0$.

The inequality restrictions imposed in Propositions 2 and 3 differ precisely because the reference and final utilities differ in their preference intensity with respect to risk aversion. Assuming that v and u are risk averse, the inequalities in part (a) of Proposition 3 imply those in part (a) of Proposition 2, since in both instances we then have $r_u > r_v > 0$ and therefore $r_u(r_u - r_v) > r_v(r_u - r_v) > 0$, while those in part ($b$) of Proposition 2 imply those in part (b) of Proposition 3, as in that case we have $0 > r_u(r_u - r_v) > r_v(r_u - r_v)$. Thus, the inverse ψ being a downside risk-loving transformation implies that φ is a downside risk-averse transformation, but not vice versa, while φ being a downside risk-loving transformation implies that ψ is a downside risk-averse transformation, but not vice versa.

For the special case in which the reference utility v is risk neutral, Propositions 2 3 imply the following paradox, described in the introduction, after exploiting the relations $d_v = p_v r_v$ and $d_u = p_u r_u$.

Corollary 2 Given $u = \varphi(i)$ and $i = \psi(u)$, we have

(a) $r_\varphi > 0$ and $d_\varphi > 0$ if and only if $r_u > 0$ and $p_u > 0$;

(b) $r_\psi < 0$ and $d_\psi < 0$ if and only if $r_u > 0$ and $p_u - 3r_u > 0$.

Part (a) states that $u = \varphi(i)$ is more downside risk-averse than i if and only if $r_u > 0$ and $p_u > 0$, while part (b) states that $i = \psi(u)$ is less downside risk averse than u. Clearly, the difference between the characterizing measure conditions for the two parts is traceable to the fact that Propositions 2 and 3 are not equivalent. Moreover part (b) is the stronger condition, since the transformation conditions $r_\psi < 0$ and $d_\psi < 0$ imply $r_\varphi > 0$ and $d_\varphi > 0$.

However, part (b) can also be viewed as an alternative to part (a) as a definition of greater downside risk aversion. Exploring this avenue reveals several logical inconsistencies, among them that compensated increases in downside risk for utility $u = \varphi(i)$ are not necessarily liked by utility i, which is neutral to all changes in risk

Keenan & Snow (2023)].[3] Here we take a complementary tack, and observe that the necessary and sufficient conditions for $d_\varphi > 0$ can be written as

$$D_u - D_v + 3r_u(r_u - r_v) = d_u - d_v - 3r_v(r_u - r_v) > 0. \tag{9}$$

Given $r_v \geq 0$ and $r_\varphi > 0$, the inequality condition $d_u - d_v > 0$ is necessary for $d_\varphi > 0$, and is also sufficient for the special case in which $v = i$ leading to the characterization $r_u > 0$ and $p_u > 0$ in part (*a*) of Corollary 2, indicating direction of risk preference for *u* with respect to risk aversion and prudence. Since $i = \psi(u)$ is less downside risk averse than *u* implies that $u = \varphi(i)$ is more downside risk averse than *i*, the conditions $r_u > 0$ and $p_u - 3r_u > 0$ with intensity of preference aversion with respect to downside risk. These observations lead us to examine prudence and downside risk aversion with respect to direction and intensity.

5 Comparative Statics for Prudence and Downside Risk Aversion

Whereas we have identified risk-averse utility functions with dislike of mean preserving spreads in the distribution for income, Arrow (1965) and Pratt (1964) link risk aversion to a positive risk premium whose magnitude increases with greater risk aversion. Insofar as the premium approach to characterizing either direction or intensity of risk preference relies on the absence of risk as the benchmark, this approach is not applicable beyond the second order. In this section, we contrast prudence and downside risk aversion with respect to the direction of preference imparted by a transformation of

[3] A compensated increase in downside risk is a shift the distribution for income *y* that induces an increase in downside risk for utility $u(y)$, and must compensate for the decline in expected utility experienced by a downside risk-averse utility.

risk-neutral utility and intensity of preference as reflected in characteristic comparative statics predictions.

While part (*a*) of Corollary 2 shows that downside risk aversion implies positive prudence when the reference utility is risk neutral, the implication is not valid when the reference is downside risk averse in the sense that $p_v - 3r_v$ is positive.

Proposition 4 [Keenan & Snow (2010)] Given $u = \varphi(v)$, for all increasing transformations φ, we have

(a) $r_\varphi > 0$ and $d_\varphi > 0$ implies $p_u > p_v$ if and only if $p_v - 3r_v \leq 0$;

(b) $r_\varphi > 0$ and $p_u > p_v$ implies $d_\varphi > 0$ if and only if $p_v - 3r_v \geq 0$.

When the reference is risk neutral, these necessary and sufficient conditions are satisfied, and we have $p_u > 0$ in both cases. Thus, introducing downside risk aversion introduces prudence, although the converse is not true, since risk-loving and downside risk-loving utility displays positive prudence.

When the reference utility is risk averse, one finds that transformations increasing prudence also increase risk aversion. As shown by Kimball, just as an increasing, concave transformation of (an increasing) utility function increases risk aversion, an increasing, convex transformation of a decreasing marginal utility, $u' = \xi(v')$, increases prudence. Given $r_v > 0$, calculation shows $r_u = r_v(\xi' v' / \xi)$ and $\xi' v' / \xi > 1$, implying that risk aversion increases. Thus, greater preference intensity for the preference directions $r_u > 0$ and $p_u > 0$ is identified with greater risk aversion and greater prudence.

The contrast between greater prudence and greater downside risk aversion is illustrated by the simple two-period saving problem with time-separable preferences described by Eeckhoudt et al (2005). The decision maker chooses saving s to maximize expected utility given by $\int v(\bar{y}-s)+v(\bar{y}+s+\varepsilon)\,dF(\varepsilon)$, where utility present and future utility functions are the same, both the interest rate and the subjective rate of time preference are set equal zero, $\bar{y}$ is the sure value of endowed income in both periods, and F is the cumulative distribution function for a zero-mean additive risk to future income, ε. Analysis of this problem by Leland (1968) showed that the introduction of an additive risk to future income increases saving by a "precautionary" amount if the decision maker exhibits risk aversion and downside risk aversion, that is, if $r_u > 0$ and $d_u > 0$, implying positive risk aversion and positive prudence, $r_u > 0$ and $p_u > 0$. Optimal saving in the absence of risk is equal to zero, and therefore in the presence of risk, optimal saving is entirely precautionary.

Kimball (1990) introduced the measure of prudence to characterize the precautionary motive in saving, in direct parallel with the characterization direction and intensity for risk aversion. Optimal precautionary saving for v, denoted by s_v, satisfies the first-order condition, $-v'(\bar{y}-s_v)+\int v'(\bar{y}+s_v+\varepsilon)\,dF=0$, or equivalently

$$v'(\bar{y}-s_v)=v'(\bar{y}+s_v-\theta_v), \tag{10}$$

where θ_v is the prudence premium for v. By analogy with risk aversion and the risk premium, the prudence premium reflects the direction and intensity of the prudence measure p_v. Solving for s_v yields $s_v=\theta_v/2$. Hence, precautionary saving is positive and increases with positive and increasing prudence.

Although the response of saving to the introduction of future-income risk is dictated by the direction and intensity of prudence, the same is not true for increases in an existing risk. Assume that an increase in the shift parameter γ for the distribution function $F(\varepsilon,\gamma)$ induces a mean preserving spread, denoted by $F_\gamma(\varepsilon,\gamma)$. The effect on precautionary saving is given by $ds_v / d\gamma = (\partial\theta_v / \partial\gamma)/2$, where equation (10) yields

$$\begin{aligned} \partial\theta_v / \partial\gamma &= -\int v' \, dF_\gamma / \hat{v}'' \\ &= -\int v''' \int^{\varepsilon} F_\gamma \, d\tau \, d\varepsilon / \hat{v}'', \end{aligned} \tag{11}$$

where the second line follows using integration by parts twice, and $\hat{v}'' = v''(\bar{y} + s_v - \theta_v)$. For the increase in saving to be greater for utility *u* than for *v*, we must have

$$\frac{\partial\theta_u}{\partial\gamma} - \frac{\partial\theta_v}{\partial\gamma} = \int \left(p_u \frac{u''}{\hat{u}''} - p_v \frac{v''}{\hat{v}''} \right) \int^{\varepsilon} F_\gamma \, d\tau \, d\varepsilon > 0, \tag{12}$$

where $\hat{u}'' = u''(\bar{y} + s_u - \theta_u)$. However, greater prudence for *u* than for *v* is not sufficient for this inequality, and therefore does not imply that a greater increase in precautionary saving in response to an increase in future-income risk.

A contrasting thought experiment introduced by Crainich & Eeckhoudt (2008) yields a complementary but distinct comparative statics prediction concerning the change in the interest rate required to maintain optimal saving equal to zero when future-income risk is introduced. Let m_v denote the compensating (gross) interest rate for *v* under which optimal saving is equal to zero when future income risk is present, defined by the first-order condition

$$-v'(\bar{y}) + m_v \int v'(\bar{y} + \varepsilon) \, dF(\varepsilon,\gamma) = 0. \tag{13}$$

In the absence of risk, $m_v = 1$. Assume that $F_\gamma(\varepsilon, \gamma)$ denotes a simple mean preserving spread with single crossing at $\varepsilon = 0$.[4] When initially there is no risk, F_γ represents an introduction of risk, and otherwise F_γ represents an increase in risk with a positive cumulative increase in probability below the mean balanced by a cumulative reduction above the mean. The effect of an increase in γ on the compensating interest rate

$$\begin{aligned} dm_v / d\gamma &= \int v'' F_\gamma \, d\varepsilon / \bar{v}' \\ &= -\int v''' \int^\varepsilon F_\gamma \, d\tau \, d\varepsilon / \bar{v}' \end{aligned} \tag{14}$$

is obtained from equation (13) using integration by parts twice, where $\bar{v}' = v'(\bar{y})$. Since the partial integrals on the third line are non-negative, the compensating interest rate falls below one when risk is introduced if $d_v > 0$, which is implied by $r_v > 0$ and $p_v - 3r_v > 0$.

Replacing v with $u = \varphi(v)$ yields

$$\begin{aligned} dm_u / d\gamma &= \int u'' F_\gamma \, d\varepsilon / \bar{\varphi}' \bar{u}' \\ &= \int (\varphi''' v^3 + 2\varphi'' v'' v') \int^\varepsilon F_\gamma \, d\tau \, d\varepsilon / \bar{\varphi}' \bar{v}' + \int \varphi' v'' F_\gamma \, d\varepsilon / \bar{\varphi}' \bar{v}', \end{aligned} \tag{15}$$

where the second line is obtained using integration by parts, and $\bar{\varphi}' = \varphi(v(\bar{y}))$. Since the partial integrals are nonnegative, the first integral is positive if v is risk averse and the transformation satisfies $r_\varphi > 0$ and $d_\varphi > 0$. Hence, we have $dm_u / d\gamma < dm_v / d\gamma$, and the compensating interest rate falls, if the final integral in equation (15) is at least as great as equation (14), that is, if

$$\int v'' (1 - \varphi' / \bar{\varphi}') F_\gamma \, d\varepsilon \geq 0 . \tag{16}$$

---

[4] A simple, or single mean preserving spread satisfies a single crossing property such that

As a simple mean preserving spread with single crossing at zero, F_γ changes sign from positive to negative at zero as ε increases, while $\varphi'' > 0$ implies that $1 - \varphi' / \overline{\varphi}'$ behaves in the opposite manner. Hence, if φ is downside risk averse, then utility $u = \varphi(v)$ requires a greater decline in the interest rate than v for saving to remain constant after the introduction of, or a simple increase in, future-income risk.

Proposition 5 Given $u = \varphi(v)$ and a single-crossing increase in future-income risk induced by $d\gamma > 0$, we have,

(a) $dm_u / d\gamma < dm_v / d\gamma$ if $r_\varphi > 0$ and $d_\varphi > 0$;

(b) $dm_u / d\gamma > dm_v / d\gamma$ if $r_\varphi < 0$ and $d_\varphi < 0$.

Part (a) shows that greater downside risk aversion implies a stronger reaction to to introductions or simple increases in future-income risk as measured by the decline in the interest rate required to maintain saving constant [Keenan & Snow (2016)]. Part (b) follows since φ is reversible, and demonstrates the reverse, that less downside risk aversion implies a weaker reaction to these increases in income risk. Thus, in contrast with greater prudence, greater downside risk aversion holds unambiguous comparatives statics implications for some mean preserving spreads of an existing risk as well as for introduction of risk into initially riskless saving decisions.

$F_\gamma(\varepsilon, \gamma) \geq [=](\leq) 0$ as $\varepsilon < [=](>) 0$ [Rothschild & Stiglitz (1970].

6 Conclusions

We characterize downside risk preference in expected utility theory with respect to direction and with respect to intensity for both greater and less aversion in terms of the risk preferences utility transformations. Downside risk preference is inherited from a transformation of risk-neutral utility, with the direction of third-order preference indicated by the sign of the third derivative. However, to obtain a strict partial ordering of utility functions by intensity of third-order risk preference requires specifying the transformations' second-order risk preferences as well as their third-order preferences. An ordering by greater third-order risk aversion is obtained when transformations are risk averse and downside risk averse, and the ordering is representable in terms of inequality restrictions on the utility measures of risk aversion r_V and prudence p_V. Moreover, reversing the preference directions from averse to loving yields an ordering by less third-order risk aversion. In particular, positive and increasing (decreasing) values for the measures r_V and $p_V - 3r_V$ are sufficient conditions for greater (less) downside risk aversion. Finally, we show that, the decline in the interest rate needed to maintain constant precautionary saving increases with greater downside risk aversion with the introduction of a zero-mean income risk or with a simple increase in risk with single crossing at zero.

Reversibility unlocks the paradox outlined in the introduction. Transformations that increase downside risk aversion are reversible, since the restrictions on r_V and p_V that characterize these transformations are reversed when love replaces aversion. For the same reason, their inverse transformations are reversible, but the reference utility *v* and final utility *u* switch roles thereby altering the characterizing inequality restrictions on the utility measures. As a consequence, we find that a transformation introducing downside risk averse preference necessarily introduces positive prudence, $p_u > 0$, while its inverse eliminates downside risk averse preference only if $p_u - 3r_u > 0$, the difference being an artifact of the switch in direction to and from utility u.

References

Arrow, K. J. (1965) Aspects of the Theory of Risk Bearing (Helsinki: Yrjo Jahnsson Lectures).

Crainich, D. & Eeckhoudt, L. (2008) On the Intensity of Downside Risk Aversion, Journal of Risk and Uncertainty 36, pp. 267-276.

Eeckhoudt, L. (2012) Beyond Risk Aversion: Why, How and What's Next, Geneva Risk and Insurance Review 73, pp. 141-155.

Eeckhoudt, L., Gollier, C., & Schlesinger, H. (2005) *Economic and Financial Decisions under Risk* (Princeton: Princeton University Press).

Keenan, D. & Snow, A. (2002) Greater Downside Risk Aversion, Journal of Risk and Uncertainty 24, pp. 267-277.

Keenan, D. & Snow, A. (2009) Greater Downside Risk Aversion in the Large, Journal of Economic Theory 144, pp. 1092-1101.

Keenan, D. & Snow, A. (2010) Greater Prudence and Greater Downside Risk Aversion, Journal of Economic Theory 145, pp. 2018-2026.

Keenan, D. & Snow, A. (2016) Strong Increases in Downside Risk Aversion, Geneva Risk and Insurance Review 41, pp. 149-161.

Keenan, D. & Snow, A. (2022) Reversibly Greater Downside Risk Aversion, Geneva Risk and Insurance Review 47, pp. 327-338.

Keenan, D. & Snow, A. (2023) Full Downside Risk Aversion, Department of Economics, University of Georgia.

Kimball, M. (1990) Precautionary Saving in the Small and in the Large, Econometrica 58, pp. 53-73.

Leland, H. E. (1968) Saving and Uncertainty: Precautionary Demand for Saving, Quarterly Journal of Economics 82, pp. 465-473.

Liu, L. & Meyer, J. (2012) Decreasing Absolute Risk Aversion, Prudence and Increased Downside Risk Aversion, Journal of Risk and Uncertainty 44 (2012), pp. 243-260.

Menezes, C., Geiss, C., & Tressler, J. (1980) Increasing Downside Risk, American Economic Review 70, pp. 921-932.

Pratt, J. (1964) Risk Aversion in the Small and in the Large, Econometrica 32, pp. 122-136.

Rothschild, M., & Stiglitz, J. (1970) Increasing Risk I: A Definition, Journal of Economic Theory 2, pp. 225-243.

The many faces of multivariate risk-taking

Risk apportionment for desirable and undesirable attributes

Richard Peter* Béatrice Rey†

June 6, 2022

Abstract

Many decisions under risk involve alternatives with multiple and possibly non-financial attributes. In this paper, we characterize risk apportionment preferences in a bivariate setting. We distinguish between desirable and undesirable attributes and show how to adapt the theory to obtain consistent results. We extend the definitions of correlation aversion, cross-prudence and cross-temperance in terms of simple lotteries to the case of undesirable attributes, provide a general characterization based on signs of cross-derivatives of the utility function, and discuss specific multivariate models for applications. Our results show how to unlock the powerful machinery of risk apportionment in the many situations in which decision-makers face undesirable attributes.

Keywords: Multivariate risk · risk apportionment · undesirable attributes · higher-order risk effects · correlation aversion · cross-prudence · cross-temperance

JEL-Classification: D11 · D81

*Corresponding author. University of Iowa, Department of Finance, E-Mail: richard-peter@uiowa.edu.

†Universite de Lyon and CNRS, GATE Lyon Saint-Etienne, E-Mail: rey-fournier@gate.cnrs.fr.

1 Introduction

Risk attitudes are without doubt a crucial determinant of economic and financial decision-making. Many decisions under risk involve more than a single attribute (see Keeney et al., 1993). Treatment decisions have financial consequences but also affect people's health. After filing one's taxes, people may be uncertain about the size of the refund but also about how long they will have to wait to obtain it. In times of a public health crisis, policymakers have to weigh the economic costs of containment measures against the loss of life. These examples also illustrate that some attributes are desirable like money, consumption or health, while other attributes are undesirable like waiting time, costs or the number of fatalities. The distinction between desirable and undesirable attributes plays a key role in our paper.

While many economists have traditionally been thinking of risk attitude as risk-averse or risk-loving, so-called higher-order risk attitudes are receiving increased attention. In the early models of precautionary saving by Leland (1968), Sandmo (1970) and Drèze and Modigliani (1972), which were later revisited by Kimball (1990), a third-order attitude called prudence guarantees that income risk leads to precautionary saving. A fourth-order attitude called temperance ensures less risk-taking in the presence of greater background risk (Kimball, 1993). Although first received with some skepticism, the notions of prudence and temperance have now been widely accepted in the economic analysis of decision-making under risk.[1]

More generally, Ekern (1980) defines the notion of Kth-degree risk aversion as an aversion to Kth-degree risk increases where K is an integer. While general, his integral conditions lack intuition and it remains unclear how to test for higher-order risk attitudes in the data. A breakthrough came with the impactful works of Eeckhoudt and Schlesinger (2006), Eeckhoudt et al. (2007) and Eeckhoudt et al. (2009), who provide a simple and intuitive way of understanding higher-order risk preferences via risk apportionment. Two basic types of apportionment preferences arise from their analysis, "combining good with bad" and "combining good with good and bad with bad" (Deck and Schlesinger, 2014).

In this paper, we provide new results on multivariate risk-taking. We use the powerful tools of risk apportionment but explicitly distinguish between desirable and undesirable attributes in our analysis. Our first contribution is to revisit the concepts of correlation aversion, cross-prudence and cross-temperance (see Eeckhoudt et al., 2009). As Deck and Schlesinger (2014) say, "restricting any analyses within economic applications to only the first four orders seems a reasonable approximation." We start with such an approximation and provide definitions of correlation aversion, cross-prudence and cross-temperance in terms of simple lotteries. When one or both attributes are undesirable, some of these definitions need to be adjusted, and we explain how and why. In the expected utility model, these simple lottery preferences pin

[1] Noussair et al. (2014) summarize the many ways in which prudence and temperance affect economic behavior. Applications include auctions, bargaining, ecological discounting, precautionary saving, investment, rent-seeking, and prevention.

down the sign of specific cross-derivatives of the utility function. The approach with simple lotteries has the advantage that it remains valid even when expected utility falls short from a descriptive standpoint (see Starmer, 2000).[2]

Our second contribution is to provide a general characterization of risk apportionment preferences in the bivariate setting under expected utility. We revisit the univariate case and explain how to accommodate undesirable attributes. All we need to do is to adjust the "seed lotteries," and the rest of Eeckhoudt and Schlesinger's (2006) risk apportionment theory stays intact. For an undesirable attribute, decision-makers who prefer to combine good with bad have all subsequent derivatives of the utility function negative. Decision-makers who prefer to combine good with good and bad with bad have subsequent derivatives alternating in sign but starting with a negative instead of a positive (Ebert, 2020). For a desirable attribute, combining good with bad is characterized by alternating signs whereas combining good with good and bad with bad is characterized by a consistent positive sign. A reversal occurs when going from a desirable to an undesirable attribute.

Once we have the univariate apportionment lotteries in place, we characterize risk apportionment preferences across attributes. We use Eeckhoudt et al.'s (2009) approach of apportioning Ekern (1980) risk increases and determine the signs of successive cross-derivatives of the utility function in three cases. We consider two desirable attributes, one desirable and one undesirable attribute, and two undesirable attributes. The orders of the risk changes play different roles in the three cases, and these roles are determined by the apportionment preferences on the individual attributes. In the expected utility model, it is easy to show that all three cases can be reconciled with each other. Hence, our results are fully consistent.

Our third contribution is to relate our findings to popular multivariate models. We discuss multiplicatively separable utility and equivalent monetary utility. In the separable case, the apportionment preference across attributes is very easy to characterize. If the component utility functions have the same sign, the decision-maker prefers to combine good with good and bad with bad. If they have opposite signs, she prefers to combine good with bad. This insight allows us to construct any of the eight combinations of risk apportionment preferences studied in this paper for applications.

We proceed as follows. Section 2 outlines the model, defines risk apportionment, and revisits the single-attribute case. Section 3 defines correlation aversion, cross-prudence and cross-temperance while distinguishing between desirable and undesirable attributes. Section 4 provides the general theory. Section 5 connects the apportionment preference across attributes to signs of cross-derivatives of the utility function and reconciles all three cases. Section 6 relates our analysis to Gollier's (2021) generalized risk apportionment theory. Section 7 presents specific multivariate models and shows how to implement different combinations of apportionment preferences in applications. A final section concludes.

[2] One such example is the recent work by Eeckhoudt et al. (2020) who characterize risk apportionment over a single attribute in Yaari's (1987) dual theory of choice under risk.

2 The model

2.1 Preliminaries

We analyze bivariate preferences. Our analysis can be extended to higher dimensions by fixing all but two of the attribute levels. Let (x, y) denote a nonnegative vector of attributes with $x \in [0, \overline{x}]$ and $y \in [0, \overline{y}]$. The domain of attribute bundles is then given by $\mathcal{D} = [0, \overline{x}] \times [0, \overline{y}]$.[3] Previous literature has mainly focused on the case that x and y both represent desirable attributes, for example, if we interpret x as consumption or final wealth and y as health or quality of life. We refer to this case as **DD** where **D** is shorthand for "desirable." We revisit this case as a benchmark and provide some extensions. The decision-maker (DM) is better off when x increases, when y increases, or when both increase.

We can also consider situations in which x is desirable and y is undesirable. For example, x can be a monetary payoff and y the time it takes to receive it (waiting time). DMs prefer higher values of x but lower values of y. This setting is studied in Ebert (2020) under the assumption that only y is uncertain and x is deterministic. Households face uncertainty regarding the value of their investments, for which they prefer higher over lower outcomes, and at the same time uncertainty over potential losses arising from auto and home ownership or legal liability. For the second attribute, they clearly prefer lower over higher outcomes. We label this case as **DU** where **U** abbreviates "undesirable." The ordering assumption that the first attribute is desirable and the second one undesirable is without loss.

Finally we consider the case that both x and y are undesirable and label it as **UU**. For the sake of example, imagine a policymaker in times of a public health crisis who considers the stringency of lockdown measures. These measures affect the economy, potentially resulting in unemployment and loss of livelihood, but they curb the spread of infectious diseases, thus mitigating the number of hospitalizations and fatalities. If x denotes unemployment and y the number of fatalities, then both are undesirable because lower outcomes are preferred over higher ones for each of the two attributes.

2.2 Risk apportionment

Eeckhoudt and Schlesinger (2006) develop the notion of risk apportionment and Eeckhoudt et al. (2009) apply it to stochastic dominance. Consider a DM who faces two independent stochastic changes that are unfavorable but unavoidable. If the DM would rather be exposed to the two changes in separate states, she exhibits a preference for combining good with bad. More formally, she prefers the 50-50 lottery that allocates one of the changes to one state and the other change to the other state over the 50-50 lottery that allocates both changes to the same state. The preferred lottery combines a relatively good outcome with a relatively bad outcome in each state whereas the dispreferred lottery has both good outcomes in the same

[3] A bounded domain avoids issues of sign permanence, see Scott and Horvath (1980) and Menegatti (2001).

state and both bad outcomes in the other state. We refer to this preference as combining good with bad or a preference for harms disaggregation, in short d for “disaggregate.”

Some DMs may have the reverse preference and rather face the two unfavorable changes in the same state than in different states. They prefer the 50-50 lottery that allocates both changes to the same state over the 50-50 lottery that allocates one of the changes to one state and the other change to the other state. Consequently, these DMs prefer to combine the two relatively good outcomes in one state and the two relatively bad outcomes in the other state instead of combining relatively good with relatively bad outcomes in the same state. We refer to this preference as combining good with good and bad with bad or a preference for harms aggregation, and abbreviate it with a for “aggregate.”[4]

In the univariate context, DMs who always prefer to disaggregate harms are called mixed risk-averse whereas DMs who always prefer to aggregate harms are called mixed risk-loving. Mixed risk aversion was first introduced by Caballé and Pomansky (1996) and Brockett and Golden (1987) whereas mixed risk lovers have not received much attention until recently, see Crainich et al. (2013) and Ebert (2013). In a laboratory experiment, Deck and Schlesinger (2014) provide evidence that the behavior of subjects classified as risk-averse is indeed consistent with mixed risk aversion while the behavior of subjects classified as risk-loving is consistent with mixed risk loving. Haering et al. (2020) confirm this dichotomy in different countries and with high stakes, and show that it is strengthened when lotteries are displayed in compound form instead of reduced form.[5]

In the bivariate context, DMs have an apportionment preference pertaining to each attribute individually and also an apportionment preference across attributes. Imagine a DM who always prefers to disaggregate harms. We label this preference as *dd-d*, where the first letter refers to the preference on the first attribute, the second letter to the preference on the second attribute, and the third letter to the preference across attributes. The example of wealth and health illustrates that a focus on *dd-d* is too narrow. A common combination of assumptions is risk aversion over wealth, risk aversion over health, and correlation loving over wealth and health. In our notation, this corresponds to *dd-a* because a DM with such a preference prefers to disaggregate harms pertaining to either wealth or either health but combines good with good and bad with bad when it comes to a sure reduction in wealth and a sure reduction in health. Following this reasoning, there is a total of eight possible combinations, *dd-d*, *dd-a*, *da-d*, *da-a*, *ad-d*, *ad-a*, *aa-d* and *aa-a*. We take these combinations as the primitive in our paper. We hope that our research will stimulate future empirical studies to investigate the relative prevalence of these preferences in the data.

[4] Of course, a DM's preference may not follow either of the two patterns, or one of the patterns in some cases and the other one in other cases. In this paper, we focus on preferences that are consistent.

[5] Bleichrodt and van Bruggen (2021) find a reflection effect for higher-order risk preferences similar to the reflection effect identified by Kahneman and Tversky (1979). Behavior in their experiment is, in general, not consistent with a preference for combining good with bad or good with good and bad with bad.

2.3 Relation to utility of univariate apportionment preferences

Before we consider multivariate risks, we recollect the relation to utility of attribute-specific risk apportionment preferences. To interpret the signs of successive derivatives of the utility function in terms of higher-order risk attitudes, we recall Ekern's (1980) definition of risk increases. Let $K \in \mathbb{N}$ be a whole number and let W_1 and W_2 be two random variables with values in $[0, \overline{w}]$. Denote by $F_1^{(1)}$ and $F_2^{(1)}$ their respective cumulative distribution functions. For $k \in \mathbb{N}$, define the functions $F_1^{(k)}$ on $[0, \overline{w}]$ recursively by setting $F_1^{(k+1)}(w) = \int_0^w F_1^{(k)}(t)\,\mathrm{d}t$ for $w \in [0, \overline{w}]$, and likewise for W_2. We state the following definition.

Definition 1 (Ekern,1980)**.** W_2 has more Kth-degree risk than W_1 if:

(i) $F_1^{(k)}(\overline{w}) = F_2^{(k)}(\overline{w})$ for all $k = 1, \ldots, K$,

(ii) $F_1^{(K)}(w) \leq F_2^{(K)}(w)$ for all $w \in [0, \overline{w}]$.

Condition (i) ensures that W_1 and W_2 have the same first $(K-1)$ moments. Condition (ii) implies that the Kth moment is larger for W_2 than for W_1 when sign adjusted by $(-1)^K$. Well-known special cases include first-order stochastic dominance for $K = 1$, a mean-preserving increase in risk for $K = 2$ (see Rothschild and Stiglitz, 1970), a mean-variance-preserving increase in downside risk for $K = 3$ (see Menezes et al., 1980), and a mean-variance-skewness-preserving increase in outer risk for $K = 4$ (see Menezes and Wang, 2005).

Under expected utility a unambiguous preference over Kth-degree risk increases pins down the sign of the Kth derivative of the utility function. We formalize this in the next result.

Lemma 1. *Let $q : [0, \overline{w}] \to \mathbb{R}$ be a real-valued function that is K times continuously differentiable. The following two conditions are equivalent.*

(i) For all pairs (W_1, W_2) such that W_2 has more Kth-degree risk than W_1, we have $\mathbb{E}q(W_1) \geq \mathbb{E}q(W_2)$.

(ii) For all $w \in [0, \overline{w}]$, we have $(-1)^{K+1} q^{(K)}(w) \geq 0$.

Ekern (1980) showed that (ii) implies (i). Following the argument in Denuit et al. (1999), Jouini et al. (2013) also prove the reverse implication. Ekern (1980) calls DMs who dislike any increase in Kth-degree risk Kth-degree risk-averse. Analogously, we call DMs who like any increase in Kth-degree risk Kth-degree risk-loving. When preferences have an expected-utility representation with a smooth utility function, we can connect the DM's apportionment preference on the individual attributes to the notion of Kth-degree risk attitudes.

Let $u(x, y)$ represent the DM's preferences and consider the **DD** case, $u^{(1,0)} \geq 0$ and $u^{(0,1)} \geq 0$. As shown in Eeckhoudt and Schlesinger's (2006) main theorem, if the DM prefers to combine good with bad on x, then $(-1)^{M+1} u^{(M,0)} \geq 0$ for all $M \geq 1$. She is then Mth-degree risk-averse on the first attribute at all orders M. This holds for *dd-d*, *dd-a*, *da-d* and *da-a* DMs. If the DM prefers to combine good with good and bad with bad on x instead, then

$u^{(M,0)} \geq 0$ for all $M \geq 1$, see Deck and Schlesinger (2014). She is then Mth-degree risk-averse on the first attribute for all M that are odd, and Mth-degree risk-loving on the first attribute for all M that are even. This holds for *ad-d*, *ad-a*, *aa-d* and *aa-a* DMs. The same applies, *mutatis mutandis*, to the DM's apportionment preference over y.

Let us now move to the **DU** case, $u^{(1,0)} \geq 0$ and $u^{(0,1)} \leq 0$. The signs of the unidirectional derivatives of u regarding the first attribute are unaffected when going from **DD** to **DU**. For the second attribute, we follow in Eeckhoudt and Schlesinger's (2006) and Deck and Schlesinger's (2014) footsteps. If the DM prefers to combine good with bad on y, we now find $u^{(0,N)} \leq 0$ for all $N \geq 1$. She is then Nth-degree risk-loving on the second attribute when N is odd and Nth-degree risk-averse on the second attribute when N is even. This holds for *dd-d*, *dd-a*, *ad-d* and *ad-a* DMs. If the DM prefers to combine good with good and bad with bad on y instead, we have $(-1)^{N+1}u^{(0,N)} \leq 0$ for all $N \geq 1$. She is always Nth-degree risk-loving on the second attribute. This holds for *da-d*, *da-a*, *aa-d* and *aa-a* DMs.

The signs of $u^{(0,N)}$ coincide with Ebert's (2020) results who studied the first four orders in the context of discounting. We show in Appendix A.1 how to obtain all signs via Eeckhoudt et al.'s (2009) approach of apportioning risk increases. As in the **DD** case, DMs agree on odd-order risk preferences but disagree on even-order risk preferences. However, odd-order risk increases on y switch from being favorable to unfavorable and from unfavorable to favorable when going from the **DD** case to the **DU** case. The reason is that lower values of y are preferred over higher ones when y is undesirable. When combined with the DM's apportionment preference, this affects her higher-order risk attitude at all odd orders.[6]

3 Lottery preference and relation to utility: Correlation aversion, cross-prudence and cross-temperance

3.1 Two desirable attributes (case DD)

We begin with the **DD** case so that both x and y are desirable. For each attribute, an unfavorable change is then a reduction or a sure loss. We have seen in Section 2.3 that a DM is averse to mean-preserving spreads on a particular attribute if her apportionment preference on that attribute is combining good with bad. If it is combining good with good and bad with bad instead, the DM likes mean-preserving spreads on that attribute. Zero-mean risks can thus be unfavorable or favorable changes compared to the status quo depending on the DM's apportionment preference.

We now characterize the DM's apportionment preference across attributes with the help of simple lotteries as in Eeckhoudt et al. (2007). For positive constants $k > 0$ and $\ell > 0$,

[6] Menegatti and Peter (2021) observe a similar reversal when comparing the comparative statics of a risky benefit with that of a risky cost. Courbage and Rey (2016) notice a reversal as well when looking at the effect of changes in risky health losses on decision thresholds for preventive treatment.

a DM is called *correlation averse* if she prefers the lottery $[(x-k,y);(x,y-\ell)]$ over the lottery $[(x-k,y-\ell);(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $x-k\geq 0$ and $y-\ell\geq 0$, and *correlation loving* if she always has the reverse preference. When both attributes are desirable, correlation aversion is consistent with a preference to disaggregate harms across attributes whereas correlation loving represents a desire to aggregate harms across attributes. In terms of our risk apportionment taxonomy, *dd-d*, *da-d*, *ad-d* and *aa-d* DMs are correlation averters whereas *dd-a*, *da-a*, *ad-a* and *aa-a* DMs are correlation lovers. The apportionment preference on the individual attributes plays no role for this classification.

Let $\widetilde{\varepsilon}$ be an arbitrary zero-mean risk on x. Eeckhoudt et al. (2007) call a DM *cross-prudent in y* if she prefers the lottery $[(x+\widetilde{\varepsilon},y);(x,y-\ell)]$ over the lottery $[(x+\widetilde{\varepsilon},y-\ell);(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $\text{Supp}[x+\widetilde{\varepsilon}]\subseteq[0,\overline{x}]$ and $y-\ell\geq 0$, and *cross-imprudent in y* if she always has the reverse lottery preference.[7] There are now two ways to interpret this lottery preference. If the DM prefers to disaggregate harms on x, then the zero-mean risk $\widetilde{\varepsilon}$ is a harm relative to zero and cross-prudence in y represents a preference to disaggregate harms across attributes. This is Eeckhoudt et al.'s (2007) interpretation. If, however, the DM prefers to aggregate harms on x, then $\widetilde{\varepsilon}$ is preferred over zero and cross-prudence in y is a preference to aggregate harms across attributes. So *dd-d*, *da-d*, *ad-a* and *aa-a* DMs are cross-prudent in y whereas *dd-a*, *da-a*, *ad-d* and *aa-d* DMs are cross-imprudent in y. What matters is whether the apportionment preference on x is aligned with the apportionment preference across attributes or not. The apportionment preference on the second attribute y is irrelevant because a sure loss of ℓ on the second attribute is always unfavorable in the **DD** case.

Now let $\widetilde{\delta}$ be an arbitrary zero-mean risk on y. Eeckhoudt et al. (2007) call a DM *cross-prudent in x* if she prefers the lottery $[(x,y+\widetilde{\delta});(x-k,y)]$ over the lottery $[(x-k,y+\widetilde{\delta});(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $\text{Supp}[y+\widetilde{\delta}]\subseteq[0,\overline{y}]$ and $x-k\geq 0$, and *cross-imprudent in x* if she always has the reverse lottery preference. There are again two ways to interpret this lottery preference. The first one by Eeckhoudt et al. (2007) relies on the zero-mean risk $\widetilde{\delta}$ being a harm relative to zero. The second one is for the case that $\widetilde{\delta}$ is preferred over zero. For cross-prudence in x, what matters is the apportionment preference on y relative to the apportionment preference across attributes. If they are aligned, we obtain cross-prudence in x, which is the case for *dd-d*, *da-a*, *ad-d* and *aa-a* DMs. If they are not aligned, we obtain cross-imprudence in x, which is the case for *dd-a*, *da-d*, *ad-a* and *aa-d* DMs. For cross-prudence in x, the apportionment preference on x does not play a role because a sure loss of k on the first attribute is always bad in the **DD** case.

Let $\widetilde{\varepsilon}$ be an arbitrary zero-mean risk on x, let $\widetilde{\delta}$ be an arbitrary zero-mean risk on y, and let $\widetilde{\varepsilon}$ and $\widetilde{\delta}$ be independent. Eeckhoudt et al. (2007) call a DM *cross-temperate* if she prefers the lottery $[(x+\widetilde{\varepsilon},y);(x,y+\widetilde{\delta})]$ over the lottery $[(x+\widetilde{\varepsilon},y+\widetilde{\delta});(x,y)]$ for all $(x,y)\in\mathcal{D}$ such

[7] We adopt Eeckhoudt et al.'s (2007) notation and let $\text{Supp}[x+\widetilde{\varepsilon}]$ denote the support of the probability distribution function associated with the random variable $x+\widetilde{\varepsilon}$. We assume that the realizations of $\widetilde{\varepsilon}$ are between $-x$ and $\overline{x}-x$ almost surely to remain in the domain of the first attribute.

that Supp$[x+\tilde{\varepsilon}] \subseteq [0,\overline{x}]$ and Supp$[y+\tilde{\delta}] \subseteq [0,\overline{y}]$, and *cross-intemperate* if she always has the reverse lottery preference. The interpretation of this lottery preference now depends on all three apportionment preferences. If the DM prefers to disaggregate harms on x and y individually, the lottery preference is consistent with a desire to disaggregate harms across attributes. This is Eeckhoudt et al.'s (2007) interpretation. However, this lottery preference is also consistent with disaggregating harms across attributes when the DM prefers to aggregate harms on x and y individually. The difference is that the harm is now not to receive the zero-mean risk. Let the individual apportionment preferences on x and y not be aligned, and assume the DM prefers to disaggregate harms on x but prefers to aggregate harms on y. Then, $\tilde{\varepsilon}$ is a harm relative to zero but $\tilde{\delta}$ is preferred over zero. The lottery preference for cross-temperance can now be understood as a preference to aggregate harms across attributes. In summary, when the individual apportionment preferences are aligned, cross-temperance represents a desire to disaggregate harms across attributes. When, however, the individual apportionment preferences are not aligned, it represents a desire to aggregate harms across attributes. So *dd-d*, *da-a*, *ad-a* and *aa-d* DMs are cross-temperate whereas *dd-a*, *da-d*, *ad-d* and *aa-a* DMs are cross-intemperate.

When the DM's preferences can be represented with a bivariate utility function $u(x,y)$, Eeckhoudt et al. (2007) show that correlation aversion, cross-prudence in x and y, and cross-temperance can be characterized via the sign of specific cross-derivatives of the utility function. We use $u^{(M,N)}(x,y)$ to denote $\partial^{M+N}u(x,y)/\partial^M x \partial^N y$ for $M \geq 0$ and $N \geq 0$ with $M+N \geq 1$. In the **DD** case, we have $u^{(1,0)} \geq 0$ and $u^{(0,1)} \geq 0$ because both x and y are desirable. Based on the above discussion and with Eeckhoudt et al.'s (2007) Proposition 1, we can then sign specific cross-derivatives of the utility function for the various underlying apportionment preferences. We summarize our findings in the following proposition.

Proposition 1. *Consider the case of two desirable attributes (case* ***DD****).*

(i) DMs with apportionment preferences dd-d, da-d, ad-d or aa-d have $u^{(1,1)} \leq 0$ *(correlation aversion), DMs with apportionment preferences dd-a, da-a, ad-a or aa-a have* $u^{(1,1)} \geq 0$ *(correlation loving).*

(ii) DMs with apportionment preferences dd-d, da-d, ad-a or aa-a have $u^{(2,1)} \geq 0$ *(cross-prudence in y), DMs with apportionment preferences dd-a, da-a, ad-d or aa-d have* $u^{(2,1)} \leq 0$ *(cross-imprudence in y).*

(iii) DMs with apportionment preferences dd-d, da-a, ad-d or aa-a have $u^{(1,2)} \geq 0$ *(cross-prudence in x), DMs with apportionment preferences dd-a, da-d, ad-a or aa-d have* $u^{(1,2)} \leq 0$ *(cross-imprudence in x).*

(iv) DMs with apportionment preferences dd-d, da-a, ad-a or aa-d have $u^{(2,2)} \leq 0$ *(cross-temperance), DMs with apportionment preferences dd-a, da-d, ad-d or aa-a have* $u^{(2,2)} \geq 0$ *(cross-intemperance).*

Table 1 in the appendix collects these signs and organizes them according to the DM's apportionment preference. Only DMs who prefer to disaggregate harms, on the individual attributes as well as across attributes, are correlation averse, cross-prudent in x and y, and cross-temperate. As soon as one of the apportionment preferences changes, at least some of the signs flip. For example, we obtain correlation loving, cross-imprudence in x and y, and cross-intemperance when the DM prefers to combine good with bad on the individual attributes but prefers to aggregate harms across attributes.

3.2 One desirable and one undesirable attribute (case DU)

We now turn to the **DU** case so that x is desirable and y is undesirable. A sure reduction is still an unfavorable change when applied to the first attribute but it is now a favorable change when applied to the second attribute. In fact, a sure *increase* is now an unfavorable change of the second attribute. As explained in Section 2.3, if the DM prefers to combine good with bad on the second attribute, she is averse to mean-preserving spreads on the second attribute and zero-mean risks are unfavorable compared to the status-quo. Conversely, if the DM prefers to combine good with good and bad with bad on y, she likes mean-preserving spreads on y and zero-mean risks are favorable changes relative to the satus-quo.

Taking into account that y is undesirable, we now define a DM to be *correlation averse* if she prefers the lottery $[(x-k, y); (x, y+\ell)]$ over the lottery $[(x-k, y+\ell); (x, y)]$ for all $(x, y) \in \mathcal{D}$ such that $x-k \geq 0$ and $y+\ell \leq \overline{y}$, and *correlation loving* if she always has the reverse preference. As in the **DD** case, correlation aversion represents a preference to disaggregate harms across attributes wheres correlation loving represents a desire to aggregate harms across attributes. What has changed is the definition of a harm on the second attribute because y is now undesirable. In terms of our classification *dd-d*, *da-d*, *ad-d* and *aa-d* DMs are correlation averters whereas *dd-a*, *da-a*, *ad-a* and *aa-a* DMs are correlation lovers.

Let $\widetilde{\varepsilon}$ be a zero-mean risk on the first attribute and let ℓ be a sure increase of the second attribute. We now call a DM *cross-prudent in* y if she prefers the lottery $[(x+\widetilde{\varepsilon}, y); (x, y+\ell)]$ over the lottery $[(x+\widetilde{\varepsilon}, y+\ell); (x, y)]$ for all $(x, y) \in \mathcal{D}$ such that $\text{Supp}[x+\widetilde{\varepsilon}] \subseteq [0, \overline{x}]$ and $y+\ell \leq \overline{y}$, and *cross-imprudent in* y if she always has the reverse lottery preference. We can interpret this lottery preference in two ways. If the DM prefers to disaggregate harms on x, the zero-mean risk is bad relative to zero and cross-prudence in y represents a preference to disaggregate harms across attributes. If the DM prefers to combine good with good and bad with bad on x instead, the zero-mean risk is preferred over zero and cross-prudence in y is consistent with a preference to aggregate harms across attributes. So *dd-d*, *da-d*, *ad-a* and *aa-a* DMs are cross-prudent in y whereas *dd-a*, *da-a*, *ad-d* and *aa-d* DMs are cross-imprudent in y. As in the **DD** case, the alignment between the apportionment preference on x and the apportionment preference across attributes matters. The apportionment preference on y does not matter for cross-prudence in y because a sure increase of ℓ is always unfavorable.

Now let $\widetilde{\delta}$ be a zero-mean risk on y. We call a DM *cross-prudent in* x if she prefers the lottery $[(x-k,y);(x,y+\widetilde{\delta})]$ over the lottery $[(x-k,y+\widetilde{\delta});(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $\mathrm{Supp}[y+\widetilde{\delta}]\subseteq[0,\overline{y}]$ and $x-k\geq 0$, and *cross-imprudent in* x if she always has the reverse lottery preference. It is now the apportionment preference on y and the apportionment preference across attributes that matter. If they are aligned, we find cross-prudence in x, if they are not, we obtain cross-imprudence in x. So *dd-d*, *da-a*, *ad-d* and *aa-a* DMs are cross-prudent in x whereas *dd-a*, *da-d*, *ad-a* and *aa-d* DMs are cross-imprudent in x. The apportionment preference on x does not play a role because a sure loss of k is always unfavorable.

Let $\widetilde{\varepsilon}$ and $\widetilde{\delta}$ be two independent zero-mean risks on x and y. We call a DM *cross-temperate* if she prefers the lottery $[(x+\widetilde{\varepsilon},y);(x,y+\widetilde{\delta})]$ over $[(x+\widetilde{\varepsilon},y+\widetilde{\delta});(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $\mathrm{Supp}[x+\widetilde{\varepsilon}]\subseteq[0,\overline{x}]$ and $\mathrm{Supp}[y+\widetilde{\delta}]\subseteq[0,\overline{y}]$, and *cross-intemperate* if she always has the reverse lottery preference. If the DM prefers to combine good with bad on x and y individually, this lottery preference is consistent with combining good with bad across attributes. The same holds if the DM prefers to combine good with good and bad with bad on x and y individually. The stated lottery preference is to rather face the undesirable changes in separate states (i.e., not getting the zero-mean risks) instead of taking the chance to face them together. Let the apportionment preferences on x and y not be aligned, and say the DM prefers to disaggregate harms on x but prefers to aggregate harms on y. In this case, $\widetilde{\varepsilon}$ is a harm relative to zero but $\widetilde{\delta}$ is preferred over zero. The lottery preference for cross-temperance can now be interpreted as a preference to aggregate harms across attributes. In summary, *dd-d*, *da-a*, *ad-a* and *aa-d* DMs are cross-temperate whereas *dd-a*, *da-d*, *ad-d* and *aa-a* DMs are cross-intemperate.

Our classification in the **DU** case is thus identical to the one in the **DD** case. We achieved this by adjusting the defining lottery preferences. When we represent preferences with a bivariate utility function, this affects some of the signs of the cross-derivatives as follows.

Proposition 2. *Consider the case in which the first attribute is desirable and the second attribute is undesirable (case **DU**).*

(i) DMs with apportionment preferences dd-d, da-d, ad-d or aa-d have $u^{(1,1)}\geq 0$ (correlation aversion), DMs with apportionment preferences dd-a, da-a, ad-a or aa-a have $u^{(1,1)}\leq 0$ (correlation loving).

(ii) DMs with apportionment preferences dd-d, da-d, ad-a or aa-a have $u^{(2,1)}\leq 0$ (cross-prudence in y), DMs with apportionment preferences dd-a, da-a, ad-d or aa-d have $u^{(2,1)}\geq 0$ (cross-imprudence in y).

