MACROECONOMIC RISK CHAIR
NEWSLETTER N°13

The sixth annual conference of the Macroeconomic Risk Chair was held on October 16, 2023 at the Paris School of Economics, on the topic “Structural changes and their implications for macroeconomic risks, dynamics and policies”. Following this conference, we had the opportunity to interview **Chad Jones (Stanford University)** about his recent research on **artificial intelligence.**

Furthermore, the chair organized its sixth annual lecture on June 1st, 2023, with **Olivier Jeanne (Johns Hopkins University)** as a special guest speaker to discuss his last article on **the sustainability of sovereign debt.**

This newsletter includes a first interview with Chad Jones and a second one with Olivier Jeanne, as well as a brief description of their research papers.

The AI Dilemma: Growth vs. Existential Risk

On September 16, 2023, **Chad Jones** (Stanford University) gave a keynote lecture on the theme of artificial intelligence. Following this lecture, we had the opportunity to interview him about his research.

Artificial intelligence (AI) has the potential to bring us into a new age of amazing economic growth, "delivering" as much change as electricity and internet did in the past. At the same time, however, it also carries risks not unlike those posed by nuclear weapons. As a matter of fact, these two scenarios are likely correlated: **it is precisely the state of the world in which AI is extremely sophisticated** (and can therefore change all of our lives) **that is also the more likely to pose a dangerous threat to humanity.**

In this paper, Chad Jones asks under what conditions should we, as a society, decide to "stop the AI experiment"? He does so by developing a model that shines light on the main trade-off lying behind such decision: does the utility deriving from AI-induced growth in consumption outweigh the potential existential threats to humanity?

> It is precisely the state of the world in which AI is extremely sophisticated that is also the more likely to pose a dangerous threat to humanity.

In the model, the answer to such trade-off fundamentally depends on two components: first, our assumptions about preferences and the utility function; and second, the type of improvements that AI would bring. In fact, the author finds that we can establish a threshold for what the society could consider as "acceptable risk" and that this threshold depends on the utility function, the AI-induced growth rate, and on the mortality rate.

Under a standard specification for preferences, the assumption we make about the curvature of the utility function (i.e. aversion to risk) is crucial: under log-utility large consumption gains would be worth gambles that involve as much as a 1-in-3 chance of extinction; on the other hand, increasing the curvature of utility (i.e. making society more averse to risk) would lead to much more conservative "bets" independently of how much faster AI would make the economy grow.

However, the conclusion drastically changes if AI can deliver significant improvements in mortality (e.g. by discovering new ways to fight cancer or slow down aging). The intuition behind this difference is very simple: due to the curvature of the utility function, additional consumption is not particularly valuable. In fact, as long as γ (the risk-aversion parameter) is greater or equal than 2, even a so-called "singularity" (infinite consumption in a short amount of time) would not significantly change the society's willingness to engage in the AI bet, which would remain well below a 1 in 10 chance of extinction. On the other hand, however, mortality improvements are not the same as additional consumption: longer lifespans imply not just more consumption, but more utility (the same as an extra year of life is worth more as an additional year worth of consump- tion). **As a consequence, in the presence of mortality improvements, the existential risk cutoffs are much higher (on the order of 25-30%).**

In fact a similar - albeit opposite - result applies also if society were to discount the future at a lower rate (i.e. by putting additional weight on the future), what the author defines as "longtermism": **because by engaging in the AI "bet" we are also putting at risk the very existence of future generations**, lowering the discount rate means we should be even less willing to risk it all.

While we still do not have a definitive answer on whether we should stop the AI experiment or not - and we will probably never have it - this paper provides a useful framework to think about the trade-offs that we, as a society, face when thinking about the possible gains from developing ever more powerful AI systems.

Interview: Chad Jones

In your lecture, you talked about the opportunities (and risks) of advances in artificial intelligence (AI). Intuitively, it feels natural to consider AI as different from other forms of automation. What makes AI so special? What sets it apart from other forms of automation?

