Transcript of 2023 Thomas Laubach Research Conference Open Remarks and Session #1
May 19, 2023

TREVOR REEVE. Good morning, everyone. It's great to see so many familiar faces. My name is Trevor Reeve, and on behalf of the Division of Monetary Affairs at the Federal Reserve board, I would like to welcome you to the inaugural Thomas Laubach Research Conference. As a highly influential researcher and policy advisor, Thomas Laubach was a beloved figure in the economics and central banking communities. Thomas served as the Director of Monetary Affairs here at the Fed and was a trusted advisor to Fed Chairs and the Federal Open Market Committee. And to many of us, he was a close colleague and a dear friend. Today, we are honored to have several current and former Federal Reserve policy makers and senior officials with us along with many other distinguished guests and friends. I'm grateful that all of you have taken time out of your busy schedules to be with us today, and I would like to extend a special welcome to Thomas' family for being here on this special occasion. Thomas is widely known for his influential research including his innovative work on the natural rate of interest, which we will hear more about soon. Thomas saw the value of bringing together researchers and practitioners, to further our understanding of the economy and to improve the conduct of monetary policy. In this spirit, today's conference pays tribute to Thomas' lasting contributions and inspires us to carry on his legacy. In conjunction with this conference and to advance Thomas' passion for promoting research excellence, the Division of Monetary Affairs established the Thomas Laubach Research Award. The annual award recognizes important contributions to the monetary policy literature from Federal Reserve System staff. Yesterday, we were pleased to celebrate the inaugural recipient of the award, Kurt Lunsford from the Federal Reserve Bank of Cleveland. Kurt's paper, Policy Language, and Information Effects in the Early Days of Federal Reserve Forward Guidance, published in the American Economic Review,

provides new insights on how the nature of the FOMC's forward guidance shapes private sector expectations and, hence, the effectiveness of monetary policy communications. Kurt, congratulations once again. Let's give a round of applause.

[Applause]

TREVOR REEVE. A few housekeeping items before we get started with our agenda. The conference is being recorded and is being live streamed on the Board's website, federalreserve.gov, as well as the Board's YouTube channel. And as I'm sure you've noticed, the press is here as well, so please bear that in mind. We will be taking questions during a few of our sessions today, and we'll have microphones around the room to facilitate that process. We will not be taking any questions online, however. If you do have any questions about anything related to today's conference or just have questions about where to go, please find any of my colleagues around the room who are wearing a white lanyard. Lastly, I want to offer my sincere gratitude to the presenters and panelists for their participation in today's conference and to my team and partners from around the board who have helped make today's event possible and help make it a success. So thank you, all. And before we move on, we have a message from someone knows firsthand the iconic nature and contributions of Thomas Laubach.

JANET YELLEN. Good morning. I regret that other business prevents me from being with you, but I did want to add my voice to those who today are remembering Thomas Laubach and his contributions as a researcher and public servant. The establishment of this conference is a very worthy tribute to a remarkable economist and a remarkable person. I had the good fortune of working closely with Thomas first during my time as Vice Chair. Back then I was looking for an advisor to help me develop and communicate some of my thinking about monetary policy. He came highly, in fact, lavishly recommended. I knew his work, of course, on inflation targeting

and the natural rate of interest, but I didn't yet really know him. The person who showed up in my office was bursting with ideas, suggestions, and enthusiasm. As an accomplished economist, he had already made important contributions and begun a career as an academic, which he surely would have continued to great success, but Thomas had also spent time at the Fed, and he longed to be back in the arena doing the important work the Fed does on behalf of the public.

In working with Thomas, I found that he possessed a rare combination of penetrating, insightful analysis, superlative technical skill, and the ability to explain complex concepts to anyone. One of our projects was to advance the work on monetary policy communication that began in earnest under Chairman Bernanke. The FOMC had adopted a 2 percent inflation target and a statement laying out the longer-run goals of monetary policy. And it might have seemed like the revolution in communications was over, but I didn't think so, and one of the first speeches the Thomas worked on with me sent the very clear message that attention to communications would remain a top priority. In a speech that was about the need for policy makers to speak clearly, Thomas's ability to write clearly and explain things like optimal control in an accessible way was of great value to me and to the Fed

Later as the Director of the Division of Monetary Affairs, Thomas helped me as Chair in the task of making the difficult judgments about when to lift rates off the effective lower bound and in beginning normalization of the balance sheet. I know he played a big role also in the huge amount of research and other work that went into the FOMC's monetary policy framework review. Many of you here at this first Laubach conference had the benefit that I did knowing and working alongside Thomas, but let me close by addressing those of you who didn't have that privilege and might be wondering about the special qualities that have earned Thomas the esteem that so many of us hold him in. And to you I would say that Thomas Laubach was a

brilliant thinker, a bottomless reservoir of economic knowledge with a broad range of skills of the highest caliber and with the perspective and wisdom about how to apply it all that seemed far beyond his years. But I'd also say that just as important as these advantages was the incredible enthusiasm and energy he brought to his work and to his co-workers. Economic policy making often requires making difficult, sometimes even agonizing choices, but economics was never the dismal science for Thomas. He so loved his work that he could make the most challenging task sheer pleasure, even fun. Of all his abilities, that was Thomas' greatest gift and one that he generously shared. Thank you for the opportunity to say a few words and congratulations to the researchers presenting today. Have a great conference.

[Applause]

TREVOR REEVE. Thank you and apologies for the technical glitches, but I think I think the message came across. As we were thinking about the agenda -- and special thanks to Secretary Yellen for providing those kind remarks. As we were thinking about the agenda for this conference, we knew there was one person who had to be a part of it, and we were thrilled that he generously agreed to deliver our keynote address. John Williams became the 11th President and Chief Executive Officer of the Federal Reserve Bank of New York in June of 2018, and in that capacity, he served as the Vice Chair of the Federal Open Market Committee. Before coming to the Federal Reserve Bank of New York, he held multiple roles at the Federal Reserve Bank of San Francisco, including President and CEO as well as Executive Vice President and Director of Research. And I would be remiss not to mention that John began his career as an economist here at the Board. So, with no further ado, it's a pleasure to welcome President John Williams.

[Applause]

JOHN WILLIAMS. Well, thank you, Trevor, and good morning, everyone. It's a special privilege to participate in this conference honoring my dear friend, colleague, and co-author, Thomas Laubach. Now many of you know Thomas personally, and all of you know him through his work. He was a brilliant economist, trusted colleague, and an inspiring leader. We miss Thomas' dedication, his good cheer and sense of humor, and his impact in everyone he touched endures. Now our collaboration started in the fall of 2000 with a shared interest in figuring out how to measure the elusive natural rate of interest, otherwise known as R-star. Now at the time, neither of us had any idea that our research on this topic would continue for 20 years. Now Kathryn Holston, who's here today, and I are carrying on this work. It's one of the many ways that we can honor Thomas and his memory. Now, because of the unprecedented and pandemic-related shocks to the economy, we paused our regular publication of R-star estimates back in late 2020, and this conference is the perfect venue to announce that we are resuming them. But first, I want to take some time to share R-star's origin story, if you will, which is really how I came to know Thomas. Okay, before I go any further, I have to give the standard Fed disclaimer that everything I say reflects my own views and not necessarily those of the Federal Open Market Committee, the Federal Reserve Bank of New York, or anyone in the Federal Reserve System.