(iii) DMs with apportionment preferences dd-d, da-a, ad-d or aa-a have $u^{(1,2)}\geq 0$ (cross-prudence in x), DMs with apportionment preferences dd-a, da-d, ad-a or aa-d have $u^{(1,2)}\leq 0$ (cross-imprudence in x).

(iv) DMs with apportionment preferences dd-d, da-a, ad-a or aa-d have $u^{(2,2)}\leq 0$ (cross-temperance), DMs with apportionment preferences dd-a, da-d, ad-d or aa-a have $u^{(2,2)}\geq 0$ (cross-intemperance).

Table 2 in the appendix organizes these signs by the DM's apportionment preferences. In Proposition 2(i), the signs for correlation aversion and correlation loving are flipped compared to Proposition 1(i). A lottery preference of $[(x-k,y);(x,y+\ell)]$ over $[(x-k,y+\ell);(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $x-k\geq 0$ and $y+\ell\leq\overline{y}$ can be equivalently stated as a lottery preference of $[(x-k,y'-\ell);(x,y')]$ over $[(x-k,y');(x,y'-\ell)]$ for all $(x,y')\in\mathcal{D}$ such that $x-k\geq 0$ and $y'-\ell\geq 0$. This is a simple change of variables by setting $y'=y+\ell$. But then we know from Eeckhoudt et al. (2007) that this lottery preference is equivalent to $u^{(1,1)}\geq 0$. Likewise, the signs for cross-prudence in y and cross-imprudence in y are flipped in Proposition 2(ii) compared to Proposition 1(ii). The same change of variables shows that a lottery preference of $[(x+\widetilde{\varepsilon},y);(x,y+\ell)]$ over $[(x+\widetilde{\varepsilon},y+\ell);(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $\text{Supp}[x+\widetilde{\varepsilon}]\subseteq[0,\overline{x}]$ and $y+\ell\leq\overline{y}$, is equivalent to a lottery preference of $[(x+\widetilde{\varepsilon},y'-\ell);(x,y')]$ over $[(x+\widetilde{\varepsilon},y');(x,y'-\ell)]$ for all $(x,y')\in\mathcal{D}$ such that $\text{Supp}[x+\widetilde{\varepsilon}]\subseteq[0,\overline{x}]$ and $y-\ell\geq 0$. Per Eeckhoudt et al. (2007), this is equivalent to $u^{(2,1)}\leq 0$. The signs for cross-prudence in x and cross-temperance are the same in Proposition 2 as in Proposition 1.

Our discussion leading up to Proposition 2 shows that, conceptually, nothing has changed when we adjust the characterizing lottery preferences accordingly. When we move to the utility representation, the signs of some cross-derivatives are flipped whereas others remain unchanged. Take correlation aversion as an example. In the **DD** case, both attributes are desirable. It is "riskier" to face a situation in which either both attributes are high or both are low at the same time instead of a situation in which low values of one attribute are compensated by high values of the other one. Intuitively, correlation averters should avoid positive correlation and seek negative correlation to hedge their bets. Now consider the **DU** case with one attribute being desirable and the other one undesirable. It is now better to face a situation in which either both attributes are high or both are low than a situation with one low and the other one high. When both are high, high values of the undesirable attribute are compensated by high values of the desirable attribute. When both are low, low values of the desirable attribute are compensated by low values of the undesirable attribute. Correlation averters now achieve hedging by avoiding negative correlation and seeking positive correlation.

3.3 Two undesirable attributes (case UU)

In a next step, we look at the **UU** case in which both x and y are undesirable. Then, a sure reduction of either attribute is a favorable change for the DM whereas a sure increase in either attribute is an unfavorable change. As shown in Section 2.3, a preference to combine good with bad on x implies that the introduction of a zero-mean risk on x is an unfavorable change compared to the status-quo while it is a favorable change if the DM prefers to combine good with good and bad with bad on x. The same holds for attribute y.

When both attributes are undesirable, we define a DM to be *correlation averse* if she prefers the lottery $[(x+k,y);(x,y+\ell)]$ over the lottery $[(x+k,y+\ell),(x,y)]$ for all $(x,y)\in\mathcal{D}$ such that $x+k\leq\overline{x}$ and $y+\ell\leq\overline{y}$, and *correlation loving* if she always has the reverse preference. Yet again correlation aversion represents a preference to disaggregate harms across attributes and correlation loving is consistent with aggregating harms across attributes. A harm is now a sure increase for either attribute. Only the apportionment preference across attributes matters so that *dd-d*, *da-d*, *ad-d* and *aa-d* DMs are correlation averters whereas *dd-a*, *da-a*, *ad-a* and *aa-a* DMs are correlation lovers. The apportionment preference on the individual attributes is irrelevant at this stage.

Let $\widetilde{\varepsilon}$ be a zero-mean risk on the first attribute and let ℓ be a sure increase of the second attribute. We call a DM *cross-prudent in y* if she prefers the lottery $[(x+\widetilde{\varepsilon}, y); (x, y+\ell)]$ over the lottery $[(x+\widetilde{\varepsilon}, y+\ell); (x, y)]$ for all $(x, y) \in \mathcal{D}$ such that $\text{Supp}[x+\widetilde{\varepsilon}] \subseteq [0, \overline{x}]$ and $y+\ell \leq \overline{y}$, and *cross-imprudent in y* if she always has the reverse lottery preference. This lottery preference has two interpretations. If the DM prefers to disaggregate harms on x, the zero-mean risk is a harm relative to zero and cross-prudence in y is consistent with harms disaggregation across attributes. If the DM prefers to aggregate harms on x instead, the zero-mean risk on x is preferred over zero and cross-prudence in y represents a desire to aggregate harms across attributes. So *dd-d*, *da-d*, *ad-a* and *aa-a* DMs are cross-prudent in y whereas *dd-a*, *da-a*, *ad-d* and *aa-d* DMs are cross-imprudent in y. As before, the alignment between the apportionment preference on x and the apportionment preference across attributes matters while the apportionment preference on y plays no role.

Let $\widetilde{\delta}$ be a zero-mean risk on the second attribute and k be a sure increase of the first attribute. We call a DM *cross-prudent in x* if she prefers the lottery $[(x+k, y); (x, y+\widetilde{\delta})]$ over the lottery $[(x+k, y+\widetilde{\delta}); (x, y)]$ for all $(x, y) \in \mathcal{D}$ such that $x+k \leq \overline{x}$ and $\text{Supp}[y+\widetilde{\delta}] \subseteq [0, \overline{y}]$, and *cross-imprudent in x* if she always has the reverse lottery preference. What matters is the alignment between the apportionment preference on y and the apportionment preference across attributes. We have cross-prudence in x when both are aligned and cross-imprudence in x when they are not. As a result, *dd-d*, *da-a*, *ad-d* and *aa-a* DMs are cross-prudent in x whereas *dd-a*, *da-d*, *ad-a* and *aa-d* DMs are cross-imprudent in x. The apportionment preference on x does not matter.

Consider two independent zero-mean risks, $\widetilde{\varepsilon}$ and $\widetilde{\delta}$, one on the first attribute x and the other one on the second attribute y. We call a DM *cross-temperate* if she prefers the lottery $[(x+\widetilde{\varepsilon}, y); (x, y+\widetilde{\delta})]$ over the lottery $[(x+\widetilde{\varepsilon}, y+\widetilde{\delta}); (x, y)]$ for all $(x, y) \in \mathcal{D}$ such that $\text{Supp}[x+\widetilde{\varepsilon}] \subseteq [0, \overline{x}]$ and $\text{Supp}[y+\widetilde{\delta}] \subseteq [0, \overline{y}]$, and *cross-intemperate* if she always has the reverse lottery preference. If the DM prefers to disaggregate harms on x and y individually, or if she prefers to aggregate harms on x and y individually, the lottery preference represents combining good with bad across attributes. If the DM's apportionment preferences on the individual attributes are not aligned, the lottery preference for cross-temperance is consistent with combining good with good and bad with bad across attributes. So *dd-d*, *da-a*, *ad-a* and *aa-d* DMs are cross-temperate whereas *dd-a*, *da-d*, *ad-d* and *aa-a* DMs are cross-intemperate.

The classification in the **UU** case is identical to the classification in the other two cases. We achieve this by adjusting the defining lottery preferences, specifically in those cases where a harm is now a sure increase of the attribute, not a sure reduction. In terms of the utility representation, the signs of some cross-derivatives are affected by these adjustments as follows.

Proposition 3. *Consider the case of two undesirable attributes (case **UU**).*

(i) DMs with apportionment preferences dd-d, da-d, ad-d or aa-d have $u^{(1,1)} \leq 0$ (correlation aversion), DMs with apportionment preferences dd-a, da-a, ad-a or aa-a have $u^{(1,1)} \geq 0$ (correlation loving).

(ii) DMs with apportionment preferences dd-d, da-d, ad-a or aa-a have $u^{(2,1)} \leq 0$ (cross-prudence in y), DMs with apportionment preferences dd-a, da-a, ad-d or aa-d have $u^{(2,1)} \geq 0$ (cross-imprudence in y).

(iii) DMs with apportionment preferences dd-d, da-a, ad-d or aa-a have $u^{(1,2)} \leq 0$ (cross-prudence in x), DMs with apportionment preferences dd-a, da-d, ad-a or aa-d have $u^{(1,2)} \geq 0$ (cross-imprudence in x).

(iv) DMs with apportionment preferences dd-d, da-a, ad-a or aa-d have $u^{(2,2)} \leq 0$ (cross-temperance), DMs with apportionment preferences dd-a, da-d, ad-d or aa-a have $u^{(2,2)} \geq 0$ (cross-intemperance).

Table 3 in the appendix organizes these signs according to the DM's apportionment preference. In Proposition 3(i), the signs for correlation aversion and correlation loving are flipped compared to Proposition 2(i) and are thus identical to Proposition 1(i). A preference of $[(x+k,y);(x,y+\ell)]$ over $[(x+k,y+\ell),(x,y)]$ for all $(x,y) \in \mathcal{D}$ such that $x+k \leq \overline{x}$ and $y+\ell \leq \overline{y}$ is equivalent to a preference of $[(x',y'-\ell);(x'-k,y')]$ over $[(x',y'),(x'-k,y'-\ell)]$ for all $(x',y') \in \mathcal{D}$ such that $x'-k \geq 0$ and $y'-\ell \geq 0$. Mathematically, this is a simple change of variables by letting $x' = x+k$ and $y' = y+\ell$. It renders the exact same characterizing lottery preference as in the **DD** case, which is why we find $u^{(1,1)} \leq 0$ for correlation aversion. The lotteries for cross-prudence in y are identical in the **DU** and **UU** cases, the lotteries for cross-prudence in x are different because the first attribute is desirable in the **DU** case but undesirable in the **UU** case. The cross-temperance lotteries are identical in all three cases.

Conceptually, we obtain the same classification in terms of the DM's underlying apportionment preference but only because we adjusted some of the characterizing lotteries. Consider correlation aversion again. In the **DD** and the **UU** case, correlation aversion obtains for $u^{(1,1)} \leq 0$. The underlying economic intuition is different. In both cases correlation averters hedge their bets by avoiding positive correlation and seeking negative correlation. In the **DD** case, positive correlation is unappealing because low values of one attribute tend to occur with low values of the other attribute. Negative correlation insulates the DM against this because low values of one attribute tend to be compensated by high values of the other one. For **UU**, positive correlation is unappealing because high values of one attribute tend to occur with high values of the other attribute. Negative correlation now helps because high values of one attribute tend to be compensated by low values of the other one. Even though a negative sign on $u^{(1,1)}$ characterizes correlation aversion in both cases, the roles of "high" and "'low" for the attribute values are reversed due to their different effects on the DM's welfare. While this is obvious for correlation attitude, the extension to higher orders is not immediate. The simple example of correlation attitude also shows that signing cross-derivatives and deriving economic intuition are two separate steps.

4 The general theory

4.1 Univariate risk apportionment

Eeckhoudt and Schlesinger (2006) define risk apportionment of any order via a specific lottery preference. Take the first attribute x and assume it is desirable. Let $\{\widetilde{\varepsilon}_i\}$ be an indexed set of zero-mean nondegenerate random variables, $i = 1, 2, 3, \ldots$, that are all mutually independent, and let k be a positive constant. Define $A_1 = [-k]$, $A_2 = [\widetilde{\varepsilon}_1]$, and $B_1 = B_2 = [0]$. Let $\mathrm{Int}(z)$ denote the greatest-integer function. For $M \geq 3$, define the univariate lotteries

$$\begin{aligned} A_M &= [B_{M-2} + 0;\ A_{M-2} + \widetilde{\varepsilon}_{\mathrm{Int}(M/2)}], \\ B_M &= [A_{M-2} + 0;\ B_{M-2} + \widetilde{\varepsilon}_{\mathrm{Int}(M/2)}]. \end{aligned}$$

A DM then prefers to combine good with bad on the first attribute if she prefers the lottery $[(x+B_M, y)]$ over the lottery $[(x+A_M, y)]$ for all $(x, y) \in \mathcal{D}$ and such that $\mathrm{Supp}[x+A_M] \subseteq [0, \overline{x}]$ and $\mathrm{Supp}[x+B_M] \subseteq [0, \overline{x}]$. She prefers combining good with good and bad with bad on the first attribute if she always has the reverse lottery preference. This is Eeckhoudt and Schlesinger's (2006) Definition 5 of risk apportionment of order M applied to the first attribute.[8]

Now assume that the first attribute is undesirable. When comparing A_1 and B_1, we now see that A_1 is preferred over B_1 because the DM appreciates a sure reduction of an undesirable attribute. To rectify this and maintain the iterative definition of higher-order risk preferences, all we need to do is to replace $A_1 = [-k]$ with $A_1 = [+k]$. For an undesirable attribute, a sure increase is now a harm relative to $B_1 = [0]$.

We can then proceed in a similar way regarding the second attribute y. Assume first that y is desirable. Let $\{\widetilde{\delta}_j\}$ be an indexed set of zero-mean nondegenerate random variables, $j = 1, 2, 3, \ldots$, that are all mutually independent and also mutually independent of the $\{\widetilde{\varepsilon}_i\}$. Let ℓ be a positive constant. Define $C_1 = [-\ell]$, $C_2 = [\widetilde{\delta}_1]$, and $D_1 = D_2 = [0]$. For $N \geq 3$, define the univariate lotteries

$$\begin{aligned} C_N &= [D_{N-2} + 0;\ C_{N-2} + \widetilde{\delta}_{\mathrm{Int}(N/2)}], \\ D_M &= [C_{N-2} + 0;\ D_{N-2} + \widetilde{\delta}_{\mathrm{Int}(N/2)}]. \end{aligned}$$

8 For ease of exposition, we take some liberty with the notation. The distribution of the lottery $[(x + B_M, y)]$ is the one that is induced by the distribution of the lottery B_M. In other words, the lottery $[(x + B_M, y)]$ has the outcome $(x + b, y)$ with probability $\mathbb{P}(B_M = b)$ for all $b \in \mathrm{Supp}[B_M]$.

A DM prefers to combine good with bad on the second attribute if she prefers the lottery $[(x, y+D_N)]$ over the lottery $[(x, y+C_N)]$ for all $(x, y) \in \mathcal{D}$ such that $\text{Supp}[y+C_N] \subseteq [0, \overline{y}]$ and $\text{Supp}[y+D_N] \subseteq [0, \overline{y}]$. She prefers combining good with good and bad with bad on the second attribute if she always has the reverse lottery preference. This is Eeckhoudt and Schlesinger's (2006) Definition 5 of risk apportionment of order N applied to the second attribute. If y is undesirable instead, we need to replace $C_1 = [-\ell]$ with $C_1 = [+\ell]$ to keep the iterative definition of higher-order risk preferences intact.

The main advantage of Eeckhoudt and Schlesinger's (2006) risk apportionment approach is its simplicity and elegance. Specifically, higher-order risk preferences can be defined purely based on a simple lottery preference and no particular representation of preferences is used. This is what some refer to as "model-free" even though, of course, preferences themselves are an economic model of choice under risk and reduction of compound lotteries is implicit in the risk apportionment literature.

4.2 Risk apportionment across attributes

Building on these univariate risk apportionment lotteries, we can now define risk apportionment across attributes. For $M, N \geq 1$, we say that preferences satisfy risk apportionment of order (M, N) if the DM prefers the lottery $[(x + B_M, y + C_N); (x + A_M, y + D_N)]$ over the lottery $[(x+B_M, y+D_N); (x+A_M, y+C_N)]$ for all $(x, y) \in \mathcal{D}$ such that $\text{Supp}[x+A_M] \subseteq [0, \overline{x}]$, $\text{Supp}[x + B_M] \subseteq [0, \overline{x}]$, $\text{Supp}[y + C_N] \subseteq [0, \overline{y}]$ and $\text{Supp}[y + D_N] \subseteq [0, \overline{y}]$. If the DM always has the reverse lottery preference, we say that preferences exhibit anti-risk apportionment of order (M, N). Eeckhoudt and Schlesinger (2006) introduce a terminology for preferences consistent with risk apportionment and Deck and Schlesinger (2014) use the qualifier "anti" for the reverse preference.[9]

We can easily connect this to the analysis of correlation aversion, cross-prudence in x and y, and cross-temperance in Section 2. Take $M = N = 1$ and consider A_1, B_1, C_1 and D_1. We have $B_1 = [0]$ and $D_1 = [0]$; furthermore, we have $A_1 = [-k]$ and $C_1 = [-\ell]$ in the **DD** case, $A_1 = [-k]$ and $C_1 = [+\ell]$ in the **DU** case, and $A_1 = [+k]$ and $C_1 = [+\ell]$ in the **UU** case. Therefore, risk apportionment of order $(1, 1)$ is characterized as follows:

$$\begin{cases} [(x, y-\ell); (x-k, y)] \succsim [(x, y); (x-k, y-\ell)], & \text{in case of } \mathbf{DD}, \\ [(x, y+\ell); (x-k, y)] \succsim [(x, y); (x-k, y+\ell)], & \text{in case of } \mathbf{DU}, \\ [(x, y+\ell); (x+k, y)] \succsim [(x, y); (x+k, y+\ell)], & \text{in case of } \mathbf{UU}. \end{cases}$$

These are the lottery preferences we used in Section 2 to characterize correlation aversion in each case, with the reverse preference characterizing correlation loving. Now take $M = 2$ and

[9] The distribution of the lottery $[(x + B_M, y + C_N)]$ is the one that is induced by the joint distribution of (B_M, C_M). Due to independence, the lottery has outcome $(x+b, y+c)$ with probability $\mathbb{P}(B_M = b)\mathbb{P}(C_N = c)$ for all $b \in \text{Supp}[B_M]$ and all $c \in \text{Supp}[C_N]$.

$N = 1$ and consider A_2, B_2, C_1 and D_1. We have $A_2 = [\widetilde{\varepsilon}_1]$, $B_2 = [0]$ and $D_1 = [0]$; when y is desirable, we have $C_1 = [-\ell]$, when y is undesirable, we have $C_1 = [+\ell]$. Risk apportionment of order $(2, 1)$ is then characterized via the following lotteries:

$$\begin{cases} [(x, y-\ell); (x+\widetilde{\varepsilon}_1, y)] \succsim [(x, y); (x+\widetilde{\varepsilon}_1, y-\ell)], & \text{in case of } \mathbf{DD}, \\ [(x, y+\ell); (x+\widetilde{\varepsilon}_1, y)] \succsim [(x, y); (x+\widetilde{\varepsilon}_1, y+\ell)], & \text{in case of } \mathbf{DU} \text{ or } \mathbf{UU}. \end{cases}$$

These are the lottery preferences we used in Section 2 to characterize cross-prudence in y, with the reverse preference characterizing cross-imprudence in y. For $M = 1$ and $N = 2$, consider A_1, B_1, C_2 and D_2. We have $B_1 = [0]$, $C_2 = [\widetilde{\delta}_1]$ and $D_2 = [0]$; when x is desirable, we have $A_1 = [-k]$, when x is undesirable, we have $A_1 = [+k]$. Risk apportionment of order $(1, 2)$ is then characterized via the following lotteries:

$$\begin{cases} [(x, y+\widetilde{\delta}_1); (x-k, y)] \succsim [(x, y); (x-k, y+\widetilde{\delta}_1)], & \text{in case of } \mathbf{DD} \text{ or } \mathbf{DU}, \\ [(x, y+\widetilde{\delta}_1); (x-k, y)] \succsim [(x, y); (x-k, y+\widetilde{\delta}_1)], & \text{in case of } \mathbf{DU}. \end{cases}$$

These are the lottery preferences we used in Section 2 to characterize cross-prudence in x, with the reverse preference characterizing cross-imprudence in x. Finally, for $M = 2$ and $N = 2$, consider A_2, B_2, C_2 and D_2, that is, $A_2 = [\widetilde{\varepsilon}_1]$, $B_2 = [0]$, $C_2 = [\widetilde{\delta}_1]$ and $D_2 = [0]$. The distinction between **DD**, **DU** and **UU** is now irrelevant. We can then always characterize risk apportionment of order $(2, 2)$ via the following lottery preference:

$$\left\{ \; [(x, y+\widetilde{\delta}_1); (x+\widetilde{\varepsilon}_1, y)] \succsim [(x, y); (x+\widetilde{\varepsilon}_1, y+\widetilde{\varepsilon}_1)], \quad \text{in case of } \mathbf{DD}, \mathbf{DU} \text{ or } \mathbf{UU}. \right.$$

This is the lottery preference we used in Section 2 for cross-temperance, with the reverse preference characterizing cross-intemperance.

Correlation aversion corresponds to risk apportionment of order $(1, 1)$, correlation loving to anti-risk apportionment of order $(1, 1)$, cross-prudence in y to risk apportionment of order $(2, 1)$, cross-imprudence in y to anti-risk apportionment of order $(2, 1)$, cross-prudence in x to risk apportionment of order $(1, 2)$, cross-imprudence in x to anti-risk apportionment of order $(1, 2)$, cross-temperance to risk apportionment of order $(2, 2)$, and cross-intemperance to anti-risk apportionment of order $(2, 2)$. When an attribute flips from desirable to undesirable, all we need to do is swap out the seed lottery and the iterative process and associated taxonomy stays fully intact. Specifically, if attribute x is undesirable, we need to replace $A_1 = [-k]$ with $A_1 = [+k]$. If attribute y is undesirable, we need to replace $C_1 = [-\ell]$ with $C_1 = [+\ell]$.

The lottery preference of $[(x+B_M, y+C_N); (x+A_M, y+D_N)]$ over $[(x+B_M, y+D_N); (x+A_M, y+C_N)]$ extends the notions of correlation aversion, cross-prudence in x and y, and cross-temperance to higher orders. Eeckhoudt et al. (2007) mention in their Footnote 12 that such an extension is possible but do not carry it out. What's more, their analysis focuses exclusively on the **DD** case. We show that, by defining the seed lotteries A_1 and C_1 accordingly, the entire risk apportionment machinery can also be applied to the **DU** and the **UU** case. As in the univariate analysis, no particular representation of preferences is necessary to define risk apportionment and anti-risk apportionment of order (M, N). Characterizations of these lottery preferences can be explored outside the confines of the expected-utility model. If the expected utility theorem holds, the stated lottery preference can be characterized by signing the corresponding cross-derivative of the utility function. We will provide this characterization in the next section based on risk apportionment via stochastic dominance.

5 Relation to utility - The general case

5.1 Two desirable attributes (case DD)

An alternative to the risk apportionment lotteries in Eeckhoudt and Schlesinger (2006) is the apportionment of risks via stochastic dominance in Eeckhoudt et al. (2009). Consider the four mutually independent random variables X_1, X_2, Y_1 and Y_2. Let X_2 have more Mth-degree risk than X_1, and Y_2 have more Nth-degree risk than Y_1. In the spirit of Eeckhoudt et al. (2009), we can then assess the DM's preference over the lotteries $[(X_1, Y_2); (X_2, Y_1)]$ and $[(X_1, Y_1); (X_2, Y_2)]$. The first lottery combines low Mth-degree risk on x with high Nth-degree risk on y, and high Mth-degree risk on x with low Nth-degree risk on y. The second lottery combines low Mth-degree risk on x with low Nth-degree risk on y, and high Mth-degree risk on x with high Nth-degree risk on y. When the DM always prefers the first lottery over the second one, we obtain $(-1)^{M+N+1}u^{(M,N)} \geq 0$ from Lemma 1. If she always has the reverse lottery preference instead, we obtain $(-1)^{M+N+1}u^{(M,N)} \leq 0$ from Lemma 1. We show this formally in Appendix A.2.

While the lottery preference of $[(X_1, Y_2); (X_2, Y_1)]$ over $[(X_1, Y_1); (X_2, Y_2)]$ pins down the sign of $(-1)^{M+N+1}u^{(M,N)}$ unambiguously, the interpretation of this sign depends on the DM's apportionment preference on the individual attributes, her apportionment preference across attributes, as well as on the parity of the orders (i.e., whether M and N are odd or even). The following result organizes the signs by the DM's risk apportionment preference.

Theorem 1 (Case **DD**). *Consider the case of two desirable attributes and let $M, N \geq 1$.*

(i) DMs with apportionment preference dd-d have $(-1)^{M+N+1}u^{(M,N)} \geq 0$, DMs with apportionment preference dd-a have $(-1)^{M+N+1}u^{(M,N)} \leq 0$.

(ii) DMs with apportionment preference da-d have $(-1)^{M}u^{(M,N)} \geq 0$, DMs with apportionment preference da-a have $(-1)^{M}u^{(M,N)} \leq 0$.

(iii) DMs with apportionment preference ad-d have $(-1)^{N}u^{(M,N)} \geq 0$, DMs with apportionment preference ad-a have $(-1)^{N}u^{(M,N)} \leq 0$.

(iv) DMs with apportionment preference aa-d have $u^{(M,N)} \leq 0$, DMs with apportionment preference aa-a have $u^{(M,N)} \geq 0$.

Appendix A.3 provides the proof. Proposition 1 is a special case of Theorem 1. When the DM prefers to disaggregate harms on x and y (*dd-a* and *dd-a*), her apportionment preference across attributes depends on whether the total order $M+N$ is odd or even. For $M+N$ odd, a positive sign on $u^{(M,N)}$ indicates a preference to disaggregate harms across attributes and a negative sign a preference to aggregate harms across attributes. When $M+N$ is even, the interpretation of the signs flips. The total order $M+N$ is decisive because DMs who prefer to combine good with bad on x and y individually will always view $[(X_1, Y_2); (X_2, Y_1)]$ as the lottery that combines good with bad across attributes and $[(X_1, Y_2); (X_2, Y_1)]$ as the lottery that combines good with good and bad with bad across attributes.

When the DM prefers to disaggregate harms on x but aggregate harms on y (*da-d* and *da-a*), her apportionment preference across attributes depends on the parity of M, the order of the risk change on the first attribute. Similarly, when she prefers to aggregate harms on x but disaggregate harms on y (*ad-d* and *ad-a*), the parity of N is decisive, the order of the risk change on the second attribute. For these DMs, lottery $[(X_1, Y_2); (X_2, Y_1)]$ is not always the one that combines good with bad across attributes relative to lottery $[(X_1, Y_2); (X_2, Y_1)]$. As we move up the orders, high Nth-degree risk on y is a good thing for *da-d* and *da-a* DMs when N is even, and high Mth-degree risk on x is a good thing for *ad-d* and *ad-a* DMs when M is even. This alternating pattern causes $(-1)^{N+1}$ to cancel from the condition for *da-d* and *da-a* DMs, and $(-1)^{M+1}$ to cancel from the condition for *ad-d* and *ad-a* DMs.

Finally, when the DM prefers to aggregate harms on x and on y (*aa-d* and *aa-a*), the parity of M, N and $M+N$ are all irrelevant because a negative sign on $u^{(M,N)}$ always indicates a preference to disaggregate harms across attributes and a positive sign on $u^{(M,N)}$ always represents a preference to aggregate them. The alternating pattern for each attribute implies that now both orders vanish and the sign of $u^{(M,N)}$ alone determines the DM's apportionment preference across attributes.

5.2 One desirable and one undesirable attribute (case DU)

In a next step, we provide the signs of the cross-derivatives of the utility function when one attribute is desirable and the other one is undesirable. Section 2.3 provides the link between risk apportionment preferences on individual attributes and the signs of the unidirectional derivatives of the utility function. When looking at cross-derivatives, the **DU** case shows a different pattern than the **DD** case. Here is our result.

Theorem 2 (Case **DU**). *Consider the case in which the first attribute is desirable and the second attribute is undesirable, and let $M, N \geq 1$.*

(i) DMs with apportionment preference dd-d have $(-1)^M u^{(M,N)} \leq 0$, DMs with apportionment preference dd-a have $(-1)^M u^{(M,N)} \geq 0$.

(ii) DMs with apportionment preference da-d have $(-1)^{M+N+1} u^{(M,N)} \leq 0$, DMs with apportionment preference da-a have $(-1)^{M+N+1} u^{(M,N)} \geq 0$.

(iii) DMs with apportionment preference ad-d have $u^{(M,N)} \geq 0$*, DMs with apportionment preference ad-a have* $u^{(M,N)} \leq 0$*.*

(iv) DMs with apportionment preference aa-d have $(-1)^N u^{(M,N)} \leq 0$*, DMs with apportionment preference aa-a have* $(-1)^N u^{(M,N)} \geq 0$*.*

Appendix A.4 gives the proof. As in the **DD** case, we obtain different criteria on the utility function depending on the DM's apportionment preference on the individual attributes, her apportionment preference across attributes, and the parity of the risk changes. How these criteria are assigned to the DM's apportionment preference has changed. When the DM prefers to disaggregate harms on x and on y (*dd-d* and *dd-a*), her apportionment preference across attributes now depends on the parity of M and not the parity of $M+N$ as was the case under **DD**. When M is odd, a positive sign on $u^{(M,N)}$ now indicates a preference to disaggregate harms across attributes and a negative sign a preference to aggregate harms across attributes. When M is even, the interpretation of the signs flips. Comparing Theorems 1(i) and 2(i), the two criteria are different when N is odd and identical when N is even.

When the DM prefers to disaggregate harms on x and aggregate harms on y (*da-d* and *da-a*), her apportionment preference across attributes depends on the parity of the total order $M+N$ in the **DU** case. For **DD**, the parity of M was decisive. Yet again, the two criteria differ for N odd and coincide for N even. When the DM prefers to aggregate harms on x and disaggregate harms on y (*ad-d* and *ad-a*), a positive sign on $u^{(M,N)}$ indicates a preference to disaggregate harms across attributes in the **DU** case. For **DD**, the parity of N was critical. The two criteria differ for N odd and coincide for N even. Finally, when the DM prefers to aggregate harms on x and on y (*aa-d* and *aa-a*), her apportionment preference across attributes depends on the parity of N in the **DU** case. For **DD**, a negative sign on $u^{(M,N)}$ always indicated a preference to disaggregate harms across attributes. As before the two criteria differ for N odd and coincide for N even.

This observation extends the comparison of Propositions 1 and 2. For correlation aversion and cross-prudence in y, the signs flip when going from **DD** to **DU** because we have $N = 1$, an odd number. For cross-prudence in x and cross-temperance, the signs stay the same because we have $N = 2$, an even number. To understand why the parity of N determines whether the sign on the cross-derivative needs to be flipped or not, we examine Nth-degree risk attitudes in the **DD** and **DU** case. Consider DMs who prefer to disaggregate harms on y. For **DD**, they are Nth-degree risk-averse for all $N \geq 1$. For **DU**, they are Nth-degree risk-loving for N odd and Nth-degree risk-averse for N even. They agree on the fact that even-order risk increases on y are unfavorable but disagree on odd-order risk increases. DMs who prefer to aggregate harms on y are Nth-degree risk-averse for N odd and Nth-degree risk-loving for N even in the **DD** case. For **DU**, they are always Nth-degree risk-loving. Yet again, they agree that even-order risk increases on y are favorable but disagree on odd-order risk increases. Regardless of whether the apportionment preference on y is combining good with bad or combining good with good and bad with bad, the signs of the cross-derivative remain unchanged when N is even but need to be flipped when N is odd.

5.3 Two undesirable attributes (case UU)

Finally, we consider the **UU** case with two undesirable attributes, $u^{(1,0)} \leq 0$ and $u^{(0,1)} \leq 0$. Section 2.3 provides the signs of the unidirectional derivatives of the utility function depending on the DM's risk apportionment preferences regarding the individual attributes x and y. We will now look at the signs of the cross-derivative of the utility function. The next result summarizes how the DM's risk apportionment preference determines these signs.

Theorem 3 (Case **UU**). *Consider the case of two undesirable attributes and let $M, N \geq 1$.*

(i) DMs with apportionment preference dd-d have $u^{(M,N)} \leq 0$, DMs with apportionment preference dd-a have $u^{(M,N)} \geq 0$.

(ii) DMs with apportionment preference da-d have $(-1)^N u^{(M,N)} \geq 0$, DMs with apportionment preference da-a have $(-1)^N u^{(M,N)} \leq 0$.

(iii) DMs with apportionment preference ad-d have $(-1)^M u^{(M,N)} \geq 0$, DMs with apportionment preference ad-a have $(-1)^M u^{(M,N)} \leq 0$.

(iv) DMs with apportionment preference aa-d have $(-1)^{M+N+1} u^{(M,N)} \geq 0$, DMs with apportionment preference aa-a have $(-1)^{M+N+1} u^{(M,N)} \leq 0$.

Appendix A.5 provides the proof. The comparison between Theorems 2 and 3 follows along the same lines as the comparison between Theorems 1 and 2. When going from **DU** to **UU**, the two criteria are different for M odd and identical for M even. The reason is that DMs agree on whether an Mth-degree risk increase on x is favorable or unfavorable for M even but disagree when M is odd.

To compare Theorems 1 and 3, we start with a DM who prefers to disaggregate harms on x and y (*dd-d* and *dd-a*). For **DD**, her apportionment preference across attributes depends on the parity of $M + N$ while for **UU**, it is determined by the sign of $u^{(M,N)}$ directly and the parity of neither M, N nor $M + N$ matter. The two criteria differ for $M + N$ odd and coincide for $M + N$ even. When the DM prefers to disaggregate harms on x and aggregate harms on y (*da-d* and *da-a*), her apportionment preference across attributes depends on the parity of N in the **UU** case and on the parity of M in the **DD** case. When the DM prefers to aggregate harms on x and disaggregate harms on y (*ad-d* and *ad-a*), her apportionment preference across attributes depends on the parity of M in the **UU** case and on the parity of N in the **DD** case. In both cases the two criteria differ for $M + N$ odd and coincide for $M + N$ even. Finally, when the DM prefers to aggregate harms on x and y (*aa-d* and *aa-a*), her apportionment preference across attributes depends on the parity of $M + N$ in the **UU** case and is determined by the sign of $u^{(M,N)}$ in the **DD** case. Yet again, the two criteria differ for $M + N$ odd and coincide for $M + N$ even.

This observation extends the comparison of Propositions 1 and 3. For cross-prudence in x and cross-prudence in y, the signs flip when going from **DD** to **UU** because we have $M + N = 1 + 2 = 2 + 1 = 3$, an odd number. For correlation aversion and cross-temperance, the signs stay the same because we have $M+N = 1+1 = 2$ and $M+N = 2+2 = 4$, two even numbers. To explain why it is now the parity of the total order that determines whether the criterion needs to be adjusted, we examine the DM's Mth- and Nth-degree risk attitudes in the **DD** and **UU** cases. Consider a DM who prefer to disaggregate harms on x and y (*dd-d* and *dd-a*). For **DD**, she is Mth-degree risk-averse for all $M \geq 1$ and Nth-degree risk-averse for all $N \geq 1$. For **UU**, she is Mth-degree risk-loving for M odd, Mth-degree risk-averse for M even, Nth-degree risk-loving for N odd, and Nth-degree risk-averse for N even. When both M and N are even, the two DMs agree that an Mth-degree risk increase on x and an Nth-degree risk increase on y are both unfavorable. When both M and N are odd, the two risk increases are unfavorable in the **DD** case but favorable in the **UU** case. Given that two reversals occur when going from **DD** to **UU**, they cancel each other out and no adjustment to the sign of the cross-derivative is necessary.[10] When M is even and N odd or when M is odd and N even, only one of the risk increases becomes favorable when moving from **DD** to **UU**, and the sign of the cross-derivative flips. The reasoning is analogous for the other apportionment preferences on individual attributes.

5.4 A simple mathematical reconciliation

We will now show the consistency between the criteria stated in Theorems 1, 2 and 3 directly. Take the case of **DU** and let preferences be represented by utility function $u(x, y)$ for $(x, y) \in \mathcal{D} = [0, \overline{x}] \times [0, \overline{y}]$. We have $u^{(1,0)} \geq 0$ and $u^{(0,1)} \leq 0$. Define utility function $v(x, y) = u(x, \overline{y} - y)$ for $(x, y) \in \mathcal{D}$. Obviously, we have $v^{(1,0)} = u^{(1,0)} \geq 0$ and $v^{(0,1)} = -u^{(0,1)} \geq 0$ so that utility function v represents the **DD** case. More generally, we find that $v^{(M,N)} = (-1)^N u^{(M,N)}$.

As a consequence, when going from Theorem 1 to Theorem 2, we need to multiply each of the criteria by $(-1)^N$. For example, *dd-d* is characterized by $(-1)^{M+N+1} u^{(M,N)} \geq 0$ in the **DD** case. Multiplying by $(-1)^N$ yields $(-1)^{M+1} u^{(M,N)} \geq 0$ or, equivalently, $(-1)^M u^{(M,N)} \leq 0$, the criterion for *dd-d* in the **DU** case. Similarly, *da-d* is characterized by $(-1)^M u^{(M,N)} \geq 0$ in the **DD** case. Multiplying by $(-1)^N$ yields $(-1)^{M+N} u^{(M,N)} \geq 0$ or, equivalently, $(-1)^{M+N+1} u^{(M,N)} \leq 0$, the criterion for *da-d* in the **DU** case.

We also show the consistency between Theorems 1 and 3. When utility function u represents preferences in case of **UU**, we define $v(x, y) = u(\overline{x} - x, \overline{y} - y)$ and obtain a utility function for the **DD** case. We have $v^{(M,N)} = (-1)^{M+N} u^{(M,N)}$. Therefore, when going from Theorem 1 to Theorem 3, each of the criteria needs to be multiplied by $(-1)^{M+N}$. Finally, if u represents preferences in case of **UU**, then $v(x, y) = u(\overline{x} - x, y)$ is a utility function for the

[10] More specifically, a lottery that allocates the risk increases to different states combines good with bad for either DM, the only difference being that the labels "good" and "bad" need to be switched for both risk increases when going from **DD** to **UU**.

DU case. We have $v^{(M,N)} = (-1)^M u^{(M,N)}$ so that each of the criteria needs to be multiplied by $(-1)^M$ when going from Theorem 2 to Theorem 3.

Mathematically, it is easy to see the equivalence between the criteria in Theorems 1 to 3 even though nothing is learned about the underlying economic intuition. Some problems are more naturally formulated in terms of undesirable attributes, and our results show how to make the entire arsenal of the risk apportionment literature available in those situations. While the mathematical equivalence of Theorems 1 to 3 is easy to see if a utility representation exists, we emphasize that the concepts of correlation aversion, cross-prudence, cross-temperance and their higher-order extensions can be defined with the help of simple lotteries, and thus do not require the existence of a utility representation. As we showed in Section 4, an appropriate adjustment to the seed lotteries ensures that the definitions stay intact when going from the **DD** case to the **DU** and **UU** cases. The definition based on simple lotteries allows researchers to utilize multivariate risk preferences outside the narrow confines of the expected-utility model and regardless of whether desirable or undesirable attributes are studied. This flexibility broadens the scope of our results significantly.

6 Gollier's (2021) generalized risk apportionment theory

Recently, Gollier (2021) provides a generalization of Eeckhoudt et al.'s (2009) risk apportionment approach. He assumes that the Mth-degree riskiness of X is uncertain and that the Nth-degree riskiness of Y is uncertain. In his model, X is parameterized by random variable Θ, and Y is parameterized by random variable Ψ. Then, for realizations $\theta_2 > \theta_1$ of Θ, $X(\theta_2)$ has more Mth-degree risk than $X(\theta_1)$, and for realizations $\psi_2 > \psi_1$ of Ψ, $Y(\psi_2)$ has more Nth-degree risk than $Y(\psi_1)$. The uncertainty over the riskiness of X and Y is represented by a joint distribution function for (Θ, Ψ).

In the original approach by Eeckhoudt and Schlesinger (2006) and Eeckhoudt et al. (2009), Θ and Ψ are both limited to a support of $\{1, 2\}$ so that the level of riskiness can either be low or high. Furthermore, state probabilities are equal because only 50-50 lotteries are considered. Thirdly, only perfect negative or perfect positive correlation between Θ and Ψ are considered. Gollier's (2021) contribution is to show that all these appendages can be removed. He accomplishes this by utilizing the following notion of dependence.

Definition 2 (Tchen et al.,1980; Epstein and Tanny,1980)**.** For two pairs of random variables (Θ_1, Ψ_1) and (Θ_2, Ψ_2) with joint cumulative distribution functions H_1 and H_2, we say that (Θ_2, Ψ_2) is more concordant than (Θ_1, Ψ_1) if H_1 and H_2 have the same marginal distributions and $H_2(\theta, \psi) \geq H_1(\theta, \psi)$ for all (θ, ψ) in the relevant domain.

Tchen et al. (1980) use the term concordance for this change in the joint distribution. Epstein and Tanny (1980) show in their Theorem 1 that, for discrete random variables, an increase in concordance is obtained as a sequence of correlation-increasing transformations. An increase in concordance implies higher correlation, a higher Kendall's τ, and a higher

Spearman's ρ, see Tchen et al. (1980). We use the notion of concordance and apply it to the uncertainty over the riskiness of X and Y.

Definition 3 (Gollier,2021)**.** Let Θ be an index of the Mth-degree riskiness of X and Ψ be an index of the Nth-degree riskiness of Y. Then, $(X(\Theta_2), Y(\Psi_2))$ is an (M, N)-degree risk increase over $(X(\Theta_1), Y(\Psi_1))$ if (Θ_2, Ψ_2) is more concordant than (Θ_1, Ψ_1).

Gollier (2021) goes on to show that a change in the joint distribution of (X, Y) is an (M, N)-degree risk increase if and only if it reduces the expectation of $u(X, Y)$ for any utility function u whose (M, N) cross-derivative has the same sign as $(-1)^{M+N+1}$.[11] We can use this result and our Theorems 1 to 3 to assess a DM's attitude towards increases in (M, N)-degree risk. We call a DM (M, N)-degree risk-averse if she dislikes any increase in (M, N)-degree risk and (M, N)-degree risk-loving if she appreciates any increase in (M, N)-degree risk. We formulate our results as corollaries and dissociate the three cases for readability.

Corollary 1 (Case **DD**)**.** *Consider the case of two desirable attributes and let $M, N \geq 1$.*

(i) DMs with dd-d (dd-a) are (M, N)-degree risk-averse (risk-loving).

(ii) DMs with da-d (da-a) are (M, N)-degree risk-averse (risk-loving) for N odd and (M, N)-degree risk-loving (risk-averse) for N even.

(iii) DMs with ad-d (ad-a) are (M, N)-degree risk-averse (risk-loving) for M odd and (M, N)-degree risk-loving (risk-averse) for M even.

(iv) DMs with aa-d (aa-a) are (M, N)-degree risk-loving (risk-averse) for $M + N$ odd and (M, N)-degree risk-averse (risk-loving) for $M + N$ even.

The only risk apportionment preference that implies (M, N)-degree risk aversion throughout is *dd-d* with *dd-a* yielding universal (M, N)-degree risk loving. In all other cases, the DM's (M, N)-degree risk attitude flips as we look at different orders, and either the parity of N, M or $M + N$ is decisive. Let us look at the **DU** case next.

Corollary 2 (Case **DU**)**.** *Consider the case in which the first attribute is desirable and the second attribute is undesirable, and let $M, N \geq 1$.*

(i) DMs with dd-d (dd-a) are (M, N)-degree risk-loving (risk-averse) for N odd and (M, N)-degree risk-averse (risk-loving) for N even.

(ii) DMs with da-d (da-a) are (M, N)-degree risk-loving (risk-averse).

[11] In his Theorem 1, Gollier (2021) uses the sign criterion on the cross-derivative of the utility function to define an (M, N)-degree risk increase and then shows the equivalence to an increase in the concordance between the index of the Mth-degree riskiness of X and the index of the Nth-degree riskiness of Y. Given the equivalence, we go the reverse route, which makes it easier to connect our results to his.

(iii) *DMs with ad-d (ad-a) are (M, N)-degree risk-averse (risk-loving) for $M + N$ odd and (M, N)-degree risk-loving (risk-averse) for $M + N$ even.*

(iv) *DMs with aa-d (aa-a) are (M, N)-degree risk-loving (risk-averse) for M odd and (M, N)-degree risk-averse (risk-loving) for M even.*

Now the only risk apportionment preferences that implies (M, N)-degree risk aversion throughout is *da-a*. When the DM prefers to combine good with good and bad with bad on y, she is Nth-degree risk-loving on the second attribute for all N. If she prefers to aggregate harms across attributes, higher concordance between the Mth-degree riskiness of x and the Nth-degree riskiness of y makes her worse off. She would rather face low Mth-degree risk on x together with high Nth-degree risk on y (two good things) or high Mth-degree risk on x together with low Nth-degree risk on y (two bad things) instead of low Mth-degree risk on x (a good thing) together with low Nth-degree risk on y (a bad thing) or high Mth-degree risk on x (a bad thing) toghether with high Nth-degree risk on y (a good thing). While both *dd-d* DMs in the **DD** case and *da-a* DMs in the **DU** case are consistently (M, N)-degree risk-averse, the reasons for their preference are quite different.

Corollary 3 (Case **UU**)**.** *Consider the case of two undesirable attributes and let $M, N \geq 1$.*

(i) *DMs with dd-d (dd-a) are (M, N)-degree risk-loving (risk-averse) for $M + N$ odd and (M, N)-degree risk-averse (risk-loving) for $M + N$ even.*

(ii) *DMs with da-d (da-a) are (M, N)-degree risk-averse (risk-loving) for M odd and (M, N)-degree risk-loving (risk-averse) for M even.*

(iii) *DMs with ad-d (ad-a) are (M, N)-degree risk-averse (risk-loving) for N odd and (M, N)-degree risk-loving (risk-averse) for N even.*

(iv) *DMs with aa-d (aa-a) are (M, N)-degree risk-averse (risk-loving).*

It is now DMs with *aa-d* who are always (M, N)-degree risk-averse. Yet again the intuition for this preference differs from the previous discussion. When the DM prefers to combine good with good and bad wit bad on x and y individually, she is Mth-degree risk-loving on x and Nth-degree risk-loving on y for all M and N. If she prefers to disaggregate harms across attributes, higher concordance between the Mth-degree riskiness of x and the Nth-degree riskiness of y makes her worse off. The DM would rather face high Mth-degree risk on x (a good thing) together with low Nth-degree risk on y (a bad thing) or low Mth-degree risk on x (a bad thing) together with high Nth-degree risk on y (a good thing) instead of high Mth-degree risk on x and high Nth-degree risk on y (two good things) or low Mth-degree risk on x and low Nth-degree risk on y (two bad things). At the surface, the resulting lottery preference is the same as for *dd-d* DMs in the **DD** case and *da-a* DMs in the **DU** case. How we obtain this lottery preference is entirely different.

Table 4 in the appendix provides a compact overview of the (M, N)-degree risk attitudes implied by different apportionment preferences in the three cases. We fully agree with Gollier's (2021) conclusion that a DM's (M, N)-degree risk attitude can be characterized without knowledge of any of her lower-degree risk attitudes by signing $u^{(M,N)}$. Our point here is that, if one imposes a consistent apportionment preference on the individual attributes and across attributes, this implies specific (M, N)-degree risk attitudes. Corollaries 1 to 3 detail what these (M, N)-degree risk attitudes are. Our results also highlight that the underlying reasons for a particular (M, N)-degree risk attitude can vary considerably within each case and across cases. By making the apportionment preference explicit, we can uncover these reasons and provide economic intuition.[12]

7 Some special multivariate models

7.1 Multiplicative separability

The utility function is multiplicatively separable if we can write it as $u(x, y) = v(x)z(y)$ for univariate utility functions v and z. Bleichrodt and Quiggin (1999) use this utility function to assess the consistency of quality-adjusted life years with life-cycle preferences when both consumption and health are arguments of the utility function. One might think that the separability assumption is constraining and restricts the types of risk apportionment preferences one can model. Our next result shows that this is not the case. To the contrary, multiplicatively separable utility is quite flexible and can be used to model any of the eight combinations of apportionment preferences discussed in this paper.