Absolutely. AI does feel different from tractors and airplanes, and I think it is different in a fundamental way. In particular, earlier forms of automation were quite narrow and focused, whereas AI may be the most "general" of the "general purpose technologies" that we have ever seen. It is conceivable that some future – perhaps distant future – AI could substitute for people in all or nearly all tasks we perform. It is that breadth that makes AI fundamentally different.

The video replay of Chad Jones' lecture is available online.

Your paper is titled "The AI Dilemma: Growth vs. Existential Risk". Let us first focus on the "existential risk" part. In popular culture such risk is often associated with some form of "robot revolution" which is really only one of several possibilities. Which scenarios would you consider the most likely to pose an existential threat and what should be done to avoid such scenarios?

By their nature, these existential risk scenarios typically sound quite far-fetched and like a science fiction movie. But that is not true of all such scenarios, including ones that may be closer to the present rather than further away. For example, imagine a future ChatGPT-7 that could distill all of scientific knowledge and generalize from that base. A "bad actor" could query the AI for instructions on cheap and easy ways to kill massive numbers of people. The AI might suggest biotech modifications to a virus like the one associated with Covid-19 to make it even more deadly. Or perhaps some kind of genetically-modified bacteria that does great harm.

One analogy that I've heard from AI experts and that I find helpful is this: if Magnus Carlsen and I play chess, I can be 100% certain that I will be checkmated, even if I cannot describe in advance exactly how the checkmate will occur!

In March 2023, an open letter calling for a pause in AI development made the headlines worldwide, with signatories including some of the most prominent AI names. That letter went mostly unanswered and companies working on AI are still moving forward, possibly at an even faster pace. What makes it so hard for companies to slow down and should something be done about it?

I suspect that several things make it difficult to get a coordinated slowdown in AI research in order to focus on AI safety. First, while many experts are worried about existential risk, many others are not. This allows the experts who are not worried to justify going ahead with the research. From a policy standpoint, us non-experts might therefore reduce our estimate of the existential risk probability accordingly, but a 5% chance of something going bad instead of a 10% chance strikes me as still being a problem we should worry about.

Second, there are very real financial gains to be won before any existential risk gets realized. If AI really is a more significant technology than even electricity or the internet, then companies that develop AI may earn astounding profits. These profits provide a strong incentive to push AI research forward. A company may rationalize that others are pursuing AI research no matter what I do, so whatever existential risk problems exist are beyond my control. So I may as well respond to the financial rewards.

Finally, there are geopolitical dimensions to the problem. Just as the US raced Nazi Germany to develop the atomic bomb during World War II, "friendly" governments may rightly worry that even if they pause their AI research, the "unfriendly" actors will not. Any pause may therefore simply serve to allow others to catch up and achieve these potentially dangerous tools first.

We have talked about the "existential risk" part of the dilemma, the other part being "growth". This partly refers to the many exciting possibilities that AI development could bring us. Which, of the many great promises of AI, do you consider more realistic?

> If AI can also invent new technologies and new ideas, this could accelerate living standards around the world.

As we discussed earlier, what makes AI different from other forms of automation is that it could potentially do any task a person could do, and perhaps even better and with more intelligence. One of the most economically significant tasks people do is invent new ideas that make all of us better off. Think about the invention of electricity or antibiotics or computers or air conditioning. If AI can also invent new technologies and new ideas, this could accelerate living standards around the world: future discoveries made by AI could solve climate change and end poverty and hunger, for example. Those are noble goals!

One of the things you discuss in your paper is the notion of a "singularity", could you explain what it is and why you think it is important when talking about AI development?

A singularity is one of the science-fiction concepts associated with AI. In particular, imagine AI can discover new ideas and new technologies. Once we invent a single AI, we could replicate it across thousands, millions, or hundreds of millions of computer systems, each working on discovering new ideas. And some of those ideas will be how to create smarter or more efficient AIs, so we get even more AIs to hunt for new ideas. The idea of a singularity is that all ideas that could ever be discovered could get discovered in a historically-short period of time, such as a few decades. So economic growth rates themselves could increase to faster and faster rates.