The idea of a natural rate of interest has been around for over a century, when the Swedish economist Knut Wicksell wrote about it back in 1898. Now early on, it was recognized that it's not something that can be directly observed or measured. As the economist John H. Williams wrote in 1931. and I love this quote and thanks to Athanasios Orphanides for finding this, "The natural rate is an abstraction. Like faith it is seen by its works. One can only say that if the bank policy succeeds in stabilizing prices, the bank rate must have been brought in line with

the natural rate, but if it does not, it must not have been." Now this topic gained renewed relevance in 1993 following John Taylor's famous description of a monetary policy rule that incorporated an assumption of a natural rate of 2 percent. But policy makers and natural questions, pardon the pun, were is 2 percent the right number? Does it change over time, and how would we know? In fact, those are the very questions that then Federal Reserve Governor Larry Meyer posed to the Board staff back in 2000, and that's what brought Thomas and me together. So in July of that year, I returned to the Board from a stint at the Council of Economic Advisors shortly after Thomas had joined the Board staff from the Kansas City Fed. And we hit it off immediately. It just so happened that both of us were actually thinking about this natural rate of interest question, and since Governor Meyer had raised the subject, we moved quickly to develop an approach to answer his questions. Although we didn't know it then, we were following exactly the problem that was set out by the other John Williams back in 1931. Simply put, Thomas had a hammer, and we found a nail, new nail. That hammer was a common filter, which he had used in earlier research, and the common filter is about inferring the behavior of an object from its effects on other objects. Of course, the nail was R-star.

On December 14th, 2000, after just a few short months of working together, we wrote our ideas and results in a memo to the Board of Governors. The memo started with a bold declaration, and remember we were relatively junior staff economists, and this memo was a first report on a broader project to study alternative definitions and estimates of the equilibrium real rate R-star and to evaluate the usefulness and the conduct of monetary policy. And so, the Laubach Williams, or LW model, was born. Starting in May of 2001, our estimates and others made their way into the staff Blue Book prepared for the FOMC. Looking back and going through these documents from 2000, I am struck by how quickly the ideas came together, but

importantly I must acknowledge the unwavering support that we received from senior leadership at the Board to bring new research to policy makers. Many of those senior leaders are in this room, so I thank you for that. Within a year of writing the memo, we put out an academic version of the paper, and then interestingly in that first working paper, we actually did have a section on time varying R-star's implications for monetary policy, but for reasons that only the editor can understand, and for all those who write research papers, we've all been there, the editor decided to cut that section, leaving it for future research. [laughter] Okay. So fast forward a decade, and an entirely new question about R-star arose. Why had estimates of R-star fallen so low? In fact, when you go back to our original work, we were talking about why is R-star so high. This is a period of rapid productivity growth in the late '90s, early 2000s. So now the question was the opposite. It's a topic Thomas and I explored in our Redux paper, and then in 2017 Kathryn, Thomas, and I expanded the set of economies for which we estimated R-star. Indeed, before the pandemic, historically low estimates of R-star characterize advanced economies across the globe, and this is illustrated by the Holston Laubach Williams or HLW estimates of R-star in the United States and the Euro area. This is using data through 2019 in figure 1. Ah, it worked. So the blue line is the Euro area, the I think black line is the United States, and this is going through the end of 2019. Both of the estimates of this, and as you can see in the figure, were about a half a percent, far lower than the estimates from preceding decades.

One of the features of the LW and HLW models is that both are designed to be very flexible and let the data speak in measuring changes in the natural rate of interest and other variables. But even this flexibility has its limits, and the economic turmoil brought on by the COVID-19 pandemic went far beyond what the models were originally designed for. In particular, the pandemic violated two key assumptions about the nature of the macroeconomic

disturbances underlying the models. First, the Kalman filter statistical method that really serves as the model's workhorse. It assumes random disturbances to the economy following normal or bell-shaped distribution. Relative to historical experience, COVID-19 represented an extremely rare tail event in terms of the effects on the economy. Second, the models assumed that these disturbances that are serially uncorrelated over time, which is at odds with the sequence of shutdowns and reopenings that we saw with COVID-19. So the highly unusual nature of the effects of COVID-19 is illustrated, now I'm going to get into a little technical thing here by what's called the model auxiliary residuals, and this is for the output gap or the deviation of GDP from its potential or natural level. So here, again, I'm going to show estimates going back to ending -- well going back to, for the United States and Euro area in figure 2. So the way to think about these residuals, which are shown by the squiggly lines here, is these are really related to the difference between the data, the GDP data, and what the model would predict for that data, and the dashed lines in the figure represent or indicate two standard deviations. So as you can see over history, the residuals from the model bounce around, but generally within two standard deviations, as you would expect. But then when you go to the pandemic period of 2020, for the United States, the residuals are as large as 15 standard deviations, and for the Euro area, they can exceed 20 standard deviations in some period. So you can go on your calculator and see that in a purely statistical sense, the probability of such an event occurring even once, much less twice, is infinitesimally small given the history of the behavior of GDP over most of our sample. So after the onset of the pandemic, we suspended publication of the R-star estimates due to the extreme volatility and the elevated uncertainty about how the pandemic would evolve. So I'm pleased to report, as I said earlier, that starting today, I think starting now, we are relaunching a regular publication of both the LW and the HLW estimates of R-star and everything that comes with that

on our website. We have the estimates along with all model documentation are available in the New York Fed R-Star website and will be updated each quarter, as we did in the past.