Proposition 4. *Let the utility function be multiplicatively separable, $u(x, y) = v(x)z(y)$, for univariate utility functions v of the first attribute and z of the second attribute. In each of the three cases,* ***DD***, ***DU*** *or* ***UU***, *we find the following:*

- *If sgn(v) = sgn(z), the utility function can accommodate dd-a, da-a, ad-a and aa-a. The DM always prefers to aggregate harms across attributes.*
- *If sgn$(v) \neq$ sgn(z), the utility function can accommodate dd-d, da-d, ad-d and aa-d. The DM always prefers to disaggregate harms across attributes.*

Appendix A.6 provides the proof. For multiplicatively separable utility, the DM's apportionment preference across attributes is simply determined by the signs of the univariate utility functions v and z. If the two signs are the same, either both positive or both negative, the DM necessarily prefers to combine good with good and bad with bad across attributes. If

[12] Of course, DMs with a given (M, N)-degree risk attitude may not belong to any of the eight groups of apportionment preferences considered in this paper. In the univariate context, some prudent DMs are risk-averse, some prudent DMs are risk-loving, and some prudent DMs may be neither risk-averse nor risk-loving.

the two signs are different with one being positive and the other one negative, the DM prefers to combine good with bad across attributes.

We can use Proposition 4 to construct any of the eight combinations of apportionment preferences studied in this paper. Consider the univariate utility function $v(x)$ for $x \in [0, \overline{x}]$. If attribute x is desirable and the DM prefers to combine good with bad on x, then v is mixed risk-averse, $(-1)^{M+1}v^{(M)} \geq 0$ for all $M \geq 1$. The class of utility functions with harmonic absolute risk aversion (HARA) provides specific examples. Let

$$v(x) = \begin{cases} \zeta \cdot \left(\eta + \frac{x}{\gamma}\right)^{1-\gamma}, & \text{for } \gamma \neq 1, \\ \zeta \cdot \log(\eta + x), & \text{for } \gamma = 1, \end{cases}$$

with $\eta > 0$, $\gamma > 0$, and $\zeta > 0$ for $\gamma \leq 1$ and $\zeta < 0$ for $\gamma > 1$. Then, v is mixed risk-averse and increases from $\underline{v} = v(0)$ to $\overline{v} = v(\overline{x})$. If v is negative, then $\hat{v}(x) = v(x) - \underline{v} + 1$ is mixed risk-averse and positive. If v is positive, then $\hat{v}(x) = v(x) - \overline{v} - 1$ is mixed risk-averse and negative. In general, a utility function displays mixed risk aversion if and only if it is the mixture of negative exponential functions (see Caballé and Pomansky, 1996).

Let $v(x)$ be a mixed risk-averse utility function for a desirable attribute $x \in [0, \overline{x}]$. Then, $\check{v}(x) = -v(\overline{x} - x)$ is a mixed risk-loving utility function for the desirable attribute x. Indeed, $\check{v}^{(M)}(x) = (-1)^{M+1}v^{M}(\overline{x} - x) \geq 0$ for all $M \geq 1$. Similarly, define $\breve{v}(x) = v(\overline{x} - x)$. This utility function satisfies $\breve{v}^{(M)}(x) = (-1)^{M}v^{(M)}(\overline{x} - x) \leq 0$ for all $M \geq 1$, and thus represents a preference for combining good with bad for an undesirable attribute x. Along the lines of Ebert (2020), we label this as anti-mixed risk-loving. If we set $\dot{v}(x) = -v(x)$, utility function $\dot{v}(x)$ satisfies $(-1)^{M+1}\dot{v}^{(M)}(x) = (-1)^{M}v^{(M)}(x) \leq 0$ for all $M \geq 1$, and thus represents a preference for combining good with good and bad with bad for an undesirable attribute x. Following Ebert (2020), we call this anti-mixed risk-averse.[13]

We can thus use the large class of mixed risk-averse utility functions to construct mixed risk-loving, anti-mixed risk-averse, and anti-mixed risk-loving utility functions. Shifting a utility function up or down ensures the desired sign. This puts us in a position to construct any of the eight combinations of apportionment preferences with the help of Proposition 4. To economize on space, we carry this out in Appendix B. Applied decision theorists can use this as a toolbox to construct utility functions with desired properties.

7.2 Equivalent monetary utility

We only consider a simple case of equivalent monetary utility, which is $u(x, y) = v(x + Ay)$ for $A \neq 0$. In this case, parameter A measures the marginal rate of substitution of attribute

[13] For an undesirable attribute, combining good with bad is characterized by a consistent negative sign while combining good with good and bad with bad is characterized by alternating signs, starting with a negative one. The pattern is thus reversed compared to the case of a desirable attribute, for which good with bad has the alternating sign pattern, and good with good and bad with bad a consistent positive sign.

y for attribute x. We assume for simplicity that A is constant and does not depend on the levels of x and y. Our next result shows that this restricts the DM's apportionment preference considerably.

Proposition 5. *Consider an equivalent monetary utility function with a constant marginal rate of substitution between attributes. In each of the three cases,* ***DD***, ***DU*** *or* ***UU***, *the utility function can accommodate either dd-d or aa-a.*

Appendix A.7 states the proof. In other words, equivalent monetary utility with a constant marginal rate of substitution imposes a strong consistency assumption on the DM's risk apportionment preference. She either prefers to combine good with bad on the individual attributes as well as across attributes, or she prefers to combine good with good and bad with bad on the individual attributes as well as across attributes. The other six combinations of apportionment preferences are excluded per assumption with this specification. This illustrates clearly that some simplifying assumptions that are sometimes made for convenience or tractability, can have far-reaching economic implications.[14]

8 Conclusion

Risk apportionment has revolutionized our understanding of higher-order risk preferences and accelerated their use in economics and finance. In this paper, we advanced the theory of risk apportionment for multivariate risks along several dimensions. We defined the concepts of correlation aversion, cross-prudence and cross-temperance in terms of simple lotteries when one or both attributes are undesirable. We characterized risk apportionment preferences across attributes by signing cross-derivatives of the utility function. We related our results to popular multivariate models and explained how to construct any of the eight combinations of apportionment preferences studied in this paper. It is our hope that these tools will help improve our understanding of risk-taking behavior in the many situations in which people face several attributes, some of which may be undesirable.

[14] Equivalent monetary utility is more flexible if we allow the marginal rate of substitution to depend on the levels of the attributes. We leave it for future research to determine how flexible this specification can be.

References

Bleichrodt, H. and Quiggin, J. (1999). Life-cycle preferences over consumption and health: when is cost-effectiveness analysis equivalent to cost–benefit analysis? *Journal of Health Economics*, 18(6): 681–708.

Bleichrodt, H. and van Bruggen, P. (2021). Reflection effect for higher-order risk preferences. *Review of Economics and Statistics (forthcoming)*.

Brockett, P. L. and Golden, L. L. (1987). A class of utility functions containing all the common utility functions. *Management Science*, 33(8): 955–964.

Caballé, J. and Pomansky, A. (1996). Mixed risk aversion. *Journal of Economic Theory*, 71(2): 485–513.

Courbage, C. and Rey, B. (2016). Decision thresholds and changes in risk for preventive treatment. *Health Economics*, 25(1): 111–124.

Crainich, D., Eeckhoudt, L., and Trannoy, A. (2013). Even (mixed) risk lovers are prudent. *American Economic Review*, 103(4): 1529–35.

Deck, C. and Schlesinger, H. (2014). Consistency of higher order risk preferences. *Econometrica*, 82(5): 1913–1943.

Denuit, M., De Vylder, E., and Lefevre, C. (1999). Extremal generators and extremeal distributions for the continuous s-convex stochastic ordering. *Insurance: Mathematics and Economics*, 24(3): 201–217.

Drèze, J. H. and Modigliani, F. (1972). Consumption decisions under uncertainty. *Journal of Economic Theory*, 5(3): 308–335.

Ebert, S. (2013). Even (mixed) risk lovers are prudent: Comment. *American Economic Review*, 103(4): 1536–37.

Ebert, S. (2020). Decision making when things are only a matter of time. *Operations Research*, 68(5): 1564–1575.

Eeckhoudt, L., Rey, B., and Schlesinger, H. (2007). A good sign for multivariate risk taking. *Management Science*, 53(1): 117–124.

Eeckhoudt, L. and Schlesinger, H. (2006). Putting risk in its proper place. *American Economic Review*, 96(1): 280–289.

Eeckhoudt, L., Schlesinger, H., and Tsetlin, I. (2009). Apportioning of risks via stochastic dominance. *Journal of Economic Theory*, 144(3): 994–1003.

Eeckhoudt, L. R., Laeven, R. J., and Schlesinger, H. (2020). Risk apportionment: The dual story. *Journal of Economic Theory*, 185: 1–27.

Ekern, S. (1980). Increasing nth degree risk. *Economics Letters*, 6(4): 329–333.

Epstein, L. G. and Tanny, S. M. (1980). Increasing generalized correlation: a definition and some economic consequences. *Canadian Journal of Economics*, 13(1): 16–34.

Gollier, C. (2021). A general theory of risk apportionment. *Journal of Economic Theory*, 192: 105189.

Haering, A., Heinrich, T., and Mayrhofer, T. (2020). Exploring the consistency of higher order risk preferences. *International Economic Review*, 61(1): 283–320.

Jouini, E., Napp, C., and Nocetti, D. (2013). Economic consequences of nth-degree risk increases and nth-degree risk attitudes. *Journal of Risk and Uncertainty*, 47(2): 199–224.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2): 263–292.

Keeney, R. L., Raiffa, H., and Meyer, R. F. (1993). *Decisions with multiple objectives: Preferences and value trade-offs*. Cambridge University Press.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica*, 58(1): 53–73.

Kimball, M. S. (1993). Standard risk aversion. *Econometrica*, 61(3): 589–611.

Leland, H. E. (1968). Saving and uncertainty: The precautionary demand for saving. *The Quarterly Journal of Economics*, 82(3): 465–473.

Menegatti, M. (2001). On the conditions for precautionary saving. *Journal of Economic Theory*, 98(1): 189–193.

Menegatti, M. and Peter, R. (2021). Changes in risky benefits and in risky costs: A question of the right order. *Management Science*, 68(5): 3625–3634.

Menezes, C., Geiss, C., and Tressler, J. (1980). Increasing downside risk. *American Economic Review*, 70(5): 921–932.

Menezes, C. F. and Wang, X. H. (2005). Increasing outer risk. *Journal of Mathematical Economics*, 41(7): 875–886.

Noussair, C. N., Trautmann, S. T., and Van de Kuilen, G. (2014). Higher order risk attitudes, demographics, and financial decisions. *The Review of Economic Studies*, 81(1): 325–355.

Rothschild, M. and Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3): 225–243.

Sandmo, A. (1970). The effect of uncertainty on saving decisions. *The Review of Economic Studies*, 37(3): 353–360.

Scott, R. C. and Horvath, P. A. (1980). On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, 35(4): 915–919.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, 38(2): 332–382.

Tchen, A. H. et al. (1980). Inequalities for distributions with given marginals. *The Annals of Probability*, 8(4): 814–827.

Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica*, 55(1): 95–115.

A Proofs

A.1 Signs of $u^{(0,N)}$ for all $N \geq 1$ in the DU case

Lemma 2. *Consider the* ***DU*** *case so that attribute y is undesirable and let preferences be represented by a smooth utility function $u(x, y)$.*

(i) If the DM prefers to combine good with bad on y, then $u^{(0,N)} \leq 0$ for all $N \geq 1$.

(ii) If the DM prefers to combine good with good and bad with bad on y, then $(-1)^{N+1}u^{(0,N)} \leq 0$ for all $N \geq 1$.

Proof. We show the two statements by mathematical induction. For $N = 1$, $u^{(0,1)} \leq 0$ holds by assumption because y is undesirable.

Now assume the statements are true for a given $N \geq 1$. Let Y_1, Y_2, Y_1' and Y_2' be four mutually independent random variables with Y_2 having more first-degree risk than Y_1, and Y_2' having more Nth-degree risk than Y_1'. The DM prefers Y_2 over Y_1 because of $u^{(0,1)} \leq 0$. If N is odd and she prefers to combine good with bad, she is Nth-degree risk-loving and thus prefers Y_2' over Y_1'. Combining good with bad implies that she also prefers the 50-50 lottery $[(x, Y_2 + Y_1'); (x, Y_2' + Y_1)]$ over the 50-50 lottery $[(x, Y_2 + Y_2'); (x, Y_1 + Y_1')]$ because the first lottery combines high first-degree risk (a good thing) with low Nth-degree risk (a bad thing) and high Nth-degree risk (a good thing) with low first-degree risk (a bad thing) whereas the second lottery combines high first-degree risk with high Nth-degree risk (two good things) and low first-degree risk with low Nth-degree risk (two bad things). If the DM always has said lottery preference, we know from Eeckhoudt et al. (2009) that $(-1)^{N+2}u^{(0,N+1)} \geq 0$. For N odd, this simplifies to $u^{(0,N+1)} \leq 0$ as claimed in statement (i).

If N is odd and the DM prefers to combine good with good and bad with bad, she is also Nth-degree risk-loving but now prefers the 50-50 lottery $[(x, Y_2 + Y_2'); (x, Y_1 + Y_1')]$ over the 50-50 lottery $[(x, Y_2 + Y_1'); (x, Y_2' + Y_1)]$ because of combining good with good and bad with bad. We know from Eeckhoudt et al. (2009) that this lottery preference is characterized by $(-1)^{N+2}u^{(0,N+1)} \leq 0$, as claimed in statement (ii).

If N is even and the DM prefers to combine good with bad, she is Nth-degree risk-averse and thus prefers Y_1' over Y_2'. Combining good with bad now implies that she always prefers the 50-50 lottery $[(x, Y_2+Y_2'); (x, Y_1+Y_1')]$ over the 50-50 lottery $[(x, Y_2+Y_1'); (x, Y_2'+Y_1)]$ because the first lottery combines high first-degree risk (a good thing) with high Nth-degree risk (a bad thing) and low first-degree risk (a bad thing) with low Nth-degree risk (a good thing) whereas the second lottery combines high first-degree risk with low Nth-degree risk (two good things) and high Nth-degree risk with low first-degree risk (two bad things). It follows from Eeckhoudt et al. (2009) that $(-1)^{N+2}u^{(0,N+1)} \leq 0$, which simplifies to $u^{(0,N+1)} \leq 0$ because N is even. This verifies statement (i) .

If N is even and the DM prefers to combine good with good and bad with bad, she is Nth-degree risk-loving and thus prefers Y_2' over Y_1'. Combining good with good and bad with bad leads to a lottery preference of $[(x, Y_2+Y_2'); (x, Y_1+Y_1')]$ over $[(x, Y_2+Y_1'); (x, Y_2'+Y_1)]$, which is characterized by $(-1)^{N+2}u^{(0,N+1)} \leq 0$, as claimed in statement (*ii*). So if statements (*i*) and (*ii*) are true for a given $N \geq 1$, they also hold for $N+1$, which completes the proof. $\square$

A.2 Sign of $u^{(M,N)}$ based on Eeckhoudt et al.'s (2009) approach

Assume the DM prefers lottery $[(X_1, Y_2); (X_2, Y_1)]$ over lottery $[(X_1, Y_1); (X_2, Y_2)]$ for all sets of four mutually independent random variables X_1, X_2, Y_1 and Y_2 such that X_2 has more Mth-degree risk than X_1, and Y_2 has more Nth-degree risk than Y_1. In terms of expected utility, the DM's lottery preference reads

$$\tfrac{1}{2}\mathbb{E}u(X_1, Y_2) + \tfrac{1}{2}\mathbb{E}u(X_2, Y_1) \geq \tfrac{1}{2}\mathbb{E}u(X_1, Y_1) + \tfrac{1}{2}\mathbb{E}u(X_2, Y_2),$$

which is equivalent to

$$\mathbb{E}u(X_2, Y_1) - \mathbb{E}u(X_2, Y_2) \geq \mathbb{E}u(X_1, Y_1) - \mathbb{E}u(X_1, Y_2).$$

Define auxiliary function $v(x) = \mathbb{E}u(x, Y_1) - \mathbb{E}u(x, Y_2)$; the last inequality can then be rewritten as $\mathbb{E}v(X_2) \geq \mathbb{E}v(X_1)$. If this inequality holds for every Mth-degree risk increase from X_1 to X_2, it follows from Lemma 1 that $-v$ must be Mth-degree risk-averse, that is, $(-1)^M v^{(M)}(x) \geq 0$. Using the definition of v, this is equivalent to

$$(-1)^M \mathbb{E}u^{(M,0)}(x, Y_1) \geq (-1)^M \mathbb{E}u^{(M,0)}(x, Y_2).$$

Per Lemma 1, this inequality holds for every Nth-degree risk increase from Y_1 to Y_2 if and only if $(-1)^M u^{(M,0)}$ is Nth-degree risk-averse in y, that is, if and only if

$$(-1)^{M+N+1}u^{(M,N)} \geq 0.$$

A.3 Proof of Theorem 1

Let $\mathcal{L}_1 = [(X_1, Y_2); (X_2, Y_1)]$ be the lottery where the Mth-degree risk increase on x and the Nth-degree risk increase on y occur in different states, and $\mathcal{L}_2 = [(X_1, Y_1); (X_2, Y_2)]$ be the lottery where they occur in the same state. For *dd-d* and *dd-a* DMs, both risk increases are unfavorable changes so that lottery $\mathcal{L}_1$ represents combining good with bad across attributes whereas lottery $\mathcal{L}_2$ represents combining good with good and bad with bad across attributes. A universal preference of $\mathcal{L}_1 \succsim \mathcal{L}_2$ is equivalent to $(-1)^{M+N+1}u^{(M,N)} \geq 0$ whereas a universal preference of $\mathcal{L}_2 \succsim \mathcal{L}_1$ is equivalent to $(-1)^{M+N+1}u^{(M,N)} \leq 0$. This shows (*i*).

For *da-d* and *da-a* DMs, the Mth-degree risk increase on x is always an unfavorable change. However, the Nth-degree risk increase on y is an unfavorable change when N is odd and a favorable change when N is even. The lottery $\mathcal{L}_1$ then represents combining good with bad when N is odd. It represents combining good with good and bad with bad when N is even. So a preference to disaggregate harms across attributes leads to $\mathcal{L}_1 \succsim \mathcal{L}_2$ for N odd and to $\mathcal{L}_2 \succsim \mathcal{L}_1$ for N even. In terms of the utility function, this means $(-1)^{M+N+1}u^{(M,N)} \geq 0$ for N odd and $(-1)^{M+N+1}u^{(M,N)} \leq 0$ for N even. When N is odd, $(-1)^{M+N+1} = (-1)^M$, and when N is even, $(-1)^{M+N+1} = (-1)^{M+1}$. So the criterion on the utility function can be consolidated to $(-1)^M u^{(M,N)} \geq 0$ for *da-d* DMs, and to $(-1)^M u^{(M,N)} \leq 0$ for *da-a* DMs. This proves *(ii)*. Result *(iii)* follows with the same argument replacing M by N.

For *aa-d* and *aa-a* DMs, the Mth-degree risk increase on x is an unfavorable change when M is odd and a favorable change when M is even. Likewise, the Nth-degree risk increase on y is an unfavorable change when N is odd and a favorable change when N is even. Lottery $\mathcal{L}_1$ thus represents combining good with bad when both M and N are odd or when both M and N are even. When M is odd and N is even or M is even and N is odd, lottery $\mathcal{L}_1$ corresponds to combining good with good and bad with bad. This implies $(-1)^{M+N+1}u^{(M,N)} \geq 0$ when both M and N are odd or both are even, which can be simplified to $u^{(M,N)} \leq 0$. For M odd and N even or M even and N odd, we obtain $(-1)^{M+N+1}u^{(M,N)} \leq 0$, which can also be simplified to $u^{(M,N)} \leq 0$. So regardless of the parity of M and N, *aa-d* DMs have $u^{(M,N)} \leq 0$ whereas *aa-a* DMs have $u^{(M,N)} \geq 0$

A.4 Proof of Theorem 2

Let X_1, X_2, Y_1 and Y_2 be four mutually independent random variables with X_2 having more Mth-degree risk than X_1 and Y_2 having more Nth-degree risk than Y_1. Let $\mathcal{L}_1 = [(X_1, Y_2); (X_2, Y_1)]$ be the lottery where the Mth-degree risk increase on x and the Nth-degree risk increase on y occur in different states, and $\mathcal{L}_2 = [(X_1, Y_1); (X_2, Y_2)]$ be the lottery where they occur in the same state. For *dd-d* and *dd-a* DMs, the Mth-degree risk increase on x is always an unfavorable change whereas the Nth-degree risk increase on y is a favorable change when N is odd and an unfavorable change when N is even. So *dd-d* DMs prefer $\mathcal{L}_2$ over $\mathcal{L}_1$ when N is odd, leading to $(-1)^{M+N+1}u^{(M,N)} \leq 0$, while they prefer $\mathcal{L}_1$ over $\mathcal{L}_2$ when N is even, leading to $(-1)^{M+N+1}u^{(M,N)} \geq 0$. Taking the parity of N into account, the condition on the utility function can be condensed to $(-1)^M u^{(M,N)} \leq 0$ for *dd-d* DMs, and to $(-1)^M u^{(M,N)} \geq 0$ for *dd-a* DMs.

For *da-d* and *da-a* DMs, the Mth-degree risk increase on x is always an unfavorable change and the Nth-degree risk increase on y is always a favorable change. Therefore, *da-d* DMs have a preference of $\mathcal{L}_2$ over $\mathcal{L}_1$ because $\mathcal{L}_2$ combines low Mth-degree risk on x (a good thing) with low Nth-degree risk on y (a bad thing) and high Mth-degree risk on x (a bad thing) with high Nth-degree risk on y (a good thing) whereas $\mathcal{L}_1$ combines low Mth-degree risk on x with high Nth-degree risk on y (two good things) and high Mth-degree risk on x with low Nth-degree risk on y (two bad things). This leads to $(-1)^{M+N+1}u^{(M,N)} \leq 0$ for *da-d* DMs and to $(-1)^{M+N+1}u^{(M,N)} \geq 0$ for *da-a* DMs.

For *ad-d* and *ad-a* DMs, the Mth-degree risk increase on x is an unfavorable change when M is odd and a favorable change when M is even whereas the Nth-degree risk increase on

y is a favorable change when N is odd and an unfavorable change when N is even. For M odd, *ad-d* DMs then prefer $\mathcal{L}_2$ over $\mathcal{L}_1$ when N is odd, and $\mathcal{L}_1$ over $\mathcal{L}_2$ when N is even. This leads to $(-1)^M u^{(M,N)} \leq 0$, which can be further simplified to $u^{(M,N)} \geq 0$ because M is odd. When M is even instead, *ad-d* DMs prefers $\mathcal{L}_1$ over $\mathcal{L}_2$ when N is odd, and $\mathcal{L}_2$ over $\mathcal{L}_1$ when N is even. This yields $(-1)^M u^{(M,N)} \geq 0$, which is also equivalent to $u^{(M,N)} \geq 0$ because M is even. Therefore, *ad-d* DMs have $u^{(M,N)} \geq 0$ and *ad-a* DMs have $u^{(M,N)} \leq 0$.

For *aa-d* and *aa-a* DMs, the Mth-degree risk increase on x is an unfavorable change when M is odd and a favorable change when M is even whereas the Nth-degree risk increase on y is always a favorable change. So for M odd, *aa-d* DMs prefers $\mathcal{L}_2$ over $\mathcal{L}_1$, leading to $(-1)^{M+N+1}u^{(M,N)} \leq 0$. This can be simplified to $(-1)^N u^{(M,N)} \leq 0$. When M is even, *aa-d* DMs prefer $\mathcal{L}_1$ over $\mathcal{L}_2$, leading to $(-1)^{M+N+1}u^{(M,N)} \geq 0$. This can also be simplified to $(-1)^N u^{(M,N)} \leq 0$. For *aa-a* DMs, matters are reversed so that $(-1)^N u^{(M,N)} \geq 0$.

A.5 Proof of Theorem 3

Let X_1, X_2, Y_1 and Y_2 be four mutually independent random variables with X_2 having more Mth-degree risk than X_1 and Y_2 having more Nth-degree risk than Y_1. Let $\mathcal{L}_1 = [(X_1, Y_2); (X_2, Y_1)]$ be the lottery where the Mth-degree risk increase on x and the Nth-degree risk increase on y occur in different states, and $\mathcal{L}_2 = [(X_1, Y_1); (X_2, Y_2)]$ be the lottery where they occur in the same state. For *dd-d* and *dd-a* DMs, the Mth-degree risk increase on x is a favorable change when M is odd and an unfavorable change when M is even, and the same is the case for the Nth-degree risk increase on y. In this case, a *dd-d* DM prefers lottery $\mathcal{L}_1$ over lottery $\mathcal{L}_2$ when M and N are both odd or both even, and has the reverse lottery preference otherwise. This implies $(-1)^{M+N+1}u^{(M,N)} \geq 0$ when M and N are both odd or both even, which simplifies the condition to $u^{(M,N)} \leq 0$, and $(-1)^{M+N+1}u^{(M,N)} \leq 0$ when M is odd and N even or M is even and N odd, which again simplifies to $u^{(M,N)} \leq 0$. The lottery preference is reversed for *dd-a* DMs, which leads to $u^{(M,N)} \geq 0$.

For *da-d* and *da-a* DMs, the Mth-degree risk increase on x is a favorable change when M is odd and an unfavorable change when M is even but the Nth-degree risk increase on y is always a favorable change. Consequently, *da-d* DMs prefer $\mathcal{L}_1$ over $\mathcal{L}_2$ when M is odd and $\mathcal{L}_2$ over $\mathcal{L}_1$ when M is even. This leads to $(-1)^{M+N+1}u^{(M,N)} \geq 0$ for M odd and to $(-1)^{M+N+1}u^{(M,N)} \leq 0$ for M even. The condition then simplifies to $(-1)^N u^{(M,N)} \geq 0$ for *da-d* DMs. Reversing the lottery preferences shows $(-1)^N u^{(M,N)} \leq 0$ for *da-a* DMs. Result (*iii*) follows with the same argument replacing M by N.

For *aa-d* and *aa-a* DMs, the Mth-degree risk increase on x and the Nth-degree risk increase on y are both always favorable changes. Hence, a *aa-d* DM always prefers $\mathcal{L}_1$ over $\mathcal{L}_2$, implying $(-1)^{M+N+1}u^{(M,N)} \geq 0$, whereas a *aa-a* DM always prefers $\mathcal{L}_2$ over $\mathcal{L}_1$, implying $(-1)^{M+N+1}u^{(M,N)} \leq 0$

A.6 Proof of Proposition 4

For multiplicatively separable utility, we have $u^{(M,N)} = v^{(M)} \cdot z^{(N)}$ for $M, N \geq 1$, and in particular $u^{(M,0)} = v^{(M)} \cdot z$ and $u^{(0,N)} = v \cdot z^{(N)}$ for the unidirectional derivatives. Let us start with the **DD** case and assume $v(x) > 0$ for $x \in [0, \overline{x}]$ and $z(y) > 0$ for $y \in [0, \overline{y}]$ so that $\text{sgn}(v) = \text{sgn}(z)$. Combining good with bad on x is equivalent to $(-1)^{M+1}u^{(M,0)} \geq 0$ for all $M \geq 1$ or $(-1)^{M+1}v^{(M)} \geq 0$ for all $M \geq 1$. Combining good with good and bad with bad on x is equivalent to $u^{(M,0)} \geq 0$ for all $M \geq 1$ or $v^{(M)} \geq 0$ for all $M \geq 1$. Combining good with bad on y is equivalent to $(-1)^{N+1}u^{(0,N)} \geq 0$ for all $N \geq 1$ or $(-1)^{N+1}z^{(N)} \geq 0$ for all $N \geq 1$. Combining good with good and bad with bad on y is equivalent to $u^{(0,N)} \geq 0$ for all $N \geq 1$ or $z^{(N)} \geq 0$ for all $N \geq 1$.

When the DM prefers to combine good with bad on both attributes individually, we obtain

$$(-1)^{M+N+1}u^{(M,N)} = (-1) \cdot \underbrace{(-1)^{M+1}v^{(M)}}_{\geq 0} \cdot \underbrace{(-1)^{N+1}z^{(N)}}_{\geq 0} \leq 0,$$

which characterizes *dd-a* according to Theorem 1(i). When she prefers to combine good with bad on x but good with good and bad with bad on y, we obtain

$$(-1)^{M}u^{(M,N)} = (-1) \cdot \underbrace{(-1)^{M+1}v^{M}}_{\geq 0} \cdot \underbrace{z^{N}}_{\geq 0} \leq 0,$$

which characterizes *da-a* according to Theorem 1(ii). When she prefers to combine good with good and bad with bad on x but good with bad on y, we obtain

$$(-1)^{N}u^{(M,N)} = (-1) \cdot \underbrace{v^{(M)}}_{\geq 0} \cdot \underbrace{(-1)^{N+1}z^{N}}_{\geq 0} \leq 0,$$

which characterize *ad-a* according to Theorem 1(iii). When she prefers to combine good with good and bad with bad on both attributes individually, we obtain

$$u^{(M,N)} = \underbrace{v^{(M)}}_{\geq 0} \cdot \underbrace{z^{(N)}}_{\geq 0} \geq 0,$$

which characterize *aa-a* according to Theorem 1(iv). Regardless of her apportionment preference on the individual attributes, she always prefers to aggregate harms across attributes.

Consider now that $v(x) > 0$ for $x \in [0, \overline{x}]$ but $z(y) < 0$ for $y \in [0, \overline{y}]$ so that $\text{sgn}(v) \neq \text{sgn}(z)$. Combining good with bad on x is now equivalent to $(-1)^{M+1}v^{(M)} \leq 0$ for all $M \geq 1$, and combining good with good and bad with bad on x is now equivalent to $v^{(M)} \leq 0$ for all $M \geq 1$. The signs of higher-order derivatives of utility function z are as in the case of both utility functions positive. As a result, all signs of the cross-derivatives flip and we now obtain *dd-d*, *da-d*, *ad-d* and *aa-d*. The DM now prefers to disaggregate harms across attributes. Similarly, if $v(x) < 0$ for $x \in [0, \overline{x}]$ and $z(y) > 0$ for $y \in [0, \overline{y}]$, combining good with bad on

y is equivalent to $(-1)^{N+1}z^{(N)} \leq 0$ for all $N \geq 1$, and combining good with good and bad with bad on y is equivalent to $z^{(N)} \leq 0$ for all $N \geq 1$. The signs of higher-order derivatives of utility function z are as in the case of both utility functions positive. Yet again, all signs of the cross-derivatives flip compared to the case with $\text{sgn}(v) = \text{sgn}(z)$, and we thus obtain *dd-d*, *da-d*, *ad-d* and *aa-d*. If both $v(x) < 0$ for $x \in [0, \overline{x}]$ and $z(y) < 0$ for $y \in [0, \overline{y}]$, the signs of higher-order derivatives of both utility functions flip and we obtain the same signs as in the case with $\text{sgn}(v) = \text{sgn}(z)$. In other words, we find *dd-a*, *da-a*, *ad-a* and *aa-a*.

The **DU** and **UU** cases follow a similar logic. We briefly look at the **DU** case. When $v(x) > 0$ for $x \in [0, \overline{x}]$ and $z(y) > 0$ for $y \in [0, \overline{y}]$, combining good with bad on x is equivalent to $(-1)^{M+1}v^{(M)} \geq 0$ for all $M \geq 1$, and combining good with good and bad with bad on x is equivalent to $v^{(M)} \geq 0$ for all $M \geq 1$. Combining good with bad on y is equivalent to $u^{(0,N)} \leq 0$ for all $N \geq 1$ or $z^{(N)} \leq 0$ for all $N \geq 1$. Combining good with good and bad with bad on y is equivalent to $(-1)^{N+1}u^{(0,N)} \leq 0$ for all $N \geq 1$ or $(-1)^{N+1}z^{(N)} \leq 0$ for all $N \geq 1$. Using Theorem 2, we find

$$(-1)^M u^{(M,N)} = (-1) \cdot \underbrace{(-1)^{M+1}v^{(M)}}_{\geq 0} \cdot \underbrace{z^{(N)}}_{\leq 0} \geq 0$$

for *dd-a*,

$$(-1)^{M+N+1} u^{(M,N)} = (-1) \cdot \underbrace{(-1)^{M+1}v^{(M)}}_{\geq 0} \cdot \underbrace{(-1)^{N+1}z^{(N)}}_{\leq 0} \geq 0$$

for *da-a*,

$$u^{(M,N)} = \underbrace{v^{(M)}}_{\geq 0} \cdot \underbrace{z^{(N)}}_{\leq 0} \leq 0$$

for *ad-a*, and

$$(-1)^N u^{(M,N)} = (-1) \cdot \underbrace{v^{(M)}}_{\geq 0} \cdot \underbrace{(-1)^{N+1}z^{(N)}}_{\leq 0} \geq 0$$

for *aa-a*. The DM always prefers to aggregate harms across attributes regardless of her apportionment preference on the individual attributes. When the sign of v switches from positive to negative, all signs of the higher-order derivatives of z flip and so do the signs of the cross-derivatives. If instead the sign of z switches from positive to negative and the sign of v is positive, all signs of the higher-order derivatives of v flip and so do the signs of the cross-derivatives. Regardless, as soon as $\text{sgn}(v) \neq \text{sgn}(z)$, we have *dd-d*, *da-d*, *ad-d* or *aa-d*, and the DM prefers to disaggregate harms across attributes. When both v and z are negative, the two sign reversals cancel each other out and we are back to *dd-a*, *da-a*, *ad-a* or *aa-a*, as in the case of both v and z positive.

A.7 Proof of Proposition 5

Let $u(x, y) = v(x + Ay)$ and consider the **DD** case first. The first attribute is desirable so that $v' \geq 0$. For the second attribute to be desirable, we then have $A \geq 0$ (except in the uninteresting case of $v' = 0$). If the DM prefers to combine good with bad on x, we obtain $(-1)^{M+1}v^{(M)} \geq 0$ for all $M \geq 1$. This implies $(-1)^{N+1}u^{(0,N)} = (-1)^{N+1}A^N v^{(N)} \geq 0$ for all $N \geq 1$ so that she prefers to combine good with bad on y. Furthermore, $(-1)^{M+N+1}u^{(M,N)} = (-1)^{M+N+1}A^N v^{(M+N)} \geq 0$ for all $M, N \geq 1$ so that she prefers to combine good with bad across attributes according to Theorem 1(i). The DM's preference is thus *dd-d*. If she prefers to combine good with good and bad with bad on x instead, we obtain $v^{(M)} \geq 0$ for all $M \geq 1$, which implies $u^{(0,N)} = A^N v^{(N)} \geq 0$ for all $N \geq 1$, and $u^{(M,N)} = A^N v^{M+N} \geq 0$ for all $M, N \geq 1$. From Theorem 1(iv), her preference is then *aa-a*.

In the **DU** case, we have $v' \geq 0$ and $A \leq 0$. If the DM prefers to combine good with bad on x, we obtain $(-1)^{M+1}v^{(M)} \geq 0$ for all $M \geq 1$. This implies $u^{(0,N)} = A^N v^{(N)} = (-1)(-A)^N(-1)^{N+1}v^{(N)} \geq 0$ for all $N \geq 1$ so that she prefers to combine good with bad on y. In addition we find $(-1)^M u^{(M,N)} = (-1)^M A^N v^{(M+N)} = (-1)(-A)^N(-1)^{M+N+1}v^{(M+N)} \leq 0$ so that she prefers to combine good with bad across attributes according to Theorem 2(i). The DM's preference is *dd-d*. If she prefers to combine good with good and bad with bad on x instead, we have $v^{(M)} \geq 0$ for all $M \geq 1$, which implies $(-1)^{N+1}u^{(0,N)} = (-1)(-A)^N v^{(N)} \leq 0$ for all $N \geq 1$, and $(-1)^N u^{(M,N)} = (-A)^N v^{(M+N)} \geq 0$ for all $M, N \geq 1$. Using Theorem 2(iv), her preference is then *aa-a*.

In the **UU** case, we have $v' \leq 0$ and $A \geq 0$. If the DM prefers to combine good with bad on x, we have $v^{(M)} \leq 0$ for all $M \geq 1$. This implies $u^{(0,N)} = A^N v^{(N)} \leq 0$ for all $N \geq 1$ so that she prefers to combine good with bad on y. We obtain $u^{(M,N)} = A^N v^{(M+N)} \leq 0$ for all $M, N \geq 1$ so that she prefers to combine good with bad across attributes, see Theorem 3(i). Her preference is *dd-d*. If she prefers to combine good with good and bad with bad on x instead, we have $(-1)^{M+1}v^{(M)} \leq 0$ for all $M \geq 1$. This implies $(-1)^{N+1}u^{(0,N)} = A^N(-1)^{N+1}v^{(N)} \leq 0$ for all $N \geq 1$ and $(-1)^{M+N+1}u^{(M,N)} = A^N(-1)^{M+N+1}v^{(M+N)} \leq 0$. According to Theorem 3(iv), the DM's preference is then *aa-a*.

In either one of the three cases, we either find *dd-d* or *aa-a* for monetary equivalent utility $u(x, y) = v(x + Ay)$ with a constant marginal rate of substitution between attributes.

order	***dd-d* and *dd-a***	***dd-d***	***dd-a***
$M+N=2$	$u^{(2,0)} \leq 0, u^{(0,2)} \leq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \geq 0$	$u^{(2,1)} \geq 0$, $u^{(1,2)} \geq 0$	$u^{(2,1)} \leq 0$, $u^{(1,2)} \leq 0$
$M+N=4$	$u^{(4,0)} \leq 0, u^{(0,4)} \leq 0$	$u^{(3,1)} \leq 0$, $u^{(2,2)} \leq 0$, $u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0$, $u^{(2,2)} \geq 0$, $u^{(1,3)} \geq 0$

order	***da-d* and *da-a***	***da-d***	***da-a***
$M+N=2$	$u^{(2,0)} \leq 0, u^{(0,2)} \geq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \geq 0$	$u^{(2,1)} \geq 0$, $u^{(1,2)} \leq 0$	$u^{(2,1)} \leq 0$, $u^{(1,2)} \geq 0$
$M+N=4$	$u^{(4,0)} \leq 0, u^{(0,4)} \geq 0$	$u^{(3,1)} \leq 0$, $u^{(2,2)} \geq 0$, $u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0$, $u^{(2,2)} \leq 0$, $u^{(1,3)} \geq 0$

order	***ad-d* and *ad-a***	***ad-d***	***ad-a***
$M+N=2$	$u^{(2,0)} \geq 0, u^{(0,2)} \leq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \geq 0$	$u^{(2,1)} \leq 0$, $u^{(1,2)} \geq 0$	$u^{(2,1)} \geq 0$, $u^{(1,2)} \leq 0$
$M+N=4$	$u^{(4,0)} \geq 0, u^{(0,4)} \leq 0$	$u^{(3,1)} \leq 0$, $u^{(2,2)} \geq 0$, $u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0$, $u^{(2,2)} \leq 0$, $u^{(1,3)} \geq 0$

order	***aa-d* and *aa-a***	***aa-d***	***aa-a***
$M+N=2$	$u^{(2,0)} \geq 0, u^{(0,2)} \geq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \geq 0$	$u^{(2,1)} \leq 0$, $u^{(1,2)} \leq 0$	$u^{(2,1)} \geq 0$, $u^{(1,2)} \geq 0$
$M+N=4$	$u^{(4,0)} \geq 0, u^{(0,4)} \geq 0$	$u^{(3,1)} \leq 0$, $u^{(2,2)} \leq 0$, $u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0$, $u^{(2,2)} \geq 0$, $u^{(1,3)} \geq 0$

Table 1: All signs up to order 4 for the **DD** case with two desirable attributes x and y, $u^{(1,0)} \geq 0$ and $u^{(0,1)} \geq 0$. Our classification distinguishes whether the DM prefers to disaggregate (in short: d) or aggregate (in short: a) harms on the first attribute (first letter), on the second attribute (second letter), and across attributes (third letter). Correlation aversion, cross-prudence in x and y, and cross-temperance are highlighted in blue, correlation loving, cross-imprudence in x and y, and cross-intemperance are highlighted in green. The signs are collected in Proposition 1.

order	***dd-d* and *dd-a***	***dd-d***	***dd-a***
$M+N=2$	$u^{(2,0)} \leq 0, u^{(0,2)} \leq 0$	$u^{(1,1)} \geq 0$	$u^{(1,1)} \leq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \geq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \leq 0$
$M+N=4$	$u^{(4,0)} \leq 0, u^{(0,4)} \leq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \geq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \leq 0$
order	***da-d* and *da-a***	***da-d***	***da-a***
$M+N=2$	$u^{(2,0)} \leq 0, u^{(0,2)} \geq 0$	$u^{(1,1)} \geq 0$	$u^{(1,1)} \leq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \leq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \geq 0$
$M+N=4$	$u^{(4,0)} \leq 0, u^{(0,4)} \geq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \geq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \leq 0$
order	***ad-d* and *ad-a***	***ad-d***	***ad-a***
$M+N=2$	$u^{(2,0)} \geq 0, u^{(0,2)} \leq 0$	$u^{(1,1)} \geq 0$	$u^{(1,1)} \leq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \geq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \leq 0$
$M+N=4$	$u^{(4,0)} \geq 0, u^{(0,4)} \leq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \geq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \leq 0$
order	***aa-d* and *aa-a***	***aa-d***	***aa-a***
$M+N=2$	$u^{(2,0)} \geq 0, u^{(0,2)} \geq 0$	$u^{(1,1)} \geq 0$	$u^{(1,1)} \leq 0$
$M+N=3$	$u^{(3,0)} \geq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \leq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \geq 0$
$M+N=4$	$u^{(4,0)} \geq 0, u^{(0,4)} \geq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \geq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \leq 0$

Table 2: All signs up to order 4 for **DU** with a desirable attribute x, $u^{(1,0)} \geq 0$, and an undesirable attribute y, $u^{(0,1)} \leq 0$. Our classification distinguishes whether the DM prefers to disaggregate (in short: d) or aggregate (in short: a) harms on the first attribute (first letter), on the second attribute (second letter), and across attributes (third letter). Correlation aversion, cross-prudence in x and y, and cross-temperance are highlighted in blue, correlation loving, cross-imprudence in x and y, and cross-intemperance are highlighted in green. The signs are collected in Proposition 2.

order	***dd-d* and *dd-a***	***dd-d***	***dd-a***
$M+N=2$	$u^{(2,0)} \leq 0, u^{(0,2)} \leq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \leq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \leq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \geq 0$
$M+N=4$	$u^{(4,0)} \leq 0, u^{(0,4)} \leq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \geq 0$

order	***da-d* and *da-a***	***da-d***	***da-a***
$M+N=2$	$u^{(2,0)} \leq 0, u^{(0,2)} \geq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \leq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \geq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \leq 0$
$M+N=4$	$u^{(4,0)} \leq 0, u^{(0,4)} \geq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \geq 0$

order	***ad-d* and *ad-a***	***ad-d***	***ad-a***
$M+N=2$	$u^{(2,0)} \geq 0, u^{(0,2)} \leq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \leq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \leq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \geq 0$
$M+N=4$	$u^{(4,0)} \geq 0, u^{(0,4)} \leq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \geq 0$

order	***aa-d* and *aa-a***	***aa-d***	***aa-a***
$M+N=2$	$u^{(2,0)} \geq 0, u^{(0,2)} \geq 0$	$u^{(1,1)} \leq 0$	$u^{(1,1)} \geq 0$
$M+N=3$	$u^{(3,0)} \leq 0, u^{(0,3)} \leq 0$	$u^{(2,1)} \geq 0,\ u^{(1,2)} \geq 0$	$u^{(2,1)} \leq 0,\ u^{(1,2)} \leq 0$
$M+N=4$	$u^{(4,0)} \geq 0, u^{(0,4)} \geq 0$	$u^{(3,1)} \leq 0,\ u^{(2,2)} \leq 0,\ u^{(1,3)} \leq 0$	$u^{(3,1)} \geq 0,\ u^{(2,2)} \geq 0,\ u^{(1,3)} \geq 0$

Table 3: All signs up to order 4 for **UU** with two undesirable attributes x and y, $u^{(1,0)} \geq 0$ and $u^{(0,1)} \geq 0$. Our classification distinguishes whether the DM prefers to disaggregate (in short: d) or aggregate (in short: a) harms on the first attribute (first letter), on the second attribute (second letter), and across attributes (third letter). Correlation aversion, cross-prudence in x and y, and cross-temperance are highlighted in blue, correlation loving, cross-imprudence in x and y, and cross-intemperance are highlighted in green. The signs are collected in Proposition 3.

Case	app. pref.	(M, N)-deg. risk att.	app. pref.	(M, N)-deg. risk att.	condition
DD	*dd-d*	averse	*dd-a*	loving	
	da-d	averse loving	*da-a*	loving averse	if N odd if N even
	ad-d	averse loving	*ad-a*	loving averse	if M odd if M even
	aa-d	loving averse	*aa-a*	averse loving	if $M + N$ odd if $M + N$ even
DU	*dd-d*	loving averse	*dd-a*	averse loving	if N odd if N even
	da-d	loving	*da-a*	averse	
	ad-d	averse loving	*ad-a*	loving averse	if $M + N$ odd if $M + N$ even
	aa-d	loving averse	*aa-a*	averse loving	if M odd if M even
UU	*dd-d*	loving averse	*dd-a*	averse loving	if $M + N$ odd if $M + N$ even
	da-d	averse loving	*da-a*	loving averse	if M odd if M even
	ad-d	averse loving	*ad-a*	loving averse	if N odd if N even
	aa-d	averse	*aa-a*	loving	

Table 4: Attitudes towards an increase in the (M, N)-degree riskiness of (X, Y) organized by the DM's apportionment preference. The table distinguishes the different combinations of apportionment preferences, states the implied (M, N)-degree risk attitude, and provides a condition if required. The first, second and third panel summarizes the results of Corollary 1, 2 and 3 for the case **DD**, **DU** and **UU**, respectively.

B Construction of multiplicatively separable utility functions with desired risk apportionment preferences

Table 5 shows how to construct any of the eight combinations of apportionment preferences (rows) in any of the three cases (columns) when the utility function is multiplicatively separable, $u(x, y) = v(x)z(y)$. The acronym "mra" stands for mixed risk-averse, "mrl" for mixed risk-loving, "amra" for anti-mixed risk-averse, and "amrl" for an anti-mixed risk-loving. Each cell contains two possibilities depending on the signs of the factor utility functions v and z.

	DD	**DU**	**UU**
dd-d	$v > 0$ amra & $z < 0$ mra $v < 0$ mra & $z > 0$ amra	$v > 0$ amra & $z < 0$ amrl $v < 0$ mra & $z > 0$ mrl	$v > 0$ mrl & $z < 0$ amrl $v < 0$ amrl & $z > 0$ mrl
dd-a	$v > 0$ mra & $z > 0$ mra $v < 0$ amra & $z < 0$ amra	$v > 0$ mra & $z > 0$ amrl $v < 0$ amra & $z < 0$ mrl	$v > 0$ amrl & $z > 0$ amrl $v < 0$ mrl & $z < 0$ mrl
da-d	$v > 0$ amra & $z < 0$ mrl $v < 0$ mra & $z > 0$ amrl	$v > 0$ amra & $z < 0$ amra $v < 0$ mra & $z > 0$ mra	$v > 0$ mrl & $z < 0$ amra $v < 0$ amrl & $z > 0$ mra
da-a	$v > 0$ mra & $z > 0$ mrl $v < 0$ amra & $z < 0$ amrl	$v > 0$ mra & $z > 0$ amra $v < 0$ amra & $z < 0$ mra	$v > 0$ amrl & $z > 0$ amra $v < 0$ mrl & $z < 0$ mra
ad-d	$v > 0$ amrl & $z < 0$ mra $v < 0$ mrl & $z > 0$ amra	$v > 0$ amrl & $z < 0$ amrl $v < 0$ mrl & $z > 0$ mrl	$v > 0$ mra & $z < 0$ amrl $v < 0$ amra & $z > 0$ mrl
ad-a	$v > 0$ mrl & $z > 0$ mra $v < 0$ amrl & $z < 0$ amra	$v > 0$ mrl & $z > 0$ amrl $v < 0$ amrl & $z < 0$ mrl	$v > 0$ amra & $z > 0$ amrl $v < 0$ mra & $z < 0$ mrl
aa-d	$v > 0$ amrl & $z < 0$ mrl $v < 0$ mrl & $z > 0$ amrl	$v > 0$ amrl & $z < 0$ amra $v < 0$ mrl & $z > 0$ mra	$v > 0$ mra & $z < 0$ amra $v < 0$ amra & $z > 0$ mra
aa-a	$v > 0$ mrl & $z > 0$ mrl $v < 0$ amrl & $z < 0$ amrl	$v > 0$ mrl & $z > 0$ amra $v < 0$ amrl & $z < 0$ mra	$v > 0$ amra & $z > 0$ amra $v < 0$ mra & $z < 0$ mra

Table 5: Construction of multiplicatively separable utility functions with desired apportionment preferences in the three cases. The acronyms "mra", "mrl", "amra" and "amrl" abbreviate mixed risk-averse, mixed risk-loving, anti-mixed risk-averse and anti-mixed risk-loving utility functions, respectively. Together with the sign of v and z, this establishes the univariate apportionment preferences, see Section 2.3. The apportionment preference across attributes follows from the alignment or misalignment of the signs of v and z, see Proposition 5.