In fact, one of the findings in my paper that surprised me is that the possibility of such singularities actually isn't that important to thinking about the costs and benefits of AI. Instead of imagining ever-accelerating rates of economic growth, it is sufficient to just consider a single increase, for example the possibility that rates of growth might rise from 2% per year to something like 10% per year. The cost-benefit calculations look very similar in this high-growth scenario to what might happen with a singularity.

Given the large incentives companies have to continue pushing forward AI development, and given the risks we discussed, do you believe governments should step up and intervene? What kind of policies do you think might be more successful at tilting the scale towards "growth" as opposed to "existential risk"?

I think these are fantastic questions and are exactly the right ones to be asking. Part of the reason I wrote this paper was to help me and other economists to start thinking about them. Through research and through public discussions among the experts, I hope we will figure out the answers to these questions. But I myself do not feel like I have a strong sense of what the right answers are at the moment. I do believe that the stakes involved are so high that these questions are among the most important questions we should be thinking about over the next five to ten years. Research on AI safety is fundamentally important.

Part of the promise of AI is that it provides a way to create incredibly rapid growth in the amount of research effort, by having AI discover new ideas.

One of the points you made in your talk, which is related to your research body more broadly, is that the growth rate of the economy is determined by ideas. In fact, in one of your recent papers you argue that "ideas, and the exponential growth they imply, are getting harder to find". In your world of ideas, growth is given by the product of research productivity and number of researchers. How do you think AI will change this world? Will it only increase research productivity or might it also increase the number of researchers by making research more "democratic"?

Yes. There is a precise sense in which economic growth is the product of research productivity and the amount of research effort we undertake. What we showed in that earlier paper on "Are Ideas Getting Harder to Find?" was that the answer is "Yes." Research productivity in many different settings is declining. These settings range from agriculture to medical innovation to Moore's Law itself (the fact that the density of computer chips is doubling roughly every two years). If research productivity is declining, then the only way we have achieved relatively constant exponential growth rates historically is through increasing the number of researchers over time: exponential growth in the number of researchers has delivered exponential growth in living standards through the new ideas that they discover. Part of the promise of AI is that it provides a way to create incredibly rapid growth in the amount of research effort, by having AI discover new ideas.

If ideas are indeed harder to find, one would expect the price of these ideas to increase. As you also explored in an earlier paper with Jihee Kim, the data seem to suggest that the increasing levels of inequality observed in the last decades is partly due to successful entrepreneurs who develop these ideas. Are these two faces of the same coin?

In the paper with Jihee Kim, we point out the possibility that the rise in inequality around the world is indeed associated with changing technology. Consider Sam Walton, the founder of WalMart in the 1960s. Walton built his fortune by gradually opening more and more Walmarts, first throughout Arkansas, then throughout the country, and finally throughout the world. But this was a slow gradual process. In contrast, because of the rise in information technology and the internet, an inventor who creates a new smartphone application can become a billionaire in a matter of months instead of decades.

This is exactly what we saw with ChatGPT for example. Information technology allows successful entrepreneurs to become very rich very quickly, and this has contributed to a rise in inequality. Importantly, of course, there is a sense in which this kind of inequality could be good instead of bad: if someone becomes wealthy because they invent the Covid-19 vaccine, that raises inequality but makes the world a better place. This is very different from someone becoming wealthy because they steal from many others.

In a recent paper Moll, Rachel, and Restrepo connect the rise in automation to the rise in inequality. However, if yours is a world of ideas, theirs would be a "world of goods". In the chapter you wrote with Philippe Aghion and Ben Jones in "The Economics of Artificial Intelligence", you consider both aspects and their consequences for growth. Are there potentially important interactions across these two worlds and, in particular, could such interactions also be important for inequality?

The Moll, Rachel, and Restrepo paper provides a very elegant model in which rising automation leads to an increase in the capital share of factor payments and a rise in wealth inequality. The automation they consider is the automation of tasks in producing goods, as you noted. In the paper with Philippe and Ben, we studied this kind of automation as well as the automation of tasks in the production of ideas.