So we had to address these two violations of the model that I mentioned, and they were caused by the pandemic. So to do that, we made two modifications to the estimation to both models, and this is detailed and described in detail in a paper measuring the natural rate of interest after COVID-19 that we posted this morning, again, on the New York Fed's R-star web page. And the goal here was to modify the model in ways that addressed the issues caused by the pandemic but maintain the basic structure in intention and flexibility of the original models. So the first thing we did is we incorporated a persistent supply shock that captures the effects of the pandemic that reduce the economy's potential. Think of a variable that's affecting shutdowns and the effects of the pandemic on the supply side of the economy. We measure that shock using the Oxford COVID-19 government response stringency index for each economy, and this index aggregates measures of government containment and shutdown policies, and we have it for each of the economies that we study. Now, they decided to stop publishing this index at the end of last year, always a source -- when you're writing research, suddenly the data change on you -- so what we're doing is we're assuming that for each economy that this index declines smoothly from where it was at the end of last year to zero over '23 and 2024. Second, we used a statistical procedure that accounts for outliers in the model's estimation. Specifically, from the second quarter of 2020, really when the pandemic hit, through the fourth quarter of 2022, we estimated a time-varying process for the variance of the shocks that hit the various economies. And this procedure intrinsically places lower weight on periods where there are huge, large outliers. Our results show that outliers were particularly large in 2020, and you saw that in the figure I've shown here. By comparison, the estimated magnitude of outliers are relatively modest in 2021

and 2022, and our estimation is consistent with that. Starting in 2023, so in this year, we assumed that the distributions of the shocks to our models are no longer affected by the pandemic. So if you will, we're moving back to the standard historical view of the variance of the shocks.

So now that we have the modified versions of our models, we're able to -- we've estimated them through the end of last year using published data, and we have three key findings. First the modified estimation procedure yields results that are overall quite similar to those from the original models during the pre-pandemic period. In other words, our goal of maintain, you know, making modifications to deal with the pandemic but maintaining the basic structure of the models was successful. Second, the current estimates of R-star are similar to those estimated directly before the pandemic. I'll show a picture on that. And third, the estimates of the natural level of output at the end of 2022 are much lower than the models predicted before the pandemic. So the current HLW estimates of R-star in the United States are shown in figure 3. For comparison, we also show the estimates using a version of the model that we didn't put in these adjustments to take in account of COVID or outliers, and we hold the parameters' values fixed at estimates using data through the end of 2019. So that's the blue line. So this tells you what the HLW model for the U.S. would look like if we froze the parameters at the 2019 Q form level values and then just rolled the model estimation forward. The gold line is our current version of the model. So what's interesting here, again, as I said, the two sets of estimates are quite similar up through 2019. There are some differences, but they're very similar. They differ sharply during the acute period of the pandemic, however, when the estimates from the unmodified model shown by the blue line exhibit large swings due to the presence of the sizable outliers I mentioned. What happens is these outliers and GDP and inflation then feed into the estimates of R-star and other variables.

Now what is interesting in this case for the U.S. is that at the end of the sample, the fourth quarter of last year, the estimates of R-star are actually quite close to each other. That is, again, I would emphasize the modifications we made to the model do not interfere with letting the data speak. So based on the new R-star estimates for Canada, the Euro area, and the United States, we don't see -- we see no signs of a significant reversal of the decline in R-star estimates that was evident in prior decades. In fact, in all three economies, the HLW R-star estimates in 2022 are within two tenths of a percentage point of the corresponding estimates of 2019. But the largest difference between the models estimates pre- and post-pandemic relate to the level of each country's potential output. So figure 4 compares the model's measure the natural level of output based on estimates using data through the fourth quarter of 2019, and that's the blue line, that's the top line. At the end of 2022, the COVID-adjusted level of potential output, which is shown by the bottom black line, is a little over 4 percent below what the pre-pandemic projection for the fourth quarter of 2022, and about half of that or nearly half of that shortfall is explained by the COVID shock measure with the remainder -- and the remainder of permanent change in the natural level of output, which is shown by the gold line. So in summary, according to our model estimates, the main longer term consequence from the pandemic period is a reduction in potential output, but the [inaudible] in R-star appears to be relatively modest. Importantly, there's no evidence that the era of very low natural rates of interest has ended. So these estimates indicate that R-star today is about where it was before the pandemic, but where is it headed in the future? Of course, this is impossible to know with any certainty, but one way to gauge how forecasters perceive the future of R-star is to use forecast data to estimate our model. So figure 5 shows the HLW-based estimates of R-star in the U.S. through the end of 2024. We have a crystal ball here, you know. It uses published data for the first quarter of 2023, which we didn't include

in the earlier charts, but we do have, you know, GDP, inflation data for the first quarter of 2023, and then it extends that data using the most recent blue chip forecast for inflation, GDP, and interest rates as proxies for data from the second quarter of this year through the fourth quarter of next year. So what you see here is really interesting is the resulting estimate of R-star just on the data through the first quarter of this year is about half a percent and then subsequently falls to slightly below zero. So evidently the value of R-star implied by private forecasts is, if anything, even lower than today's estimate. Now time will tell whether this turns out to be the case or not.

So let me conclude by sharing something I learned from Thomas and I hold dear. It's actually hidden in the last line of the abstract from our first published paper. Now it says "estimates of the natural rate of interest however are very imprecise and subject to real-time measurement error." Now to some people that might sound like a pretty negative statement about your research, but it's also what makes empirical research so exciting and so challenging. There's no one right answer. There's no single way to view the world. Empirical research is a process of continuous learning and adaptation fueled by perseverance. People come up with different approaches, and we learn from all of them. So while we started with the problem of how to estimate R-star, it didn't end with that first board memo. Rather, it opened the door to new questions, which led to further challenges and even more research, and this is the legacy of Thomas Laubach, and it's what brings us together today. So thank you.

[Applause]

JOHN WILLIAMS. So I think -- I think I got about 10 minutes to answer any questions from the audience that you may have and see if anyone is courageous enough to ask the first question. We have microphone runners in the room. This can't be a shy group. I mean, Kurt, you

should -- ah, there we go, Don. After I was praising you for your support of Thomas and me and this work, so here we go. [laughter]

DON. So I -- I'd like to take your focus back to the global financial crisis and that drop in R-star. That's really the most striking thing on the chart, goes from two and a half to a half, basically, and then doesn't recover. So I admit, my forecast in 2009 was that it would -- R-star would recover just about back to where it was. So -- and we faced a decade of discussions of secular stagnation and whatnot, but what do you think really -- what do you think happened? Why did -- why did we get such a sharp cutback in growth, really, that showed up in R-star in that period.