Consolidation of the US property and casualty insurance industry: Is climate risk a causal factor for mergers and acquisitions?*

Georges Dionne
Akouété Fenou
Mohamed Mnasri

Canada Research Chair in Risk Management and Finance Department
HEC Montréal

February 28, 2023

Abstract

This report analyzes the difference between mergers and acquisitions (M&As) of target insurers in the US life and non-life insurance sectors. We first document M&A transactions in the US insurance market between 1990 and 2021 and select the M&A transactions related to US target insurers. We then study the evolution of the life and non-life insurance sectors over time in order to determine whether there are parallel trends between the evolution of M&As of target insurers in these two sectors over time. We empirically test the difference between the M&As of the life and non-life insurance sectors by employing a natural experiment method and verify whether climate risk has been a causal factor in the observed difference in mergers and acquisitions between the two sectors after 2012. Our results do not support a causal link between climate risk and M&As during the period of analysis. Insurers choose other diversification sources of capital, including reinsurance, premium management, CAT bonds, and better capital management under stronger risk regulation.

Keywords: Mergers and acquisition, US insurance industry, property and casualty insurance, life insurance, health insurance, climate risk, capital management, reinsurance, ILS, CAT bonds, premium management, risk regulation.

* *Research financed by the SCOR Foundation for Science and the Canadian SSHRC. Only the authors of the paper are responsible of its content. They thank Claire Boisvert, Denise Desjardins and Mohamed Jabir for their excellent contribution.*

1. Introduction

The main objective of this study is to measure a causality effect of climate risk on property and casualty (P&C) insurance industry consolidation. More generally, we examine how catastrophic events may have affected industry resilience by focusing on M&As in the US insurance industry.

The proponents of diversifying risk portfolios via M&A argue that acquisitions between different industries allow the acquiring insurer to benefit from economies of scope and scale through the joint use of customer databases, managerial expertise, and brand name. In addition, diversified transactions are expected to reduce acquirers' risk because this allows them to operate in a broader range of insurance lines and to better diversify extreme risks. By contrast, the proponents of focusing transactions within the same industry (or business line) argue that insurers are better off when they concentrate on their core business. It is not clear that such concentration is always beneficial in presence of climate risk.

In both cases transactions are also likely to be initiated by managers wishing to protect their human capital or increase their private benefits (Amihud and Lev, 1981; Jensen, 1986). Such behavior could be very risky for poorly diversified acquirers.

We have not found studies linking catastrophic risks to M&As in the insurance industry. Cummins and Weiss (2004), Cummins and Xie (2008) and Boubakri et al. (2008) analyze M&As in the insurance industry. They do not focus on catastrophic or climate risks, and their methodology is not up-to-date because they do not perform a causality analysis on the effect of different factors on M&As. One way to extend this literature is to investigate how climate risk events might be causal variables in explaining M&As. Difference-in-differences analysis is a methodology that can be applied by using insurers in activities less exposed to climate risk events as a control group and insurers in more climate-exposed activities as a treatment group. For example, insurers in the life insurance industry can be considered less exposed to climate risks than P&C insurers.

There are two major difficulties associated with isolating climate risk events as a causal effect on M&As during our period of analysis (1990 to 2021). The first is separating M&As from the varied alternative sources of capital consolidation that the insurers can use to protect themselves from natural catastrophes. Dionne and Desjardins (2022) show that US property and casualty insurers significantly increased their capital over recent years. They also identify various potential sources of capital, such as reinsurance, M&As, premium management, capital regulation, and insurance-linked securities (ILS).

The second difficulty is identifying factors other than climate risk events that may have affected M&As during the period of analysis. Notably, our period of analysis contains the 2007–2009 financial crisis. The US insurance industry was affected by this crisis, albeit less significantly than banks. Market conditions were difficult after the crisis, particularly for the life insurance industry. Premium growth was low, as were interest rates. Moreover, new federal regulations for capital were introduced, particularly in and after 2012. These new regulations have affected the level of capital and introduced some uncertainty in the markets regarding M&As.

Our results do not support a causal link between climate risk and M&As during the period of analysis. Insurers choose other diversification activities, including reinsurance, premium management, catastrophe bonds, and better capital management under stronger risk regulation.

The rest of the paper is organized as follows. Section 2 presents a literature review on M&As in the insurance industry. Section 3 describes the evolution of M&As in the US insurance industry from 1990 to 2001. Section 4 documents natural weather disasters during the same period. Section 5 analyzes the impact of markets conditions and regulation on M&A after 2012. Section 6 proposes an analysis of the parameters for a DID analysis, while Section 7 describes the DID analysis. Section 8 discusses the results. Section 9 concludes. A robustness analysis is presented in the Online appendices along with additional results and literature review.

2. Literature review

Usually, bidders initiate M&A transactions only when they anticipate that these activities will create value for their shareholders. Thus, studying the impact of such deals on bidders' performance is of particular interest, especially for intra-industry transactions, because these are most likely to be driven by synergies, and hence, create value. The empirical literature shows that acquiring insurers in the US insurance industry experience greater efficiency and higher profitability three years after the M&A (Cummins et al., 1999; Cummins and Xie, 2008; Boubakri et al. 2008).

Among insurers' economic rationales for these operations are a desire to increase their geographical reach and product range (Amel et al., 2004) and to benefit from economies of scale and scope (Cummins et al., 1999). Further, insurers may initiate these transactions to benefit from financial synergies (Chamberlain and Tennyson, 1998) or to reduce their riskiness and/or improve the amount/timing of their cash flow streams (Cummins and Weiss, 2004). Estrella's (2001) findings refute the risk-reduction argument from transactions between different industries. Indeed, the article shows that the median failure probability resulting from combinations of two property-casualty firms is lower than that resulting from a combination of a property-casualty firm and a bank holding company.

The financial literature also suggests that M&A transactions may destroy rather than create value, especially if these transactions are motivated by managerial hubris, that is, where managers are more interested in maximizing the size of their business empires than in returning cash to shareholders (Roll, 1986; Denis and McConnell, 2003). Hence a negative impact on the bidders' firm value could be observed. For such behavior to be constrained, effective governance mechanisms must be put in place, such as 1) a strong board with competent independent directors, and 2) a legal environment that offers strong protection to minority shareholders. The legal environment relates not only to investor protection but also to transparency and overall quality of accounting standards, which were all recently shown by Rossi and Volpin (2004) and Moeller and Schlingemann (2005) to be significant determinants of M&A (see also Boubakri et al., 2008). Asymmetric information between acquiring firms on particular targets can also affect M&A activities by modifying the premiums of different deals (Dionne et al., 2015; Betton et al., 2009; Brockman and Yan, 2009).

Akhigbe and Madura (2001) report a positive and significant abnormal return for acquiring insurers and conclude that this favorable valuation effect is driven by the similarity of services provided by both the acquirer and the acquired. In other words, standardization in their products makes the merger of operations easier for both parties. Interestingly, Akhigbe and Madura (2001) document a higher positive and significant market effect for acquirers that are non-life insurers. Floreani and Rigamonti (2001) also report a positive and significant valuation effect for the bidder, following M&A transactions involving pure insurance partners. This market valuation is positive but slightly lower when the target firm is publicly traded. However, only transactions involving insurers buying insurers seem to create value for the bidder. Indeed, Cummins and Weiss (2004) report a small negative valuation effect on the bidder’s shares following transactions that do not involve pure insurance partners.

Additionally, cross-border transactions may generate a higher positive valuation effect for the bidder because they are perceived to lead to a geographic expansion of its market. The results of Floreani and Rigamonti (2001) support this argument. Specifically, they demonstrate that transactions involving insurance partners that are both located in the European Union countries are not welcomed by the financial market. On the other hand, cross-border transactions may also destroy value for the bidder because they are more difficult to manage (Cummins and Weiss, 2004)—a result not supported by Floreani and Rigamonti (2001). In the Online appendix 1, we present a detailed analysis of various contributions on the insurance industry.

3. M&A transactions related to US target insurers from 1990 to 2021

From the SDC database, we identify 3,198 M&A transactions related to US target insurers from 1990 to 2021. Data are annual observations as of December 31 of each year.

Figure 1 identifies the two main waves of target insurer M&As recorded in the US insurance industry over the past 32 years. There was strong M&A growth until the years 1997 to 1999, when the market reached its first peak since 1990.

Figure 1: Histogram of the annual number of M&A transactions related to US target insurers from 1990 to 2021

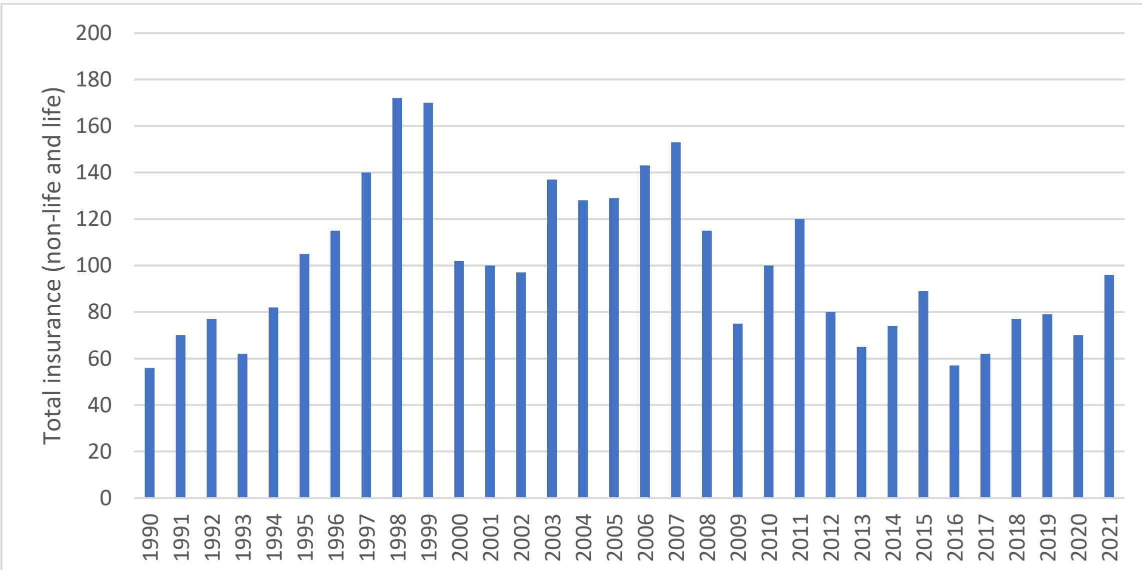

Data source: SDC database.

After a sharp decline in 2000, the M&A market resumed growth in 2003, and reached its second peak in 2007. Each of these wave years has more than 120 annual transactions. The two peaks correspond to periods of economic expansion. The wave recorded around 1997-1999 represents the largest of the US insurance industry during the period of analysis. The record years of 1998 and 1999 have not been broken since then. In fact, this period corresponds to the internet and new technologies growth of the years 1998-2000. The years of the second largest wave of M&As correspond to the economic expansion period before the financial crisis that began in August 2007.

Figure 2 depicts three peaks of M&As across all industries in the US (1998, 2007, and 2017) during the same period. As documented above, only two waves of M&As occurred in the US insurance industry during that period. Since the 2007 peak, the M&A market has exhibited an overall downward trend throughout the US insurance industry (life and non-life combined). By comparison, the all-industry M&A market resumed its overall upward trend after a short decline during the financial crisis, from 2007 to 2009, and reached a new peak in 2017. Figure 2 suggests that the post-2007 period is marked by a shift behavior of insurers across the US insurance industry, which may be explained by changes in industry regulation after the 2007-2009 financial crisis, market conditions, and climate risk.

Figure 2: M&A trends in the US insurance industry (total M&A for non-life and life targets, left) and for all industries in the US (right), 1990 to 2021

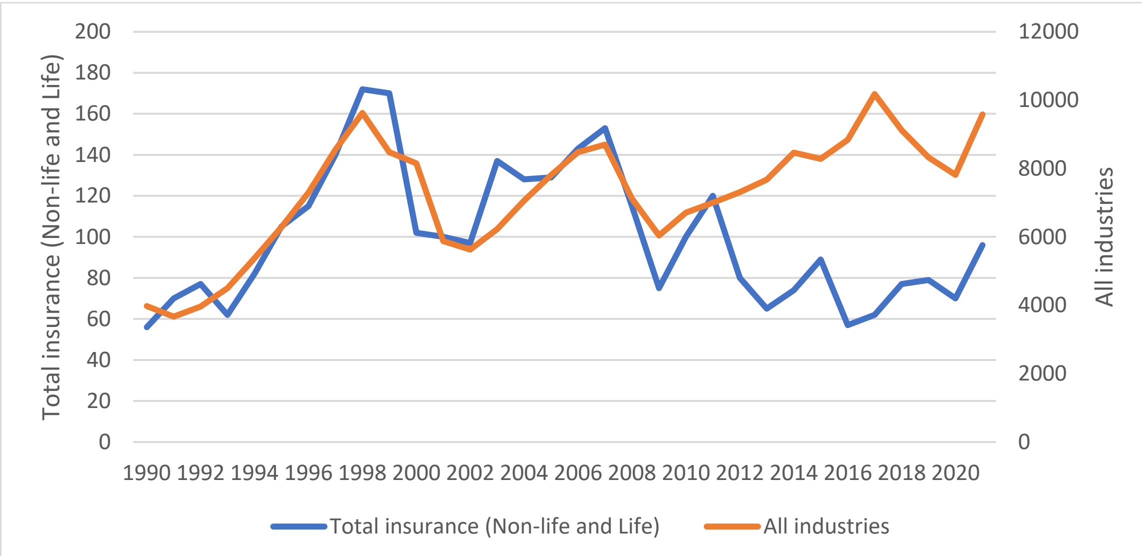

Data source: SDC database.

Figure 3 presents the evolution of the numbers of M&As in the three insurance lines and Table 1 summarizes their main statistics. Property and casualty insurers and health insurers appear to be more similar than with life insurers. We also observe the large reduction in M&As in the life sector after 2011. In this paper, we consider that the US insurance industry consists of two main lines of business: life insurance, and non-life insurance that includes property and casualty insurance and health insurance.[1] Given that the two main lines of insurance can be affected differently by climate risk, market conditions, and insurance regulation, we have plotted the M&A transactions recorded in each of these two lines in order to analyze their behavior in relation to the target insurer M&A phenomenon. Figure 4 shows the evolution of M&As in each of the two main US insurance lines and that of the US insurance industry as a whole over the period of 1990 to 2021.

[1] We perform a robustness analysis in Online appendix 2 by merging health insurers with life insurers.

Figure 3: MA trends of target insurers by the three insurance sectors
in the US, 1990 to 2201

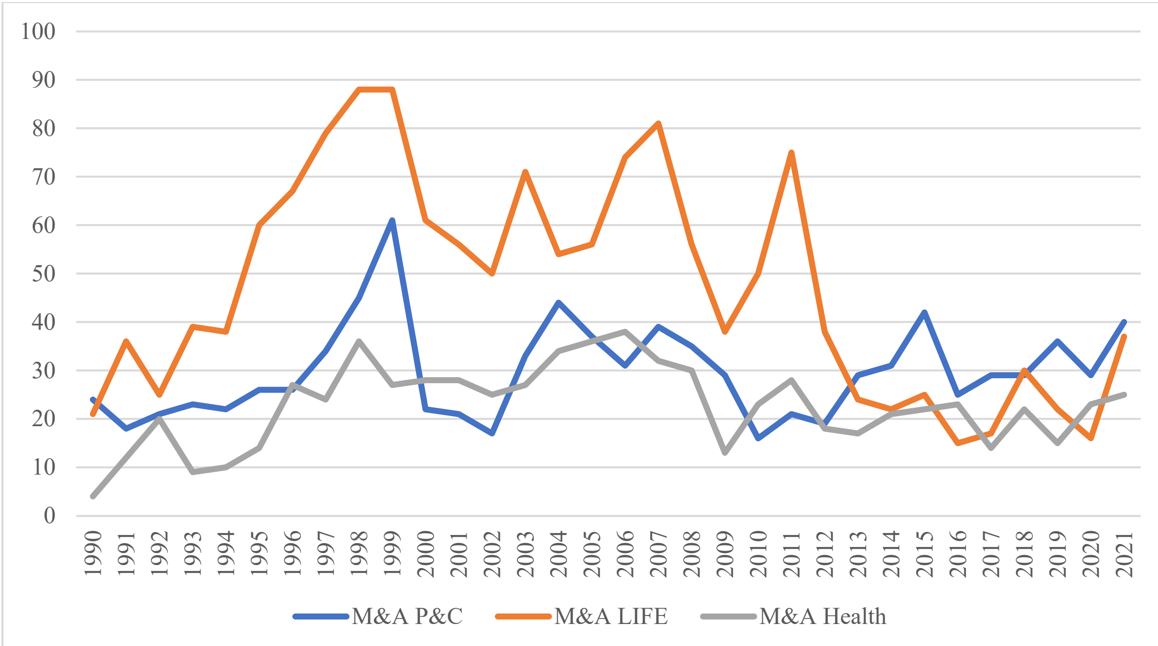

Data source: SDC database.

Table 1: Mean and standard deviation of the M&A in each sector

Period	1990-2021		1990-2012		Post-2012	
Annual number of MA	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
P&C sector	29.813	9.822	28.870	10.981	32.222	5.761
Life sector	47.156	22.598	56.565	19.294	23.111	7.079
Health sector	22.656	8.407	23.609	9.524	20.222	3.898

We observe, in Figure 4, that the evolution of M&As of target insurers in the life insurance sector seems to mirror the evolution of M&As of target insurers observed in the entire US insurance industry. More importantly, we confirm the strong decrease in mergers and acquisitions in the life insurance industry after 2012 while this activity seems more stable in the non-life insurance sector during the same period.

Figure 4: M&A trends for target insurers by the two major insurance lines
(life or non-life, left) and the overall US insurance industry (right), 1990 to 2021

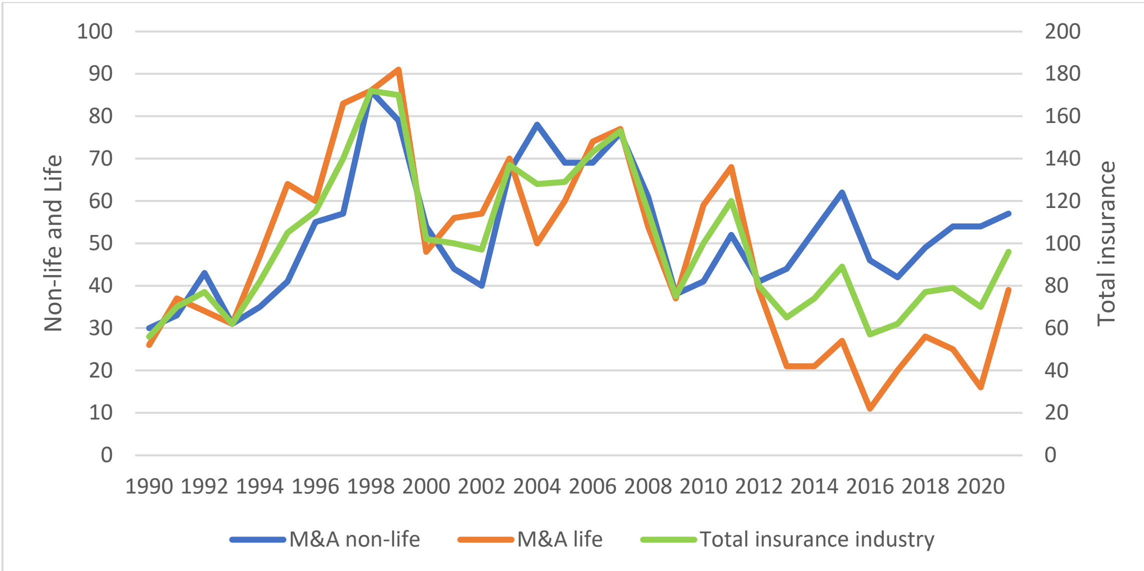

Data source: SDC database.

Figure 5: M&A trends of target insurers by the two main insurance sectors
(life and non-life) in the US, 1990 to 2021

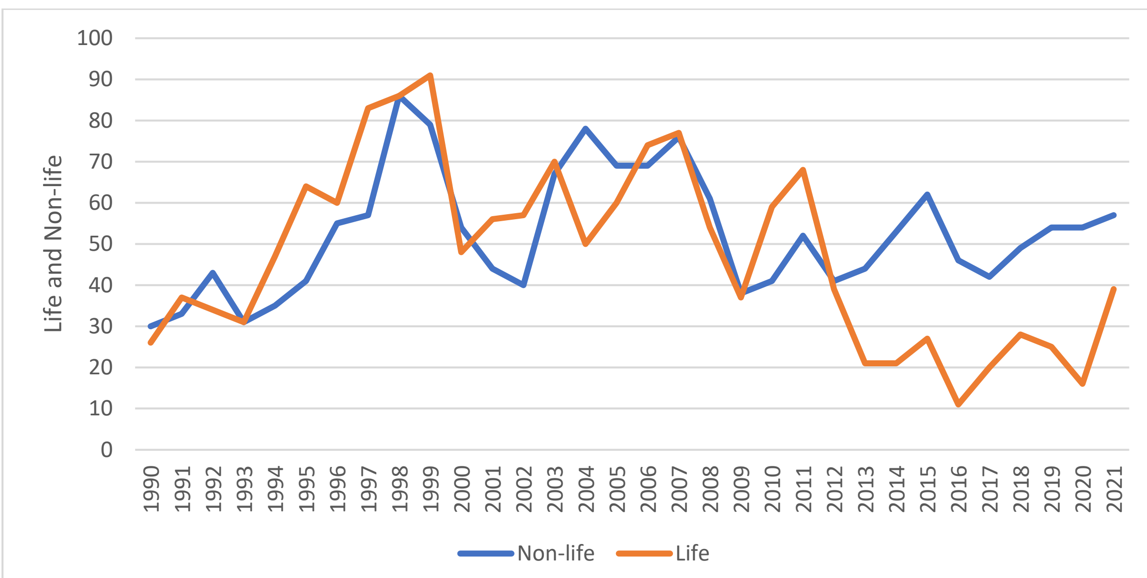

Data source: SDC database.

Figure 5 shows a parallel time trend in the evolution of target insurer M&As for life and non-life insurance from 1990 until 2009 and even 2012. This result suggests that the evolution of target insurer M&As in the non-life insurance sector is almost identical to that observed in the life insurance sector during this period. The parallel trends observed between the two groups started to disappear after 2009. The difference is more pronounced after 2012. Based on Figure 5, we retain the years 2009 and 2012 as potential candidates for the treatment date in our analysis with the difference-in-differences (DID) method. The choice of the treatment date for our DID method thus seems ambiguous. We will use a statistical approach, applied to time series, to validate the year that best suits our data.

It is worth trying to understand the divergence in the temporal trends in M&As observed between our two groups. It is possible that the temporal trends in M&As observed between our two groups cease being parallel in 2009 or 2012 owing to series of natural disaster events in the US or to the relative change in the regulation and market conditions of the two industries after the 2007-2009 financial crisis. To analyze these possible causes, we will first describe the evolution of the number and the severity of natural disaster events occurring in the US from 1990 to 2021.

4. Analysis of the evolution of natural weather disasters events from 1990 to 2021

4.1. General statistics

The year 2011 will remain etched in the memory of insurers and reinsurers. It generated losses of exceptional magnitude, particularly in Japan, Thailand, New Zealand, Australia and the US. In other words, 2011 was a year of huge losses both globally and nationally (speaking of the US).

Globally, the last few decades have seen an increase in extreme weather-related events that have fueled the rise in the number of claims paid by insurers. Figure 6 shows three major peaks in the insured losses paid by insurers worldwide. The first largest peak in claims costs was in 2017. The year 2011 represents the second largest peak in the cost of claims borne by insurers worldwide. The year 2005 represents the third highest peak in insured losses. Looking only at the period prior to 2017, 2011 is the worst year for claims over the period of 1990 to 2017.

Figure 6: Insured losses (in million $) from natural disaster events worldwide, 1990 to 2020

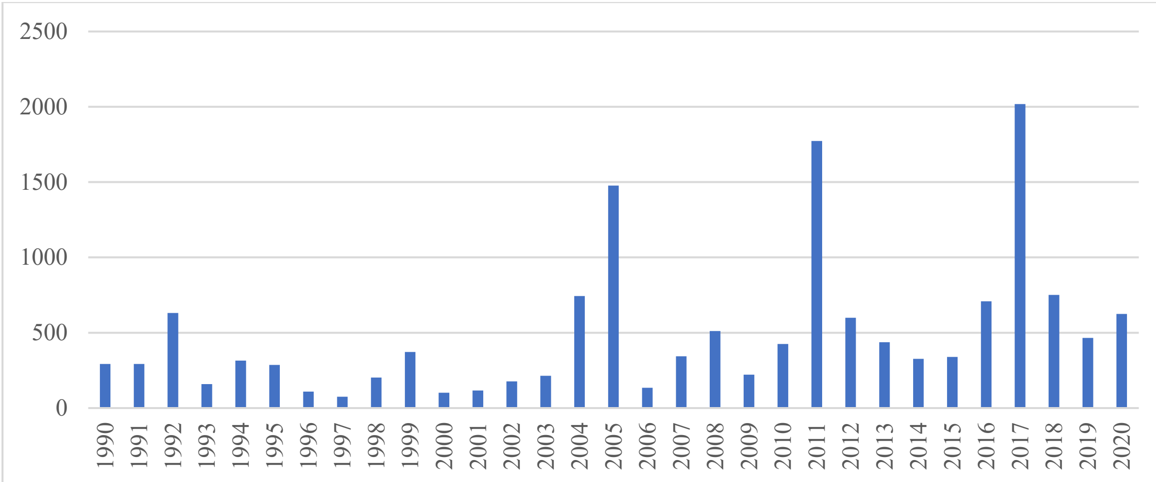

Data source: Our World in Data. Insured losses: property damage and business interruption, excluding liability and life damage.

Figure 7 indicates that 2011 represents the third deadliest year due to natural disasters in the US. This 2011 record can be linked to the exceptional series of severe tornadoes that occurred that year in the Midwestern US. The most catastrophic year was 2005, the year Katrina struck. Figure 8 shows that 2011 is the year with the first highest number of injuries and deaths from natural disasters after 1998, the year of Hurricane Georges. Finally, the figure indicates a decrease in total casualties after 2011. Bear in mind that when a single natural catastrophe event affects a large number of policyholders, it increases claims costs on the one hand and management expenses (operating costs) on the other, putting upward pressure on the combined ratio and other financial ratios of insurers.

Figure 7: Numbers of injuries (left) and deaths (right)
from natural disasters observed in the US, 1990 to 2021

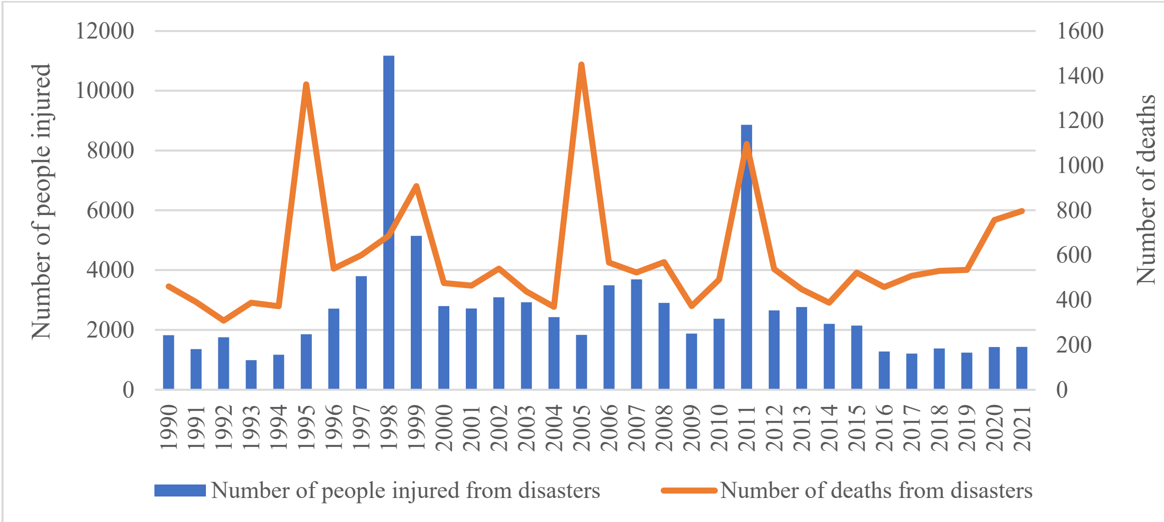

Data source: NOAA Weather Related Fatality and Injury Statistics. People injured or killed by natural disasters are not necessarily insured.

Figure 8: Total casualties (injuries and deaths)
from natural disasters in the US, 1990 to 2021

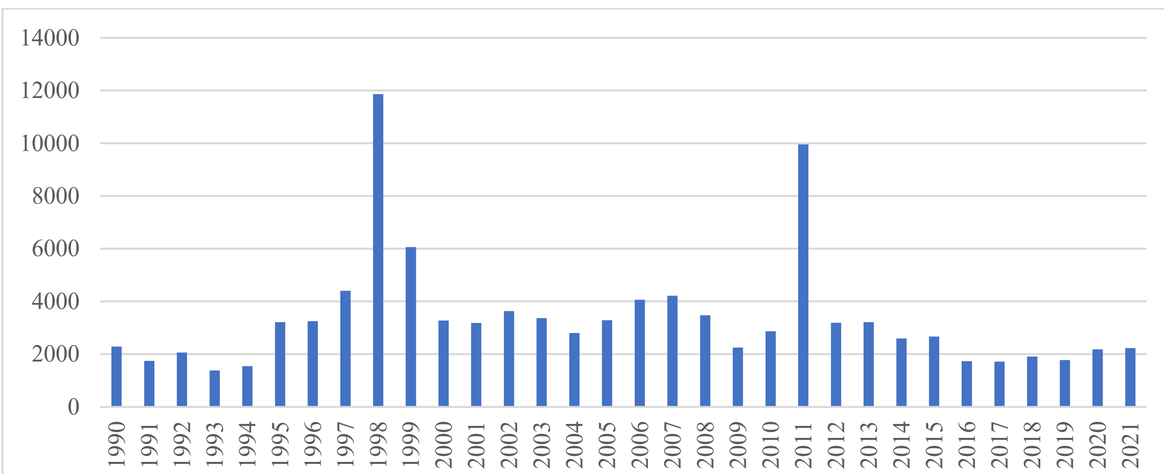

Data source: NOAA Weather Related Fatality and Injury Statistics. People injured or killed by natural disasters are not necessarily insured.

4.2. Our data

We now present the definition of the three main variables used in the following analysis of the US insurance industry. The data for the first two measures of weather disasters are from the VERISK database. Our first variable is the annual number of natural weather disaster events that cause insured losses to the insurance industry of $25 million or more which is the VERISK threshold to document a catastrophe. Events that meet or exceed this threshold are considered natural disasters, given the magnitude of the loss costs incurred by insurers. Our second variable measures the total annual insured losses from natural weather disaster events that cause losses of $25 million or more to the insurance industry. Finally, our third variable measures the number of natural disaster casualties. It represents the sum of the annual number of deaths and injuries caused by natural disaster events. The data for the number of natural disaster casualties were obtained from the National Oceanic and Atmospheric Administration (NOAA) website. Figure 9 shows the evolution of the number of natural weather disaster events occurring in the US from 1990 to 2021, as reported by VERISK. They cover hurricane, tropical storm, wildland fire, wind and thunderstorm, and winter storm.

Figure 9: Number of natural disaster events in the US, 1990 to 2021

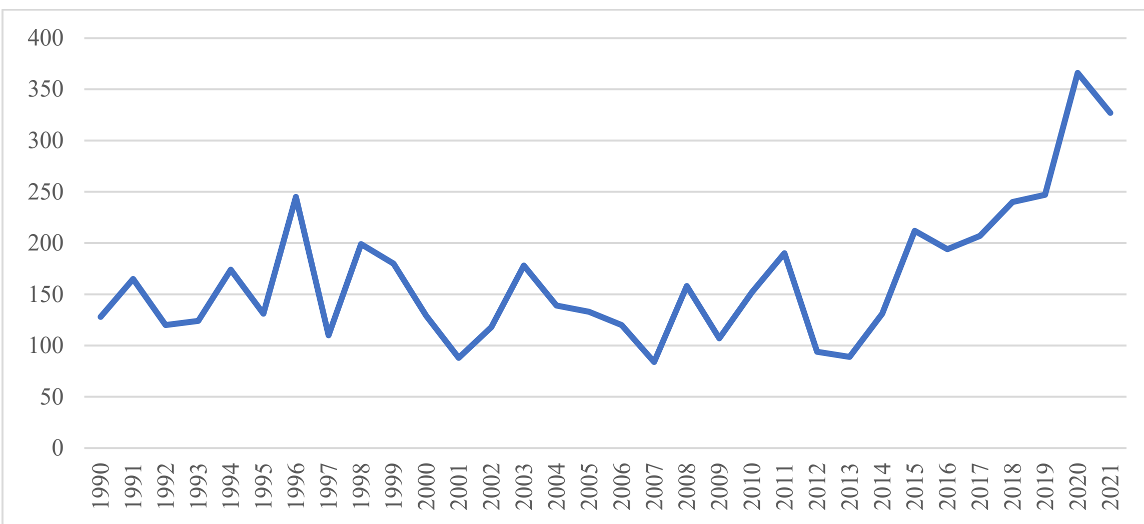

Data source: VERISK database.

Note: An "ISOnet PCS Loss Event" means an event occurring within the Service Area to which ISO assigns a serial number, based on ISO's judgment that the event is likely to cause $25,000,000 or more in total insured property losses within such Service Area and is likely to affect a significant number of property and casualty insurance policy holders and property and casualty insurance companies.

Figure 9 shows that there have been significant variations in the number of weather disaster events in recent years with an upward trend in the post-2013 period. The year 2013 is the turning starting point for this increase in the numbers. The increase in disaster weather events observed after 2013 could be attributed to variation in climate change.[2] This phenomenon may have posed a real threat to the American insurance market because of some extreme natural disaster events it has caused in the US. As can be seen in Figure 9, the number of natural disaster events has reached extremes over the last five years (2017 to 2021). Arguably, the insurance industry can be weakened by the increase in extreme natural disaster events because of the high claims costs they incur, particularly after 2017.

Our data indicates an average number of 241 natural disaster events per year during the post-2013 period, compared with 140 from 1990 to the end of 2013.[3] This analysis was limited to the number of events. It may be more appropriate to consider the losses in the insurance industry. Figure 10 relates annual numbers of natural disasters events and annual insured losses. See Appendix A3 for different correlation results. These results do not support any causality link.

Figure 10: Number of natural disaster events (left) and insured losses (right) linked to these natural disaster events observed in the US, from 1990 to 2020

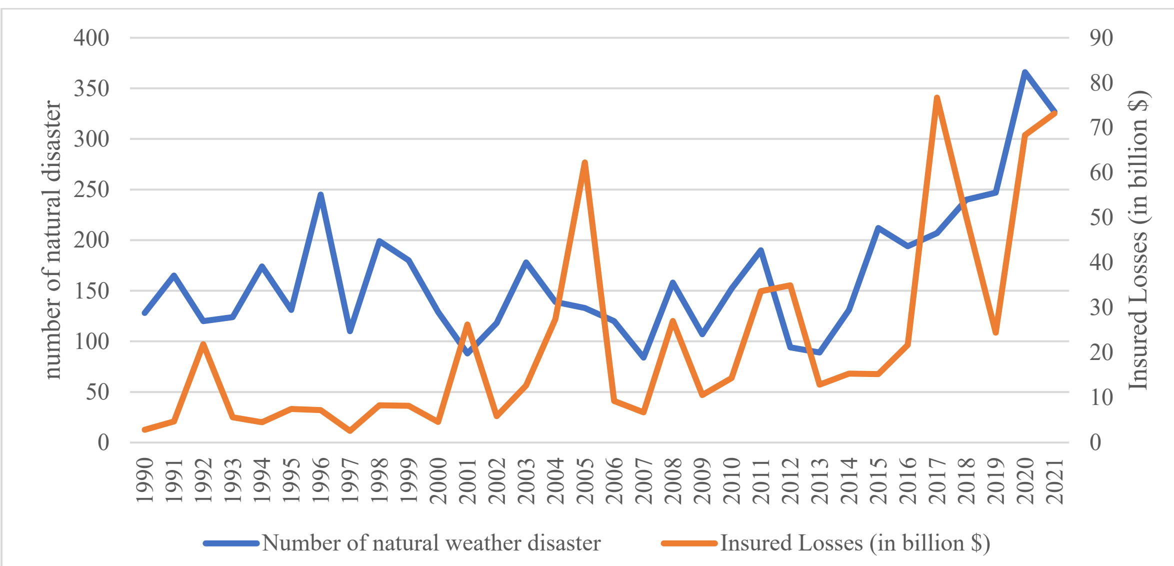

Data source: VERISK database.

[2] Many references consider weather and climate risks to be synonymous. In this study, as in Dionne and Desjardins (2022), we use the NASA (2005) definitions of climate and weather. The main difference between the two definitions is time. Weather is atmospheric conditions over a short period of time, while climate covers a long period of time. Climate change is related to changes in average daily weather.

[3] The corresponding numbers for the period post-2012 and before are respectively 233 and 142.

4.3. Comparative analysis of the evolution of M&As and insured natural disaster losses

Figure 11 shows a link between insured losses from natural wealth disasters and the number of M&As per year in the non-life insurance sector. This link seems to confirm graphically the hypothesis that the number of target insurer M&As is an increasing function of the insured losses from natural disasters variable, particularly after 2012.

Given that the post-2012 period marked by the resurgence of natural disaster events coincides with the period of the loss of parallel trends observed between our two groups identified graphically (see Figure 5), we can assume that the upsurge in natural disaster weather events observed after the year 2012 may have caused the difference in the number of M&As of target insurers in the non-life insurance sector compared with the number of M&As of target insurers in the life insurance sector observed after 2012. We will consequently select target insurers in the non-life insurance sector as organizations affected by the increase in natural disaster events observed during the post-2012 period, as our potential treatment group for our DID analysis between the M&As of target insurers in the life and non-life insurance sectors in the US.

Figure 11: Comparison of M&A trends in the non-life insurance sector (left) and observed insured losses from weather events (right) during the period 1990 to 2021

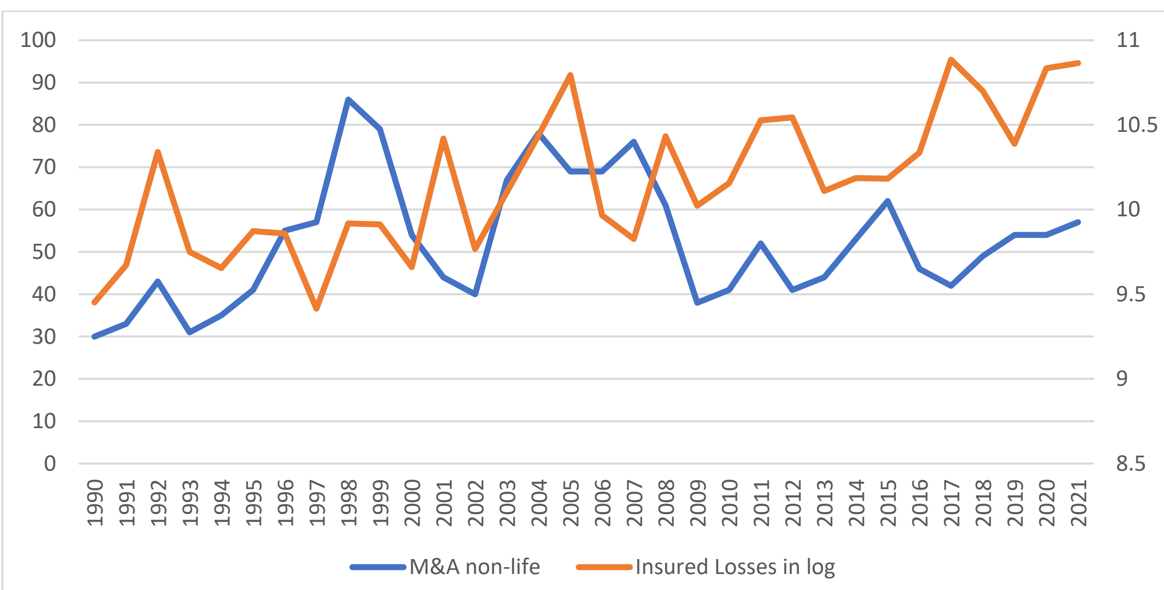

Data sources: SDC database and VERISK database.

According to a study published by Atlas Magazine, the emergence of new hazard detection technologies and the generalization of anti-seismic construction standards, especially in developed countries, have significantly limited the number of natural disaster casualties in the world. This information seems relevant to explain the relatively stable level of casualties observed after the year 2012 (Figure 8) despite the upsurge in extreme natural events compared with the period of 1990 to 2012.

The capacity of new hazard detection technologies to warn residents of potential extreme natural events enables these individuals to leave their areas of residence when natural disasters occur, which limits the number of deaths and injuries. However, even if residents are warned about the possibility of an extreme natural disaster, they cannot take real estate such as houses and buildings with them when they evacuate the area. In other words, insured losses are still potentially present in the non-life insurance sector despite the advent of new hazard detection technologies. The direct consequence of this would be an increase in insured losses associated with extreme natural disasters, which would increase the claims costs paid by non-life insurers, thereby worsening their financial performance and potentially increasing the number of M&As.

We have shown above that the upsurge in natural disaster events observed after 2012 has led to increased growth in insured losses from natural disasters for non-life insurers (Figure 10). We have also shown that the number of natural disaster casualties remains relatively stable despite the upsurge in extreme natural events observed in the post-2012 period (Figure 8).

As to which event may have produced an exogenous change in treatment that further increased the number of M&As for target insurers in the non-life insurance sector relative to the life insurance sector, our analysis indicates that the upsurge in natural disaster events observed in the post-2012 period may represent a causal shock on M&As in the non-life sector.

After having motivated our first theoretical hypothesis graphically and statistically, we will analyze a second potential causal factor explaining the difference in M&As between life and non-life sectors after 2012.

5. Impact of market conditions and regulation on M&As after 2012[4]

5.1. Markets conditions and regulation

In the preceding sections, we emphasized climate risk as motivating the difference between the life and non-life insurance industries in the evolution of M&A after 2012. In this section, we document potential alternative economic explanations of this difference before proceeding to the formal DID analysis.

Another catastrophe in the US economy in recent years was the 2007–2009 financial crisis. Although this crisis affected banks more significantly, it also disrupted the insurance industry. It took many years for the US insurance industry to recover. Moreover, the insurance industry was subject to new federal regulations in the years following the crisis. In these years, economic growth was slow due to a lack of liquidity in the US economy, partly explained by the strong new banking regulation. In particular, the secondary market for bond trading was out of liquidity. Interest rates were very low for investments, and the European economy was in distress. These facts seem to have affected the life insurance industry more strongly than the P&C insurance industry.

The year 2012 was an active one for life insurance M&As, with 39 transactions, as shown in Figure 5. The aggregate deal value involving US targets for the year was about $4.2 billion, which is higher than the $775 million in 2011, but significantly less than the $21.6 billion reported in 2010 (59).[5] This can be explained by AIG's activity of selling firms following the financial crisis (Mayer Brown, 2013). This decrease was mainly due to the need for acquirers to maintain capital under new regulatory capital requirements and to the uncertainty around the impact of Solvency II in Europe.

Acquisition activity in the property-casualty sector was significantly lower in 2012 than in 2011. The announced aggregate US deal value for 2012 (39) was approximately $6 billion,

[4] This section is based on many reports from industry, including the annual reports of Mayer Brown and documents from KPMG. The SDC database is also used to document the annual numbers of mergers and acquisitions.

[5] Numbers in parentheses are observations on the number of mergers and acquisitions, as illustrated in Figure 5.

down from approximately $10 billion in 2011 (68). Moreover, 2012 was characterized by small and medium-sized deals under $500 million (Mayer Brown, 2013). P&C activity was driven primarily by geographic or product expansions, as well as by runoff transactions involving insurers deciding to exit some lines of business.

The year 2013 was characterized by the continued decline in deal activity in the US life insurance M&A market (transactions involving US targets), as compared to 2010, in terms of deal values and numbers (21 instead of 59). Deal value in the life sector was $3.2 billion, compared to $4.2 billion in 2012. Continued macroeconomic uncertainty presented challenges for product sales in this industry, and low interest rates continued to create challenges for long-term investment returns in bonds. Regulatory changes, such as the NAIC's Own Risk Solvency Assessment (ORSA, adopted in 2012, effective in January 2015) and the international accounting convergence project contributed to a climate of caution among buyers and sellers in the M&A markets. To increase shareholder value, insurers tended to use excess capital for share repurchases and dividend distributions rather than M&A activity. ORSA represented a major regulatory change in the insurance industry. Insurers must now use market value information instead of accounting values to compute economic capital. It represented an additional source of uncertainty, because many insurers had to learn about capital computation with market information.

Acquisition activity in the P&C sector was stable in 2013 compared to 2012, despite generally favorable market valuations on companies' balance sheets in a year marked by few large catastrophe losses. Major runoff acquisition specialists continued to be active acquirers in the global P&C sector. Many P&C companies were still overcapitalized. Some companies were returning capital in the form of stock buybacks and dividends, but high stock prices made stock buybacks expensive.

At the NAIC's Summer 2013 National Meeting, the Solvency Modernization Initiative (SMI) Task Force adopted a white paper: the US National State-band System of Insurance Financial Regulation and the Solvency Modernization Initiative (NAIC, 2013). The white paper also highlighted the importance of the national state-based system of insurance regulation, instead of state only regulation as before the financial crisis.

In addition, regulatory scrutiny of M&As in the two areas may have had a slight negative effect on capital management, thus limiting M&As: the restrictive use of captives for reserve financing and additional requirements for approval of acquisitions raised difficulties in making acquisitions (Mayer Brown, 2014).

Acquisition activity in the P&C sector was lower in 2013 than in 2012, continuing the trend from the prior year (21 instead of 39). This occurred despite generally favorable market valuations and significant cash balances on P&C companies' balance sheets in a year marked by few large catastrophe losses. Since catastrophe losses had been relatively modest, many P&C companies remained overcapitalized. M&A was not considered an important activity for consolidation during these years.

The number of US life insurance M&A deals in 2014 was down for the third straight year, but overall, the deal value on announced transactions was $8 billion in 2014, more than double the total for 2013 (Mayer Brown, 2015). There were 53 announced M&A deals involving property and casualty companies (Figure 5). The year was again characterized by small- and medium-sized deals.