> The risks associated with AI are very real and therefore should be taken seriously to ensure that the best-case scenario comes about.

There are indeed important interactions between these forms of automation, both for economic growth and for inequality. Speaking somewhat loosely, the more things can be automated – either in producing goods or ideas – the better it is for productivity and GDP overall. On the other hand, as Pascual Restrepo has pointed out in a series of papers with Daron Acemoglu, it is possible for this automation to increase inequality. Intuitively, the workers whose jobs are automated away now have to compete with cheap machines rather than just with other workers, and those cheap machines can lead to lower wages for some types of workers. This is one of those tricky instances in economics where the overall size of the "pie" can be made larger, but some people may get a smaller slice. One of the key lessons of economics is that we want to make the size of the pie as large as possible, but economic policies become of central importance in situations like this where some people can be made worse off by the changes.

To conclude, are you more excited or more worried about these recent advances in AI technology?

I am an optimist by nature. There are many problems we face as a society where new ideas could be incredibly useful. New technologies could drive the price of clean energy to zero, solve our climate change problems, cure heart disease and cancer, and lead to cheap, healthy, plentiful foods to end world hunger. The history of the past century is one where new discoveries have helped us make progress on all of these fronts. In the best-case scenario, AI could continue this progress and spread it throughout the world. There is, therefore, much to be excited about. However, the risks associated with AI are very real and therefore should be taken seriously to ensure that this best-case scenario comes about.

Chad Jones is The STANCO 25 Professor of Economics at Stanford Graduate School of Business and a research associate of the National Bureau of Economic Research. Professor Jones has been honored as a member of the American Academy of Arts and Sciences, a fellow of the Econometric Society, and a co-editor of *Econometrica*. He is currently the area coordinator for the economics group at Stanford GSB.

Professor Jones is the author of numerous research papers as well as two textbooks, *Introduction to Economic Growth* (2013) and *Macroeconomics* (2020).

Annual Macroeconomic Risk Chair Lecture: Should Central Banks backstop Government Debt?

On June 1st, 2023, **Olivier Jeanne** (Johns Hopkins University) gave a lecture on the theme of central banks and governement debt. Following this lecture, we had the opportunity to interview him about his research.

After the Great Financial Crisis in 2008, the Euro-Area experienced a sovereign debt crisis to which the European Central Bank (ECB hereafter) answered with the lender in last resort (LOLR hereafter) tool. By intervening at the peak of the crisis, the ECB de facto started backstopping euro governments debt. The use of LOLR instruments was motivated by the admitted existence of multiple equilibria: the use of LOLR would make the economy shift from a bad to a good equilibrium without dealing with costly inflation. However, when debt is not sustainable, an alternative to LOLR may be financial repression. **To what extend financial repression could be used as a tool when the alternative is government default?** What would be the optimal way to implement financial repression?

To answer those questions, Olivier Jeanne develops a dynamic model in which a government faces different policy choices to stabilize its debt level: it can do fiscal adjustments, implement financial repression, or default. Fiscal adjustment is the least costly option, but takes time to be implemented. It makes the use of financial repression or default credible as they can be triggered at any time even if they are more costly. Financial repression remains less costly than a government default. The model shows how to optimally design financial repression to circumvent incentives to default and its consequences.

Without financial repression, default occurs about every twenty years at equilibrium, which reduces output by 2% on average. **When financial repression is implemented as a last resort instrument, it is triggered only when debt reaches almost 400% of GDP, which occurs very rarely in numerical simulations.** Yet, welfare gains arising from the simple fact that the government may use financial repression are substantial, equal to a 1% permanent increase in consumption. In that respect, financial repression may be a useful instrument to manage public debt.

Interview: Olivier Jeanne

Your paper "Whatever it takes: Government Default VS Financial Repression" discusses how central banks could use financial repression as a tool to backstop optimally government debt. What's financial repression in practice? How does that tool differ from other instruments a central bank could use to backstop government debt?