JOHN WILLIAMS. Yeah, so it's a, you know, it's a, I think, hugely important question and a very hard one to answer. I will mention that, you know, in both of our models we do link the estimate of R-star [inaudible] trend output growth along with another unobserved variable, and so what you see, as you mentioned, is that we also saw a significant slowdown in trend GDP growth. So the two are happening. They're not exactly the same. We saw both R-star declining both for trend growth reasons and for other reasons. So I do think there was a period where we had gone from a very, a relatively high estimate of trend GDP growth with them you know, the dot com boom and the internet and that to a more normal or slower level of trend growth. And then I think a lot of the research has highlighted some longer term trends, which are happening in the background, demographic trends, perhaps trends around the demand for safe assets, which will probably even, will likely reinforce during that period, which would tend to push estimates of R-star down. So I think there's a lot of reasons we now understand years later why the trend of trend growth and the trend of R-star have moved down, but it doesn't answer the question you started with is why has it happened relatively suddenly. And I think that my own kind of thinking

about this is, you know, these are -- it's a model that's trying to separate the transitory or the cyclical from long-term, highly persistent, and permanent, but of course there's a lot going on that's highly persistent in the economy. So if you think about it, the period of the '90s and into the 2000s was a period, a very strong economy, very strong aggregate demand. We have the technology boom and productivity boom followed by the, you know, the rise in demand for housing. So we had some pretty big rightward shifts of the IS curve, if you will. And then after the financial crisis, we see the IS curve, again, using that lingo move far to the left. So I think the model is getting pulled by that switch of the IS curve. I do think the underlying structural factors were actually moving more slowly. It's just that, you know, when the housing kind of bubble, if you will, crashed, then we had the crisis, you know, it happened much more abruptly in the data, so I think a lot of the research that's looked at these drivers of productivity or trend GDP growth in R-star would say that it was actually the thing was moving over many years and slowly it shows up in our model estimates as more abrupt because of some of the other macro factors that were happening. But that's my, you know, one of the great, you know, one of the things that some colleagues kid me about is they start talking about models as if they were people. And so I'm getting a little into doing that here.

AUDIENCE MEMBER. John, thank you, thank you very much. Interesting paper. To follow up on Don's question, my question, so your forecast comes as a bit of a surprise, or your forecast conditioned on private forecasts. Number one, what is the key driver there. Number two, if we coming out of this business cycle, we have an AI boom with substantial pickup in productivity growth, could you see R-star rising back to previous levels?

JOHN WILLIAMS. Yeah. So the second question, absolutely. I mean this is a model where the estimates of, if you will, G, which is the underlying trend growth of the economy, and

Z, which is this other factor driving R-star are treated as random walks. So, if there is a new boom in productivity or a shift in in labor supply or other factors affecting the supply side of the economy or the natural rate of interest, clearly, those could surprise us. So I think that that's absolutely possible. Going back to the first question, I did look at that. So what's going on, and again, we're just using the Blue Chip forecasts. These are just, you know, this is not, you know, my particular view or something. We're just using that data and running the estimation procedure. So I did look at how the model -- again, I'm going to talk about the model, you know, like this oracle. So I went into the room and talked to the HLW model, and basically what's happening is you're seeing a pretty significant shift in the shocks hitting output and inflation from very positive shocks in the model in '21 and 2022. Basically the model is seen, you know, as the economy bounced back, this U.S economy bounced back and inflation took off, a good part of that, according to the model, is driven by shocks. And then if you, again, infer this from the Blue Chip forecast, those shocks are actually reversing. So this is a bit of a reversal and maybe some of the -- in the view of the forecast, the pandemic shocks. Now, in the end of the day, I would argue that the -- I wouldn't take too literally the fact that R-star goes up or down by a small amount. I think the thing that's clear in these forecasts is given the path for real interest rates, given the path for output, they basically see inflation coming down relatively quickly, output getting, you know, output gaps getting somewhat smaller, and that happening consistent with a very low, you know, natural way of interest. Who has the mic? David.

DAVID. Your estimate of R-star being as low as it is raises a question, well, why is, why is inflation as high as it is? Is that simply to anthropomorphize the model, is the model telling us, therefore, be patient? It's lags in the macro system, and this will all work out?

Monetary policy is approximately restrictive enough to bring inflation down, or what is the model telling us when you go into the room with it?

JOHN WILLIAMS. Well the model does have built into it, long not variable, they're fixed, relatively long lags of monetary policy, and the transmission in the model is pretty standard. Real interest rate gaps feed into output gaps, output gaps then feed into inflation, and there's quite a bit of inflation inertia in the model. So there is that kind of inertial process. So what is happening in 2021 and '22? What's the model telling us is that it's saying that there were significant increases -- or positive demands to output gap that were not explained by monetary policy or shifts in the output gap not explained by monetary policy. By the end of the sample, the monetary policy, the very low real interest rates of the past two years did contribute to even higher output gap, as you would expect, and the inflation is in part driven by the very significant positive output gaps. You can see all of this in our paper that we posted today, but is -- in terms of what the output gap is, along with significant inflation shocks. So part of it is driven by, you know, shocks to demand outside of monetary policy. Some of it's monetary policy, but a lot of it, again, is some shocks to inflation. So in terms of how to answer both of your questions, you know, I'm not going to opine on what monetary policy should or should not do, but I would say that the, you know, in this model clearly how you view how these shocks evolve is very important, because they are a significant driver of what's happened, at least according to this model, and I'll just reiterate what I said at the end of my prepared remarks is there's lots of different models, lots of different ways to slice and dice this data to understand this. I know we had Bill.

BILL. So, this is working? Yes. So I wanted to urge you to talk a little bit further into the future. So we're all asking questions about kind of what's going to happen to R-star, but over

the next decade or something like that, do you have any thoughts about where R-star is going? So productivity was raised earlier, but they're also -- debt to GDP is now really high relative to where it was pre-pandemic. There's been deglobalization. There's potentially shifts in the demographic effects on R-star. Charles Goodheart has a story about how for a while demographics were pushing interest rates down, now they're going to be pushing interest rates up. Do you want to offer any speculation on kind of where is R-star going to be 10 years from now?

JOHN WILLIAMS. Well, I think it will depend on all the factors that you described and others. Demographics, I think, the evidence from the research on this is pretty convincing. I actually would point to some other research that highlights birth rates and longevity beyond just the baby boom kind of dynamic. So I think there's some global factors. I always, you know, emphasize R-star is essentially a global concept. There are global factors around demographics, you know, you think about China and other countries, they're probably going to push R-star down. Obviously, there's a possibility the technology is going to transform this. Now, which way it pushes R-star is kind of fascinating. If you go back to some of the discussion years ago about secular stagnation, some of those theories were really based on how technology, if you will, the production function has changed, and maybe that is actually reduces the demand for capital rather than increases the demand for capital. So I think these are really interesting questions. It depends on how technology changes, the form it takes will affect the influence on R-star and other variables. The other thing that tends to come up in these conversations is really about climate policies and investments in green energy, and again, I think that depends on kind of the details. How is it financed? How big is the net investment increase into green investment versus reducing investments in fossil fuel? So I think that, you know, to me what these models and the

work of actually many people in this room and others teaches is they give us a framework to think about this is really about global, you know, supply and demand of savings and investment, and as we think through, you know, deglobalization, we think through some of these other things, it does give you an apparatus to separate between supply, demand, and between, you know, what are transitory and permanent. I do, so if you would ask me kind of where my, kind of how I view 10 years from now, I think it's these demographic and other factors are going to be, I would say, more of a negative on R-star net over the next 10 years. I do think this issue of fiscal debt, debt to GDP, clearly that's research that Thomas himself had made really important contribution to. It would, all its equal, push estimates of R-star around the world up. Now the real question is we watch the dated, how big are these various factors and how do they all fit together. The reason we really wanted to get, you know, these estimates back in running is it's one useful way to gauge these various influences. Obviously, there's lots of other models and approaches. Krishna? I think this is the last question based on at least my watch. Yeah. Yeah. Like Faith, you know, like I said with R-star, you just have to believe that the –

KRISHNA. So as you know, one alternative way of thinking about some of these issues is to frame them in terms of say a 10-year R-star as opposed to an overnight R-star, and in thinking about the sort of difference between framing things that way, you've obviously got the evolution of an equilibrium term premium to be thinking about in this process as well. So do you have any thoughts as to how that side of things may evolve? What would be sort of implied consistent with your estimates of the short, neutral rate?