Insured losses from natural catastrophes fell significantly in 2014, according to research from Swiss Re's Sigma (2015), as reported in Mayer Brown (2015). The global insured losses for 2014 fell by 24% to $34 billion, compared to $45 billion the previous year. The number of life insurance M&A transactions involving US targets was on the rise in 2015 after falling in each of the previous two years. The number of annual P&C insurance M&A transactions in 2015 was up for the third straight year, increasing from 44 to 62. The overall deal value on announced transactions was also up, from approximately $12 billion in 2014 to $48 billion in 2015. The year 2015 saw a number of very large transactions being announced, as buyers increasingly sought scale, diversification, and market access (Mayer Brown, 2016).

The number and size of life insurance M&A deals was very low in 2016 (only 11), compared to 2015 (27). The slowdown in activity was due to a number of obstacles facing the US life industry, including low life insurance policy sales, continued profit pressure in investments arising from the low interest-rate environment, and regulatory-change uncertainty.

The number of M&A transactions involving P&C insurance targets decreased in 2016 to 45, as compared to 62 in 2015, according to data compiled from the SDC database. The 2016 P&C insurance segment was again characterized by small and medium-sized transactions, with more than 75% of all announced deals valued below $200 million. The growing need for capital expenditure for investments, to support new digital and high-tech business models demanded that smaller and mid-sized companies look to M&As as an option for continued growth. Insurers worked to adapt to technological growth. For example. developments in insurtech continued to be important in 2016, with significant deals and expansion across product lines and markets. Moreover, in 2016, regulators took significant steps to enhance the regulation of insurers' data practices. Cybersecurity became a new priority for regulators (Mayer Brown, 2017).

In January 2017, the US and Europe announced an agreement regarding international insurance groups doing business in the US and the EU, to enhance regulatory certainty for insurers and reinsurers operating in both places. Meanwhile, the number of M&A transactions involving P&C insurance targets continued to decrease in 2017, to about 42, as compared to 46 in 2016 (SDC database). Overall, the deal value on transactions in 2017 was down to $7.5 billion, compared to $12 billion in 2016 (Mayer Brown, 2017).

With excess capital, more insurers saw themselves as buyers rather than sellers, which pushed the valuation levels of target companies upwards. Insurers in the P&C market appeared more likely to allocate their excess capital to investments in technology and marketing. Consequently, instead of buying competitors, insurers were more likely to make acquisitions of insurtech enterprises to improve their diversification.

The number and size of life insurance M&A deals involving US targets were up in 2017 (20), compared to 2016 (11). According to the SDC database, 2017 saw several large deals take place. The continued low-interest-rate environment, combined with the significant amount of capital available for deployment into the life and annuity sector led to a number of large annuity transactions in 2017. The year 2017 was notable for the occurrence of a number of catastrophic events, including hurricanes Harvey and Irma and wildfires in California all of which caused losses for several outstanding catastrophe bonds. The availability of this financial market protection in a year with significant catastrophe losses illustrates the robust nature of the insurance market and its critical importance in providing the resources needed to pay claims (Dionne and Desjardins, 2022).

The number of M&A transactions in 2018 involving P&C insurance targets rose to 60, compared to 42 in 2017, according to data compiled by the SDC database. The $32 billion in aggregate transaction value ranks as the most active year for P&C M&As since 2015. It should be noted that approximately two-thirds of that amount is attributable to two very large acquisitions. As in the previous years, small and medium-sized transactions of deals valued below $500 million represented more than 70% of transactions (Mayer Brown, 2019).

Despite around $80 billion of catastrophe losses in 2018, which followed on record catastrophe losses in 2017, the P&C industry continued to be regarded as overcapitalized. Other key factors limiting the increase in P&C M&As included federal tax reform and continued inbound interest from international acquirers seeking a meaningful presence in the US market (Mayer Brown, 2019). Established players were pursuing strategic investments in insurtech businesses.

Issuance of RWI policies continued to be important in the Americas, predominantly in the US. RWI is a form of insurance policy that is purchased in connection with an M&A transaction that protects the insured party (almost always the buyer) against financial loss arising from an unanticipated or unknown breach of certain conditions in the purchase agreement. While there are no market studies that provide reliable figures on the numbers of RWI policies written each year, data from several market studies suggest that numbers have doubled every two years since 2013. The year 2018 also saw the first transfer of pure wildfire risk to the capital markets. Two California utility providers sponsored a catastrophe bond covering third-party liability losses due to wildfires caused by their respective infrastructure. Demand for reinsurance remained high following the ongoing capital requirements of the Solvency II regime, which made reinsurance attractive.

One of the consequences of the 2007–2009 financial crisis was a decision by the federal government to revisit the regulatory system in the McCarran-Ferguson Act. The Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank) gave increased systemic risk regulatory authority to the Federal Reserve. In addition, Dodd-Frank also created a Federal Insurance Office within the Department of the Treasury to establish greater uniformity among the states with regard to excess and surplus insurance and reinsurance lines.

The development of the COVID-19 pandemic in the first quarter of 2020 created uncertainty regarding all aspects of the insurance business. This resulted in a halt in insurance P&C transactions in the US, as insurers and investors reevaluated their strategic plans. Despite of this first quarter slowdown, an increase in industry M&As from the third quarter of 2020 resulted in deal-making in 2020 whose value exceeded that of 2019 (Mayer Brown, 2021).

The year 2020 has been described as the Year of the SPAC.[6] According to SPAC Insider, 248 special purpose acquisition corporations (SPACs) completed their initial public offerings (IPOs), raising over $83 billion. The recent rise of the SPAC has had an important effect on the US IPO market and, to a lesser extent, the US IPO market for insurance companies. In 2020, three SPACs completed IPOs, with a stated focus on the insurance (including insurtech) industry.

During 2020, US jurisdictions began revising their laws and regulations governing credit for reinsurance to implement the amendments to the NAIC Credit for Reinsurance Model Law and Model Regulation adopted in 2019. Those amendments were designed to satisfy the requirements of the bilateral agreement on insurance and reinsurance between the US and EU.

Climate risk and sustainability were established as a key theme of the IAIS (International Association of Insurance Supervisors) strategy for 2020–2024. Included in this strategy is its partnership with the United Nations Environmental Programme's Sustainable Insurance

[6] A SPAC is a newly formed company with no assets or operations, also known as a blank check company.

Forum. The IAIS is one of the first global standard-setting bodies to adopt policy to guide its performance in terms of environmental issues: incorporating risks from climate change into their governance frameworks, risk management processes, and business strategies.

5.2. Use of ILS for catastrophes losses

The use by insurers and reinsurers of insurance-linked securities (ILS) as a supplemental source of capital for their protection continued after 2012. The capital markets have become a critical component of managing catastrophe risk for a growing number of insurers and reinsurers, although the relative magnitude is still low compared to the total capital available in the industry (Dionne and Desjardins, 2022).

The catastrophe bond market was quite strong in 2013, with a total of $7.5 billion of new catastrophe bonds issued, the second highest annual issuance volume in market history. As of December 31, 2013, there was $20.3 billion of catastrophe bonds outstanding. US catastrophe risks (particularly US wind) continued to dominate, representing approximately 51% of outstanding bonds (Mayer Brown, 2014).

In 2017, the ILS market solidified its importance as a critical component of the global reinsurance market, representing almost 20% of dedicated reinsurance capacity. There was a $31.0 billion total aggregate principal amount of risk-linked securities outstanding, almost 20% higher than the amount at the end of 2016 (Mayer Brown, 2018).

In 2020, the volume issued was the largest in market history, beating the record level of 2018. The total aggregate principal amount of risk-linked securities outstanding of $46.4 billion represented a yearly growth of approximately $5.7 billion. It should be mentioned that the total capital of the US insurance industry was about $1.1 trillion in 2020 (Dionne and Desjardins, 2022).

Reinsurance and premium growth are other sources of capital in the P&C insurance industry (Dionne and Desjardins, 2022). We shall look at these sources of capital later on. In the next section, we continue our statistical analysis of M&As.

6. Validation of the selected treatment date and the presence of parallel trends

In our DID approach, we propose that the increase in natural disaster events observed in the post-2012 period could be a cause of the difference in the number of M&As of target insurers in the non-life insurance sector, relative to the number of M&As of target insurers in the life insurance sector. The varied changes in regulations and economic conditions in the insurance industry during the post-2012 period could also be a cause. These new regulations were motivated by the 2007–2009 financial crisis. Very low interest rates significantly affected the benefits of the insurance industry, particularly in the life insurance industry. Looking at these two potential causes, it appears that a shock event occurred in the years preceding 2013 that might have caused an exogenous change in the treated units that increased the difference in the number of M&As of the treatment group relative to the control group. In short, we consider the increase in natural disaster losses observed after 2012 as a situation that induced an exogenous variation in the treated units (target non-life insurers) that maintained the number of M&As of target insurers in the non-life insurance sector (treatment group), compared to those in the life insurance sector (control group), which decreased significantly during the post-2012 period.

Based on an analysis of Figure 5, we have identified two years in which the parallel trends observed between our two groups began to disappear: 2009 and 2012. However, our analysis of Figure 10 allows us to propose that it was the insured losses from natural disaster events observed after the year 2012 that likely caused the increase in the number of M&As of target insurers in the non-life insurance sector, compared to the number of M&As of target insurers in the life insurance sector, observed in the post-2012 period. Therefore, we can define our treatment effect as a positive difference between the average number of M&As per year of target insurers in the non-life insurance sector and the average number of M&As of target insurers in the life insurance sector. Alternatively, market conditions and variations in the regulation of the insurance industry may also explain the difference observed in Figure 5. The following analysis is independent of the two potential causes.

6.1. Validation of the choice of treatment date using five statistical tests

To choose the most appropriate treatment date for our data, we use a statistical approach applied to the annual data of M&As in the two insurance sectors (Berck and Villas-Boas, 2016; Imbens and Wooldridge, 2009; Roberts and Whited, 2012). We first calculate the annual difference between the number of M&As of target insurers in the non-life insurance sector versus the number of M&As of target insurers in the life insurance sector observed over our entire study period, that is 1990 to 2021. Next, we calculate the mean and median of the difference between the number of target insurer M&As in the non-life insurance sector and the number of target insurer M&As in the insurance sector over the pre-treatment period (including the year of the candidate date) and over the post-treatment period for each of our two selected candidate dates (2009 and 2012). Finally, we perform five statistical tests—the mean statistical test, the median statistical test, the distribution statistical test, the monotonicity test, and the median-criteria test—to validate the choice of treatment date. The results of the first three tests are presented in Table 2, where the differences between various statistics are presented.

Table 2: Statistical descriptions (median, mean of the number of M&As) and validation tests of the treatment date

Period	1990-2009	Post-2009	1990-2012	Post-2012	1990-2021
Median	–2	22.50	–3	29	2
Mean	–2.75	18.4167	–3.78	28.11	5.18
Student's test	–0.9864	3.3066	–1.4679	11.015	1.6014
Median test[1]	0.8238	0.0386	0.6776	0.0039	0.3771
Wilcoxon test[2]	–0.915	2.589	–1.354	2.666	1.356

[1] Sign test (Snecdecor and Cochran, 1989).
[2] Signed rank test (Wilcoxon, 1945).

6.1.1. Statistical test based on the mean (Student's test)

Our decision criterion for the choice of treatment date is to test the null hypothesis (H0) that the average number of M&As in the non-life sector and the average number of M&As in the life sector are statistically similar over the period of 1990 to the end of the candidate date (2009 or 2012) on the one hand, and, on the other hand, to test the null hypothesis

(H0) that the average number of M&As in the non-life sector and the average number of M&As in the life sector are statistically different over the post-treatment date period (post-2009 or post-2012) due to the treatment effect.

According to Table 2, the *t*-test statistic yields a value of -0.9864 over the period of 1990 to 2009 and 3.3066 over the post-2009 period. Given that the absolute *t*-test value is less than 1.96 over the period of 1990 to 2009, the null hypothesis (H0) is not rejected. In addition, because the *t*-test value is greater than 1.96 over the post-2009 period, the null hypothesis (H0) is rejected. The year 2009 is therefore retained by our *t*-test criterion as the treatment date for our DID method. Further, Table 2 shows that the *t*-test statistic yields a value of -1.4679 over the 1990 to 2012 period and 11.015 over the post-2012 period. The null hypothesis (H0) is not rejected over the 1990 to 2012 period and the null hypothesis (H0) is rejected over the post-2012 period. We can therefore conclude that the average number of M&As in the non-life sector and the average number of M&As in the life sector are statistically the same over the period of 1990 to 2012 and statistically different over the post-2012 period. Our *t*-test statistic criterion also retains the year 2012 and cannot discriminate between the two years and between the two potential interpretations.

6.1.2. Statistical test based on the median

This test was proposed by Snecdecor and Cochran (1989). Based on this test, the analyze of the null hypothesis (H0) that the difference between the median number of M&As of target non-life insurers and the median number of M&As of target life insurers is equal to 0.

Our treatment date decision criterion is to test the null hypothesis (H0) that the median number of M&As in the non-life sector and the median number of M&As in the life sector are statistically similar over the period of 1990 to the end of the candidate date (2009 or 2012) on the one hand, and, on the other hand, to test the null hypothesis (H0) that the median number of M&As in the non-life sector and the median number of M&As in the life sector are statistically different over the post-treatment date period (post-2009 or post-2012) due to the treatment effect.

Table 2 reports a *p*-value of 0.8238 over the period of 1990 to 2009 and 0.0386 over the post-2009 period. Because the *p*-value is above the critical threshold of 5%, the null hypothesis is not rejected. In addition, because the *p*-value is lower than the 5% threshold over the post-2009 period, the null hypothesis (H0) is rejected. We can therefore conclude that the median number of M&As in the non-life sector and the median number of M&As in the life sector are statistically similar over the period of 1990 to 2009 and statistically different over the post-2009 period. The year 2009 is therefore retained by our median-based statistical test as the treatment date for our DID method. Further, Table 2 shows a *p*-value of 0.6776 over the 1990 to 2012 period and 0.0039 over the post-2012 period. Because the *p*-value is greater than the 5% critical threshold, H0 is not rejected. In addition, because the *p*-value is below the 5% threshold in the post-2012 period, the null hypothesis (H0) is refuted. We can therefore conclude that the median number of M&As in the non-life sector and the median number of M&As in the life sector are statistically similar over the period of 1990 to 2012 and statistically different over the post-2012 period. Our test based on the median also retains the year 2012 and cannot discriminate between the two dates.

6.1.3. Statistical test based on distributions

This test was proposed by Wilcoxon (1945). We test the null hypothesis (H0) that the distributions of the number of M&As per year of target non-life insurers and the number of M&As per year of target life insurers are close.

According to Table 2, the Z-test statistic yields a value of -0.915 over the period of 1990 to 2009 and 2.589 over the post-2009 period. Because the Z-test value in absolute terms is less than 1.96 over the period of 1990 to 2009, the null hypothesis (H0) is not rejected. In addition, because the Z-test value is greater than 1.96 over the post-2009 period, the null hypothesis (H0) is rejected. We can therefore conclude that the distribution of the number of M&As in the non-life sector and the distribution of the number of M&As in the life sector are statistically similar over the period of 1990 to 2009 and statistically different over the post-2009 period. The year 2009 is therefore retained by our statistical test based on the distributions as the treatment date for our DID method. In contrast, Table 2 shows that the *t*-test statistic yields a value of -1.354 over the 1990 to 2012 period and 2.666 over the post-2012 period. Because the value of the Z-test statistic in absolute terms is less than 1.96 over the period of 1990 to 2012, the null hypothesis (H0) is therefore not rejected. In addition, because the Z-test value is greater than 1.96 over the post-2012 period, the null hypothesis (H0) is rejected. We can therefore conclude that the distribution of the number of M&As in the two industries are statistically similar over the period of 1990 to 2012 and statistically different over the post-2012 period. Our test of the distribution-based statistic also retains the year 2012 and cannot discriminate between the two dates.

6.1.4. Monotonicity hypothesis

We employ an additional criterion called the monotonicity hypothesis, often used in econometrics to evaluate the treatment effect. This hypothesis postulates that when there is a change, the treatment effect can go in only one direction. To choose our treatment date based on the criterion of the monotonicity assumption, we used a graphical approach based on the analysis of Figure 12.

Figure 12 clearly shows a large difference between the number of M&As of target insurers in the non-life insurance sector compared with the number of M&As of target insurers in the life insurance sector observed over the post-2012 period. Moreover, we note that our treatment effect, defined as a positive difference between the number of M&As per year of target insurers in the non-life insurance sector and the number of M&As of target insurers in the life insurance sector, is respected for each year of the post-2012 period (9 years with a positive difference versus 0 year with a negative difference). In other words, 2012 changes the treatment effect in only one direction (positive difference) for each of the years in the post-2012 period. This affirms our monotonicity hypothesis. In contrast, Figure 12 shows that the year 2009 does not cause a change in the treatment effect in a single direction for each of the years in the post-2009 period (10 years with a positive difference versus 2 years with a negative difference). As can be seen, we get a negative difference for the years 2010 and 2011 and a positive difference for each of the other years in the post-2009 period. This violates our monotonicity condition (hypothesis). To conclude, because only the year

2012 meets the monotonicity condition, we select the year 2012 as the treatment date for our DID method with the monotonicity hypothesis.

Figure 12: Evolution of the number of M&As per year in each of the two insurance sectors (non-life and life, left) and their difference (in histogram, right)

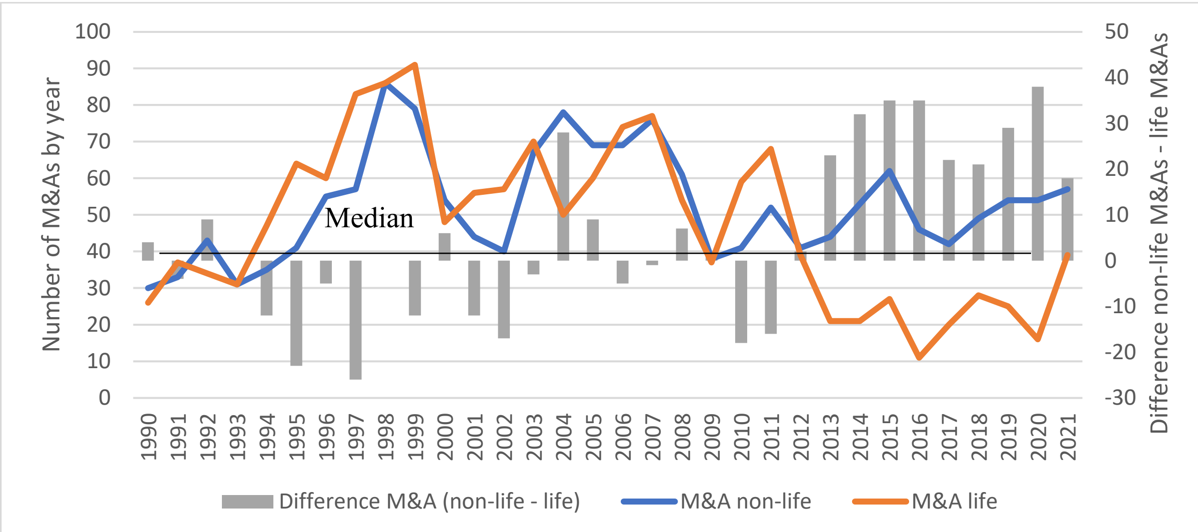

Data source: SDC database.

6.1.5. Median-criteria test of Guest (2021)

For robustness, a last statistical criterion based on the median is applied to ensure the reliability of the choice of the selected year 2012. To do this, we draw on the work of Guest (2021), who applies a median-based statistical criterion. This allows us to define a selection criterion whereby the treatment effect for each of the years in the post-treatment period (post-2009 or post-2012) is greater than the median value of the difference between the number of M&As per year of target insurers in the non-life insurance sector and the number of M&As of target insurers in the insurance sector over our entire study period (1990 to 2021), which is equal to 2 (see Table 2). This criterion supports the choice of 2012 as the treatment date for our DID method. As can be seen in Figure 12, the positive difference between the number of M&As per year of target insurers in the non-life insurance sector and the number of M&As of target insurers in the life insurance sector is greater than the median value of our entire study period (1990 to 2021) for each of the years in the post-2012 period. This is not the case for the post-2009 period, where we in fact observe a negative difference for the years 2010 and 2011, which is thus lower than the median of the entire sample. Therefore, our median-based criterion rejects the choice of the year 2009 as the treatment date for our DID method. To summarize, the statistical criterion based on the median supports the choice of the year 2012 retained by our affirmation of the monotonicity hypothesis.

6.2. Parallel trends analysis

We have just validated the choice of 2012 as the treatment year for our DID method. We will now perform a validation test for the presence of parallel trends before the end of that period. To do this, we first create 32 dummy variables for each of the years in the period of 1990 to 2021. Then, we create a dummy variable $\mathrm{Treated_i}$ with i equal to 1 for the treated group and 0 for the control group. Our Treated dummy (non-life sector) is then represented by the $\mathrm{Treated_i}$ variable. We also create 32 interaction variables between the Treated dummy and the year dummy for each year from 1990 to 2021. Finally, we regress our dependent variable, number of M&As per year and state, on our 32 $\mathrm{Treated_i} \times \mathrm{Year}$ interaction variables in each of the 51 states and in the two insurance sectors using the OLS method of estimation for panel data. With the OLS method, we capture the individual effect (state) and the time effect (year). The results are presented in Table 3 with 3,264 observations (32 × 51 × 2).

The results of our regressions validate the presence of a parallel trend before the end of 2012. As can be observed, the obtained coefficients are overall not statistically significant for the pre-treatment period (before 2013). Our F-test supports this result. It shows that the F-statistic on our $\mathrm{Treated_i} \times \mathrm{Year}$ interaction variables prior to the treatment date (1990 to 2012) is F (23, 3200) = 0.59 with a probability Prob > F = 0.9709. Given that the p-value is greater than 5%, we do not reject the null hypothesis, and we can conclude that the coefficients obtained before the treatment date are not significantly different from zero overall. In contrast, the coefficients obtained for each of the years during the post-2012 period are all statistically significant at the 1% level (except for the year 2021). Our F-test supports this result. The F-test over the post-treatment period (2013 to 2021) yields an F (9, 900) = 5.20 with Prob > F = 0.0008. Because the p-value is less than 5%, we reject the null hypothesis and can thus say that the coefficients considered as a whole are significant over the post-2012 period. These results allow us to validate our parallel trend test econometrically and thus confirm the choice of the year 2012 as the treatment year to be retained for our DID method.

Table 3: Parallel trends analysis for DID validation test

Test	Validation test		1st Robustness Test		2nd Robustness test	
Independent variable	Coefficient	Standard error	Coefficient	Standard error	Coefficient	Standard error
Treated×Year1990	2.034***	(0.242)	–		–	
Treated×Year1991	–0.0784	(0.182)	2.005***	(0.260)	–	
Treated×Year1992	0.176	(0.191)	0.176	(0.192)	2.220***	(0.260)
Treated×Year1993	0.00	(0.209)	0.00	(0.210)	0.00	(0.211)
Treated×Year1994	–0.235	(0.163)	–0.235	(0.164)	–0.235	(0.163)
Treated×Year1995	–0.451**	(0.206)	–0.451**	(0.207)	–0.451**	(0.208)
Treated×Year1996	–0.0980	(0.274)	–0.0980	(0.274)	–0.0980	(0.274)
Treated×Year1997	–0.510**	(0.239)	–0.510**	(0.239)	–0.510**	(0.239)
Treated×Year1998	0.00	(0.342)	0.00	(0.341)	0.00	(0.340)
Treated×Year1999	–0.235	(0.327)	–0.235	(0.327)	–0.235	(0.325)
Treated×Year2000	0.118	(0.239)	0.118	(0.240)	0.118	(0.240)
Treated×Year2001	–0.235	(0.224)	–0.235	(0.226)	–0.235	(0.227)
Treated×Year2002	–0.333	(0.208)	–0.333	(0.208)	–0.333	(0.208)
Treated×Year2003	–0.0588	(0.272)	–0.0588	(0.272)	–0.0588	(0.271)
Treated×Year2004	0.549**	(0.270)	0.549**	(0.271)	0.549**	(0.270)
Treated×Year2005	0.176	(0.248)	0.176	(0.249)	0.176	(0.248)
Treated×Year2006	–0.0980	(0.289)	–0.0980	(0.289)	–0.0980	(0.290)
Treated×Year2007	–0.0196	(0.320)	–0.0196	(0.321)	–0.0196	(0.320)
Treated×Year2008	0.137	(0.238)	0.137	(0.236)	0.137	(0.235)
Treated×Year2009	0.0196	(0.201)	0.0196	(0.202)	0.0196	(0.203)
Treated×Year2010	–0.353*	(0.211)	–0.353*	(0.212)	–0.353*	(0.213)
Treated×Year2011	–0.314	(0.198)	–0.314	(0.198)	–0.314	(0.198)
Treated×Year2012	0.0392	(0.203)	0.0392	(0.204)	0.0392	(0.204)
Treated×Year2013	0.451***	(0.162)	0.451***	(0.164)	0.451***	(0.165)
Treated×Year2014	0.627***	(0.182)	0.627***	(0.183)	0.627***	(0.182)
Treated×Year2015	0.686***	(0.199)	0.686***	(0.199)	0.686***	(0.198)

Treated×Year2016	0.686***	(0.188)	0.686***	(0.190)	0.686***	(0.191)
Treated×Year2017	0.431**	(0.206)	0.431**	(0.206)	0.431**	(0.207)
Treated×Year2018	0.412*	(0.211)	0.412*	(0.211)	0.412*	(0.210)
Treated×Year2019	0.569***	(0.151)	0.569***	(0.152)	0.569***	(0.153)
Treated×Year2020	0.745***	(0.182)	0.745***	(0.183)	0.745***	(0.183)
Treated×Year2021	0.353	(0.233)	0.353	(0.234)	0.353	(0.234)
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,264		3,162		3,060	
R-squared	0.631		0.628		0.630	

Robust standard errors in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

To ensure the reliability of our validation test of the choice of treatment date for our DID method, we conduct two robustness tests. The first test consists in ignoring the first year of observation: Treated×Year1990. The second test consists in ignoring the first two years of observations: Treated×Year1990 and Treated×Year1991. The results of these two robustness tests, as shown in Table 3, confirm the validation of the year 2012 as the treatment date to retain for our DID method.

7. DID analysis

In this section, we present in detail the variables of interest that we introduced into our regressions to analyze the difference between M&As in the US life and non-life insurance sectors using the DID method. The data utilized in this study come from the SDC database. The SDC database provides comprehensive quantitative and qualitative information on the characteristics of M&A transactions over the period of 1990 to 2021 in the two insurance sectors.

7.1. Description of variables

7.1.1. Natural experiment

In our econometric approach, we opted for a natural experiment methodology using the difference-in-differences estimator (DID). This estimator must separate the firms that have received a treatment (treatment group) and firms that have not received a treatment (control group).

7.1.2. Treatment group and control group variable

The purpose of our study is to determine the impact of climate risks or regulatory changes and market conditions on target insurers in the US. Because insurers in the non-life insurance sector are more exposed to climate risks and less exposed to market conditions and regulatory changes than are insurers in the life insurance sector during our period of analysis, we select insurers in the non-life insurance sector as our treatment group. We create a dichotomous variable $\mathrm{Treated_i}$ with *i* equal to 1 for the treatment group (non-life insurance sector) and 0 for the control group (life insurance sector).

7.1.3. Regression model

Based on our variables of interest, we consider the following regression model:

$$\mathrm{Nbr\ M\&A}_{\mathrm{it}} = \alpha + \delta_1 \mathrm{Treated_i} \times \mathrm{Post2012} + c_i + \eta_t + \epsilon_{it} \tag{1}$$

where:

$\mathrm{Nbr\ M\&A}_{\mathrm{it}}$: number of M&A in state *i* during year *t*;

$\mathrm{Treated_i} \times \mathrm{Post2012}$: equal to 1 for the treatment group after the treatment period and equal to 0 otherwise;

α : constant;

c_i : individual effects that exert the same influence on the state i in all periods;

η_t: temporal effects that affect all states equally in period *t*;

ϵ_{it} : standard random effects.

What interests us in equation (1) is the interaction variable $\mathrm{Treated_i}$ x $\mathrm{Post2012}$. It indicates the impact of the treatment on the insurers in the treatment group. Given that the regulation of insurance companies differs from state to state in the US, we created dummy-states variable to capture the individual effect of each state. The model assumes that the time shocks η_t affect all units in the two groups equally in period *t*. For this reason, we create dummy-periods to capture the time effect in each period. In our estimation of equation (1), we maintain the constant α since we use an estimation procedure that controls for multicolinarity. This approach is contrary to those of Dionne and Liu (2021) and Giorcelli and Moser (2020).

7.1.4. Description of targets

The targets selected for our study are US insurers that were acquired or merged during the period of 1990 to 2021. These targets operated in the life or non-life insurance sectors prior to the M&A transaction. We exclude from our sample of targets financing agency insurers or brokers with an SIC code of 6411 (Insurance Agents, Brokers and Service). The US targets selected for this study have the following SIC codes:

- 6311: Life Insurance
- 6321: Accident and Health Insurance
- 6324: Hospital and Medical Service Plans
- 6331: Fire, Marine, and Casualty Insurance
- 6351: Surety Insurance
- 6361: Title Insurance
- 6399: Insurance Carriers, Not Elsewhere Classified

Targets with the SIC codes 6321, 6324, 6331, 6351, 6361, and 6399 (Non-life Insurers) represent our treatment group, and targets with the Code 6311 (Life Insurance) represent our control group.[7]

After having presented the SIC codes of the target insurers selected for our analysis, we now document geographic information to determine the US states in which target insurers were most affected by the two waves of M&A transactions that we identified in Figure 1. Most large insurers have developed models based on geographic, seismic, and meteorological information to estimate the level of exposure to climate risks and the associated losses. In this study, we document geographic information to estimate targets' level of exposure to climate risks captured by the fixed effects. To do this, we break down the number of M&A transactions of the targets by state over the period of 1990 to 2021.

[7] In Online appendix 2, we regroup 6321 and 6324 with 6311. The statistical results remain the same but their interpretation changes.

We find that states such as California (324), Florida (288), New York (256), Texas (268), Illinois (158), Pennsylvania (155), Ohio (122), Michigan (87), Connecticut (101), New Jersey (119), Indiana (74), Massachusetts (69), Georgia (68), Maryland (68), Missouri (68), Minnesota (65), North Carolina (65), Arizona (64), and Delaware (63) each have a number of M&A transactions that exceeds the insurance industry average of 62. In other words, these regions have seen a significant number of M&A transactions over the past 30 years.

Using the distribution of the number of target M&A transactions by state shows that states can be subdivided into two groups based on whether the state is located in a coastal or a non-coastal zone. According to the National Oceanic and Atmospheric Administration (NOAA) website classification,[8] coastal zones include the following 30 states: New York, Florida, Connecticut, Pennsylvania, Texas, Illinois, California, Georgia, South Carolina, Maryland, Ohio, Virginia, Washington, Louisiana, Mississippi, New Jersey, Michigan, Alabama, North Carolina, Oregon, Maine, Massachusetts, Delaware, New Hampshire, Hawaii, Indiana, Minnesota, Wisconsin, Rhode Island and Alaska. The remaining 21 states (including District of Columbia) are located in non-coastal zones.

Figure 13 shows that all states identified as having a number of M&A transactions that exceeds the all-state average are in coastal zones except for Missouri and Arizona. In contrast, all non-coastal states have a number of M&A transactions per state that is below the all-state average except Missouri and Arizona. This distribution suggests that insurers located in coastal zones are more active in M&As. The extreme weather conditions that occur in these zones could explain this situation. Extreme weather can quickly trigger natural disaster events such as hurricanes, wildfires, tornadoes, and winter storms, and cause significant or extreme losses to insurers located in coastal zones. To summarize, insurers located in coastal zones have a higher level of exposure to climate risks than do insurers located in non-coastal zones. In our estimations, these differences will be taken into account by the fixed-effects variable.

[8] https://coast.noaa.gov/czm/mystate/.

Figure 13: Geographic distribution of the number of M&As transactions by state (1990 to 2021)

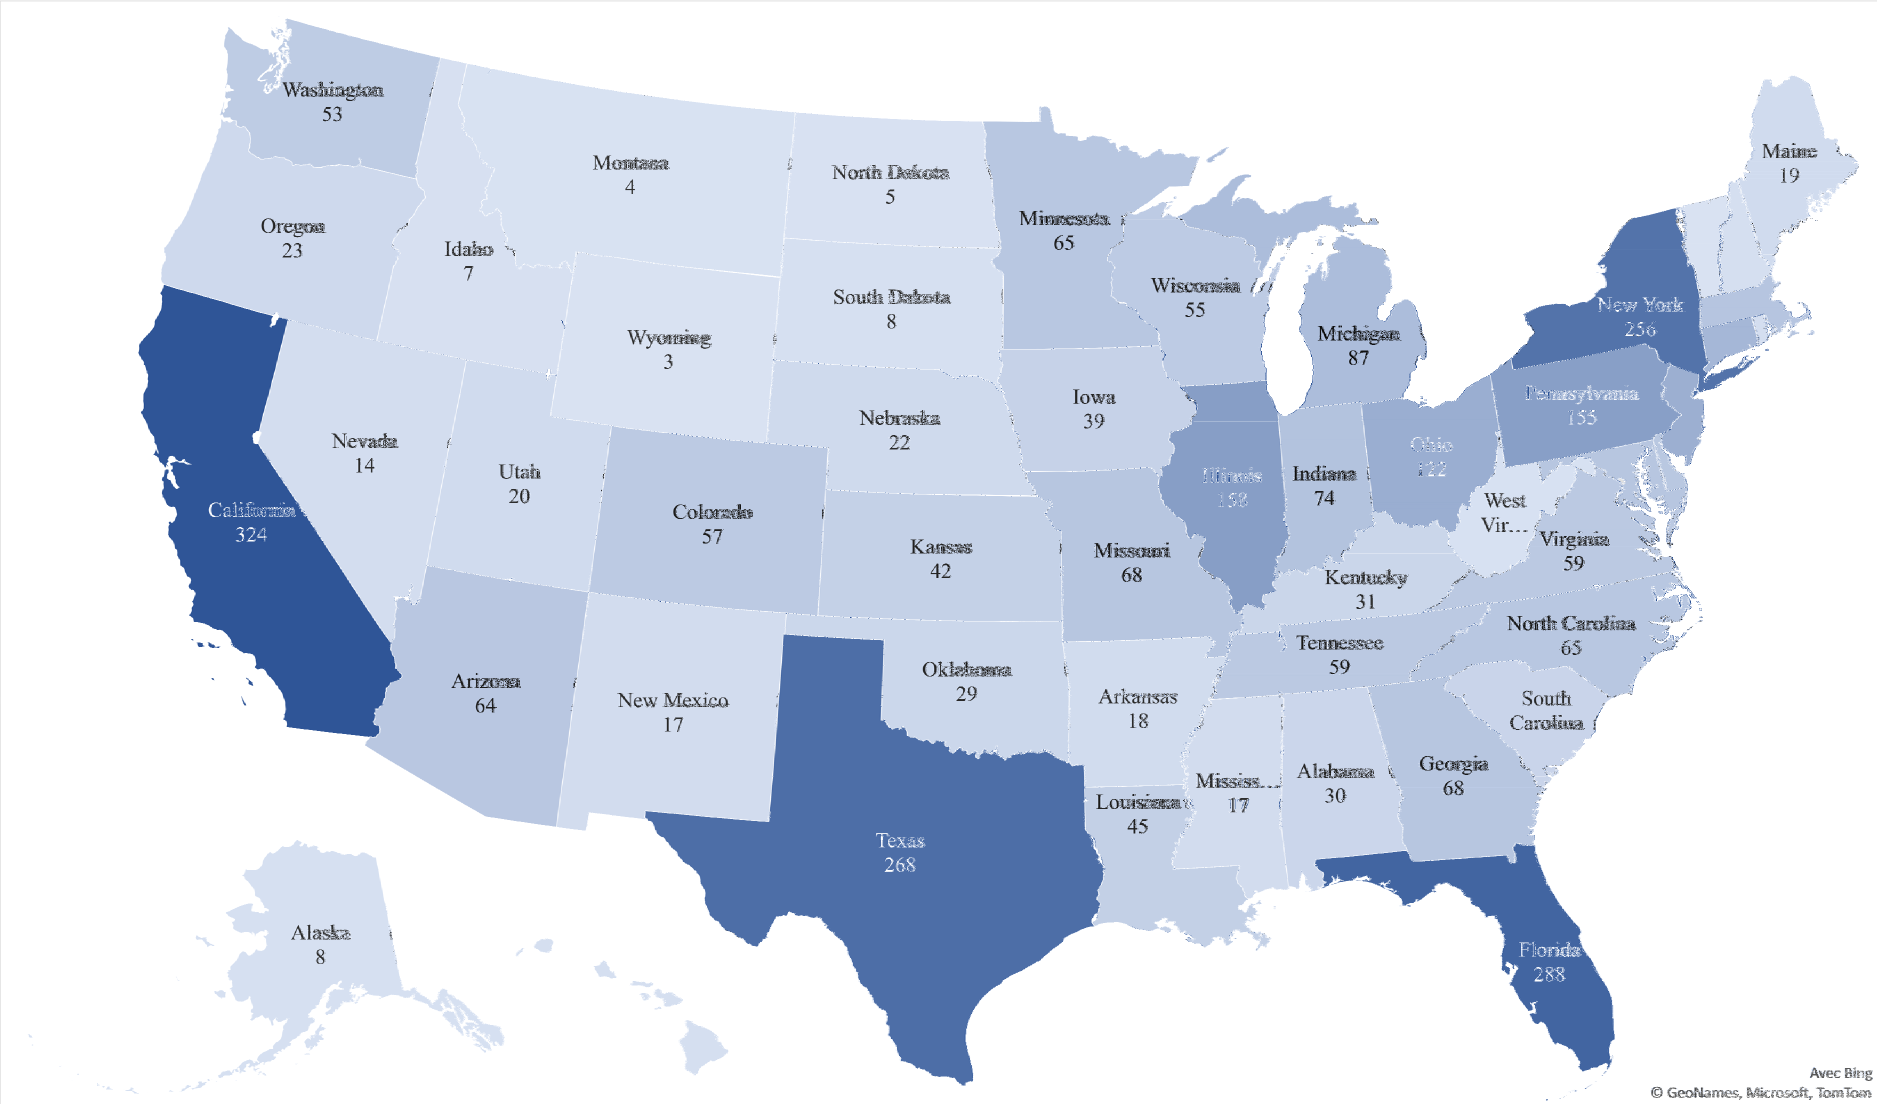

Data source: SDC database.

Additional states with numbers of MA in parentheses: South Carolina (34), Connecticut (99), Delaware (63), Maryland (67), Massachusetts (67), New Hampshire (10), New Jersey (116), Vermont (3), West Virginia (6).

The larger the number, the darker the color.

7.1.5. Description of acquirers

The acquirers are US or foreign companies that have carried out M&A transactions with the US target insurers over the period of 1990 to 2021. Based on the distribution of M&A transactions observed between 1990 and 2021, we identify two categories of transactions: inter-state transactions and intra-state transactions. According to this categorization of transactions, we determine that, over the period of 1990 to 2021, 24.14% of the M&A transactions were carried out by targets and acquirers from the same state (intra-state) and 75.86% of M&A transactions were carried out by targets with acquirers from different states (inter-state) or with foreign acquirers. Thus, this distribution suggests that acquirers have increased their geographic scope significantly over the period of 1990 to 2021.

Further, based on the distribution of M&A transactions observed between 1990 and 2021, we identify and determine the percentage of M&A transactions that occurred between targets and acquirers that operate in the same industry sector (i.e. that has the same SIC code). Our data show that 36.15% of the transactions were between targets and acquirers that have the same SIC code (concentration). In other words, 63.85% of the transactions were between targets and acquirers that have different SIC codes (diversification). This distribution suggests that acquirers have mostly opted for a management strategy based on diversification of operations rather than on concentration of operations.

7.1.6. Description of explanatory variables

Table 4 presents in detail the description of the variables we introduce into our model (1) to empirically test the difference between M&As in the US life and non-life insurance sectors by adopting the natural experiments method or the DID estimator.

We argue that the increase in natural disaster events that occurred in the post-2012 period, and especially the significant insured losses that they caused to insurers in the non-life insurance sector after 2012, seriously weakened target insurers in the non-life insurance sector. This has caused an increase in the number of M&A targets per year in the non-life insurance sector relative to the life insurance sector in the post-2012 period.

Table 4: Description of explanatory variables

Explanatory variable	Description	Expected sign
$\mathrm{Treated_i}$ (dichotomous)	$\mathrm{Treated_i}$ variable with *i* equal to 1 for the treated group (non-life insurance sector) and 0 for the control group (life insurance sector)	n.a
$\mathrm{Post2012}$ (dichotomous)	The $\mathrm{Post2012}$ variable takes the value 0 if the period is before the treatment (12-2012) and the value 1 if the period is after the treatment.	n.a
$\mathrm{Treated_i} \times \mathrm{Post2012}$ (dichotomous)	The interaction variable $\mathrm{Treated_i} \times \mathrm{Post2012}$ captures the effect of the treatment administered to the insurers in the treated group (non-life insurance sector) after the treatment.	+

We expect a positive sign of the coefficient of the variable $\mathrm{Treated_i} \times \mathrm{Post2012}$ on the number of target M&As per year. Otherwise, market conditions and changes in regulation after 2012 seem to have more negatively affected the life insurance industry. This observation may also explain a positive sign on the coefficient of the interaction variable.

7.2. Data and descriptive statistics of variables

The database used is the population of state-aggregated data on the characteristics of the target insurers' M&A transactions, observed in the two main sectors of US insurance (non-life and life) over a 32-year period and documented in the SDC database. Our data includes the 50 states of USA and the District of Columbia. This means that if a typical non-life insurance company operates across the country, it will be subject to 51 different regulations and different climate risk exposures. In order to capture the different structure of insurance companies as it often changes from state to state, we separate our data by state (51) and by year (32) according to each of our two insurance sectors. We obtain a total of 3,264 observations.

Table 5 presents the descriptive statistics of the variables related to the characteristics of M&As according to the two groups in our study sample. To compile this table, we calculate the means and standard deviations of the different variables within our two groups.

Table 5 shows that the average number of M&As per year and by state is 1.030 in the non-life insurance sector and 0.928 in the life insurance sector. In addition, the number of M&As for our two groups as a whole is 0.979 with a standard deviation of 1.634. Table 6 presents the mean and standard deviation of mergers and acquisitions by period. The mean is lower after 2012.

Table 5: Mean and standard deviation of the variables by insurance sector

Sample	Total sample (N=3264)	Non-life sector (N=1632)	Life sector (N=1632)
Dependent variable			
Number of M&As per year and by state	0.979 (1.634)	1.030 (1.662)	0.928 (1.605)
Variable of interest			
$\text{Treated}_i \times \text{Post2012}$	0.140 (0.347)	0.281 (0.449)	n.a n.a

Numbers in parentheses are standard deviations.

Table 6: Mean and standard deviation of the M&A by period

Period	1990-2021		1990-2012		Post-2012	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Number of M&As per year and by state	0.979	1.634	1.077	1.743	0.728	1.282

Table 5 indicates that the average number of M&As per year and by state observed in the non-life insurance sector over the period of 1990 to 2021 is roughly the same as that observed in the life insurance sector. To validate this observation, we statistically test the null hypothesis that the average number of M&As per year and by state in the non-life sector and the average number of M&As per year and by state in the life sector are statistically the same. Our statistical *t*-test yields a value of 1.60. Because the *t*-test value obtained is below the critical value of 1.96 (5% threshold), the hypothesis is not rejected. We can therefore conclude that the average number of M&As per year and by state in the non-life sector and the life sector are statistically the same over our entire study period, i.e. from 1990 to 2021.

7.3. Estimation results

The regression results of model (1) were obtained using the OLS method of estimation with fixed-effects. Our results presented in Table 7 indicate that the coefficient of our variable $\mathrm{Treated_i} \times \mathrm{Post2012}$ is positive and statistically significant at the 1% level. This result suggests a higher number of M&As in the treated group following the treatment date of 2012.

Table 7: Regression results for model (1) using OLS
with fixed effect on the state and on time

Dependent variable	Number of M&As per year (non-life and life)	
Independent variables	Coefficient	Standard error
$\mathrm{Treated_{NL}} \times \mathrm{Post2012}$	0.626***	0.0871
Constant	3.013***	0.226
State FE	Yes	Yes
Year FE	Yes	Yes
Observations	3,264	
R-squared	0.551	

Robust standard errors.
*** $p<0.01$.

The sign of the coefficient of the variable $\mathrm{Treated}_i \times \mathrm{Post2012}$ is as expected. This result empirically validates the assumption that the increase in natural disaster events or the variations of market conditions and in regulation that occurred during the post-2012 period may have seriously modified the insurers consolidation behavior between the two insurance sectors. These potential causes may have increased the difference of target M&As per year in the non-life insurance sector compared with the life insurance sector during the post-2012 period.

8. Financial health of US P&C insurers, 1990 to 2021

8.1. Combined ratio

Figure 14 shows the insured losses from natural disasters, while Figure 15 describes the evolution of the combined ratio. The combined ratio of the US non-life insurance industry has reached three major peaks since the 2000s. The first was in 2001 and reflects the major economic losses associated with the September 11, 2001, terrorist attack. The second peak occurred in 2005 and reflects the large economic losses associated with hurricanes Katrina, Rita, and Wilma, in 2005. Finally, the third peak was reached in 2011 and illustrates the costs of major claims generated by the exceptional series of violent tornadoes that occurred in 2011 in the US Midwest. If one considers only the level of the combined ratio attributable to natural catastrophe events in the US since the early 2000s, it is clear that 2011 was the second-most costly year for US insurers, after 2005.

Analysis of Figure 15 shows that the combined ratio for 2011 is higher than for 2017, which was a year of extremes in terms of US natural event losses, as shown in Figure 14. In other words, insured losses from natural catastrophe events in 2011 are lower than in 2017, but the combined ratio is higher.

Figure 14: Insured losses (billion $) from natural disaster events in US, 1990 to 2021

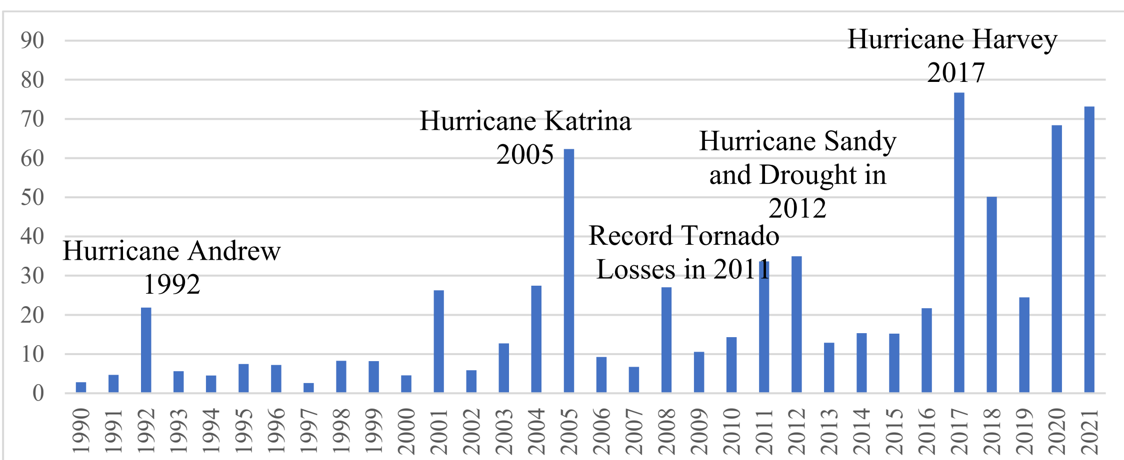

Data source: VERISK database. VERISK selects events with insured losses of $25 million and above. Insured losses: property damage and business interruption, excluding liability and life damage.