Sure! Financial repression occurs when the banking sector is required to finance the government at interest rates lower than the market rates. For instance, the central bank might lend to the government at a low interest rate and demand banks to hold equivalent amounts in reserves. Another way is to mandate commercial banks to lend directly to the government at reduced interest rates. There are many examples of such policies in advanced economies after WWII, and in less developed economies at various points in time. These policies are often associated with inflation because inflation magnifies the quasi-fiscal revenue from financial repression.

Central bank backstop policies are different. Examples of central bank backstop are, in the euro area, the outright monetary transactions (OMTs) and the transmission protection instruments (TPIs). These measures involve the central bank purchasing government debt in the market to prevent excessive default spread due to market dysfunction. Interestingly, just having the possibility of using these instruments was enough to stabilize the debt markets, and they haven't been utilized yet.

Financial repression and central bank backstop are different but they are not totally unrelated. If the central bank attempted to backstop the debt of an insolvent government, the backstop could morph into financial repression. And solvency being in the eyes of the beholder, the backstop could be a slippery slope towards financial repression and inflation. This can explain the German worries about the ECB backstop in the euro area.

Your model assumes a form of "fiscal inertia". Could you explain what does this notion of fiscal inertia capture, and why it matters?

Fiscal inertia is something we observe in many countries. It happens when the government debt is on an unsustainable path and requires a fiscal adjustment, but it doesn't happen due to political deadlock. In my model, I capture this idea by assuming that opportunities for fiscal adjustment occur infrequently.

This leads to long waves in the dynamics of government debt. We see periods when the debt-to-GDP ratio increases, followed by episodes where it decreases. These periods can be quite long. We see that in the model and in the data. To avoid default, the government needs room to borrow during the upward phase, and for that it may need the support of the central bank.

One of the main results of your paper is that the very existence of financial repression as a last resort tool can rule out defaults. Could you summarize for us the mechanism driving this result?

The key lies in the uncertainty surrounding government solvency. Let's consider a situation where the government's debt is increasing. It might appear insolvent as its debt dynamics seem unsustainable in the long run. However, the government might rectify this at any time through a fiscal adjustment. The catch is that for the government to be willing to make such an adjustment, its debt must not cross a critical threshold, as that would make the fiscal adjustment too costly.

Temporary financial repression comes into play here. It can prevent the debt from surpassing that critical threshold. Essentially, the government remains solvent as long as it can rely on financial repression as a last resort tool. This changes the entire dynamics of debt.

The alternative to financial repression is default. The crucial question is whether financial repression is a better option than default. In my calibrated model, the answer is yes, financial repression is preferable. One reason is that financial repression occurs infrequently in equilibrium, about once every two hundred years on average. On the other hand, defaults happen more frequently, roughly every twenty years, due to the volatile nature of debt dynamics under default risk.

Your paper ends with the case of financial repression in a Monetary Union. Could you summarize the main challenges faced by a monetary union to implement financial repression as an optimal policy choice?

In a monetary union, the main challenge is that the incentives for fiscal adjustments are weaker. Financial repression may sometimes be necessary but it is costly, so governments should implement a fiscal adjustment to exit financial repression whenever they can. This is the case when financial repression is implemented in one country. In a monetary union, the costs of financial repression are shared among all the members. As a result, the incentives to adjust for high-debt countries are diluted. The optimal balance between insurance and incentives is more difficult to achieve in a monetary union.

> The crucial question is whether financial repression is a better option than default. In my calibrated model, the answer is yes, financial repression is preferable.

A large part of your research has focused on prudential policies. In one of your papers, you investigate the design of ex-ante macroprudential regulation when policy makers also have access to ex-post policy instruments, such as provision of liquidity, after a financial crash. Could you tell us more about the interactions between ex-ante and ex-post policies? How does the optimal design of one affect the other? More generally, what are the main policy recommendations you've drawn from your research on this topic?