JOHN WILLIAMS. Well, I think, you know, I think it's a really important question because, you know, based on, you know, kind of how we think about the term premium and risk premium in general, if you think you're in a world, let's just postulate we're in a world of very

low R-star, then you get into, as we've, you know, studied, you know, many papers and seen in the real world, that you're at the effective lower bound more often. You're using perhaps other monetary policy tools such as QE more frequently, perhaps the downside risks, the risks are skewed in different ways. So I do think the way, at least with the lens I think about this and thinking about longer term yields or, you know, neutral rates and things, is that if you really think that the neutral radius is very low, then it does affect the distribution of outcomes and behavior in the economy, which then feeds into risk pricing. Again, it's a reason to, you know, analyze these kind of models to understand, has the world changed in the past four years, and in what ways has that changed. So that's the how I think about it. I think that were done. Thank you so much for the questions and discussion.

[Applause]

TREVOR REEVE. Thank you very much, President Williams. That was great, and I'm sure we'll get lots of downloads of the estimates over the course of the day. Now, we're on to our first session, kind of continuing the theme here. This is on consequences of declining natural rates. And it's a privilege to introduce our session Chair Annette Vissing-Jørgensen. Annette is, of course, a very well-known, preeminent economist, and we're also delighted that she is a senior advisor here in the Division of Monetary Affairs.

ANNETTE VISSING-JØRGENSEN. Thank you very much, Trevor. So today's first academic session has a star-studded lineup. Stephanie Schmitt-Grohé is one of the leading macroeconomists of her generation. She's going to speak on the topic closely related to John's keynote. Stephanie graduated in '94 from Chicago, and I actually didn't know her first job was right here in the Division of Monetary Affairs at the Fed along with a bunch of people who have since become a whole bunch of the Fed bosses. So, you know, it's good to see that people do well

starting in MA. She has since, of course, had a glorious career in academics and since 2008 has been a professor at Columbia. So the discussant, Christian Wolf, I'm pretty sure will be one of the leading macroeconomists of his generation, although he's just getting started, graduated in 2020 with his PhD from Princeton, and then after a year at Chicago got the job that everyone wants at the MIT Economics Department. So with no further ado, please go ahead Stephanie.

STEPHANIE SCHMITT-GROHÉ. Good morning. It's an enormous pleasure to be here. I knew Thomas, and you can feel this here, Thomas was really beloved. And I think he was so beloved for many reasons. Let me just highlight two of them. This guy was creative. This guy was original, and this guy had an eye for big questions. And he pursued that with a lot of love for economics. So that is impressive about him, and I guess in this audience everybody read Volcker's last book and remember that he ended with a big credo that we should invest more in schools of public service, that we should create more public servants to run this country, and if I could think of one person who maybe somebody wants to endow a Chair of public policy and public servant for it should be called the Thomas Laubach Public Policy Chair. He was somebody who dedicated his career to public service. He was an example of kindness and somebody with values and principles. So I think he's a great inspiration.

Okay. So having said that, I could equally be here going back to his economic contributions because this is the 20th year or 20th anniversary of the publication of the R-star paper of John and Thomas. So I think that paper has an importance in the literature that would induce people to have this type of conferences, conference in honor of the 20th year of the publication of this in this paper. So I could be equally here for this, and there are other contributions that Thomas made where we could have an anniversary of that, a conference. Most of us, especially looking at the academics in this room, we know that 20 years after we published

a paper and we go to Google Scholar hoping that somebody still cites our work, we see that the little piles have gone to zero, whereas the R-star paper had some impact when it came out, but the R-star paper had a big revival and still is a paper that collects a lot of citations and has a lot of impact. And so this is something rare. Not many people have several of these contributions, while also dedicating most of their work to public service. So let me see, yeah, if I can operate this. Yeah.

So this paper is squarely Thomas' paper with John on R-star. So we ask, what do long-run data tell us about R-star and Pi-star? So I hope Christian doesn't get scared, but I think he has seen me giving it under this title. This is not the title of the paper but it is the same content. So let me first tell you what we do, and then I will relate it to Thomas' paper. So we are not the first to observe this, but business cycle analysis is a story told on post-war data, and the question is whether that's the only way to tell that story or is there something we can learn from longer-term data. So what we do here is we look at learning something about R-star and Pi-star using just longer data since 1900. Marco Del Negro and his co-authors also have done this, so I don't want to claim we are the first, and we have a little bit of a different focus than much of the literature. So it was great that John went ahead of me, and he showed you a bunch of figures. And many of these figures were time series plots of R-star. So much of the profession has been interested in learning what is the time series behavior of this latent variable, how big are the error band, so how imprecise is it, and then as the questions that came here, the other question is please give us a forecast of R-star. So that is what many of these papers are interested in. We are interested in something slightly different. We are interested in what are the effects that we can estimate from long-term data of variations in R-star on the macro economy. In particular, we want to understand what if there's a change in this natural rate of interest, what happens to the level of

output. So our paper, that will become clear in a moment, is squarely focused on that, and we, as the title says, we don't just have R-star, we also have Pi-star, and so we are also interested in what are the [inaudible] of Pi-star but also what are the effects of innovations in Pi-star on the economy. So we, in this paper, we ask -- too, it's a purely empirical paper, so there's not going to be -- I can comment on how I think our empirical results inform the validity of various theories, what moves us down, what are their consequences, but I will not provide -- we are working on that, but we don't have a very good theoretical explanation for what's going on. So this is going to be a purely empirical paper. And we are asking how do natural rate shocks affect output and inflation in the long run and in the short run, and then we have a little bit more current application, namely to which extent the current spike in inflation is driven by the permanent component of inflation. And these are two completely separate questions, but they can be addressed within the same empirical framework that we apply.