Figure 15: Combined ratio US property-casualty, 1990 to 2021

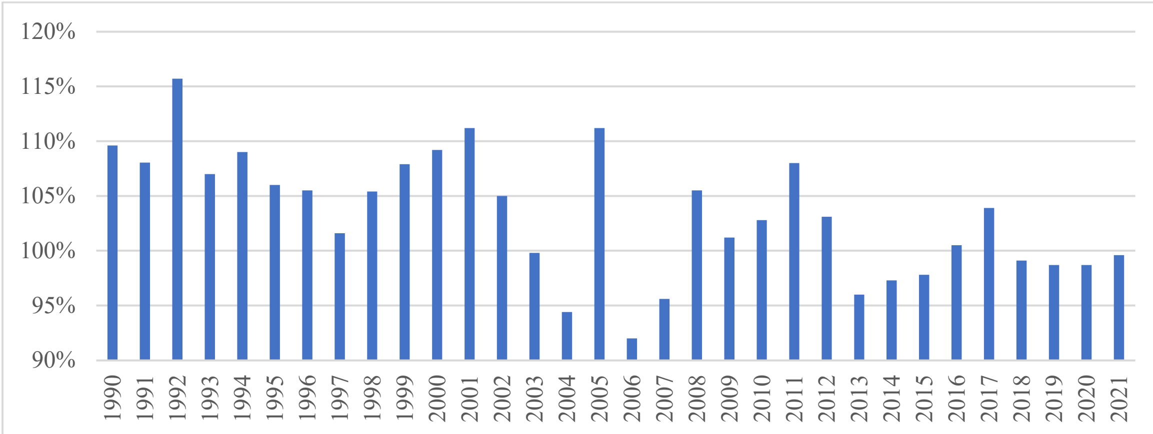

Data source: NAIC data, Federal Insurance Office, US Department of the Treasury, Annual Report on the Insurance Industry (before 2018), and Statista data. Combined ratio formula = (claims costs + management expenses) / premiums earned.

The combined ratio is affected by the claims losses variable (the combined ratio being an increasing function of insured losses). The combined ratio is also affected by the management expenses variable (the combined ratio being an increasing function of management expenses). Another variable that affects the level of the combined ratio is the premiums earned variable. As the formula noted below Figure 15 indicates, the combined ratio is a decreasing function of the premiums earned variable.

Our data from the NAIC[9] indicate that total claims costs (including those due to natural catastrophe events) in 2011 were $296 billion, as compared to $354 billion in 2017, an increase of 20% from 2011 to 2017. These loss cost figures suggest that the 2017 combined ratio level should be higher than that of 2011. In addition, management expenses in 2011 were $180 billion, versus $214 billion in 2017, for an increase of 19% from 2011 to 2017.

In other words, we should expect a higher combined ratio in 2017 than in 2011, given that the total loss costs and management expenses, which were $477 billion in 2011, rose to $568 billion in 2017, an increase of 19%. Our data, however, indicate the opposite: in Figure 15, a ratio of 108% in 2011 (the record year for natural event losses in the US) versus a ratio of 103% in 2017, equal to a 5% decrease in the combined ratio.

[9] US Property & Casualty and Title Insurance Industries – 2020 Full Year Results.

Our NAIC data also indicate that net premiums earned, which were $443 billion in 2011, grew to $550 billion in 2017, an increase of 24%. By contrast, the same data source shows that total loss costs and management, which were $477 billion in 2011, increased to $568 billion in 2017, a 19% increase. We clearly see that it is the increase in the growth of net premiums earned of 24% versus the increase in total loss costs and management expenses of 19% over the period from 2011 to 2017 (a difference of 5%) that could explain the reduction in the combined ratio level by 5% over the same period (108% in 2011 versus 103% in 2017).

8.2. ROA and asset-turnover of targets

To illustrate the very sharp deterioration in growth volume of all public non-life target insurers after the series of violent tornadoes that occurred in 2011, we use two profitability measures. The first is the return on total assets (ROA) profitability indicator and the second is the asset-turnover efficiency ratio. We use the ROA profitability indicator as a reliable instrument to measure the viability (growth) of our targets and non-life insurers. To be viable, insurers, like any other company, must generate profitability in all their businesses. They must repay their clients and creditors, satisfy their shareholders' demands, and finance their growth (on which their viability depends). Second, we use the asset-turnover ratio as another reliable measure of the viability of our non-life public targets. This ratio measures the efficiency with which a company uses its assets to produce revenue. In other words, asset-turnover measures performance in terms of return on assets.

Figure 16 and Figure 17 compare the ROA and asset-turnover efficiency ratios of a sample of M&A targets in the US non-life insurance market with those of the non-life insurance industry. The two target ratios do not look very different than those of the industry, which indicates that the financial conditions of the targets were not necessarily bad at the merger or acquisition dates. We must note that these results are limited to a sample of 224 targets that may not represent the entire industry. They do not necessarily make it possible to reach a final conclusion about the overall insurance industry.

Figure 16: Return on total assets (ROA) for a sample of non-life targets (left) and for the non-life insurance industry (right) in the US, 1990 to 2021

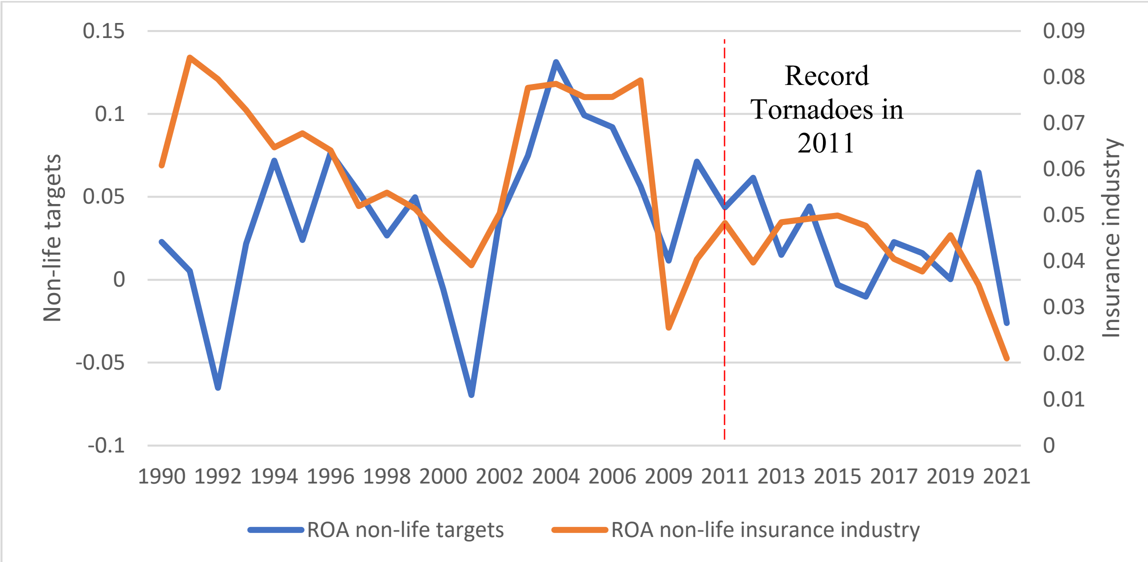

Sources: COMPUSTAT and NAIC databases.

Figure 17: Asset-turnover efficiency ratios for a sample of non-life targets (left) and non- life insurance companies (right), 1990 to 2021

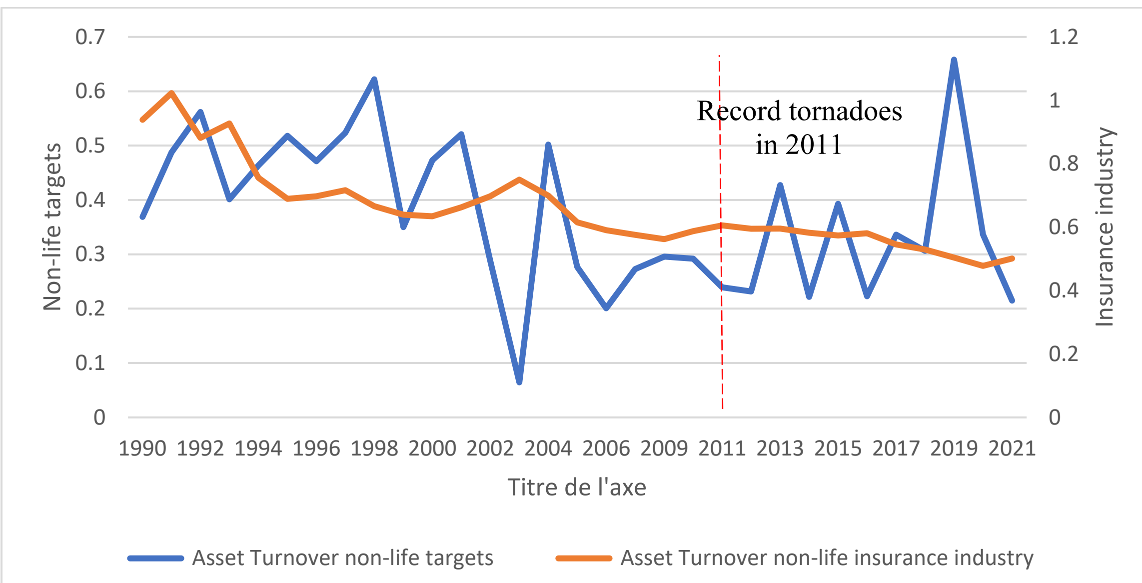

Sources: COMPUSTAT and NAIC databases.

8.3. CAT bonds

The exceptional series of severe tornadoes in 2011 also resulted in very high losses on two Mariah Re catastrophe (CAT) bonds: the Mariah Re 2010-1 CAT Bond triggered[10] on September 30, 2011; and the Mariah Re 2010-2 CAT Bond triggered on August 30, 2011. These two CAT bonds were issued in November 2010 (for Mariah Re 2010-1) and December 2010 (for Mariah Re 2010-2) by Mariah Re Ltd. They covered the risks of severe storms in the US. The losses on these two Mariah Re CAT bonds issued in 2010 represent the highest losses in the history of CAT bonds in the US. These results indicate how the utilization of ILS instruments helps the insurance industry maintain capital in years of very high losses.

8.4. World Economic Forum

The magnitude of the loss costs caused by the natural disasters in the US in 2011, to which can be added the natural disaster events that occurred internationally, notably in Japan, Thailand, New Zealand, and Australia, may have raised the collective awareness of the danger of natural (or weather) disasters, as indicated by the works from the experts of the World Economic Forum (Table 8).

The experts of the World Economic Forum show that awareness of environmental risks appeared among companies' top five concerns only starting in 2011, that is to say, after the occurrence of very large natural disasters. The analysis is based on an assessment of hazards by specialists from various sectors through a risk mapping model. Risk mapping is one of the risk management tools most widely used by companies, particularly insurers. It involves a graphic representation of a number of risks and serves to identify the threats and dangers incurred by organizations, synthesizing them in a hierarchical form. According to *Atlas* magazine (consulted on 6 December 2022), this hierarchy is based on criteria such as probability of occurrence, potential impact, and level of risk control. Further, mapping natural, economic, and social catastrophe risks enables insurance companies to better identify the threats likely to impact their business. Table 8 presents the World Economic

[10] Triggered means that the risk underlying the (CAT) bond has materialized and that the principal or capital is used to cover the insurer's loss instead of going back to the investors.

Forum's assessment of the perception (by year) of the five global risks to which companies are most sensitive, for the years 2007 to 2013.

The table shows that in 2011, the overall risk that leaders considered most worrisome for the next 10 years is meteorological catastrophes (storms, tornadoes and hurricanes). Climatological catastrophes (rain, snow, or hail) are ranked fifth, following the series of violent tornadoes in the Midwestern US and the natural and nuclear disasters in Japan and Thailand.

Table 8: Top five global risks in terms of probability of occurrence

	2007	2008	2009	2010	2011	2012	2013
1st	Breakdown of critical information infrastructure	Asset price collapse	Asset price collapse	Asset price collapse	Meteorological catastrophes	Severe income disparity	Severe income disparity
2nd	Chronic disease in developed countries	Middle East instability	Slowing Chinese economy (<6%)	Slowing Chinese economy (<6%)	Hydrological catastrophes	Chronic fiscal imbalances	Chronic fiscal imbalances
3rd	Oil price shock	Failed and failing states	Chronic disease	Chronic disease	Corruption	Rising greenhouse gas emissions	Rising greenhouse gas emissions
4th	China economic hard landing	Oil and gas price spike	Global governance gaps	Fiscal crises	Biodiversity loss	Cyber attacks	Water supply crises
5th	Asset price collapse	Chronic disease, developed world	Retrenchment from globalization (emerging)	Global governance gaps	Climatological catastrophes	Water supply crises	Mismanagement of population ageing

Economic | Environmental | Geopolitical | Societal | Technological

Source: World Economic Forum.

Figure 18: US Property catastrophe rate-on-line index (private and public insurers)

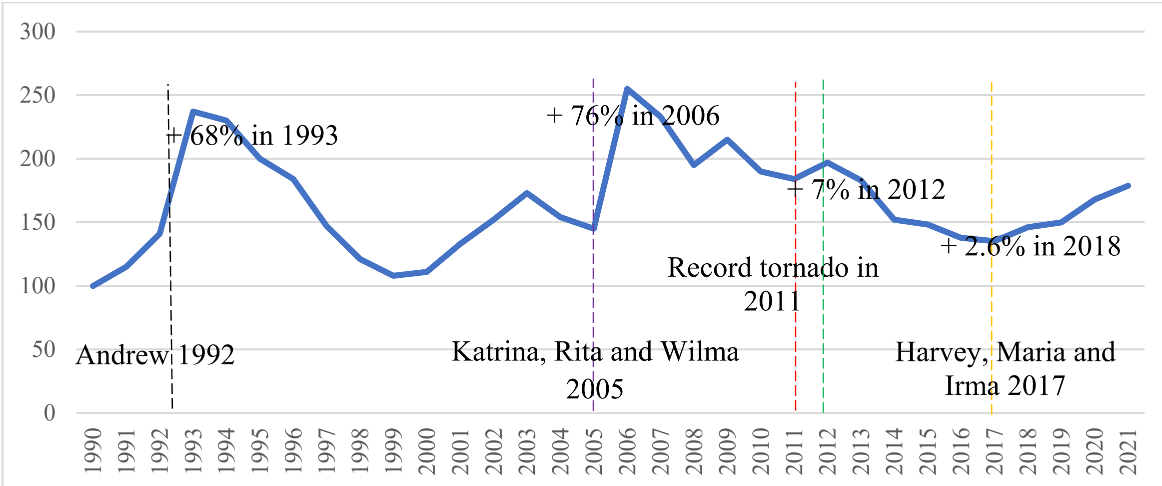

Data source: Data from Guy Carpenter, presented by Artemis.bm.

Definition: Rate-on-line index (ROL) is the ratio of premium paid to loss recoverable in a reinsurance contract. In simple terms, ROL represents the amount of money an insurer must commit to obtain reinsurance coverage. A high ROL indicates that the insurer must pay more for coverage, while a low ROL means that an insurer pays less for the same level of coverage.

8.5. ROL index

Figure 18 indicates that major disasters led to large changes in the ROL index until 2012, and small changes thereafter. This is the case, for example, with Hurricane Andrew in 1992 and Hurricane Katrina in 2005. After Andrew in 1992, the catastrophe index increased 68% in 1993. It increased 76% in 2006 after Hurricane Katrina in 2005, and by 7% in 2012 after the series of severe tornadoes in the Midwest in 2011. By contrast, Figure 18 shows very small changes in the ROL index after 2012. All ROL changes remained below the 7% mark (ROL change from 2011 to 2012) throughout the post-2012 period, even after major hurricanes Harvey, Maria, and Irma of 2017 (the year of extremes); the ROL increased by only 2.6% in 2018.

8.6. Premium earned

Premiums earned are one of the main resources available to insurers to cover loss costs. Therefore, the small changes in the ROL index observed after 2012 suggest that non-life insurers increased their level of premium collection in the post-2012 period. To verify this, we use premium earned data and calculate the market share of each of our insurance sectors (non-life and life) over the period of 2007 to 2017. We retain this period because data on premiums earned, from the Insurance Information Institute, are available only for the period of 2007 to 2017.

Figure 19: Market share of premiums earned by all non-life (left) and life insurers (right) (private and public)

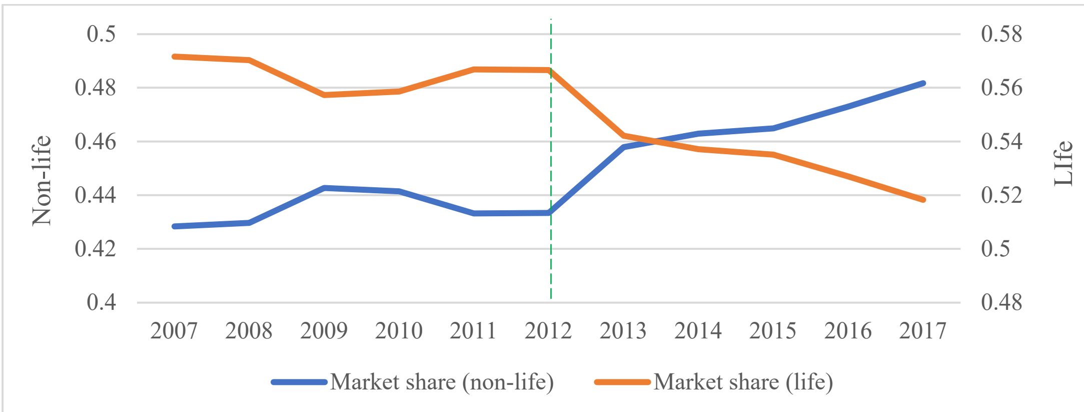

Data source: Insurance Information Institute.

Figure 19 shows that premiums earned share increased significantly in the post-2012 period in the non-life insurance sector. By contrast, premiums earned share decreased significantly during the post-2012 period in the life insurance sector. Over five years (2012 to 2017), the non-life sector's premium market share grew by 12%, while the life insurance sector's premium market share declined by 9%.

Figure 20 presents the different premium indexes during the period of analysis. Life premium growth is much lower than P&C premium growth. The P&C Homeowner's Insurance Premium Index more than doubles during the period of analysis.

The results obtained from figures 19 and 20 suggest that the recognition of natural catastrophe risk may have led insurers to readjust their pricing, to properly take climate risk into account. The net increase in the level of premiums earned in the post-2012 period illustrates this.

Figure 20: Insurance Premium Indices

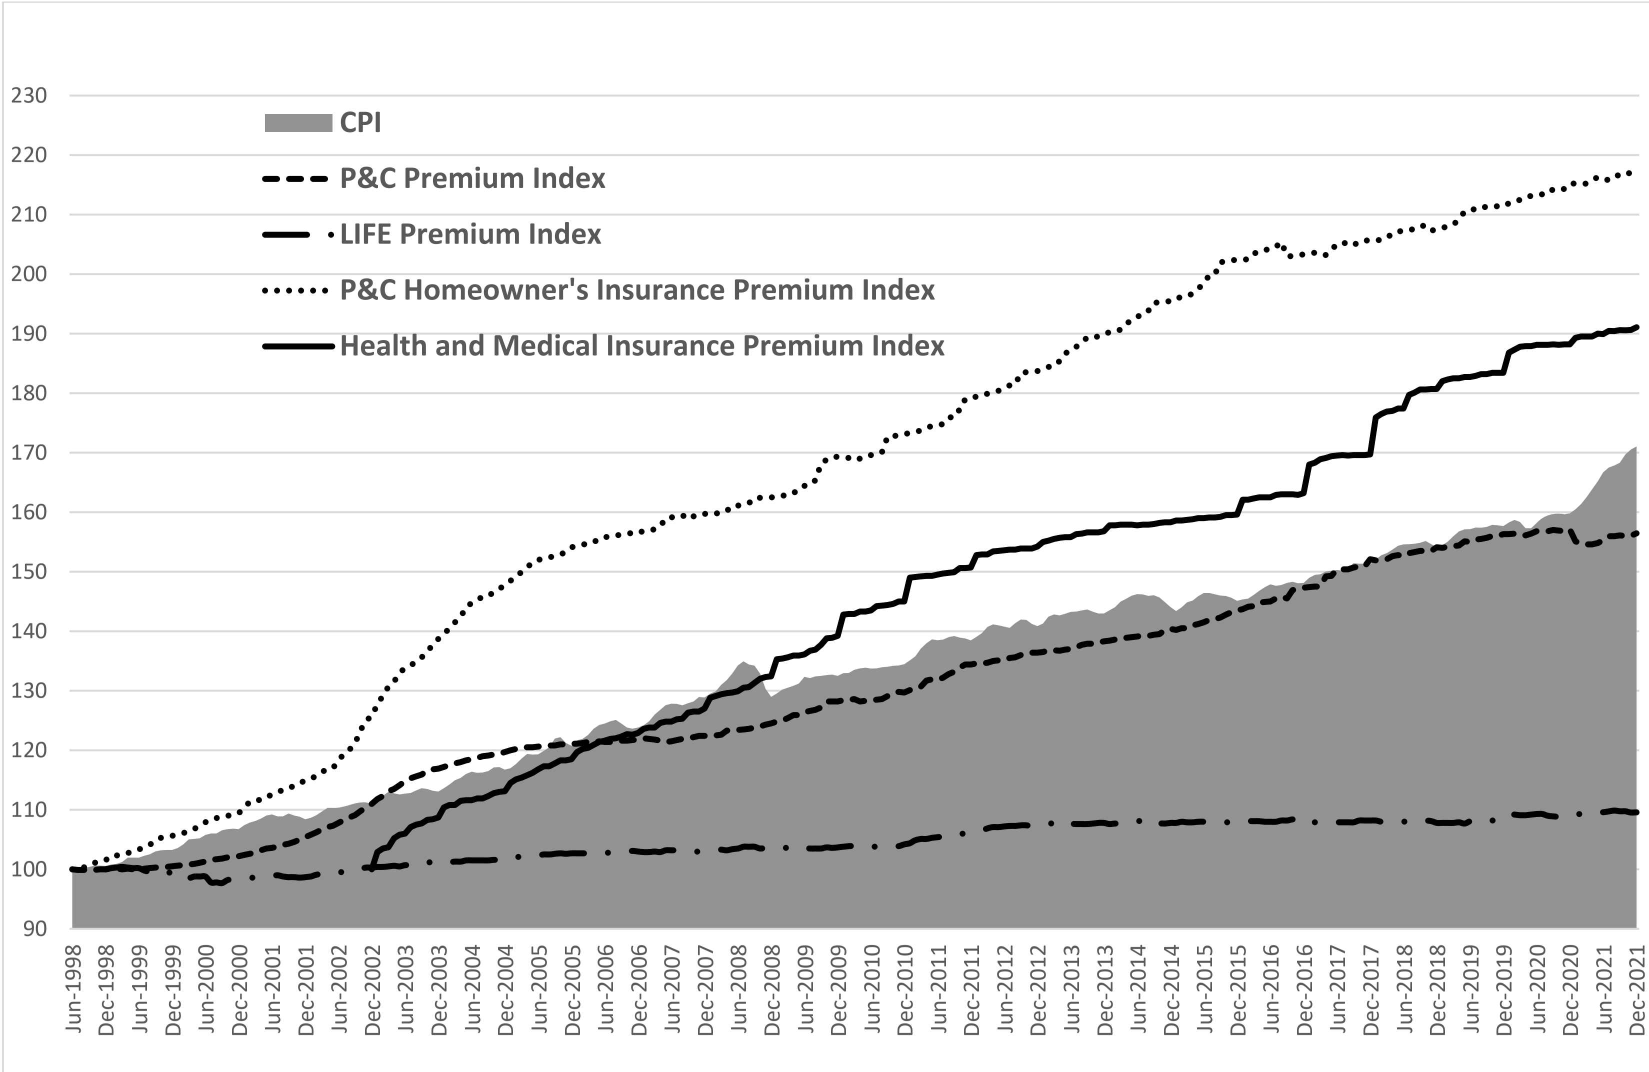

Data source: US Bureau of Labor Statistics.

8.7. Market-to-book and price/book

The results in Figure 21 suggest that there has been resilience to property damage due to natural disasters, in the non-life insurance industry in the post-2012 period, a period that was marked by sharp increases in claims costs due to natural disasters, especially starting in 2017 (the year of Harvey, Maria, and Irma). In other words, recognition of the risk of large claims from natural disasters in post-2012 allowed US non-life insurers to sufficiently cover loss costs with reserves from written premiums, allowing them to improve their financial health in the post-2012 period, as shown in Figure 21. Indeed, Figure 21 shows that the financial health (as measured by the price/book and market-to-book (MTB) indicators) of all insurers in the US non-life insurance industry improved significantly in the post-2012 period.

Figure 21: Evolution of the price/book and MTB ratios in the US non-life sector

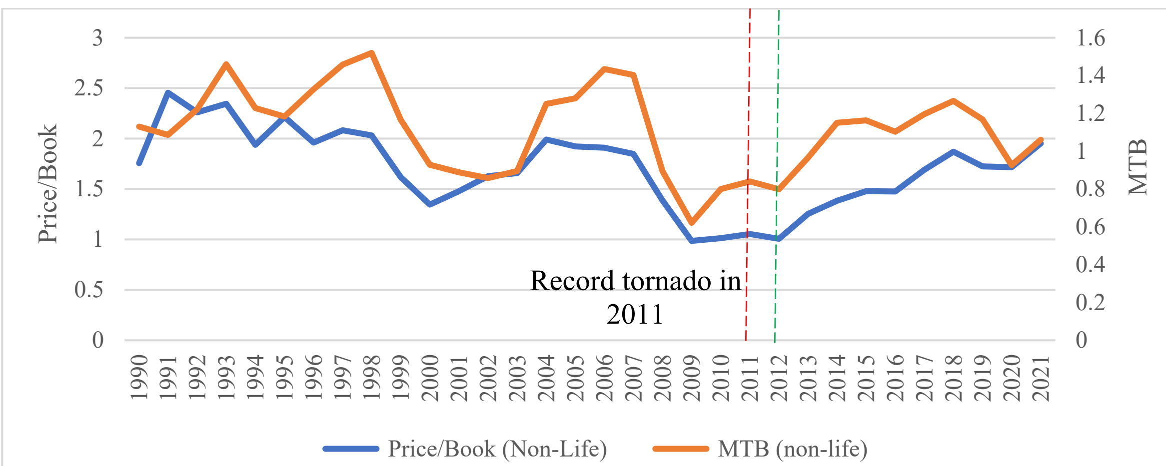

Data source: COMPUSTAT database.

8.8. ROA in both sectors

Figure 22 shows the evolution of the ROA ratio. It suggests that non-life insurers as a whole have returned to growth after the great economic recession of 2009 and the decline in 2012 caused by the Midwestern tornados in 2011 and the impact of Hurricane Sandy in 2012. By contrast, Figure 22 still points to a deterioration in organic growth across all life insurers during the same period. Figure 22 also shows a divergence in the trend between overall growth of non-life insurers and life insurers after 2012.

Figure 22: Evolution of the ROA ratio in the non-life and life insurance sectors in the US, 1990 to 2021

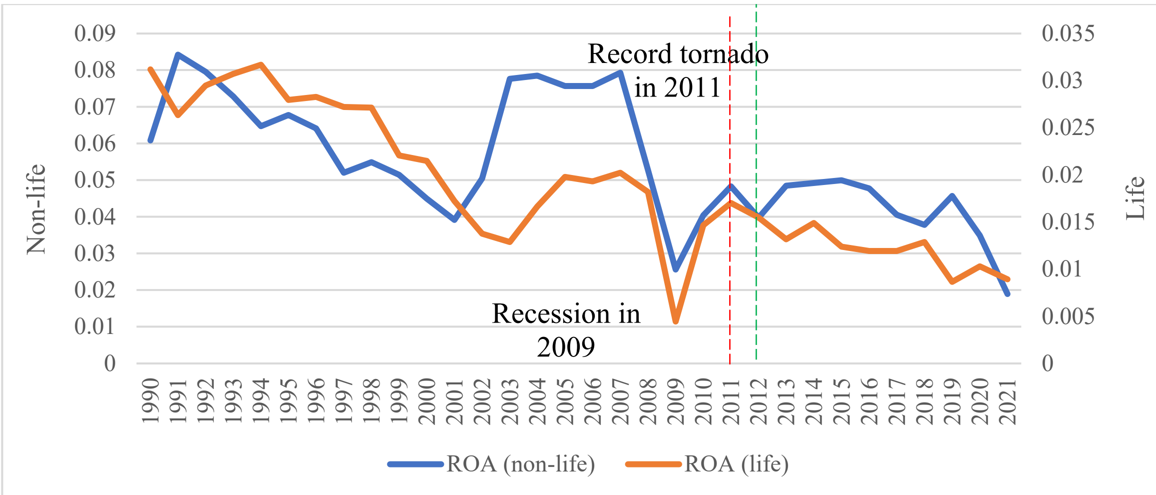

Data source: COMPUSTAT database.

Our data show, as Figure 22 indicates, that there is a clear positive difference between the ROA of the US non-life insurance industry and that of the US life insurance industry for almost every year in the post-2012 period. This difference was also observed between M&As of the US non-life insurance industry and those of the US life insurance industry, for each of the years over the same post-2012 period.

9. Conclusion and discussion

The main objective of this study is to test for the presence of a statistical link between climate risk and mergers and acquisitions (M&As) in the US property and casualty (P&C) insurance industry. The main research question is the following: is the observed increase in claims costs associated with climate risk events a causal factor for M&As growth during the 1990-2021 period? More generally, the study examines how the costs of catastrophic weather events associated with climate risk have impacted the insurance industry's resilience by affecting economic capital during the 1990–2021 period. The financial literature often describes M&As as consolidation activities in different industries.

We develop a natural experimental event study by identifying two groups of insurers that are exposed differently to climate risk events. The control group of insurers was less exposed to weather risk events, and the treatment group of insurers was more exposed to weather risk events. Life insurers were considered less exposed than P&C insurers. Our statistical results indicate that the post-2012 period was associated with a difference in M&A activity between the two insurance sectors, while both sectors had parallel trends in M&A prior to January 2013. The number of M&As was statistically higher in the P&C insurance sector than in the life insurance sector in the post-2012 period.

We faced two major difficulties isolating climate risk as having a causal effect on M&As. The first was separating M&As from other sources of capital consolidation that insurers can use to protect themselves from natural catastrophes. Dionne and Desjardins (2022) show that US P&C insurers significantly increased their capital between 1997 and 2020. These authors also identify different potential sources of capital, such as reinsurance, premium management, M&As, capital regulation, and insurance linked securities (ILS).

The second difficulty was identifying potential factors other than weather risk events that may have affected M&As in the two insurer groups in the 1990–2021 period of analysis. The US insurance industry overall was affected by the 2007–2009 financial crisis, and the life insurance industry in particular (Barnes et al., 2016). Market conditions were difficult after the crisis for the life insurance industry (NAIC, 2022; Federal Insurance Office, 2022). Premium growth was low in this line of business, and interest rates were very low in the whole economy. Different federal regulations for capital were introduced, particularly in and after 2012, to consolidate capital risk management following the financial crisis. These new regulations affected capital levels and may have introduced uncertainty into the markets about the potential future growth of M&As.

Our main results do not support a causal link between climate risk and M&As in the US insurance market during the period of analysis. We obtain a significant increase in the number of M&A events in the treatment group (target non-life insurers) compared to the control group (target life insurers) after the year 2012, but we cannot yet identify the actual cause of this result. Climate risk costs significantly increased after 2012 in the P&C insurance industry, but it is not clear that M&As were chosen to consolidate the industry. The observed difference could also be attributed to a significant reduction in M&As in the life insurance industry after 2012, which could be explained by stagnant activity growth in insurance premiums and very low interest rates in the economy.

It seems that P&C insurers choose other diversification activities, including reinsurance and premium management. ILS, including catastrophe bonds, became more popular during our period of analysis, but cannot be considered one of the main sources of capital in the US P&C insurance industry. Better capital risk management under the stronger risk regulation introduced in 2012 and following years could also have been another significant source of resilience for the P&C insurance industry. A preliminary analysis of all these potential sources of capital is presented in the appendix. It indicates that premium growth and reinsurance demand were the two main sources of capital in the P&C insurance industry during our period of analysis. Finally, our analysis of different financial indicators confirms the relative good health of P&C insurers after 2012.

Many extensions of our research are in development. Reinsurance is important to diversify climate risks around the world over time (Cummins and Weiss, 2000, 2004). It has been documented that the presence of reinsurance can affect P&C insurers' behavior (Desjardins et al., 2022). The introduction of a more active role for reinsurance in modeling insurers' capital should improve our understanding of the stability of this industry despite the increasing number and severity of climate risk events. But reinsurance capacity may have its limit, particularly with the increase of climate risk worldwide, which reduces international diversification capacities.

Our period of analysis ends with the year 2021. Many extreme events have been observed in the P&C insurance industry since 2017, which was a record year. The years 2021 and 2022 were particularly expensive and have significantly affected both the insurance and reinsurance industries. Some reinsurance companies have been downgraded by rating agencies and others have reduced their participation in the extreme weather risk market. Reinsurance premiums are very high in 2023, and insurers are also leaving the market in high-risk states such as Florida. To date, 2022 was the third-highest for total insured costs, behind 2017 and 2005, according to Aon re (2023) and Munich re (2023). Total economic losses were $165 billion in the US, with about $100 billion in insured losses for 2022. It seems that the annual $100 billion in insured losses is becoming the standard, or perhaps even a minimum! Updates of the data and analyses from this paper will be needed to take into account the new trend in the severity of catastrophic events that began in recent years.

Before 2021, many reports described the US P&C insurance industry as overcapitalized. It is not clear that this will remain true in the future, when we look at insured costs since 2017. These costs are not only high, they repeat every recent year. The years 2005 and 2011 used to be considered outliers, with a low probability of recurrence. This does not seem to be the case anymore with the recent years, as we observe the climate changing.

Finally, another issue concerns the effect of climate risk on life insurance. In a recent SCOR analysis (2022), climate change risks are related to potential life liabilities in the long run. The relevance of climate change risks for life insurance liabilities depends mainly on the insurer's location in the world. For example, the study shows that climate change could generate additional US heat mortality over a time horizon of several decades. More research on the effect of climate risk on the life insurance industry also seems necessary.

10. References

Aerts, J.C.J.H., Wouter Botzen, W.J., Emanuel, K., Lin, N., de Moel, H., Michel-Kerjan, E.O., 2014. Evaluating flood resilience strategies for coastal megacities. Science 344, 6183, 473-475.

Akhigbe, A., Madura, J., 2001. Intra-industry signals resulting from insurance company mergers. The Journal of Risk and Insurance 68, 489-506.

Amel, D., Barnes, C., Panetta, F., Salleo, C., 2004. Consolidation and efficiency in the financial sector: a review of the international evidence. Journal of Banking and Finance 28, 2493-2519.

Amihud, Y., Lev, B., 1981. Risk reduction as a managerial motive for conglomerate mergers. Bell Journal of Economics 12, 1, 605-617.

AON, 2023. Weather, climate and catastrophe insight. https://www.aon.com/weather-climate-catastrophe/index.aspx.

Atlas magazine. https://www.atlas-mag.net/en/category/regions-geographiques/monde/natural-catastrophes-115-billion-usd-of-insured-losses-in-2022. Consulted on 6 December 2022.

Baldauf, M., Garlappi, L., Yannelis, C., 2020. Does climate change affect real estate prices? Only if you believe in it. The Review of Financial Studies 33, 3, 1256-1295.

Betton, S., Eckbo, B.E., Thorburn, K., 2009. Merger negotiations and the toeholds puzzle. Journal of Financial Economics, 91, 2, 158-178.

Boubakri, N., Dionne, G., Triki, T., 2008. Consolidation and value creation in the insurance industry: The role of governance. Journal of Banking and Finance 32, 56-68.

Brockman, P., Yan, X., 2009. Block ownership and firm-specific information. Journal of Banking and Finance 33, 2, 308-316.

Chamberlain, S., Tennyson, S., 1998. Capital shocks and merger activity in the property-liability insurance industry. The Journal of Risk and Insurance 65, 563-595.

Cummins, D., Klumpes, P., Weiss, M.A., 2015. Mergers and acquisitions in the global insurance industry: Valuation effects. The Geneva Papers on Risk and Insurance - Issues and Practice, 40 (3), pp. 444-473. ISSN 1018-5895.

Cummins, D., Tennyson, S., Weiss, M., 1999. Consolidation and efficiency in the US life insurance industry. The Journal of Banking and Finance 23, 325-357.

Cummins, D., Weiss, M., 2004. Consolidation in the European insurance industry: Do mergers and acquisitions create value for shareholders? Working Paper, The Wharton Financial Institutions Center.

Cummins, D., Xie, X., 2008. Mergers and acquisitions in the US property-liability insurance industry: Productivity and efficiency effects. Journal of Banking and Finance 32, 1, 30-55.

Deloitte, 2022. 2022 insurance M&A outlook (27 pages document internet). https://www2.deloitte.com/us/en/pages/financial-services/articles/insurance-m-and-a-outlook.html.

Denis, D.K., McConnell, J.J., 2003. International corporate governance. Journal of Financial and Quantitative Analysis 38, 1-36.

Desjardins, D., Dionne, G., Koné, N., 2022. Reinsurance demand and liquidity creation: A search for bicausality. Journal of Empirical Finance 66, 137-154.

Dionne, G., Desjardins, D., 2022. A re-examination of the U.S. insurance market's capacity to pay catastrophe losses. Risk Management and Insurance Review 25, 4, 515-549.

Dionne, G., Liu, Y., 2021. Effects of insurance incentives on road safety: Evidence from a natural experiment in China. Scandinavian Journal of Economics 123, 2, 453–477.

Dionne, G., 2015. Policy Making and Climate Risk Insurability: How Can (Re)Insurers Contribute to Economic Resilience in Climate Risk Events? Working paper 15-06, Canada Research Chair in Risk Management, HEC Montréal.

Dionne, G., La Haye, M., Bergerès, A.S., 2015. Does asymmetric information affect the premium in mergers and acquisitions? Canadian Journal of Economics 48, 3, 819-852.

Estrella, A., 2001. Mixing and matching: prospective financial sector mergers and market valuation. Journal of Banking and Finance 25, 2367-2392.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3-56.

Floreani, A., Rigamonti, S., 2001. Mergers and shareholders' wealth in the insurance industry. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=267554.

Geneva Association 2022. Anchoring Climate Risk Change Risks Assessment in Core Business Decisions in Insurance, Third report of the Geneva Association task force on climate change risk assessment fo the insurance industry (Maryam Golnaraghi coordinator and lead author), 51 pages document internet. https://www.genevaassociation.org/sites/default/files/2022-10/Climate%20Risk%203_web.pdf.

Giorcelli, M., Moser, P., 2020. Copyrights and creativity: Evidence from Italian opera in the Napoleonic Age. Journal of Political Economy 128, 11, 4163-4210.

Guest, P.M., 2021. Risk management in financial institutions: A Replication. Journal of Finance 76, 5, 2689-2707.

Imbens, G.W., Wooldridge, J.M., 2009. Recent developments in the econometrics of program evaluation. Journal of Economic Literature, 47, 1, 5-86.

Issler, P., Stanton, R., Vergara-Alert, C., Wallace, N., 2020. Mortgage markets with climate-change risk: Evidence from wildfires in California. Working paper, University of California, Berkeley.

Jensen, M., 1986. Agency costs of free cash flow, corporate finance and takeovers. American Economic Review 76, 323-329.

Kessler, D., 2015. Conference at the SCOR Scientific Seminar About Climate Risks, Paris, June. https://www.scor.com/en/media/news-press-releases/scientific-seminar-about-climate-risks.

Klein, R.W., Wang, S., 2009. Catastrophe risk financing in the United States and the European Union: A comparative analysis of alternative regulatory approaches. The Journal of Risk and Insurance 76, 3, 607-637.

Klumpes, P., 2022. Consolidation and efficiency in the major European insurance markets. Working paper, Accounting Research Group, Aalborg University Business School.

KPMG, 2020. KPMG Zoom: Mergers and Acquisitions Trends – 2020. https://home.kpmg/tr/en/home/insights/2021/02/kpmg-zoom-mergers-acquisitions-trends-2020.html.

Mayer Brown. Global Insurance Industry Year in Review. Web reports 2013 to 2021. https://www.mayerbrown.com/en/information/global-insurance-industry-year-in-review.

Michel-Kerjan, E., 2012. How resilient is your country? Nature, 491, p. 497. https://www.nature.com/news/how-resilient-is-your-country-1.11861.

Mills, E., 2009. A global review of insurance industry responses to climate change. The Geneva Papers 34, 323-359.

Moeller, S., Schlingemann, F., 2005. Global diversification and bidder gains: A comparison between cross-border and domestic acquisitions. Journal of Banking and Finance 29, 533-564.

Munich re, 2023. Climate change and La Niña driving losses: the natural disaster figures for 2022. https://www.munichre.com/en/company/media-relations/media-information-and-corporate-news/media-information/2023/natural-disaster-figures-2022.html.

NAIC, 2013. The U.S. National state-based system of insurance financial regulation and the Solvency Modernization Initiative. NAIC White paper. https://content.naic.org/sites/default/files/inline-files/committees_e_us_solvency_framework.pdf.

Roberts, M.R., Whited, T.M., 2012. Endogeneity in empirical corporate finance. In: G. Constantinides (ed.), Handbook of the Economics of Finance. Vol. 2, Elsevier North-Holland, Amsterdam, 493-572.

Roll, R., 1986. The hubris hypothesis of corporate takeovers. Journal of Business 59, 197-216.

Rossi, S., Volpin, P. 2004. Cross country determinants of mergers and acquisitions. Journal of Financial Economics 74, 277-304.

SCOR, 2022. The Relevance of Climate Risk Change for Life and Health Insurance (Irene Merk, author), 10 pages, document internet. https://www.scor.com/en/expert-views/relevance-climate-change-life-insurance.

Seifert, D.L., Lindberg, D.L., 2012. Managing climate change risk: Insurers can lead the way. Risk, Hazards & Crisis in Public Policy 3, 2, article 5.

Sigma, 2015. Natural catastrophes and man-made disasters in 2014: Convective and winter storms generate most losses. https://www.swissre.com/institute/research/sigma-research/sigma-2015-02.html.

Snecdecor, G.W., Cochran, W., 1989. Statistical Methods, 8th Edition. Wiley-Blackwell.

Wagner, K., 2020. Adaptation and Adverse Selection in Markets for Natural Disaster Insurance. American Economic Journal: Economic Policy 14, 3, 380-421.

Wilcoxon, F., 1945. Individual comparisons by ranking methods. Biometrics Bulletin 1, 6, 80-83.

Appendix 1
Sources of capital in the US insurance industry

Table A1: Descriptive statistics, P&C insurance industry, 1990-2021

Variable in 10^{12} $	N	Mean	Std	Min	Median	Max	Data source
Total capital	32	0.76830	0.24772	0.36562	0.75035	1.30444	NAIC
Reinsurance demand ratio	32	0.46390	0.03451	0.40622	0.47744	0.50991	NAIC
Liquidity creation ratio	32	-0.51560	0.02989	-0.58240	-0.51357	-0.45720	NAIC
Direct premium written	32	0.60237	0.09591	0.47176	0.61269	0.79358	NAIC
Net premium written	32	0.55145	0.07755	0.44708	0.55084	0.71815	NAIC
Premiums earned	32	0.53979	0.07403	0.44336	0.53749	0.69036	NAIC
MA	32	29.8125	9.82242	16	29	61	SDC
Catastrophic losses	32	0.02769	0.02363	0.00439	0.01747	0.08644	VERISK
CAT and ILS issued	25	0.00632	0.00413	0.00133	0.00630	0.01400	Artemis
ILS issued	25	0.00062	0.00063	-0.00019	0.00041	0.00212	Artemis
CAT issued	25	0.00561	0.00359	0.00132	0.00566	0.01251	Artemis

Note: Annual values in 2021$.

Table A1 presents the data and their sources for the 1990-2021 period when there are 32 observations. The period is 1997-2021 otherwise.

Table A2: Sources of capital in the US P&C insured industry, 1997-2021 (all variables)

	With ILS		Without ILS	
Variable	Parameter	*t*	Parameter	*t*
Intercept	–2.71144**	–4.95	–2.63371**	–5.01
Reinsurance demand	2.16704*	2.47	2.12044*	2.52
Liquidity creation ratio	–3.06447**	–6.32	–2.97884**	–6.30
Post-2012	0.09938**	3.25	0.09807**	3.38
Premium earned	1.61666**	5.15	1.57369**	5.14
MA	–0.00168	–1.48	–0.00151	–1.36
Catastrophic losses	0.60532	1.29	0.62122	1.38
Catastrophe bonds and ILS	6.67811	1.39	–	–
Catastrophe bonds	–	–	8.87143	1.75
Number of observations		25		
R-squared	0.9639		0.9660	
R-squared adjusted	0.9491		0.9520	

*p<0.05; ** p<0.01

We observe in Table A2 that MA, catastrophe losses, and ILS are not statistically significant to explain the sources of capital in the P&C insurance industry. Reinsurance demand and Premium earned are important sources of capital.

Table A3: Sources of capital in the US P&C insurance industry, 1997-2021 (significant variables only)

	With ILS		Without ILS	
Variable	Parameter	*t*	Parameter	*t*
Intercept	–2.66866**	–4.70	–2.61057**	–4.82
Reinsurance demand	2.57214**	2.94	2.51303**	3.01
Liquidity creation ratio	–2.68013**	–5.96	–2.62595**	–6.03
Post-2012	0.08954**	2.87	0.08965**	3.02
Premium earned	1.46227**	5.29	1.44937**	5.44
Catastrophe bonds and ILS	9.44476*	1.99	–	–
Catastrophe bonds	–	–	11.59059*	2.33
Number of observations		25		
R-squared	0.9566		0.9592	
R-squared adjusted	0.9451		0.9484	

*p<0.10; ** p<0.01

Table A3 presents a robustness analysis of results of Table A12 when we drop non-significant variables. P&C insurers significantly increased their capital after 2012 (Post2012).

Collusion between Retailers and Customers: The Case of Insurance Fraud in Taiwan

Pierre Picard* Jennifer Wang† Kili C. Wang‡

August 7th, 2023

Abstract

We analyze how insurance distribution channels may affect fraud through claim manipulation, when car repairers may collude with policyholders. We focus attention on the Taiwan automobile insurance market with a database provided by two large Taiwanese automobile insurers. The theoretical underpinning of our analysis is provided by a model of claims fraud with collusion and audit. Our econometric analysis confirms the evidence of fraud through the postponing of claims to the end of the policy year, possibly by filing a single claim for several events. It highlights the role of car dealer agencies in the fraud process, and its change from 2010 to 2018.

*CREST-Ecole Polytechnique. Email: pierre.picard@polytechnique.edu. Pierre Picard acknowledges financial support from Investissements d'Avenir (ANR-121-IDEX-0003/Labex ECODEC).

†Department of Risk Management and Insurance, National Chenghi University, Taiwan. Email: jenwang@nccu.edu.tw.

‡Department of Insurance, Tamkang University, Taiwan. Email: kili@mail.tku.edu.tw.

1 Introduction

Vertical relationships frequently involve the outsourcing of services from upstream firms to downstream retailers. This may be at the origin of agency costs, associated with the discretion in the way retailers do their job. Such agency costs sometimes go through the collusion between retailers and customers, who exploit loopholes in the contracts between producers and customers. Discount fraud and warranty fraud are instances of such customer misbehaviors that involve collusion with retailers or frontline employees. Discount fraud exploits the special discounts that companies may offer under particular circumstances, for instance when discounted products are used for a specific purpose, e.g., educational use for softwares. Warranty fraud occurs especially when a service provider - e.g. a car repairer - replaces a defective part with a new spare part and triggers the producer's warranty, although the defective part was not original and thus was not protected by the warranty.[1]

This paper investigates another form of customer misbehavior facilitated by service providers acting on behalf of distributors: insurance fraud. Our empirical focus is on the Taiwan automobile insurance market and on the role of car dealer-owned insurance agents (DOAs) in this market. In such cases, dealers sell not only cars, but also automobile insurance to their clients, and furthermore they own an auto repair shop. Understandably, this multi-faceted activity and the long-term connection with car owners favor the creation of a mutually advantageous policyholder-DOA alliance. Concerning fraud itself, we will focus attention on two misbehaviors in the Taiwanese car insurance market. Firstly, the fact that policyholders may file small false claims by the end of the policy year if they have not receive any indemnity previously, a behavior highlighted by Li et al. (2013).

[1] See Harris and Daunt (2013) on managerial strategies under the risk of customer misbehavior. Murthy and Djamaludin (2002) survey the literature on new product warranty. Insufficient maintenance effort by buyers and inadequate behavior by retailers are at the origin of a double moral hazard problem in warranty management.