Ah, yes, our "mopping up" paper with Anton Korinek! Anton and I wrote this paper because we found that there was some confusion in the literature around the notion of "overborrowing" in credit booms. It seems obvious, at least after the fact, that there was too much mortgage lending in the U.S. before 2008. At the same time, the notion of overborrowing is not obvious from a theoretical perspective. Financial crises come from financial frictions, and financial frictions tend to reduce borrowing below the optimal level, so how can they lead to overborrowing? This is the question that we try to clarify in the paper with Anton. We use a model in which crises take the form of fire sales between banks.

The main takeaway is that the relationship between financial safety nets and overborrowing is subtle and sometimes counterintuitive. The expectation of using safety nets ex post can increase banks' borrowing and leverage ex ante. However, this increase in borrowing isn't necessarily overborrowing; it can be efficient. To the extent that financial safety nets mitigate financial frictions ex post they make borrowing less dangerous ex ante. Therefore, the increase in borrowing should be welcomed rather than resisted by regulation. We discuss some properties that safety nets should have in the paper, as badly designed safety nets can indeed lead to overborrowing and moral hazard. Overall, we feel that the financial reforms that were implemented after the crisis, such as Dodd-Frank in the U.S., may have excessively limited ex-post financial safety nets.

Another topic you've worked on is the use of capital controls by emerging economies. Could you explain what are the main benefits of using capital controls for governments? How does that instrument perform compared to an alternative one like foreign exchange interventions?

There have been debates about the use of countercyclical capital controls by emerging economies in the last fifteen years. They offer advantages by providing additional policy instruments in a world with various frictions and targets. For instance, countercyclical capital controls can help shield these economies from shocks associated with the global financial cycle.

However, there are counterarguments, such as stigma---capital controls may scare away foreign investors. But there are ways of mitigating these problems. For example, establishing a code of good practices, limiting and predicting the use of capital controls, under the supervision of organizations like the IMF. There have been efforts in this direction but they were not very successful.

Some countries have used countercyclical capital controls, like Brazil in 2009-12. But overall, emerging market economies have been using foreign exchange interventions more than capital controls. This could be because foreign exchange interventions are perceived to have less stigma and are more readily controlled by the central bank, offering greater agility in their implementation.

Recent crises, such as the Great Financial crisis or the pandemic, have modified our understanding of the role of central banks. Moving away from the strict mandate of managing prices changes, there has been multiple calls over the last few years for more active central banks. How does that challenge the independence of central banks? Can central banks remain independent while becoming more active?

The role of central banks has always been multifaceted, and this isn't a recent development. The U.S. Fed for example has a mandate involving not only inflation and unemployment but also stabilizing long-term interest rates. And circling back to your first question, I do think that central banks have a central role to play in stabilizing government debt markets.

The role of central banks has always been multifaceted, and this isn't a recent development. I do think that central banks have a central role to play in stabilizing government debt markets.

But how broad should the central banks' objectives be? Should central banks for example try to address issues related to climate change? The "all hands on deck" approach to this type of questions runs the risk of distracting central banks from their core missions, and of generating a false sense of problem resolution. There is an ongoing debate on this question. Count me in the team that thinks that central banks should focus and deliver on their core missions.

Olivier Jeanne joined the Johns Hopkins Department of Economics in September 2008, after ten years in various positions at the Research Department of the International Monetary Fund. His research spans an array of applied and theoretical topics in international and domestic macroeconomics: capital flows, exchange rate regimes and currency crises, sovereign debt and defaults, international liquidity, and monetary policy. Outside of Hopkins, he is also a research associate at National Bureau of Economic Research (NBER, Cambridge MA), a research fellow at the Center for Economic Policy Research (CEPR, London), and a nonresident senior fellow at the Peterson Institute for International Economics (Washington, D.C.).

MACROECONOMIC RISK CHAIR

Paris School of Economics
48 boulevard Jourdan
75014 Paris

www.parisschoolofeconomics.eu

Publication:
December 2023

Layout:
Samuel Chich