So before I start with the empirical model, let me explain what we find. So to our surprise a little bit or in the long run data, we find that a permanent fall in R-star, which we call the real interest rate, causes a really large, once and for all, downward shift in the level of output growth. So and the number we estimate is that if the natural rate of interest declines by one percentage point, then there was a -- I mean of course output keeps growing, but there's one shift down, some of the pictures that John had had a shift down of 8 percent. So we get a very large negative effect of this. And when we were thinking about do we know any theories that can explain that, then we only came up with secular stagnation theories, and we find that this effect is there even if we restrict attention to non-zero lower bounds period. So our main interpretation of our findings is there is a void where you can come up. Maybe somebody in the room knows the theory, but I don't know any good theory that can explain that. So that is what we learned about

how R-star interacts with trend of output. And the second answer we give, which might be comforting to this audience, is if you want to know how much of the recent increase in inflation is due to an increase in its permanent component and you follow the usual tradition of telling business cycle stories based on post-war data, you would conclude that the vast majority of the increase in inflation is due to its permanent component. But if you look at telling the story on long-term data, you would conclude that actually econometric or empirical exercise, that only a very small portion of the increase in inflation that we observe during COVID is reflecting an increase in the permanent component of inflation. So these are our two results. So our work builds on -- we cite more papers than one in this paper, but let me at least highlight how this builds on literature created by Thomas and John when they were pretty junior economists, not that many years out of graduate school, and they wrote a paper, and maybe I misinterpreted you, I think when I had asked you in 2000 do you think this is going to be a seminal paper where people will hold conferences on 20 years later, he said say, no, no, no, I don't think so. But you never know what sticks. You never know what sticks. This one stuck, and this paper created a literature on which many people have contributed. And I just want to make something clear, because this is a room full of central bankers. I just want to say what in that paper this concept of the natural rate of interest means and what it could also mean, but it's not how it's estimated in that paper, but we follow Laubach and Williams.

So there's, first let me tell you what it is not the natural rate of interest, at least in an empirical way how it is estimated here. Some people have this idea or this magic real interest rate, and if the Central Bank facing sticky prices is able to put the nominal interest rate so that it hits this magic bullet of the true real interest rate is going to get full employment and stable inflation, that is not exactly it. What is being estimated here is what is the permanent component

of the real interest rate. So that is what that is trying to be estimated. And yeah. So and they -- that concept has, you know, prevailed, and we are also following the tradition. But let me tell you -- sorry you have to now endure a model, but this is an academic conference, and let me explain a little bit. I apologize to the family. That will be very boring. [laughter] So we have a humble ambition. We want to explain the behavior of output inflation in the short-term nominal interest rate, the federal funds rate. So we have three variables. Y is the log of output, pi is inflation, and I is the nominal interest rate. And then we start making assumptions, and I'm going to point out where they actually deviate from Thomas' and John's assumption. So we define what the hat means is a stationary variable. So that is the business cycle component. So of output, we are saying output, the cyclical component of output, y hat is the difference between output and its trend, and its trend has two components, two big X's. One is technology, something technology, and then there is XR which as will become clear in maybe three, four minutes is what we call the natural rate of interest. And our paper is to a large extent concerned with estimating this coefficient little delta in front of XR. Why? Because you just have to trust me for the moment, if there's a permanent change in the level of the natural rate, which is XR, then delta tells you which way the trend level of output moves. So we will assume that the growth rates of the excess are all stationary. That is, we follow the tradition in much of time series study of GDP, the GDP growth is stationary. So output is integrated of order one, whereas the tradition in Laubach and Williams is to assume that output is integrated of order two. That is to say that the growth rate of output has a unit root. So as you can see from the first equation there, we are not assuming that, and we try to gauge a little bit for that possibility. So that is our motive for doing this.

Why is that relevant? Because people -- John can tell me if I'm putting words into their mouths or thoughts into their heads that were never there, but I think people were led to making

a relationship between permanent changes in the growth rate of output and the permanent component of the real interest rate because they just look at the Euler equation, they see on the left-hand side a growth rate of consumption and on the right-hand side they see the real interest rate, and they wanted to create a connection between these two. But I think that -- and then that leads you to a telling of a story in which there are permanent changes in the growth rate of TFP, and they are related to permanent changes in the permanent component of the real interest rate, which I think is something you can do, but it's not that clear how it matches with the fact that usually output growth is stationary, or that is to say that output is not integrated of order two. So that's an assumption we make, and we impose there, and then this allows us to -- or an implication of that is that this parameter delta, what our paper is about, is going to tell you what happens if you have a natural rate shock. The second equation tells you that we assume that the signaling component of inflation is -- or the permanent component of inflation is just its own XM shock. What does that reflect? It reflects that some long-run, non-naturality assumptions. So it's not possible in this model that permanent changes in productivity change the long run or the permanent component of inflation. You could think about XM, what people in this room do, right. So the only way out from giving that interpretation is if you think you live in a regime of fiscal dominance, which I don't think is for the United States the right paradigm, but unless you live in a regime of fiscal dominance, XM is what the Central Bank, I believe, controls. And then we have the nominal interest rate, and it has two permanent components. One comes from the long run inflation rate, which is XM, and then it has its own innovations XR, which is the natural rate of interest. You might look at that equation and say how come the coefficient in front of XM is 1. That is to say that the rate of inflation or the permanent component of the rate of inflation will not have an effect on the real interest rate, the difference between I and Pi in the long run.

That is known as the long-run Fisher effect, and we have some sensitivity analysis around that, but we do believe that in the long run the Central Bank cannot influence a real interest rate, and our result is actually robust to that. And so this is what we are trying to learn in this paper is, roughly speaking, what is delta, and what are the effects of shocks to XR, the permanent component of the real interest rate and the permanent component of the nominal interest rate.

Let me say something, I think I will relieve you from equations. It's going to be one more slide, and then we're going to go to the results. But we are going to -- as Thomas and John did, we are going to set this problem up as a problem of latent variables. So these are everything on this slide is things that the econometrician doesn't see, and he tries to learn something about, so that is to say we are setting it up as a problem of a common filter, and following the literature on DSGE models or borrowing a little bit from them and mixing it up with SVR models, we are setting up a structure in which the stationary components of output inflation, the nominal interest rate is an AR1 process. It's going to be -- or is auto regressive of first order, and this is annual data, and then we are having this big C matrix in which all the innovations to this economy can contemporaneously affect the variables of interest. So, for example, it is possible that these permanent shocks either to the natural rate of interest, to TFP, or to the permanent component of inflation, not only have long-run effects, as described in the first three lines, but they also have effects in the short run, which are captured by the matrix C. So for those in the room who are familiar with the work of Marco Del Negro, Marco Del Negro's paper would be first, third, and fifth column is a zero. It's a straight zero. So no short-run effects of the permanent component. We assume that the law of motion of these exogenous shocks, given in the second rho, has rho and psi are diagonal, and we give -- we will give priors to the values of rho and prior to the value of psi, so let me also make a relation to the, Thomas' and John's seminal paper, if you ask