Recouping a part of the insurance premium paid to presumably unfair insurers may be the psychological motivation behind this behavior. Secondly, postponing claims to the last month of the policy year and, when possible, merging two losses in a single claim. Deductibles and the bonus-malus mechanism are the underlying reasons for this second type of misbehavior. Disentangling these two types of fraud will be one of our main challenges in what follows.

An insurance market model yields the theoretical underpinnings of our analysis. The model focuses on the strategic interaction between, on the one side, policyholders who file fraudulent claims by colluding with car repairers, and, on the other side, insurers who audit claims. Auditing claims is all the more costly when the collusion between policyholders and car repairers is more difficult to detect, which is particularly the case when car repairers are sheltered by DOAs. In addition, should irregularities be detected by the insurer, the bargaining power of DOAs may allow them to deter insurers from enforcing penalties. This suggests that there are potentially two reasons for which DOAs may facilitate insurance fraud: firstly, it may be hard for insurers to establish the truth because of the risk of collusion between DOAs and policyholders, and secondly, the bargaining power of DOAs may allow them not to be penalized when fraud is detected. The outcome is a higher fraud rate when insurance is distributed by DOAs than through other channels. As we will see below, this is reinforced in the case of deductible contracts, because deductibles increase the gain that policyholders obtain from fraud, and weaken the insurers' incentives to monitor claims.

Our empirical analysis draws on a database obtained from two large insurance companies in Taiwan. One of them, company 1, provided information on the policyholders who have filed an automobile claim in 2010 or 2018, and the other one, company 2, provided information on the policyholders who have filed an automobile claim in 2010. Company 1 relied heavily on DOAs to sell policies, although the market share of this distribution channel strongly decreased from 2010 to 2018, while company 2 never used DOAs.Starting with year 2010, our results confirm that there was more fraudulent claim manipulation when insurance policies were sold through DOAs than through other distribution channels, and also that deductibles stimulated fraud.[2] We also show that the causal mechanisms on which we focus (i.e., postponing claims, and possibly filing one claim for several accidents) were related to the bonus-malus system in force in Taiwan, and also to incentives inherent in the design of deductible contracts. This will go through an approach which consists of providing indirect evidence of such misbehaviors and of its mechanisms.[3] More explicitly, we show that, in 2010, the intertemporal pattern of claims was consistent with policyholder's fraudulent behavior favored by DOAs, after controlling for other explanations, including moral hazard and the money recouping behavior highlighted by Li et al. (2013).

However, things have changed dramatically from 2010 to 2018: DOAs were used less frequently by insurers, with presumably a lower bargaining power at the claim settlement stage. In other words, in 2018 it was more difficult for DOAs to collude with their customers at the expense of the insurer. As will be shown, the role of DOAs as facilitators of insurance fraud through claim manipulation vanished in 2018, in accordance with the decrease in their bargaining power.

The paper is organized as follows. Section 2 provides further motivation for our analy-

[2]Other authors have emphasized the effect of deductibles on insurance fraud. Using data from Québec, Dionne and Gagné (2001) show that the amount of the deductible is a significant determinant of the reported loss when no other vehicle is involved in the accident which led to the claim, and thus when the presence of witnesses is less likely. Miyazaki (2009) highlights through an experimental study that higher deductibles result in a weaker perception that claim padding is an unethical behavior, and thus in a larger propensity toward fraud.

[3]Although Dionne et al. (2009a) is an exception, it is usually very difficult to use direct information on fraudulent claims to analyze insurance fraud, either because identified fraud is just the top of the iceberg, or because of insurers'reluctance to share confidential information on any fraud they are victims of. The preferred approach consists in establishing indirect evidence of fraud. For instance, Dionne and Gagné (2002) and Dionne and Wang (2013) deduce the existence of fraud in automobile theft insurance from the time pattern of claims among the twelve policy months. Pao et al. (2014) provide evidence of opportunistic theft insurance fraud by analysing the claim pattern in areas hit by a typhoon.

sis. We introduce some factual observations that should convince the reader that there is claim manipulation in the Taiwanese car insurance market, and we describe regular fraud patterns. Section 3 develops a simple theoretical model of insurance fraud that shows how these patterns are linked to specific features of insurance contracts, particularly per-claim deductibles, and to the insurance distribution channel.[4] Section 4 describes the data in more detail, it presents our econometric approach, and discusses our results about claim manipulation. We particularly highlight the changes in the fraud pattern and in the role of DOAs from 2010 to 2018. Section 5 concludes.

2 Factual background

Our investigation will be based on information yielded by two large Taiwanese insurers, refered to as companies 1 and 2, about their automobile policyholders and their claims in 2010 and 2018. In 2010, company 1 sold approximately 37% of its automobile policies through DOAs, and this share dropped to about 20% in 2018. On the contrary, company 2 never sold insurance through the DOA channel.

Insurance agents, be they DOAs or standard agents, are in charge of handling claims. This frequently involves some bargaining between the insurer whose objective is to minimize the cost of claims, and the agents, who may favour their customers, particularly when they receive sales-based commissions. In this bargaining process, DOAs take advantage of the size of their activity, and of the fact that they own the list of their customers. In particular, an insurer who discovers a claim manipulation by a DOA may be reluctant to take retaliatory actions because of this strategic advantage of DOAs, who could switch to another insurer.[5] In the case of company 1, the bargaining power of DOAs is expected

[4] This section may be skipped by readers who are mostly interested in the empirical analysis of insurance fraud.

[5] On average, Taiwanese DOAs sell more policies than other agents: three times more on average, and much more for the largest DOAs. They are independent agents, and, as emphasized by Mayers and Smith (1981), this status gives them more discretion in claim administration (e.g. they may intercede

to have decreased from 2010 to 2018, because this insurer has dramatically reduced its dependence on DOAs. The specificity of DOAs has also an informational dimension, related to the fact that they work in partnership with car repairers, both being sheltered by car dealers. This multifaceted agency relationship creates an informational advantage: establishing that a claim has been falsified is particularly difficult and costly when it has been filed through a DOA.

Our study is also related to specific forms of insurance contract manipulation in Taiwan. Li et al. (2013) have observed that a large proportion of automobile insurance claims are filed during the last months of the policy year. This is confirmed by our own database. Figure 1 presents the distribution of claims and their average cost (in hundred US dollars) in 2010 over the twelve policy months, with a striking concentration of claims and a slight decrease in the claim cost in the last months of the policy year. Li et al. (2013) interpret this phenomena as a "premium recouping effect": some policyholders without accident during the previous months would tend to file small false claims during the last month of the policy year to express their feeling that they have been unfairly treated by the insurance company.

Figure 1

Some information about insurance contracts is useful for what follows. There are three different types of automobile physical damage insurance contracts in Taiwan: types A, B and C. Types A and B contracts cover all kinds of collision and non-collision losses, with more exclusions for B than for A,[6] while type-C contracts only cover the damages incurred in a collision involving two or more vehicles. Contracts also differ in terms of

on behalf of their customers at the claim settlement stage) because they can credibly threaten to switch their business from one insurer to another. Actually, DOAs provide comparatively less stable customers to company 1 than other insurance agents, with continuation rates (i.e. the fraction of customers who continue purchasing insurance from the same insurer year on year) which are about sixty percent for DOAs and seventy to eighty percent for other insurance agents.

[6]Type B contracts cover all the areas of type-A contracts, except the non-collision losses caused by intentional damage, vandalism, and any unidentified reasons.

indemnity: Type A contracts offer low coverage with a deductible, type B contracts may be purchased with or without a deductible, and type C contracts provide full coverage without a deductible. Claims are per accident, with a specific deductible for each claim. The change in premium is ruled by a bonus-malus system when policyholders renew their contracts with the same insurance company, with a no-claim discount and an increase in premium proportional to the number of claims, without reference to their severity. The policyholders who switch to another insurance company bargain with this company about the new starting point of their bonus-malus record

In this setting, opportunist policyholders may take advantage of manipulating claims for several reasons. Firstly, according to the premium recouping interpretation of Li et al. (2013), policyholders who wrongly pretend to have incurred some small losses in order to recoup part of their insurance premium are more likely to be among the policyholders who do not plan to keep a long term relationship with the same insurance company. Intuitively, such customers feel a lower moral cost of defrauding than those who intend to keep a long-term relationship with their insurer.[7] In our empirical analysis, this will lead us to define a *Recoup Group RG* that includes the policyholders who did not renew their contract more than one year after the policy year under consideration.[8] Secondly, for two reasons, insurance contracts may also incentivize opportunistic policyholders to manipulate claims corresponding to true accidents. Indeed, the claims filed during the last month of policy year t are not registered in the bonus-malus record of year $t+1$ (they will be taken into account in the premium paid in year $t+2$), and consequently, the policyholders who plan

[7] It is well known that insurance fraud is often associated with the feeling that the insurance company is unfair; see Fukukawa et al. (2007), Miyazaki (2009) and Tennyson (1997, 2002). The premium recouping phenomenon highlighted by Li et al. (2013) could reflect a kind of resentment against insurers, particularly from policyholders who have not filed a claim during the previous months the policy year.

[8] In Taiwan, filing a claim during the last month of the policy year does not affect the policyholder through the bonus-malus system if he/she does not stay more than one year with the same insurer. Our definition of the *Recoup Group* thus corresponds to policyholders without strong attachments to their current insurer, and for whom false claims filed toward the end of the policy year have no consequence through the bonus-malus system.

to renew their contract with the same insurer may see an advantage in postponing their claim to the last policy month, in order to delay the increase in premium.[9] In addition, since the bonus-malus record depends on the number of claims and not on their severity, policyholders may benefit from filing one single claim for two accidents, should a second accident occur. This is even more profitable in the case of deductible contracts, since deductibles are per-claim. In brief, because of the bonus-malus system and of deductible contracts, postponing the first claim and merging any other accident with the first one within a single claim is a winning strategy for opportunistic policyholders.[10]

Type A and B contracts are subject to such claim manipulation, because they include coverage for losses other than those associated with the collision between two cars. There is no third-party involved in such claims and no police report. On the other hand, the claims filed for type C contracts correspond only to collisions, and they have to include a police report, which makes manipulation very unlikely. In our empirical analysis of year 2010, the set of policyholders who renewed type A or B contracts in 2011, but not in 2012, with the same insurer will be called the *Suspicious Group SG* because of this maximum incentive to manipulate the bonus-malus system, with subgroups $SG1$ and $SG2$ for no-deductible and deductible contracts, respectively.[11] In 2018, all type A or B contracts included a deductible, and thus the distinction between $SG1$ and $SG2$ is no longer appropriate for this year.[12]

One of the key insights of our analysis will be about the role of DOAs in this fraudulent claim manipulation process in 2010. Figure 2 provides a preliminary idea of this role by considering how the type of contract and the sale process (DOA or standard insurance

[9] In addition, the bonus-malus system forgives the first accident for drivers who have had no other accidents for three years, which provides an even larger manipulation gain.

[10] Since 1996, the per-claim deductible is increasing with the number of claims, which strengthen even more the incentives to manipulate claims.

[11] The bonus-malus record has a new departure point when the policyholder switch insurers. Thus, by postponing their first claim to the last policy month, such policyholders were able to fully escape the consequence of this claim on their bonus-malus record.

[12] In other words, SG is 2018 corresponds to $SG2$ in 2010.

agents) have affected the time distribution of claims during the policy year. It is striking how the claim distribution during the last policy month is peaking at the end of the year for members of $SG1$ and $SG2$ who have purchased insurance through agents sheltered by car dealers. Comparing with type C contracts used as a benchmark without claim manipulation reinforces the intuition that DOAs played an important role in this fraud process.

Figure 2

It nevertheless remains that Figure 2 does not allow us to assess whether this timing favored by DOAs resulted from the manipulation of claims corresponding to actual losses or from the behavior consisting in filing a small false claim at the end of the policy year in order to recoup some money from the insurer. However, if a substantial number of claims filed in the last policy month correspond in fact to first claims that have been postponed, possibly with the cumulated losses of two events, then such claims should be more costly than average. In other words, we should expect that the ratio of "the average cost of first claims" over "the average cost of all claims" (hereafter called the *first claim cost ratio*) should increase during this month, contrary to the premium recouping interpretation of Li et al. (2013), hence a possible way of disentangling these two interpretations. But this may be misleading if the cost of claims is affected by an intertemporal moral hazard mechanism. Indeed, if a first accident makes drivers more cautious, then one may expect that subsequent accident would tend to be less severe, hence another possible explanation for an increase in the *first claim cost ratio* by the end of the policy year. To separate claim manipulation from moral hazard, we may consider type C contracts as a benchmark, since claim manipulation is vey unlikely for such contracts.

Figure 3

Figure 3 sustains the claim manipulation hypothesis for policyholders of the $SG2$ group who have purchased insurance through a DOA in 2010: their *first claim cost ratio* strongly increases in the last month of the policy year, and this is not the case for the other groups of policyholders. This suggests that in 2010 the claim manipulation mechanism dominated the premium recouping mechanism in $SG2$ (the subgroup of policyholders who benefit the most from claim manipulation), with DOAs acting as fraud facilitators, while the reverse occurs in the other subgroups. We also observe that for the RG group, the first claim cost ratio slightly increases when insurance has been purchased through the DOA channel, while it slightly decreases otherwise. This suggests that, among RG policyholders, the claim manipulation mechanism may be stronger than the premium recouping mechanism when insurance goes through DOAs. As we will see later, things have changed from 2010 to 2018.

3 Theoretical background

The model features the non-cooperative interaction between policyholders and insurers, in a costly state verification setting.[13] Consider a population of risk-averse drivers, whose type is defined by the couple (i, h) with $i \in \{D, A\}$ and $h \in \{1, 2\}$. Index i refers to the individuals' preference for a specific distribution channel through which they purchase insurance: DOA when $i = D$ or standard insurance agents when $i = A$.[14] Index h reflects the individual's degree of absolute risk aversion: $h = 1$ corresponds to a higher absolute risk aversion than $h = 2$. Assume that drivers may have either 0,1 or 2 accidents during

[13] See Picard (1996). For the sake of brevity, several aspects of the insurance market analysis are deliberately overlooked here. This particularly concerns the way individuals choose their contract and their insurance distribution channel, depending on their risk aversion and on their intinsic preference for a specific channel.

[14] For the sake of simplicity and brevity, we do not analyze the reasons for which an individual may prefer to purchase insurance through a car dealer or through another distribution channel. Such preferences are likely to depend on many factors, such as the valuation of time saved by bundling the purchase of a new car and the taking out of an insurance policy, the repeated relationship between individuals and car dealers that also provide repair services, or the level of trust in car dealers.

the same policy year, with probability π_1 and π_2 for 1 and 2 accidents, respectively, and $\pi_1 + \pi_2 < 1$, and also that these probabilities are independent of the policyholders' type. Insurance contracts include a deductible per accident. We respectively denote d_{ih} and P_{ih} the deductible and the premium of the contract chosen by type h individuals who purchase insurance through channel i. Less risk averse individuals choose a larger deductible, and thus we have $d_{i2} > d_{i1} \geq 0$.[15]

Each accident may be severe or minor, and the corresponding claims small or large, with probability q_s or $q_m = 1 - q_s$, respectively, irrespective of the policyholder's type, and whether it is the first or second accident during the policy year. To simplify our analysis of fraud through claim manipulation, its is assumed that a large claim exactly doubles a small claim, with loss ℓ and 2ℓ, respectively. Fraud is committed by policyholders who postpone small claims till their last policy month. They will file one single large claim for two minor accidents presented as a severe accident that occured during the last policy month, should another minor accident occur later during the same policy year. Otherwise, the claim corresponding to the first minor accident will be denied because filed outside the permitted time. Fraud reduces the retained cost of the accidents by half since the deductible is paid only once. It also provides a supplementary gain through the manipulation of the bonus-malus system if the policyholder intends to stay with the same insurer at least during the next year. Fraud requires collusion with a car repairer, the policyholder and the repairer sharing the benefits according to their respective bargaining powers. If they are spotted defrauding, they have to pay a penalty (considered, for simplicity, as a fine to the government), and, in that case, the claim is fully denied.

Let us denote by α_{ih} and β_{ih} the fraud and audit mixed strategy of the policyholder and the insurer, respectively, for a population of type (i, h) individuals. α_{ih} is the probability that a type (i, h) policyholder postpones a first small claim (when the corresponding

[15] For notational simplicity, we assume that the deductible is the same whether it is the first or second claim during the policy year.

minor accident occurs before the last policy month), with the intention to file a single large claim for two accidents during the last policy month, should another minor accident occur before the end of the year. Fraud is concentrated among those policyholders who are willing to stay with the same insurer at the end of the policy year because they are the ones who benefit the most through the bonus-malus mechanism.[16] β_{ih} is the probability that a large claim (filed by a type (i, h) policyholder) is audited by the insurer.[17] Such large claims correspond either to true severe accidents or to two minor accidents that have been fraudulently aggregated and postponed to the last month). We assume that audit allows the insurer to detect with certainty whether the claim has been manipulated or not.

The expected cost of claims per type (i, h) policyholder is written as

$$C_{ih} = L - D_{ih} + FC_{ih} + AC_{ih}, \tag{1}$$

where L is the expected costs of accidents, D_{ih} is the cost retained by the policyholder (in the absence of claim manipulation), FC_{ih} is the cost of claim manipulation for the insurer and AC_{ih} is the audit cost.

L and D_{ih} are equal to the expected number of accidents per policyholder $\pi_1 + 2\pi_2$ multiplied by the weighted average loss per accident and by the deductible per accident,

[16]The policyholders who may benefit the most from defrauding through claim manipulation are those who have a first minor accident before the last month of their policy year and who do not intend to switch insurers. If these policyholders are just indifferent between defrauding and not-defrauding, as will be the case at the equilibrium of the policyholder-insurer interaction game presented in the following analysis, then the other policyholders will be detered from defrauding.

[17]Note that the degree of risk aversion is not directly observed by the insurer. However, individuals choose different contracts (i.e., different deductibles) depending on their risk aversion, and thus insurers can condition their audit probability on the level of the deductible, and thus indirectly on the policyholder's type. Note also, that the policy year and the calendar year do not coincide. The beginning of the policy year is evenly distributed over the calendar year among the policyholders. Only the first claims that correspond to (true or falsified) severe accidents are audited. For practical reasons, it is assumed that insurers audit all these claims with the same probability, whether they are filed within or outside the last month of the policy year.

respectively. This gives

$$\begin{aligned} L &= (\pi_1 + 2\pi_2)[q_s\ell + 2q_m\ell] \\ &= (\pi_1 + 2\pi_2)(2 - q_s)\ell, \end{aligned}$$

and

$$D_{ih} = (\pi_1 + 2\pi_2)d_{ih}.$$

FC_{ih} is proportional to α_{ih} but, for given α_{ih}, it decreases linearly with β_{ih}, because auditing a larger fraction of large claims reduces average indemnity payment through the detection of falsified claims. DOAs have some bargaining power with insurers and they may intercede with the insurer when a claim is denied for fraud. This intervention is successful with some probability, and thus it decreases the financial benefit drawn by the insurer from spotting a defrauding policyholder-car repairer coalition. Thus, we may write

$$FC_{ih} = \alpha_{ih}[a_1(d_{ih}) - a_2(d_{ih}, \zeta_i)\beta_{ih}], \tag{2}$$

where $a_1(d_{ih})$ and $a_2(d_{ih}, \zeta_i)$ correspond to the expected cost of fraud (in the absence of audit), and to the expected gain from claim audit. We have $a_1' > 0$ and $a_{2d}' < 0$ because the larger the deductible, the larger the financial impact of claims falsification and the smaller the gain to the insurer when a claim is denied after audit. Furthermore, ζ_i is a parameter that measures the bargaining power of distribution channel i, with $\zeta_D > \zeta_A$.[18] We have $a_{2\zeta}' < 0$ because the distribution channel's bargaining power leads to a smaller insurer's

[18] Claim manipulation, as it is described, may be committed by policyholders who intend to renew their insurance policy and who have two accidents, the first one being minor and occurring before the last month of the policy year. Thus, $a_1(d_{ih})$ and $a_{2i}(d_{ih})$ depend on the probability that a type (i, h) individual is in this situation, which depends on π_1, π_2 and q_s, but also on the timing of accidents throughout the policy year, which is left undescribed for the sake of brevity.

expected benefit when fraud is detected.

DOAs own and control their repair shop. Thus, it is assumed that auditing a claim (i.e., spending resources to discover whether a claim has been manipulated or not) is more costly when insurance has been purchased through a DOA than through a standard insurance agent, because the protection of the DOA makes the detection of the policyholder-repairer collusion more difficult. We denote c_i the audit cost when the insurance distribution channel is $i = D$ or A, with $c_D > c_A$.

Since here fraud consists in filing one single large postponed claim for two accidents, the number of large claims filed by type (i, h) policyholders is linearly increasing with α_{ih}, which allows us to write[19]

$$AC_{ih} = c_i \beta_{ih}(a_3 + a_4 \alpha_{ih}). \tag{3}$$

The insurer chooses β_{ih} in $[0, 1]$ in order to minimize C_{ih} given by (1), which implies

$$\beta_{ih} \begin{cases} = 0 \text{ if } \alpha_{ih} < \alpha^*(d_{ih}, \zeta_i, c_i), \\ \in [0,1] \text{ if } \alpha_{ih} = \alpha^*(d_{ih}, \zeta_i, c_i), \\ = 1 \text{ if } \alpha_{ih} > \alpha^*(d_{ih}, \zeta_i, c_i), \end{cases} \tag{4}$$

where

$$\alpha^*(d, \zeta, c) \equiv \frac{c a_3}{a_2(d, \zeta) - c a_4}. \tag{5}$$

with $\alpha_d^{*\prime} > 0, \alpha_\zeta^{*\prime} > 0$ and $\alpha_c^{*\prime} > 0$. Let us assume that $\alpha^*(d, \zeta, c) < 1$ for the relevant values of d, ζ, c, which means that systematic fraud would trigger the auditing of all the large claims. Depending on the bribe that they have to pay to car repairers for them to collude

[19] Here also, a_3 and a_4 depend on π_1, π_2 and q_s (but not on d_{ih}), and furthermore a_4 depends on the timing of accidents throughout the policy year.

(which is not explicitly defined here)[20], on the fine imposed on spotted defrauders, and on their degree of risk aversion, type h policyholders are willing to defraud if the probability of being caught is smaller than a threshold $\beta_h^*(P_{ih}, d_{ih}, \zeta_i) \in (0,1)$. Individuals always defraud when the audit probability is zero, and they never defraud if all large claims are audited: hence the audit probability $\beta_h^*(P_{ih}, d_{ih}, \zeta_i)$ for which they are indifferent between fraud and honesty is between 0 and 1.[21] This audit probability threshold is type dependent (hence the subscript h in the β_h^* function) because it is affected by the intrinsic risk aversion of the policyholder, but it also depends on P_{ih} because an increase in premium may affect the policyholder's risk aversion through a wealth effect,[22] and it is increasing with d_{ih} because an increase in the deductible makes fraud more attractive. Furthermore, β_h^* is increasing with ζ_i because a larger bargaining power of the agent corresponds to a larger probability of avoiding the full cancellation of the insurance payout when a fraudulent claim is detected through an audit. Thus, we have

$$\alpha_{ih} \begin{cases} = 0 \text{ if } \beta_{ih} > \beta_h^*(P_{ih}, d_{ih}, \zeta_i), \\ \in [0,1] \text{ if } \beta_{ih} = \beta_h^*(P_{ih}, d_{ih}, \zeta_i), \\ = 1 \text{ if } \beta_{ih} < \beta_h^*(P_{ih}, d_{ih}, \zeta_i). \end{cases} \tag{6}$$

A type (i,h) policyholder who has a minor accident before the last policy month and her insurer play a non-cooperative game, with strategies α_{ih} and β_{ih} respectively. Its Nash equilibrium is easily characterized. If $\alpha_{ih} = 0$, then (4) gives $\beta_{ih} = 0$, which implies

[20] We may, for instance, assume that policyholders make take it or leave it offers to car repairers. The word "bribe" refers to any form of advantage that the car dealer-repairer firm may obtain from the arrangement with the policyholder, such as the guarantee of a future car purchase.

[21] β_h^* could be defined in a more explicit way by considering the expected utility of a type h individual who has a minor accident before the last policy month, and who has to choose between two strategies: either honestly filing a small claim without delay, or postponing her claim to the last policy month in order to file a single large claim if another minor accident occurs. β_h^* is the audit probability that makes the policyholder indifferent between these two strategies.

[22] For instance, under DARA preferences, an increase in the insurance premium makes the policyholder more risk averse, and thus less prone to conclude a risky fraudulent arrangement with a car repairer. In that case, the larger the insurance premium, the lower the audit probability threshold above which fraud is deterred.

$\alpha_{ih} = 1$ from (6), hence a contradiction. Similarly, if $\alpha_{ih} = 1$, then (4) gives $\beta_{ih} = 1$, which implies $\alpha_{ih} = 0$ from (6), hence again a contradiction. Thus, $\alpha_{ih} \in (0, 1)$ and (4),(6) give $\beta_{ih} = \beta_h^*(P_{ih}, d_{ih}, \zeta_i) \in (0, 1)$ and $\alpha_{ih} = \alpha^*(d_{ih}, \zeta_i, c_i) \in (0, 1)$.

In brief, at equilibrium, the audit probability $\beta_{ih} = \beta_h^*(P_{ih}, d_{ih}, \zeta_i)$ makes the policyholder indifferent between manipulation and honesty, and the manipulation probability $\alpha_{ih} = \alpha^*(d_{ih}, \zeta_i, c_i)$ makes the insurer indifferent between auditing and not-auditing.

This leads us to simple predictions about the effect of the type of contract and distribution channel on claim manipulation. Firstly, using $\alpha_d^{*\prime} > 0$ shows that higher deductibles go along with more manipulation. Since $d_2 > d_1 \geq 0$, we have $\alpha_{i2} > \alpha_{i1}$ for $i \in \{D, A\}$. In other words, for a given distribution channel, fraud is more prevalent among type 2 than type 1 individuals. More simply, if $d_1 = 0$, we can say in a shortcut that deductibles encourage fraud. Furthermore, using $c_D > c_A, \xi_D > \xi_A$, and $\alpha_\zeta^{*\prime} > 0, \alpha_c^{*\prime} > 0$ yields $\alpha_{Dh} > \alpha_{Ah}$ for $i \in \{1, 2\}$. Put briefly, for a given type of individual, there is more fraud when insurance has been purchased through the DOA agents than through standard insurance agents, either because it is more costly to audit a claim that goes through a DOA or because DOAs have a larger bargaining power than standard insurance agents.

4 Data and testing of hypotheses

4.1 The data

The data yielded by Companies 1 and 2 provides detailed information about the policyholders, their insurance contracts and the claims they have filed. Available variables are listed in Table 1. Data was collected over the 2010-2012 and 2018-2020 periods. However, our analysis of insurance claims will be restricted to 2010 and 2018, in order to know whether policyholders subsequently renewed their contracts for less or more than one year.[23] We will start by considering year 2010 in sections 4.2 and 4.3. As previously mentioned, Company 1 has strongly reduced its dependence on DOAs from 2010 to 2018, and thus in section 4.4 comparing results obtained for years 2010 and 2018 will allow us to appraise the consequence of this strutural change.

We target the owners of private usage small sedans and small trucks with type A, B or C insurance contracts for automobile physical damage. In 2010, there was $121,952$ policyholders in the sample, and 8.10% of them filed at least one claim, which corresponds to $9,874$ observations. This subset defines our "research sample", i.e. the sub-sample of policyholders with claims.

Tables 1 and 2

The mean values of the variables in the two samples are displayed in the first two columns of Table 2, with some significant differences. In particular, the percentages of type A or B contracts, and particularly those in the suspicious groups $SG1$ and $SG2$, are much larger in the research sample. The three other columns in Table 2 separate the research sample into three subgroups, according to the insurance distribution channels (DOA in Company 1 and non-DOA in Companies 1 and 2), with significant differences in terms of gender, usage, and vehicle size. There is also a much larger proportion of new vehicles for the DOA channel, which reflects the fact that, most of the time, a DOA sells an insurance contract when the corresponding dealer sells a new car. The percentage of claims filed during the last month of the policy year, measured by the average value of dummy SC, and the share of the RG group are larger in the DOA channel than in the two other channels.

[23] In what follows, years are policy years: a contract corresponds to year 2010 if it started in 2010.

4.2 Evidence on claim manipulation

Our first step consists in testing whether in 2010 the perspective of a one-year contract renewal and the choice of a deductible contract stimulate insurance fraud by postponing claims to the last policy month, called the "suspicious period", possibly by filing one claim for two events. In other words, we wonder whether belonging to the Suspicious Group SG, and particularly subgroup $SG2$, is a factor that has stimulated insurance fraud through claim manipulation. Defining the fraud rate as the number of claims per policyholder filed during the suspicious period[24] leads us to formulate the following hypothesis.

Hypothesis 1 (H1): *The fraud rate tends to be higher in the suspicious group than in the non-suspicious group, and this is particularly the case for individuals covered by deductible contracts.*

Testing **H1** amounts to identifying whether there is a conditional dependence between belonging to the suspicious group and filing a claim within the suspicious period, respectively associated with dummies SG (or $SG1$ and $SG2$ for each subgroup) and SC. We do so through the following three Bivariate Probit models, where $\Phi(.)$ is the cumulative normal distribution function, and X is the vector of explanatory variables (with vectors of coefficients $\beta_{SC}, \beta_{SG}, ...$), including the premium amount and all the variables used in pricing and underwriting decisions.[25] In order to control for the recouping effect, dummy RG is also included in X.

Model 1:

$$\text{Prob}(SC = 1) = \Phi(X\beta_{SC} + \varepsilon) \tag{7}$$

$$\text{Prob}(SG = 1) = \Phi(X\beta_{SG} + \eta) \tag{8}$$

[24] Of course, this definition of the fraud rate does not mean that all claims filed during the suspicious period have been fraudulently manipulated.

[25] This includes all the observable characteristics of the insured (e.g., age, gender, bonus-malus coefficient, premium, etc...), the characteristics of the vehicle (e.g., age, brand, registered area, etc...) and recoup dummy RG. Hence, X includes all the variables listed in the first part of Table 1, and *logprem* and RG in the second part.

Model 2:

$$\text{Prob}(SC = 1) = \Phi(X\beta_{SC} + \varepsilon) \tag{9}$$

$$\text{Prob}(SG1 = 1) = \Phi(X\beta_{SG1} + \eta) \tag{10}$$

Model 3:

$$\text{Prob}(SC = 1) = \Phi(X\beta_{SC} + \varepsilon) \tag{11}$$

$$\text{Prob}(SG2 = 1) = \Phi(X\beta_{SG2} + \eta) \tag{12}$$

The results of these regressions are presented in Table 3, with a special interest in the residual correlation coefficient ρ. **H1** should lead to a positive conditional correlation between filing a suspicious claim and belonging to a suspicious group. More formally, the estimated residual correlation coefficients of these models $\widehat{\rho}_{SC,SG}, \widehat{\rho}_{SC,SG1}$ and $\widehat{\rho}_{SC,SG2}$ should be positive and significantly different from 0, which leads us to test for the null hypothesis $H_0 : \rho_{SC,SG} \leq 0, H_0 : \rho_{SC,SG1} \leq 0$ and $H_0 : \rho_{SC,SG2} \leq 0$, in models 1, 2 and 3 respectively.

The three estimated residual correlation coefficients are significantly positive, which allows us to reject the null hypothis in each model, and thus to state that there is a significantly positive conditional correlation between SC and $SG, SG1$ or $SG2$, in each model. In other words, in accordance with **H1**, there is a conditional dependence between belonging to the suspicious group and filing a claim within the suspicious period, whether the individual is covered by a deductible contract or not.[26]

[26] Table 3 also offers some interesting byproducts that are worth mentioning. Firstly, the policyholders from the RG group tend to file their first claims in the suspicious period, which echoes the conclusions of Li et al. (2013). Secondly, females file their first claim during the suspicious period more frequently than males, but that does not necessarily reflect a gender effect in fraudulent behavior. It is usual in Taiwan to register cars under the name of females (e.g. a wife or mother), even when the primary driver is a male, in order to benefit from cheaper insurance premiums. Hence, instead of a gender effect, the above mentioned correlation may just reflect the fact that the policyholders who carefully manage their insurance budget may also try to obtain undue advantage from their insurance company.

Table 3

When manipulation consists in postponing claims to the suspicious period, by cumulating several losses in a single claim when possible (which differs from small claims filed by the end of the policy year to recoup a part of the insurance premium), then the suspicious period should be characterized by high values of the first-claim cost ratio. This is expressed in Hypothesis 2.

Hypothesis 2: *In the suspicious group, the first-claim cost ratio is larger in the suspicious period than during the rest of the policy year.*

Hypothesis 2 is tested through the following regression:

$$clmamt = \alpha_c SC + \alpha_f first + \alpha_{fs} first * SC + \alpha_X X + e, \tag{13}$$

which is performed among the claims filed by members of $SG1$ and $SG2$ groups. This corresponds to 6,974 claims filed by 6,521 policyholders from SG1, and 695 claims filed by 647 policyholders from SG2. In these regressions, $clmamt$ is the value of the claim (in US dollars), while SC and $first$ are dummies indicating respectively that the claim was suspicious (i.e., it was filed during the last month of the policy year), and that it was the first claim of the policyholder during this policy year. Regression (13) also includes the interaction variable $first * SC$. Results are reported in Table 4.

Table 4

The estimated coefficients of the interaction variable are $\widehat{\alpha}_{fs} = -113.3$ with $p-$value 0.1627 for $SG1$, and $\widehat{\alpha}_{fs} = 1465.7$ with $p-$value lower than 0.0001 for $SG2$. This sustains Hypothesis 2 for $SG2$, but not for $SG1$, which confirms the fact that being covered by a deductible contract is a factor that stimulates fraud through claim manipulation. Hypothesis 3 focuses attention on the role of DOAs in this type of insurance fraud.

Hypothesis 3 (H3): *The fraud rate in the suspicious group is larger when insurance has been purchased through the DOA channel than through other distribution channels.*

We test **H3** by testing Bivariate Probit models 1, 2 and 3 in sub-samples that include the policyholders who purchased insurance through the DOA channel or through other distribution channels. This leads us to estimated residual correlation coefficients $\widehat{\rho}_{SC,SG}, \widehat{\rho}_{SC,SG1}$ and $\widehat{\rho}_{SC,SG2}$ in each subsample.

Tables 5,6 and 7

Detailed results are displayed in Tables 5, 6 and 7, for models 1, 2 and 3 respectively, with conclusions on residual correlation summarized as follows:[27]

		Company 1 Dealer	**Company 1 Non-dealer**	**Company 2**
Model 1	$\widehat{\rho}_{SC,SG}$	0.5393***	0.1344	0.0562
Model 2	$\widehat{\rho}_{SC,SG1}$	0.5729***	0.0916	−0.0610
Model 3	$\widehat{\rho}_{SC,SG2}$	0.7492***	−0.2020	0.2076***

Hence, when the regressions are performed in the sub-sample of policyholders who purchased coverage through the DOAs of Company 1, there is a significant positive residual correlation between SC and $SG, SG1$ or $SG2$ at the 1% threshold. This correlation vanishes in the two other sub-samples, except between SC and $SG2$ in Company 2.

4.3 Complements on the role of car dealers in claim manipulation

The previous conclusions may be reinforced by testing whether $\widehat{\rho}_{SC,SG}, \widehat{\rho}_{SC,SG1}$ and $\widehat{\rho}_{SC,SG2}$ are significantly larger among the policyholders who purchased insurance through car dealers than through other channels. To do so, we successively consider the two null

[27] *** refers to significance level at the 1% threshold.

hypotheses $H_0 : \widehat{\rho}^{D}_{SC,SG} \leq \widehat{\rho}^{ND}_{SC,SG}$ and $H_0 : \widehat{\rho}^{D}_{SC,SG} \leq \widehat{\rho}^{C2}_{SC,SG}$ in model 1, where D, ND and $C2$ refer respectively to insurance purchased from Company 1 through dealers, from Company 1 through other distribution channels, and from Company 2. We proceed in the same way for models 2 and 3, hence with $SG1$ and $SG2$ instead of SG. Results are displayed in Table 8. The two null hypotheses $\widehat{\rho}^{D}_{SC,SG} \leq \widehat{\rho}^{ND}_{SC,SG}$ and $\widehat{\rho}^{D}_{SC,SG} \leq \widehat{\rho}^{C2}_{SC,SG}$ are rejected at 1% significance level, and the conclusion is unchanged for $SG1$ and $SG2$. In other words, whatever the definition of the suspicious group ($SG, SG1$ or $SG2$), the conditional correlation between filing a suspicious claim and belonging to the suspicious group is significantly larger when contracts are sold through the car dealer associated with company 1 than through another distribution channel of company 1 or from company 2.

Table 8

Further evidence on the role of car dealers may be obtained by focusing attention on the first-claim cost ratio during the suspicious period (as in Hypothesis 2) by considering subsamples defined by the distribution channel, and by using type C contracts as a benchmark. A first-claim cost ratio during the suspicious period larger for $SG1$ or $SG2$ than for type C contracts would signal claim manipulation by members of the suspicious groups. Symmetrically, a lower first-claim cost ratio would be compatible with the premium recouping mechanism highlighted by Li et al. (2013), with small claims filed at the end of the policy year if no claim has been filed before. This leads us to consider regression (14) below, where the claim amount is the dependent variable as in regression (13). In (14), $first, SC$ and X are identical to those in regression (13), and S_1, S_2 and S_3 are dummies indicating that the policy has been purchased from Company 1 through the DOA channel, from Company 1 through another distribution channel and from Company 2, respectively. Furthermore C is a dummy indicating that the insurance policy is a type C contract, used as a benchmark without claim manipulation.

$$
\begin{aligned}
clmamt &= \alpha_c SC + \alpha_f first + \alpha_{fs} first * SC \\
&\quad + s_{SG11fs} SG1 * S_1 * first * SC \\
&\quad + s_{SG21fs} SG2 * S_1 * first * SC \\
&\quad + s_{SG23fs} SG2 * S_3 * first * SC \\
&\quad + s_{Cfs} C * first * SC + \alpha_X X + e. \qquad (14)
\end{aligned}
$$

The estimation of regression (14) shows that the null hypothesis $H_0 : s_{SG21fs} \leq s_{Cfs}$ is rejected at 1% significance level, contrary to the results obtained when s_{SG11fs} and s_{SG23fs} are compared to s_{Cfs}.[28] This means that the first-claim cost ratio is significantly higher during the last policy month when a deductible contract has been purchased from Company 1 through the DOA channel. All in all, in 2010 deductible contracts sold through DOAs have created the most favorable condition for insurance fraud through the postponing and aggregation of claims.

4.4 Smaller bargaining power for DOAs in 2018

From 2010 to 2018, Company 1 has cut almost by half the share of its automobile insurance contracts sold through car dealers. The latters became less important partners of the insurer, with presumably a lower bargaining power in the claim settlement process.[29]

To assess the consequences of this change, we have collected information about 269,475 automobile insurance contracts of type A, B and C, sold by Company 1 in 2018. The content of these contracts basically remained the same as in 2010, the only important

[28] Detailed estimation results are available from the authors upon request.

[29] In 2011, the largest car dealer group in Taiwan created its own insurance company, which induced other insurers to gradualy redirect a substantial part of their business toward other distribution channels. In the case of Company 1, the market share of DOAs decreased from 36.92% to 27.95%.

change being that in 2018 Company 1 only sold type A and B contracts with a deductible. Therefore, the Suspicious Group SG has no longer to be splitted between $SG1$ or $SG2$, and it coincides with what we called $SG2$ for year 2010. Table A1 in Appendix provides detailed information about the data. Comparing Tables 1 and A1 confirms the decrease in the proportion of contracts sold through DOAs, and other important changes including the decrease from 6.16% to 3.71% in the proportion of policyholders who filed a claim in 2010 and 2018, respectively. Figure A1 also confirms that claim rates are still higher during the last month than during the previous months of the policy year, with a large decrease in the average claim cost during the last policy month, and Figure A2 shows a decrease in the first claim cost ratio for all types of contracts, including those in SG going through DOAs, contrary to what was observed for $SG2$ in 2010.

Does this mean that the claim manipulation favored by DOAs has vanished in 2018? Formal tests have been performed to find out for sure. The results of Bivariate Probit regressions (similar to Model 1 above) are presented in Table A2. The estimated residual correlation between SG and SC is still significantly positive whatever the distribution channel, but the null hypothesis $H_0 : \rho^D \geq \rho^{ND}$ is rejected at the 1% significance threshold. In other words, the positive residual correlation between belonging to the suspicious group and filing a claim in the suspicious period still holds, which confirms claim manipulation, but the role of DOAs in this fraud process has vanished. For the sake of completeness, we have checked that the difference $\rho^D - \rho^{ND}$ has significantly decreased between 2010 and 2018, which means that the higher conditional correlation between SC and SG for the DOA channel, by comparison with other distribution channels, has significantly decreased from 2010 to 2018.[30]

We have performed a robustness check by a two-stage method in order to confirm this change from 2010 to 2018. To do so, we have created a new data set that includes $SG2$

[30] In more precise terms, the null hypothesis $H_0 : (\rho^D - \rho^{ND})_{2018} - (\rho^D - \rho^{ND})_{2010} \geq 0$ can be rejected at the 1% significance level.

and type C contracts sold by Company 1 in 2010 or 2018, with dummy y_{2018} used to indicate that the contract has been sold in 2018.[31] The first stage consists in estimating the following Probit regression:

$$\Pr[SG = 1] = \Phi(X\beta_{SG} + \eta),$$

and the estimated probability of belonging to the Suspicious Group $\widehat{SG}$ and dummy D for the DOA channel are used as explanatory variables in the second-stage regression:

$$\begin{aligned}\Pr[SC &= 1] = \Phi(\beta_{estSG}\widehat{SG} + \beta_{SG}SG + \beta_{D}D + \beta_{2018}y_{2018} \\ &+\beta_{SGD}SG * D + \beta_{SG2018}SG * y_{2018} + \beta_{D2018}D * y_{2018} \\ &+\beta_{SGD2018}SG * D * y_{2018} + X\beta_{SC} + \varepsilon)\end{aligned}$$

Results are presented in Table A3. The estimated coefficient of the triple interaction term $SG*D*y_{2018}$ is $\widehat{\beta}_{SGD2018} = -1.7265$, and it is significantly different from 0 at the 1% significance threshold. In other words, the stimulation effect of DOAs on the manipulation of claims by policyholders from the suspicious group SG has significantly decreased from 2010 to 2018.

Considering that DOAs played a crucial role in the manipulation of claims in 2010, one may wonder whether the decrease in their bargaining power has fully cancelled the fraud process in 2018, be it under the form of claim manipulation or of the premium recouping behavior. To get an idea, we have estimated regression (13) for the SG group and for the type C contracts, with the data of year 2018. Results are presented in Table A4. The estimated coefficient $\widehat{\alpha}_{fs}$ is not significantly different from 0 in the two subsets of contracts. In other words, in 2018, contrary to what occured in 2010, there was no significant change in the average amount of the claims filed during the last month of the

[31] $SG2$ and SG coincide in 2018 since Company 1 has stopped selling type A or B contracts without deductible.

policy year by comparison with previous months.

5 Conclusion

The purpose of this paper was to analyze some aspects of the policyholder-service provider coalition in insurance fraud mechanisms: how it can affect the credibility of claim auditing, and how fraudulent claim manipulation may emerge. It is a fact that the economic analysis of insurance fraud is often based on a very abstract picture of claim fraud (filing a fraudulent claim although no accident has occurred, or exagerating a claim), but in practice understanding insurance fraud often requires a much more specific analysis of the claims fraud process. The Taiwan car insurance case offers such a possibility, with fraud also taking place through the manipulation of the claim's date in order to avoid a penalty from the bonus-malus system and to reduce the burden of a second deductible, should another accident occur. The policyholders with deductible contract who intend to renew their policies (the suspicious group) have a larger propensity to defraud in that way than other policyholders

Our main focus was on the role of DOAs in this fraud process, with two specificities for this distribution channel. Firstly, the collusion between car repairers and policyholders is easier when insurance agents and car repairers are sheltered by a car dealer, and establishing claim manipulation unambiguously is more costly (i.e. the audit cost is larger) in that case. Secondly, DOAs may more easily escape penalties when fraud is detected (i.e. their bargaining power is larger at the claim settlement stage) because they can retaliate by redirecting their customers toward other insurers if the relationship with the current insurer deteriorates. Both specificities are related to the multi-faceted activities of DOAs: they sell insurance contracts, but they also work hand in hand with car repairers and car dealers. The comparison between years 2010 and 2018 suggests that reducing the dependence on car dealers has allowed Company 1 to deter claim manipulation more efficiently, because of the decrease in the bargaining power of DOAs. In other words, the role of DOAs in car insurance fraud seems to be much more related to their bargaining power, than to the difficulty for the insurer to establish that claims had been manipulated.

References

Chiappori, P.A. and Salanié, B. (2000), "Testing for asymmetric information in insurance markets", *Journal of Political Economy*, 108(1): 56–78.

Dionne, G., and Gagné, R. (2001), "Deductible contracts against fraudulent claims: evidence from automobile insurance", *Review of Economics and Statistics*, 83, 290–301.

Dionne, G. and R. Gagné R. (2002), "Replacement cost endorsement and opportunistic fraud in automobile insurance", *Journal of Risk and Uncertainty*, 24: 213-230.

Dionne, G., F. Giuliano and P. Picard (2009a), "Optimal auditing with scoring: theory and application to insurance fraud", *Management Science*, 55, 58-70.

Dionne, G., C. Gouriéroux and C. Vanasse (1997), "The informational content of household decisions with application to insurance under adverse selection", *Manuscript*. Montréal: Ecole Hautes Etudes Commerciales.

Dionne, G., C. Gouriéroux and C. Vanasse (2001), "Testing for evidence of adverse selection in the automobile insurance market : a comment", *Journal of Political Economy*, 109(2) : 444-453

Dionne G., P. St-Amour and D. Vencatachellum (2009b), "Asymmetric information and adverse selection in Mauritian slave auctions", *Review of Economic Studies*, 76 (4): 1269-1295.

Dionne, G., M. La Haye and A. S. Bergeres (2015) "Does asymmetric information affect the premium in mergers and acquisitions?", *Canadian Journal of Economics*, 48(3), 819-852.

Dionne, G. and Wang, K. C. (2013), "Does opportunistic fraud in automobile theft insurance fluctuate with the business cycle?", *Journal of Risk and Uncertainty*, 47, 67-92.

Fukukawa, K., C. Ennew and S. Diacon (2007), "An eye for an eye : investigating the impact of consumer perception of corporate unfairness on aberrant consumer behavior", in *Insurance Ethics for a more Ethical World (Research in Etical Issues in Organizations)*, edited by P. Flanagan, P. Primeaux and W. Ferguson, Vol.7, 187-221, Emerald Group Publishing Ltd.

Gouriéroux, C., A. Monfort (1995), "Statistics and econometric models", *Cambridge University Press*, 2: 458–475.

Harris, L.C. and K. Daunt (2013), "Managing customer misbehavior: Challenges and strategies", *Journal of Services Marketing*, 27(4): 281-293.

Li, C-S., C-C. Liu and S-C. Peng (2013), "Expiration dates in automobile insurance contracts : The curious case of last policy month claims in Taiwan", *Geneva Risk and Insurance Review*, 38(1): 23-47.

Mayer, D. and Jr. C. S. Smith (1981), "Contractual provisions, organizational structure, and conflict control in insurance markets", *Journal of Business*, 54: 407-434.

Miyazaki, A.D. (2009), "Perceived ethicality of insurance claim fraud : do higher deductibles lead to lower ethical standards", *Journal of Business Ethics*, 87(4):589-598.

Murthy, D.N.P. and I. Djamaludin (2002), "New product warranty: a literature review", *International Journal of Production Economics*, 79, 231-260.

Pao, T.I., L.Y. Tzeng and K.C. Wang (2014), "Typhoons and opportunistic fraud: claim patterns of automobile theft insurance in Taiwan", *Journal of Risk and Insurance*, 81(1): 91-112.

Picard, P. (1996), "Auditing claims in insurance market with fraud: the credibility issue", *Journal of Public Economics*, 63: 27-56.

Tennyson, S. (1997), "Economic institutions and individual ethics : a study of consumer attitudes toward insurance fraud", *Journal of Economics Behavior and Organization*, 32:247-265.