yourself you can see that from here, what is element rho, 5, 5, that is to say how persistent is the change in the natural rate of interest? Their paper is a paper in which that rho is a 1 by construction. In our paper that rho is estimated, and we give it a prior of 0.3, and I don't know if Roberts is in the room, but we went back knowing we would see John today and looked up what is the justification given in that paper for putting a rho of 1, and it was a -- it was actually a paper by somebody here in the Board now. Maybe comment more on that. Anyway, we are not going to go for having permanent changes in the natural rate. So we are not going to go for output being integrated of order 2. And then we also have to say something about this psi matrix, which tells you how big are the innovations to the natural rate, and it will come up in a moment how we cut -- how we put the prior on that. Again I'm interested in element 5, 5, because I want to know how big can these shocks to be the natural rate, and we will come to the result that the variations in the natural rate follow a super cycle. So we're going to say there are two super cycles -- oh, that's what our empirical work finds, since 1900, and there will be a debate about is this result influenced by our assumption how big we make psi, 5, 5, meaning how big we allow to be the standard deviation in innovations to the natural rate of interest. We all -- I mean this -- nothing is observable here. So we whip out the common filter. As John said, Thomas knew how to unleash -- to put this weapon on data, and then the Karman filter just tells you there are some things you can see, and you can tie those things in a precise way to things you cannot see, which are the observations equation, and I just want to make two comments on the three observations equations I have. You can see here that in the first line, this parameter delta that I'm interested in, which is the growth rate of GDP, will depend on delta, and so you will see how it can -- a positive delta can push down the trend path of output. And the last thing I'm going to say about three, four, and five is we do have measurement errors in here, and those measurement errors are

try to handle a little bit of the excesses of periods of big crisis. So John said they stopped publishing their R-star numbers when COVID came, and we hope that some of the unusualness of COVID is in our measurement errors, and that's why we have that. Okay. So do I need to -- no, I'm not going to -- how many more minutes left? Can I go more technical.

ANNETTE VISSING-JØRGENSEN. No, you're doing fine. You have 15.

STEPHANIE SCHMITT-GROHÉ. Fifteen. Okay, no, if I have 15 fifteen minutes, okay. So I'm not -- I'm not going to say much more technical. So let me just tell you what is the data we use, and so part of, as I said, the idea of that paper is to deviate from the usual paradigm that business cycle analysis is a story told on post-war data, as other people in the literature have also done, we're just going to go to 100 years of data. Why is that data interesting to us? Because nothing really happened to inflation since 1955 until recently. It was a total sleeper, right. So there was one big event. There was the '80s. We had some high inflation, but other than that, it was extremely boring. So the young guys working at the Board now should be grateful that the variable supposedly the Central Bank can control well, maybe not instantaneously but with time, is moving around, and so we can see how much the muscle of the Fed, at least in that regard, controlling inflation, now the Fed has many other regulatory obligations is there. So it was a bit of a sleeper from 1955 to the beginning of COVID, and there was just one big event, which was the big inflation of the 1980s, and that's it. But prior to 1955, between 1900 and 1955, you'll see data in a moment, there were very large swings in the inflation rate. You know, inflation going from +20 to -10, so there was something going on, and the question is how does a purely econometric procedure then interpret the data having seen much more volatile inflation. What we do, we do a Bayesian estimation with, as you will see will, we get a binomial distribution. So we use random work Metropolis Hastings, and we are patient, and we do it for 50 million draws. We

also tried sequential Monte Carlo but we think if you're patient and you do enough draws, you can learn something, and then so we, that's how we estimate the model. So what do we get?

So let me first show a picture that is usually at the heart of many of the R-star papers, namely what is the path of R-star. And this is not the focus of our paper, but we first want to show that we do get similar results, or I'll point out what is similar and what is different than the related literature. So there are three lines here. The solid line is the median, the posterior median of XR, which is our R-star. We only identified up to a constant, so we add a constant, so we get of 115 basis points, so we get the mean real interest rate in our sample to match our natural rate. And we also show 95 percent intervals. So what is -- let's go to say one observation Don Cohen just made. He said, wow, Asta [phonetic] measure dropped enormously during the Great Recession. That was the steepest decline. If you were to cut my picture in 1980, it is true that still what would look as the most obvious and relevant event was the sharp decline in 2008, and much of the work of Thomas and John has actually always stressed that there were very large declines in the natural rate. So that makes us feel okay. We're getting that right. We are also getting right that in general as many other authors have emphasized, there was sort of a long-run decline in R-star since the beginning of the great moderation. So since the 1980s, this measure of the permanent component of the natural rate, which is a random walk, has been going one-way traffic. So, you know, like [inaudible] is always hoping something will come up and this will turn around. This has been one-way traffic, then the narrow interest rate has declined. But the point we want to make with this graph is at the least based on our estimations is that it is not a story that the natural rate was constant from 1900 to the end of the great moderation in the 1980s, but it has fluctuated, and we call this -- there are natural rate super cycles, so we claim that there was one natural super cycle that had started in 1900 or earlier and that bottomed out

some time around the Great Depression and then recovered, and then that the second peak of the super cycle was just at the onset of the great moderation in the 1980s. In particular, a decline in the estimated natural rate of interest of about three percentage points, which is what we have seen in the last 20 years, is also something we have seen before in history. So it's not an unprecedented event. And saying something to -- there's one idea, I mean everybody, I think, is throwing up their hands and doesn't have a good explanation, I don't want to say, but they are -- people are looking for good explanations, and one story is we are -- we have been aging, and once we're aging, the number of people who are creative and productive and come up with ideas relative to the ones who are economically inactive has been going down, and as such there has been maybe a decline in the rate of innovation, I think, yes, we can maybe talk about it in that way if we just have a short run look at what happened since 1980, but if you are convinced by our story that this is a super cycle, so the natural rate going down and up and down again, then you would have to see a similar pattern in the aging in U.S. society. but the aging has been monotonically. I mean there has been constant, continuous aging. So if aging is a driver, there has to be something else. It could be aging. It could be something else. Yeah. Then the second result is we are interested in this delta. So delta tells you how much output, the trend of output changes if you have a natural rate shock. So you might think the natural rate doesn't affect the trend level of output. That would be a delta of zero. If you believe in a neoclassical model, then there's clearly a positive relationship. That is to say if the natural rate of interest goes up or stick with going down, natural interest rate goes down, the level of output will go up, right. So that's a standard thing in any neoclassical Euler equation, if you stay at that. And so, on the other hand, we have theories of secular stagnation in which a decline in the natural rate of interest is also associated with a decline in the trend path of output. So given those two views, we put a prior

that is extremely broad, so it's centered at zero and has a standard deviation of 5 percent. So that's the broken line in this graph, and then we also put our posterior -- so this posterior is bimodal, but the vast majority of the density lies to the right, meaning delta is positive, and at the mode of the posterior distribution, which is around 8, we can make the sentence if the natural rate of interest declines by one percentage point, the trend level of output goes down by eight percentage points. So that's what our estimation shows. But it also says there is a little bit of mass to the left. Can you tell me how many I have?