Tennyson, S. (2002), "Insurance experience and consumers'attitudes toward insurance fraud", *Journal of Insurance Regulation*, 21(2):35-55.

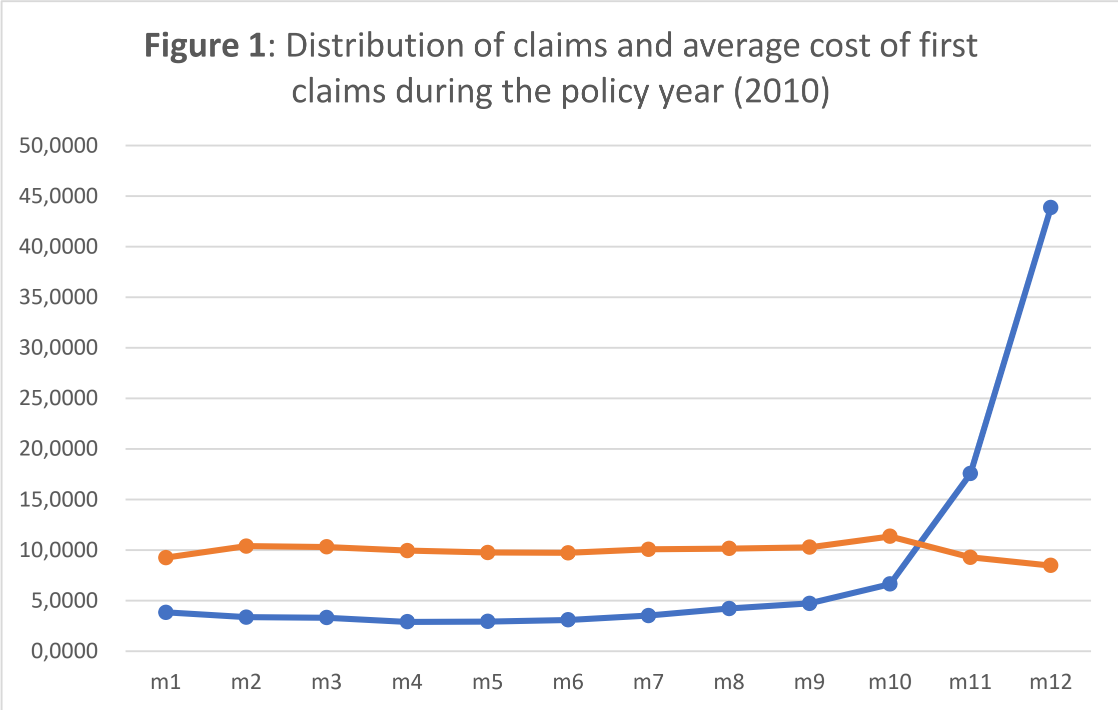

Distribution of claims (unit : %) :

Average cost of first claims (unit : USD100) :

Figure 2: Distribution of claims during the policy year (2010) for *SG1*,*SG2* according to the sale process, with type C contracts as benchmark.

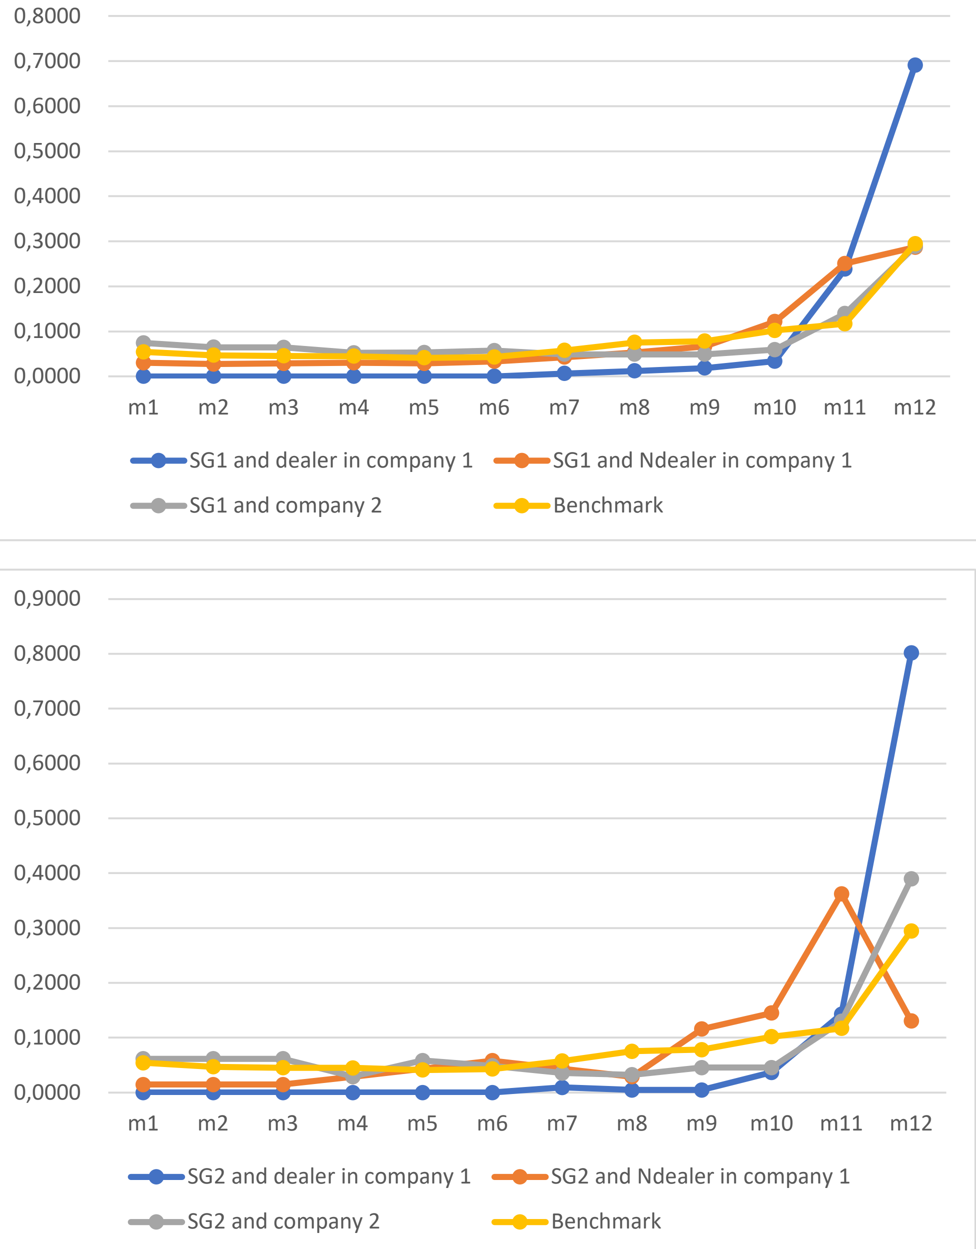

Qualifiers « dealer » and « Ndealer » refer to the cases where the insurance policy has been sold through a DOA or through a standard agent, respectively.

Figure 3: Changes in the First Claim Cost Ratio from months 1-11 to month 12 of the policy year (2010)

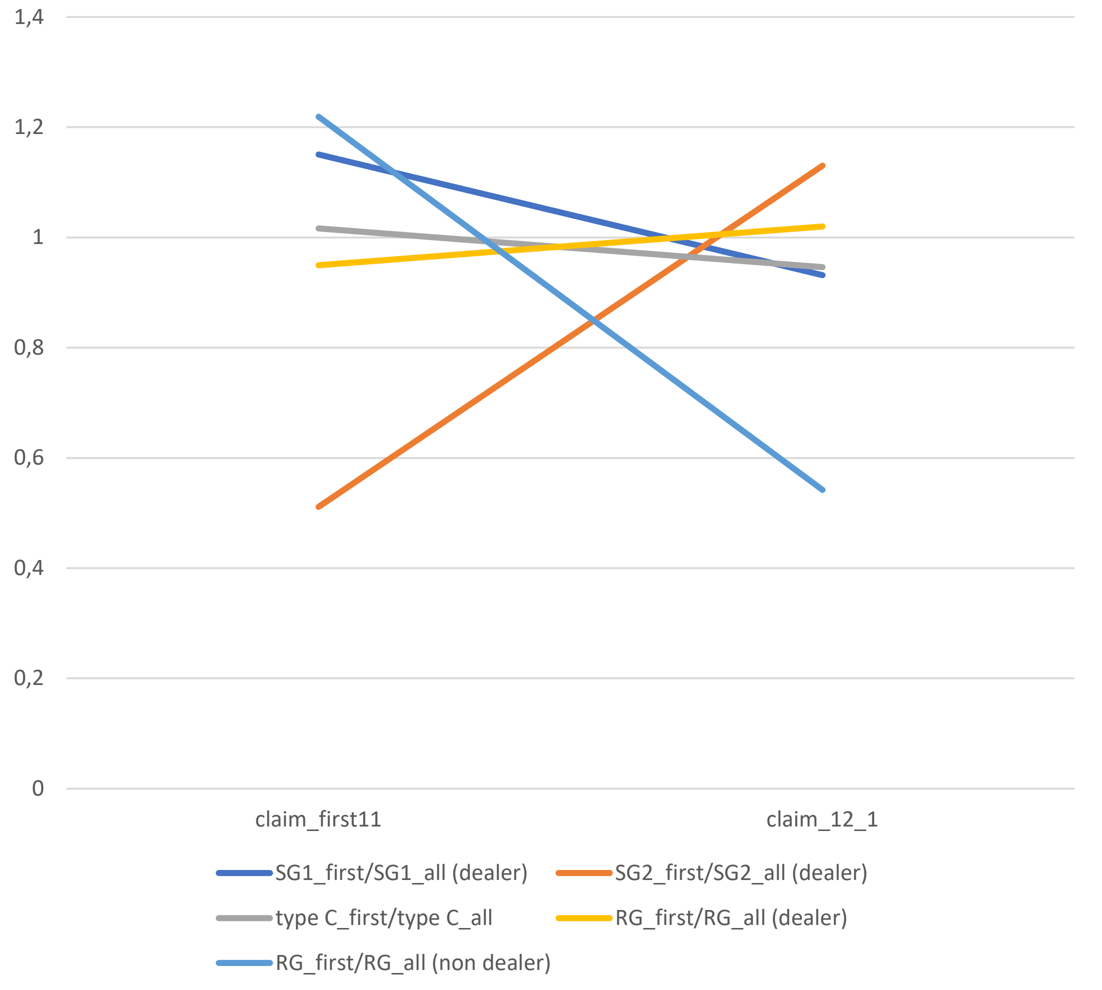

Figure A1: Distribution of claims and average cost of first claims during the policy year (2018)

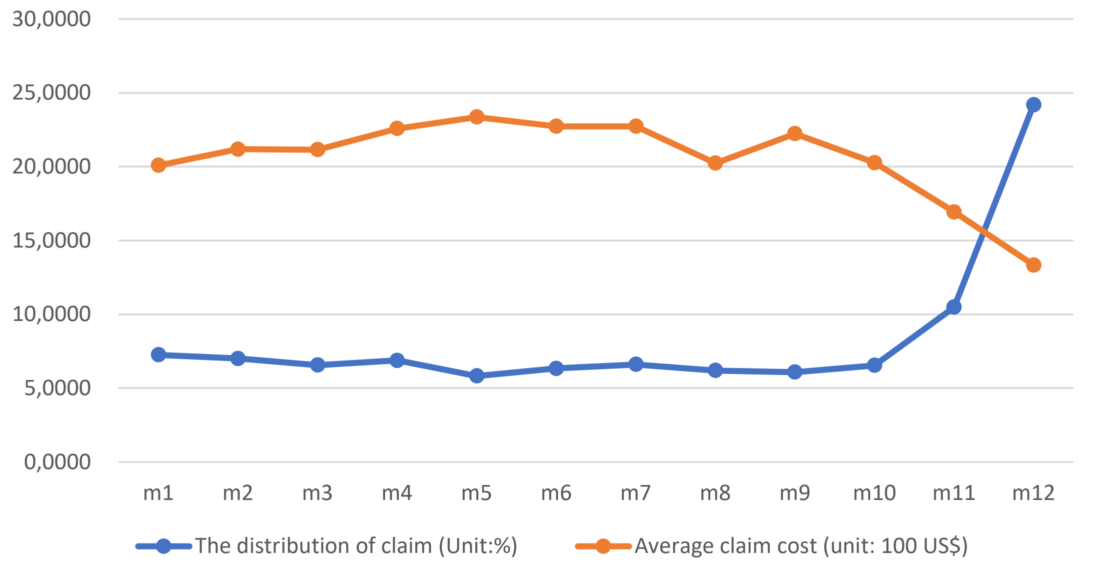

Figure A2: Changes in the First Claim Cost Ratio from months 1-11 to month 12 of the policy year (2018)

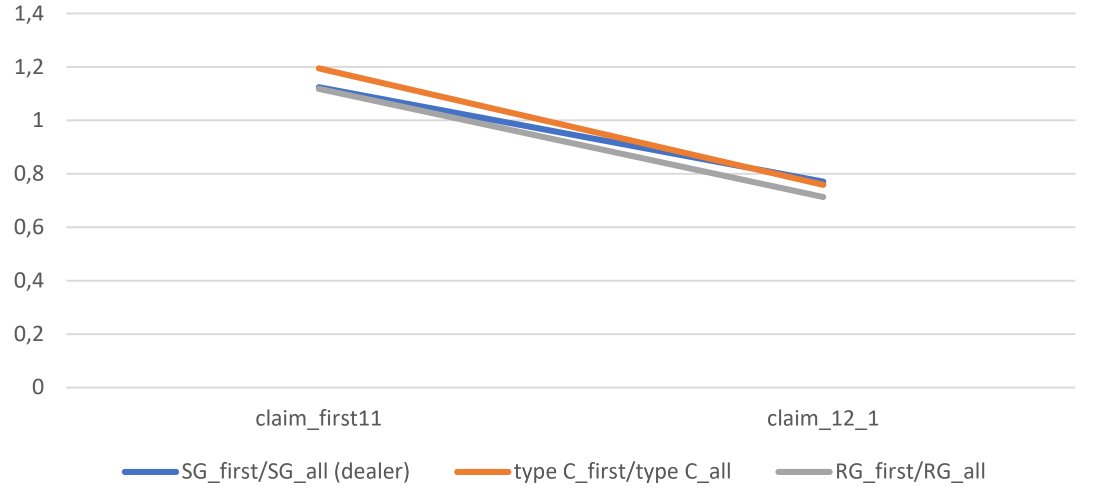

Table 1: Definition of variables

Variable	Definition
Explained variables:	
claim	Dummy variable equal to 1 when the insured has filed at least one claim during the policy year, 0 otherwise.
SC	Dummy variable equal to 1 when the insured has filed his or her first claim during the suspicious period (in the last policy month), 0 otherwise.
SG	Dummy variable equal to 1 when the insured belongs to the "suspicious group",[1] and 0 otherwise.
SG1	Dummy variable equal to 1 when the insured belongs to "suspicious group 1",[2] and 0 otherwise.
SG2	Dummy variable equal to 1 when the insured belongs to "suspicious group 2",[3] and 0 otherwise.
Explanatory variables:	
first group (underwriting and pricing variables)	
female	Dummy variable equal to 1 if the insured is a female, 0 otherwise.
age2025	Dummy variable equal to 1 if the insured is in the 20-25 age group, 0 otherwise.
age2530	Dummy variable equal to 1 if the insured is in the 25-30 age group, 0 otherwise.
age3060	Dummy variable equal to 1 if the insured is in the 30-60 age group, 0 otherwise.
ageabv60	Dummy variable equal to 1 if the insured is older than 60, 0 otherwise.
carage0	Dummy variable equal to 1 when the car is less than one year old, 0 otherwise.
carage1	Dummy variable equal to 1 when the car is two years old, 0 otherwise.
carage2	Dummy variable equal to 1 when the car is three years old, 0 otherwise.
carage3	Dummy variable equal to 1 when the car is four years old, 0 otherwise.
carage4	Dummy variable equal to 1 when the car is five years old, 0 otherwise.
veh_m	Dummy variable equal to 1 when the capacity of the insured car is between 1800 and 2000 c.c., 0 otherwise.
veh_l	Dummy variable equal to 1 when the capacity of the insured car is larger than 2000, 0 otherwise.
sedan	Dummy variable equal to 1 when the car is a sedan and is for non-commercial

[1] The "suspicious group" (*SG*) includes the individuals who renew their contract with the same insurance company for only one year. The counter group for SG includes the policyholders who do not renew their contract, or renew their contract for more than one year with the same insurance company.
[2] The "suspicious group 1" (*SG*1) are the *SG*-group policyholders who also purchased the no-deductible contracts. The counter group for *SG1* includes the policyholders with deductible contract or who are not in *SG*-group.
[3] The "suspicious group 2" (*SG*2) are the *SG*-group policyholders who also purchased the deductible contracts. The counter group for *SG2* includes the policyholders with no-deductible contract or who are not in *SG*-group.

	or for long-term rental purposes, and 0 otherwise.[4]
bonus	Bonus-malus coefficient used to calculate the premium in the current contract year. It is a multiplier on the premium. Hence, it is a discount if it is smaller than 1 and it is a penalty if it is larger than 1.
tramak_j	Dummy variable equal to 1 when the brand of the insured car is *j*, with *j*=*n*, *f*, *h*, *t*, *c*, and 0 otherwise.[5]

second group (other control variables)

logprem	Logarithm of the premium of the contract in the current contract year.
D	Dummy variable equal to 1 if the insurance contract is sold through the DOA channel of company 1, and 0 otherwise.
company2	Dummy variable equal to 1 if the insurance contract is sold by company 2, and 0 otherwise. [6]
AB	Dummy variable equal to 1if the insured is covered by a type_A or type_B contract, and 0 otherwise.[7]
RG	Dummy variable equal to 1 when the insured belongs to the "recoup group",[8] and 0 otherwise.

[4] The counter group includes the insured cars are not small sedan, for example small or large truck, cargo…etc.

[5] The insured cars in counter group for *tramak_j*, *j*= *n*, *f*, *h*, *t*, *c* , are other brands (other than Nissan, Ford, Honda, Toyota, and China.)

[6] The contracts in counter group for *D* and *company 2* are the insurance contarcts sold through the channels other than DOA of *company 1*.

[7] The contracts in counter group for *type_A* and *type_B* are type_C contracts.

[8] The "recoup group" includes the policyholders who are in "recoup group" include the ones who are covered by type A or type B contracts and who do not renew their contract or renew it for only one year.

Table 2: Sample structure (2010)

	Whole sample	Sub-sample with claim	DOA in Company 1	Non-DOA in company 1	Company 2
claim	0.8100				
SC	0.0616	0.4386	0.6628	0.2723	0.2954
RG	0.2285	0.2365	0.3165	0.2197	0.1739
AB	0.3719	0.7316	0.8979	0.6175	0.6228
C2	0.2751	0.4741	0.0000	0.0000	1.0000
SG	0.1982	0.7316	0.8979	0.6175	0.6228
SG1	0.1748	0.6670	0.8386	0.5589	0.5522
SG2	0.0234	0.0645	0.0593	0.0585	0.0706
D	0.3692	0.3978	1.0000	0.0000	0.0000
female	0.7149	0.7436	0.7758	0.7176	0.7236
age2025	0.0024	0.0022	0.0022	0.0025	0.0021
age2530	0.0313	0.0386	0.0317	0.0339	0.0456
age3060	0.8930	0.8943	0.8965	0.8872	0.8944
ageabv60	0.0734	0.0650	0.0696	0.0763	0.0580
carage0	0.1947	0.2983	0.4926	0.1383	0.1785
carage1	0.1562	0.2403	0.2387	0.2214	0.2468
carage2	0.0951	0.1010	0.0688	0.0882	0.1315
carage3	0.1175	0.1117	0.0699	0.1272	0.1425
carage4	0.1041	0.0749	0.0440	0.0941	0.0956
veh_m	0.2912	0.2589	0.2283	0.2807	0.2786
veh_l	0.2580	0.2678	0.2701	0.3070	0.2553
sedan	0.9102	0.9247	0.9612	0.8974	0.9015
lnprem	9.0442	9.5277	10.0894	9.5279	9.0563
bonus	0.8954	1.1140	0.8760	0.7154	1.4214
Observations	149452	9205	3662	1179	4364

Table 3: Conditional dependence between SC and SG

Variables	Model 1		Model 2		Model 3	
	SC	***SG***	***SC***	***SG1***	***SC***	***SG2***
RG	0.1193***	1.5444***	0.1374***	1.3111***	0.3037***	2.1846***
	[0.0332]	[0.0861]	[0.0348]	[0.0732]	[0.1076]	[0.1744]
female	-0.0138	0.1493***	0.0034	0.3138***	-0.0299	0.0263
	[0.0317]	[0.0386]	[0.0330]	[0.0393]	[0.0524]	[0.0687]
age2025	0.0281	-0.1512	0.1413	-0.1264	0.2296	-0.8458
	[0.2993]	[0.3506]	[0.3045]	[0.3507]	[0.4003]	[0.6399]
age2530	-0.2138**	-0.2505**	-0.4537***	-0.2802***	-0.2659*	-0.7943***
	[0.0884]	[0.1071]	[0.0926]	[0.1083]	[0.1464]	[0.2049]
age3060	0.0409	0.0152	-0.0292	0.0367	-0.0627	-0.2297*
	[0.0551]	[0.0680]	[0.0568]	[0.0685]	[0.0964]	[0.1213]
tramak_n	0.1442	0.3462	0.1815	0.4390*	0.5884*	0.3559
	[0.1662]	[0.2265]	[0.1774]	[0.2277]	[0.3364]	[0.3788]
tramak_f	-0.1785***	0.0285	-0.1999***	0.0592	-0.0827	0.0165
	[0.0623]	[0.0749]	[0.0656]	[0.0769]	[0.0979]	[0.1249]
tramak_h	-0.1205**	-0.2087***	-0.0468	-0.1026	-0.1615*	-0.3503***
	[0.0566]	[0.0652]	[0.0582]	[0.0664]	[0.0897]	[0.1265]
tramak_t	0.0409	0.2473***	0.0697**	0.3562***	-0.0694	-0.2450***
	[0.0317]	[0.0392]	[0.0334]	[0.0400]	[0.0563]	[0.0748]
tramak_c	-0.4594***	-0.1594*	-0.3729***	-0.2453***	-0.2362**	-0.6231***
	[0.0765]	[0.0818]	[0.0784]	[0.0829]	[0.1086]	[0.1738]
carage0	0.3822***	0.4696***	0.3307***	0.4260***	0.4480***	0.5117***
	[0.0506]	[0.0600]	[0.0536]	[0.0611]	[0.870]	[0.1023]
carage1	0.1381***	0.0837	0.1361***	0.1840***	0.0998	0.3499***
	[0.0469]	[0.0532]	[0.0492]	[0.0547]	[0.0758]	[0.0949]
carage2	0.0573	-0.0801	0.0060	0.0526	0.1059	0.1412
	[0.0553]	[0.0626]	[0.0576]	[0.0643]	[0.0865]	[0.1164]
carage3	0.0956*	0.0376	-0.0547	0.0272	0.1182	-0.1662
	[0.0529]	[0.0595]	[0.0551]	[0.0609]	[0.0793]	[0.1176]
carage4	-0.1447**	-0.1928***	-0.2558***	-0.1329*	-0.2502***	0.0632
	[0.0607]	[0.0659]	[0.0637]	[0.0681]	[0.0901]	[0.1190]
veh_m	0.0783**	-0.2005***	0.1401***	-0.2903***	0.1156*	0.1263
	[0.3456]	[0.0421]	[0.0357]	[0.0430]	[0.0596]	[0.0806]
veh_l	0.0636	-0.0568	0.0544	-0.1804***	0.1670**	0.3927***
	[0.0403]	[0.0507]	[0.0418]	[0.0519]	[0.0771]	[0.0922]

sedan	0.0685	-0.3449***	0.0246	-0.2958***	-0.0286	0.0040
	[0.0585]	[0.0695]	[0.0605]	[0.0700]	0.0951]	[0.1277]
lnprem	0.1036***	0.4852***	0.1086***	0.4849***	0.0006	0.3405***
	[0.0257]	[0.0238]	[0.0272]	[0.0245]	[0.0426]	[0.0436]
bonus	-0.4794***	-0.1689***	-0.5341***	-0.1805***	-0.1382**	0.0550
	[0.0345]	[0.0400]	[0.0371]	[0.0415]	[0.0559]	[0.0682]
ρ	0.1395***		0.0873***		0.2608***	
	[0.0319]		[0.0337]		[0.0514]	

Standard errors in brackets; ***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$

Table 4: Testing hypothesis 2 (year 2010)

	SG1		***SG2***	
	Est. coeff	***P value***	***Est. coeff***	***P value***
Intercept	-2869.3	<.0001	-4642.1	<.0001
SC	-198.9	0.0113	-742.1	0.0183
first	46.5	0.4172	-403.0	0.0630
*first*SC*	-113.3	0.1627	1465.7	<.0001
female	17.1	0.4871	-145.3	0.0853
age2025	-237.2	0.3869	621.7	0.5280
age2530	-107.3	0.1077	-442.5	0.0865
age3060	-36.9	0.3693	223.1	0.1660
tramak_n	-201.7	0.0932	-620.8	0.1240
tramak_f	-184.5	0.0002	-134.0	0.3972
tramak_h	-117.8	0.0082	-138.6	0.4594
tramak_t	-193.9	<.0001	-401.3	<.0001
tramak_c	-219.1	0.0003	-836.5	0.0006
carage0	-149.2	0.0002	-108.0	0.3834
carage1	-103.0	0.0069	-192.3	0.1308
carage2	-25.8	0.5639	-159.7	0.2931
carage3	12.9	0.7677	-192.7	0.2024
carage4	103.0	0.0409	-30.5	0.8493
veh_m	-14.2	0.5944	-151.1	0.1689
veh_l	214.9	<.0001	148.3	0.1818
sedan	269.6	<.0001	305.2	0.0785
logprem	371.2	<.0001	697.1	<.0001
bonus	48.1	0.0681	-536.6	<.0001
Adj. R^2	0.1138		0.4206	
observations	6567		633	

Table 5: Conditional dependence between *SC* and *SG* in sub-samples – Model 1 (year 2010)

	Company 1 dealer		***Company 1 non-dealer***		***Company 2***	
	SC	***SG***	***SC***	***SG***	***SC***	***SG***
RG	0.2087***	1.4486***	-0.0127	1.2273***	0.1663***	1.0515***
	[0.0490]	[0.1864]	[0.1068]	[0.1624]	[0.0560]	[0.0861]
female	0.0602	0.1776**	-0.1042	0.1454	-0.0456	0.2719***
	[0.0535]	[0.0790]	[0.0914]	[0.1007]	[0.0463]	[0.0507]
age2025	0.3982	-0.3157	-0.2333	-0.4811	0.0419	-0.3576
	[5405]	[0.6116]	[0.9014]	[0.8294]	[0.4649]	[0.4989]
age2530	-0.4620***	-0.2670	0.0615	-0.6889**	-0.3168**	-0.5582***
	[0.1456]	[0.2155]	[0.2624]	[0.3040]	[0.1310]	[0.1383]
age3060	-0.0284	0.0917	0.0773	-0.2390	-0.0226	0.0069
	[0.0856]	[0.1254]	[0.1563]	[0.1686]	[0.0869]	[0.0947]
tramak_n	-0.0162	0.0922	0.5034	0.4999	0.3149	0.2470
	[0.3307]	[0.4795]	[0.4041]	[0.5685]	[0.2239]	[0.2816]
tramak_f	-0.0451	0.1857	-0.0466	-0.0111	-0.0337	-0.1600
	[0.1189]	[0.1755]	[0.1517]	[0.1682]	[0.0906]	[0.0993]
tramak_h	-0.0381	0.0760	-0.0707	-0.0458	-0.0103	-0.3598***
	[0.1253]	[0.1689]	[0.1447]	[0.1634]	[0.0750]	[0.0805]
tramak_t	-0.1260**	0.4904***	-0.1435	0.2018*	0.0418	-0.0453
	[0.0538]	[0.0760]	[0.0935]	[0.1037]	[0.0491]	[0.0539]
tramak_c	-0.1936	0.3031	0.1873	-0.3445	-0.1218	-0.3294***
	[0.3462]	[0.4113]	[0.2121]	[0.2479]	[0.0872]	[0.0937]
carage0	-0.0363	0.4796***	-0.1329	0.6156***	0.0914	0.4867***
	[0.0975]	[0.1268]	[0.1535]	[0.1729]	[0.0809]	[0.0934]
carage1	0.0224	0.1614	-0.1268	0.4121***	0.0151	0.2196***
	[0.0943]	[0.1219]	[0.1240]	[0.1320]	[0.0695]	[0.0703]
carage2	-0.1635	-0.1748	0.2890*	0.4076**	-0.0292	0.1062
	[0.1126]	[0.1472]	[0.1509]	[0.1756]	[0.0754]	[0.0790]
carage3	0.0092	-0.0060	-0.0056	0.1812	-0.1602**	0.1101
	[0.1111]	[0.1433]	[0.1332]	[0.1457]	[0.0719]	[0.0757]
carage4	-0.2389*	-0.2915*	-0.3620**	0.2705*	-0.1641**	0.0986
	[0.1254]	[0.1540]	[0.1597]	[0.1561]	[0.0810]	[0.0848]
veh_m	-0.0862	-0.2110**	-0.1666*	-0.1394	0.1415**	0.0495
	[0.0585]	[0.0836]	[0.1006]	[0.1110]	[0.0564]	[0.0618]

veh_l	-0.1789***	-0.2330**	-0.3471***	-0.0786	0.0888	0.2930***
	[0.0689]	[0.0926]	[0.1242]	[0.1263]	[0.0828]	[0.0909]
sedan	-0.1914	-0.2939	0.0666	-0.3395*	0.0700	-0.1867**
	[0.1258]	[0.1809]	[0.1567]	[0.1755]	[0.0816]	[0.0884]
lnprem	0.2229**	0.6791***	-0.0554	0.7074***	-0.0662	0.5853***
	[0.0874]	[0.0632]	[0.1058]	[0.0752]	[0.0453]	[0.5853]
bonus	-0.2188	-1.2374***	0.4098*	-1.0545***	-0.1504*	-0.4616***
	[0.1636]	[0.1890]	[0.2408]	[0.2304]	[0.080]	[0.0955]
Constant	-4.9188*	-4.9188***	-0.2026	-5.6378***	0.1742	-4.6010***
	[0.7791]	[0.5876]	[0.9081]	[0.6592]	[0.3822]	[0.3136]
ρ	0.5393***		0.1344		0.0562	
	[0.0729]		[0.1201]		[0.0480]	

Standard errors in brackets; ***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$

Table 6: Conditional dependence between *SC* and *SG1* in sub-samples – Model 2 (year 2010)

	Company 1 dealer		***Company 1 non-dealer***		***Company 2***	
	SC	***SG1***	***SC***	***SG1***	***SC***	***SG1***
RG	0.2010***	1.5341***	-0.0221	1.1578***	0.0709	1.2021***
	[0.0501]	[0.2253]	[0.1104]	[0.1645]	[0.0599]	[0.0980]
female	0.1167**	0.2759***	-0.0620	0.1638	-0.1032**	0.2954***
	[0.0549]	[0.0824]	[0.0935]	[0.1028]	[0.0480]	[0.0537]
age2025	0.4237	-0.4132	-0.1955	-0.1888	-0.3207	-0.2477
	[0.5564]	[0.6221]	[0.8616]	[0.8311]	[0.4869]	[0.5016]
age2530	-0.4478***	-0.5637**	-0.0376	-0.4613	-0.2080	-0.4562***
	[0.1503]	[0.2213]	[0.2680]	[0.3076]	[0.1329]	[0.1452]
age3060	-0.0198	0.0052	0.0910	-0.2645	-0.0325	-0.1003
	[0.0884]	[0.1334]	[0.1599]	[0.1707]	[0.0892]	[0.0987]
tramak_n	0.3639	0.0218	0.6697*	0.3062	0.3782	0.1518
	[0.3716]	[0.4651]	[0.4006]	[0.5392]	[0.2385]	[0.2967]
tramak_f	-0.1312	0.1808	0.0723	-0.1424	-0.2251**	-0.2843***
	[0.1246]	[0.1894]	[0.1563]	[0.1727]	[0.0958]	[0.1054]
tramak_h	-0.0022	0.1606	0.0368	-0.2072	-0.0244	-0.4242***
	[0.1272]	[0.1764]	[0.1485]	[0.1682]	[0.0769]	[0.0844]
tramak_t	-0.1027*	0.5401***	-0.0549	-0.0003	0.0102	-0.0394
	[0.0572]	[0.0802]	[0.0972]	[0.1056]	[0.0505]	[0.0565]
tramak_c	-0.3570	0.1435	0.0864	-0.4660*	-0.1067	-0.4536***
	[0.3465]	[0.4135]	[0.2209]	[0.2552]	[0.0892]	[0.0979]
carage0	-0.0659	0.5069***	-0.0428	0.5662***	0.0966	0.6870***
	[0.1021]	[0.1329]	[0.1574]	[0.1766]	[0.0888]	[0.0882]
carage1	0.0044	0.2771**	-0.0676	0.4466***	0.1138	0.4128***
	[0.0987]	[0.1267]	[0.1287]	[0.1354]	[0.0749]	[0.0739]
carage2	-0.0842	0.0655	0.2540	0.4376**	0.0698	0.2247***
	[0.1171]	[0.1550]	[0.1568]	[0.1804]	[0.0793]	[0.0830]
carage3	-0.0432	0.0676	0.0953	0.2778*	0.0236	0.2375***
	[0.1142]	[0.1472]	[0.1352]	[0.1472]	[0.0745]	[0.0796]
carage4	-0.3611***	-0.1421	-0.4584***	0.3412**	-0.0691	0.1746*
	[0.1289]	[0.1602]	[0.1689]	[0.1630]	[0.0842]	[0.0896]
veh_m	0.0165	-0.2147**	-0.1595	-0.3025***	0.1506***	0.0094
	[0.0596]	[0.0864]	[0.1011]	[0.1143]	[0.0584]	[0.0652]

veh_l	-0. 1736** [0.0699]	-0.2738*** [0.0971]	-0.3776 [0.1246]	-0.2738** [0.1308]	0.1270 [0.0866]	0.2227** [0.0968]
sedan	-0.1373 [0.1279]	-0.3944** [0.1921]	0.0408 [0.1612]	-0.4320** [0.1846]	0.0674 [0.0838]	-0.3254*** [0.0924]
lnprem	0.2181** [0.0890]	0.7433*** [0.0660]	-0.0377 [0.1169]	0.7134*** [0.0781]	-0.0475 [0.0484]	0.6283*** [0.0385]
bonus	-0.1648 [0.1718]	-1.0223*** [0.1953]	0.4459 [0.2571]	-1.0166*** [0.2367]	-0.2510*** [0.0932]	-0.4993*** [0.1022]
Constant	-1.5136* [0.7873]	-5.7331*** [0.6162]	-0.4028 [1.0080]	-5.4420*** [0.6751]	0.1944 [0.4086]	-4.8641*** [0.3291]
ρ	0.5729*** [0.0699]		0.0916 [0.1362]		-0.0610 [0.0534]	

Standard errors in brackets; ***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$

Table 7: Conditional dependence between *SC* and *SG2* in sub-samples – Model 3 (year 2010)

	Company 1 dealer		***Company 1 non-dealer***		***Company 2***	
	SC	***SG2***	***SC***	***SG2***	***SC***	***SG2***
RG	0.2622	1.6147***	-0.4441	2.3649***	0.2728*	1.7827***
	[0.2004]	[0.2621]	[0.4225]	[0.4639]	[0.1594]	[0.2152]
female	-0.1535	-0.0666	-0.0555	0.3131	-0.0175	-0.0709
	[0.1249]	[0.1494]	[0.1397]	[0.2001]	[0.0669]	[0.0833]
age2025	0.2937	-0.4864	0.2417	0.4348	-0.0347	-0.2253
	[0.7007]	[0.8008]	[1.1300]	[1.0916]	[0.6076]	[0.7808]
age2530	-0.6705	-0.3919	0.6021	-0.7861	-0.3573*	-0.4311*
	[0.3223]	[0.3941]	[0.4272]	[0.7589]	[0.1828]	[0.2376]
age3060	0.0564	-0.1246	0.4362*	-0.1269	-0.1732	-0.1471
	[0.2093]	[2562]	[0.2572]	[0.3146]	[0.1268]	[0.1547]
tramak_n	0.7126	-0.1096	-0.3747	0.0300	0.2852	0.2575
	[0.6845]	[0.7266]	[1.7463]	[1.7441]	[0.4113]	[0.4415]
tramak_f	-0.1230	0.3446	-0.0555	0.0392	-0.0269	-0.0740
	[0.2342]	[0.2781]	[0.2220]	[0.2885]	[0.1283]	[0.1624]
tramak_h	-0.0649	-0.1577	0.0791	-0.5671*	0.0696	-0.2862**
	[0.2789]	[0.3410]	[0.2231]	[0.3424]	[0.1063]	[0.1434]
tramak_t	-0.4971***	-0.8100***	-0.1005	-0.1753	0.0083	-0.1858**
	[0.1352]	[0.1601]	[0.1473]	[0.1983]	[0.0727]	[0.0946]
tramak_c	0.0140	-0.0334	-0.2963	-0.2935	-0.1097	-0.8299***
	[0.5654]	[0.8117]	[0.2935]	[0.4233]	[0.1245]	[0.1970]
carage0	0.1719	0.4071*	-0.0541	0.2479	0.3881***	0.3900***
	[0.1898]	[0.2236]	[0.2660]	[0.3513]	[0.1199]	[0.1319]
carage1	-0.0991	0.0391	-0.0870	0.4110	0.0899	-0.0666
	[0.1911]	[0.2321]	[0.1933]	[0.2520]	[0.0954]	[0.1193]
carage2	-0.0875	-0.0744	0.2821	0.3666	0.0783	-0.0445
	[0.2321]	[0.2908]	[0.2407]	[0.3243]	[0.1058]	[0.1366]
carage3	0.3269	-0.1232	0.0080	-0.2482	0.2198**	-0.0513
	[0.2303]	[0.2926]	[0.2028]	[0.3401]	[0.0963]	[0.1284]
carage4	-0.2220	-0.4267	-0.3660	0.4180	0.0271	-0.0476
	[0.2400]	[0.3248]	[0.2303]	[0.2698]	[0.1083]	[0.1430]
veh_m	0.0385	0.0276	0.1005	-0.1990	0.1787**	-0.0262

	[0.1383]	[0.1680]	[0.1489]	[0.2141]	[0.0815]	[0.1075]
veh_l	0.1327	0.2949	0.1798	-0.0297	0.2780**	0.3828***
	[0.1801]	[0.1809]	[0.1944]	[0.2319]	[0.1228]	[0.1455]
sedan	-0.2460	-0.2723	0.0885	-0.6997**	0.1115	-0.3116**
	[0.2635]	[0.3166]	[0.2384]	[0.2843]	[0.1199]	[0.1471]
lnprem	0.1987	0.6067***	-0.1528	0.4173***	-0.0067	0.3557***
	[0.1692]	[0.1097]	[0.2006]	[0.1440]	[0.0713]	[0.0688]
bonus	-0.0319	-0.9141**	0.2821	-0.9554**	-0.3427**	0.0301
	[0.3338]	[0.3744]	[0.4282]	[0.4655]	[0.1385]	[0.1585]
Constant	-1.4106	-5.0827***	0.0842	-3.8736***	-0.1057	-3.6732***
	[1.4319]	[1.0026]	[1.6199]	[1.2197]	[0.5649]	[0.5400]
ρ	0.7492***		-0.2020		0.2076***	
	[0.1355]		[0.2206]		[0.0702]	

Standard errors in brackets; ***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$

Table 8: Difference of conditional dependence between *SC* and *SG/SG1/SG2* in 2010

	Model 1 *SC, SG*	**Model 2** *SC, SG1*	**Model 3** *SC, SG2*
$\rho^{D} - \rho^{ND}$	0.4049***	0.4814***	0.9513***
	[4.6650]	[5.3121]	[5.2735]
$\rho^{D} - \rho^{C2}$	0.4831***	0.6339***	0.5416***
	[7.9626]	[10.3059]	[6.0383]
$\rho^{C2} - \rho^{ND}$	-0.0782	-0.1526	0.4096***
	[-1.1193]	[-1.9339]	[3.4480]

Table A1: Sample structure in 2018

	Whole sample	**Sub-sample with claim**	**DOA**	**non-DOA**
claim	0.0371			
SC	0.0090	0.2420	0.2695	0.2275
RG	0.3400	0.3614	0.4912	0.2931
AB	0.3500	0.3708	0.5039	0.3007
SG	0.2189	0.2697	0.3505	0.2272
D	0.2795	0.3452	1.0000	0.0000
female	0.5599	0.6515	0.6541	0.6502
age2025	0.0087	0.0057	0.0061	0.0055
age2530	0.0376	0.0401	0.0344	0.0430
age3060	0.7674	0.8107	0.7968	0.8180
carage0	0.0439	0.1970	0.2246	0.1825
carage1	0.0591	0.1702	0.1705	0.1701
carage2	0.0644	0.1477	0.1424	0.1504
carage3	0.0608	0.1132	0.1106	0.1146
carage4	0.0618	0.0993	0.1004	0.0987
veh_m	0.2800	0.3320	0.3372	0.3292
veh_l	0.1612	0.1881	0.1910	0.1866
sedan	0.9719	0.9951	0.9962	0.9945
lnprem	8.9993	9.0407	9.2992	8.9045
bonus	0.9370	0.7083	0.7316	0.6960
Observations	269475	10010	3455	6555

Table A2: Conditional dependence between SC and SG (year 2018)

	dealer		Non-dealer	
	SC	***SG***	***SC***	***SG***
RG	-0.1621*	3.9976***	-0.0552	4.1421***
	[0.0833]	[0.1734]	[0.0539]	[0.1275]
female	0.1163**	0.1735	0.1878***	0.1635*
	[0.0565]	[0.1131]	[0.0379]	[0.0839]
age2025	-0.4879	1.7509***	-0.3606	-0.2547
	[0.4042]	[0.6026]	[0.2516]	[0.7147]
age2530	-0.4798***	0.5993**	-0.6081***	0.1064
	[0.1719]	[0.3043]	[0.1090]	[0.2218]
age3060	-0.0701	0.2061	-0.0548	-0.0222
	[0.0716]	[0.1511]	[0.0509]	[0.1054]
tramak_n	-0.5044	1.2453**	0.4607**	-0.2498
	[0.5860]	[0.5731]	[0.2073]	[0.5385]
tramak_f	0.3979***	0.1127	0.5452***	0.7828***
	[0.1484]	[0.2855]	[0.0839]	[0.1915]
tramak_h	0.1188	0.2980	0.2119***	-0.3979**
	[0.1440]	[0.2660]	[0.0795]	[0.1850]
tramak_t	0.5698***	0.1843	0.3826***	0.3883***
	[0.0578]	[0.1162]	[0.0393]	[0.0891]
tramak_c	0.4150	-0.4958	-0.0831	0.3025
	[0.2818]	[0.6022]	[0.1984]	[0.4272]
carage0	0.2969***	-0.9596***	0.2489***	-0.6280***
	[0.0899]	[0.1991]	[0.0581]	[0.1419]
carage1	-0.1035	-0.9074***	-0.0309	-0.3448**
	[0.0878]	[0.1917]	[0.0576]	[0.1366]
carage2	-0.2374***	-0.1850	-0.0843	-0.3009**
	[0.0891]	[0.1909]	[0.0579]	[0.1320]
carage3	-0.2848***	-0.2969	-0.3812***	-0.5707***
	[0.0980]	[0.1936]	[0.0658]	[0.1408]
carage4	-0.2814***	-0.3862*	-0.2758***	-0.9843***
	[0.1003]	[0.2060]	[0.0684]	[0.1550]
veh_m	-0.0768	-0.0678	0.1282***	0.2059**
	[0.0593]	[0.1243]	[0.0407]	[0.0956]
veh_l	-0.3323***	-0.4531***	-0.1396***	0.3258***
	[0.0739]	[0.1601]	[0.0525]	[0.1124]

sedan	-0.5778	-1.3503*	-0.4488	-1.1315**
	[0.4758]	[0.6892]	[0.2802]	[0.5508]
lnprem	0.1095**	-0.0162	-0.0536*	-0.0823
	[0.0453]	[0.0916]	[0.0291]	[0.0592]
bonus	0.9318***	0.3404	0.7223***	-1.0135***
	[0.1128]	[0.2296]	[0.0790]	[0.1906]
ρ	0.5229**		0.9277***	
	[0.2575]		[0.2172]	

Standard errors in brackets; ***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$

The difference of conditional dependence between *SC* and *SG*:

$\rho^D - \rho^{ND}$ = -0.4048 (t = -1.7524; H_0: $\rho^D \leq \rho^{ND}$ cannot be rejected)

Table A3: Comparative manipulation ability of DOAs (years 2010 and 2018)

	First stage		Second stage	
	Est. coeff.	**P value**	**Est. coeff.**	**P value**
Intercept	-37.0773	<.0001	-0.7793	0.0027
SG			-0.4212	0.0899
$\widehat{SG}$			-0.1565	0.5312
dealer			0.00618	0.9478
y2018			-0.2399	0.0005
*SG*dealer*			1.8147	<.0001
*SG*y2018*			0.7457	0.0029
*dealer*y2018*			0.0443	0.6633
*SG*dealer*y2018*			-1.7265	<.0001
recoup	16.7495	0.8673	-0.1304	0.5139
female	0.3715	0.0198	0.0640	0.0322
age2025	-4.3641	0.9975	-0.4786	0.0182
age2530	0.0644	0.9055	-0.5400	<.0001
age3060	0.6748	0.0241	-0.0957	0.0182
tramak_n	-4.9743	0.9972	0.1025	0.575
tramak_f	0.6856	0.0085	0.0525	0.4566
tramak_h	-0.3143	0.3305	0.0618	0.3446
tramak_t	-0.3328	0.0447	0.2423	<.0001
tramak_c	0.3515	0.4887	-0.2892	0.0562
carage0	0.2408	0.2913	0.4741	<.0001
carage1	0.1337	0.5468	0.1807	<.0001
carage2	-0.2893	0.2825	0.1914	<.0001
carage3	-0.3436	0.2325	0.0823	0.1051
carage4	0.3167	0.1901	0.0314	0.5563
veh_m	-0.1722	0.3223	0.0300	0.3498
veh_l	-0.1807	0.3411	0.0333	0.3941
sedan	-1.4823	<.0001	-0.1291	0.3234
logprem	3.8113	<.0001	-0.0223	0.3687
bonus	-0.8332	0.0087	0.5239	<.0001
-2logL	12270.286		11655.898	

Table A4: Testing Hypothesis 2 for year 2018

	SG		**Type C**	
	Est. coeff.	**P value**	**Est. coeff.**	**P value**
Intercept	20.8341	0.0259	30.3130	0.0009
SC	-4.8308	0.0581	-9.0929	0.0006
first	-0.8095	0.5265	-3.7417	0.0036
*first*SC*	-0.7615	0.7789	1.8227	0.5221
female	-3.7137	<.0001	-0.7140	0.4280
age2025	-10.3178	0.0970	11.9403	0.0314
age2530	-3.5532	0.1216	3.4655	0.1534
age3060	-2.5520	0.0254	-0.9612	0.4385
tramak_n	-1.7228	0.7422	-7.9312	0.1331
tramak_f	-4.9029	0.0172	-9.3758	<.0001
tramak_h	-7.9695	<.0001	-11.4974	<.0001
tramak_t	-7.0401	<.0001	-11.1717	<.0001
tramak_c	-8.6608	0.0585	-10.9192	0.0065
carage0	2.7724	0.0447	3.4785	0.0142
carage1	5.1712	<.0001	4.1478	0.0029
carage2	3.7767	0.0039	5.8271	<.0001
carage3	2.5255	0.0692	5.5491	0.0003
carage4	2.3239	0.1084	4.1172	0.0100
veh_m	5.0363	<.0001	4.5882	<.0001
veh_l	9.7725	<.0001	15.1139	<.0001
sedan	6.2842	0.3710	7.3519	0.2293
logprem	-0.6929	0.2462	-1.5460	0.0395
bonus	2.1328	0.2304	-0.0887	0.9615
$Adj.R^2$	0.0773		0.0549	
Observations	3149		7530	