ANNETTE VISSING-JØRGENSEN. Nine.

STEPHANIE SCHMITT-GROHÉ. Nine. Okay. So, okay. So then we ask -- okay that's sort of in the long run. What about in the short run? How long does it take if there's a one percentage point decline in the natural rate of interest, that the trend level of output falls? Is there something very gradual, very protracted, and so we look at the short run effect. So this is like the estimates of the C matrix I showed you. And let's start in row one, column two, where we look at the response of output to a shock that will eventually depress the natural rate of interest by one percentage point. And we basically see that the entire, I mean, after the units on the x-axis are years, and so you see within two, three years the entire transition has happened. So it's not if you have a decline in the permanent component of the natural rate of interest that it takes a very long time to lead to a decline in the trend component of output. This adjustment is very quick. It's done in one or two years. So an impact is 6 percent and then takes one or two more years until you're down to 8 percent. What about the response of the real rate? The response of the real rate, which is row one, column one, is a little bit lower. So on impact, the real rate falls only by 30 basis points, and by five years it's maybe 75 basis points, and by 10 years it's down. So the adjustment in the real interest rate is much slower, that's at least what we estimate, than in the

trend level of output. And what about inflation? I mean what does it mean for the Central Bank if there are these shocks to the natural rate of interest? By construction those type of shocks can have no effect on inflation in the long run, but they can have effects on inflation in the short run, and we estimate that a shock that lowers the natural rate of interest by one percentage point in the long run brings at least for one or two years inflation down by 25 basis points. So it is -- it smells, let's say it smells more like a demand shock. Okay.

So then let me turn to our second question, since this is a model that like in the DSG literature identifies all of the five innovations, and we do have more shocks than observables, like in a DSG model, we can also ask about Pi-star. That is to say we can look at this item, which we called XM. So here finally comes the data that I had promised. There's two lines. The red line is CPI inflation in the United States since 1900, and if you could cover half of that plot, you would see what I -- and only look at what happened since 1955, was a pretty boring story other than what happens in the '70s. But if you look at what happened between 1900 and 1955, you see that the red line has -- we had to change the scale of the graph. I had to go from minus 10 percent inflation per year to plus 15 percent. And so if you look at the first spike, that spike just coincidentally was during the Spanish Flu of 1918, but that inflation was building up a little bit earlier before the pandemic came, and you see then very shortly thereafter in 1923, the U.S. went into a big deflation, and so we had like minus 10 percent inflation. But the broader point we want to make is that is clearly across that period inflation has a different behavior. It was very volatile in the first half of the century and much less volatile in the second half of the century. So then the blue line is what we estimate as Pi-star or the permanent component of inflation, and you can see that let's go for the Spanish pandemic of 1918, inflation goes to 15 percent, and this econometric tool gives a very small increase in the permanent component, right, and attributes,

of course, everything else to non-permanent, to non-Pi-Star shocks. But if we make it to the 1970s and 1980s, which was a very different type of inflation, the inflation of the 1980s was like slow and steady, you know. It started building in the late '60s' and there was no -- kept going and going and going until the 1980s where Volcker took very -- well I don't know if it's causal, but Volcker is the one who defeated inflation and did something, and then by 1986 inflation was over. So the algorithm interprets that way, that the inflation of the '60s and '70s and then peaking in the '80s was all permanent' the red and the blue line are the same. And so as you can see to a more pressing question, this is data that runs until 2021, if you look at the increase in inflation at the very end of the sample, this algorithm says, okay, sudden spikes I've seen before. I've seen them during, you know, prior to 1955, and this algorithm interprets the increase in inflation of being basically caused by factors other than the permanent component, and inflation says the permanent component went up by 50 basis points. But, I mean, you can intuitively, I think, see if I had set that algorithm with data in the tradition of telling business cycle analysis on post-war data, I would just see one spike, and that spike was permanent. So the algorithm is going to give you that. So if you say I just instead estimate the model on the 1955 to 2021 data, so all of the parameters are re-estimated on that shorter sample, then you see that it continues to be true that the algorithm interprets the inflation of the '70s and '80s as almost entirely of a permanent nature, but for the more narrow question that I'm interested in right now, how does the algorithm interpret the current spike of inflation up until 2021, and it interprets almost the entire increased inflation is coming from the permanent component. So the point we just want to make with that is that sometimes as times change and the world becomes more volatile and we had gotten very used to the period of the great moderation, and we are now living in a more volatile world in many dimensions, that there is maybe some value at looking at longer data.

So let me conclude. We what we did in this paper, we formulated a structural model that we estimated on 122 years of data from the United States, and we concluded that natural rate shocks, a one percentage point declined in the natural rate of interest pushes the trend of output down by about 8 percent, and the puzzling part about that is that the natural rate has super cycles, makes it hard to explain this on demographic rounds, which were monotone over the century. Also, I didn't show you that when we re-estimate the model picking only periods in which we were not near the zero lower bound, we get also that there is a decrease, a one percentage point decline in the natural rate of interest lowers the trend component of output, the estimate is smaller, is more like 5 percent instead of 8 percent, but I think we lack theories for explaining that because the only theory that we, not only but one that is very prominent, is a secular stagnation theory, and that theory has a necessary ingredient that the nominal interest rate is hitting the zero lower bound. So this is opening up a theoretical challenge how to explain that. And the second point we make is with that framework we can, we believe we can also identify the permanent component of inflation, and if we look at long-term data, we think the current spike is transitory, so we go with it. I mean it depends how you call transitory, but I think what the Fed has been saying that is transitory, and they got upset if it was two or three years, but overall we would say it is transitory and not permanent. And if you had estimated that model on only post-war data, you would have arrived at the opposite conclusion. Thank you. Thank you.

[Applause]

CHRISTIAN WOLF. Okay. Okay, perfect. Well great to be here. Thanks again for the invitation, very excited to be part of this, to get to know Stephanie's paper, to get to know Thomas's work a little bit more. I mean, obviously, I also like everyone else here knew about the R-star set of papers, but I kind of did a little bit of a deeper dive now, and it's been very

enjoyable. And I also look forward to the other papers later today that touch on the other parts of his work. So what I'm going to do over the next like 15, 20 minutes is going to share my thoughts on the paper we just saw, the macroeconomic consequences of natural rate shocks, and in particular the version that I've seen was more focused on the first question. So natural rate shocks and how they're going to be moving output and inflation, a little bit less on Pi-star, so there's nothing directly going to be on Pi-star in my slides, but I do have a couple of thoughts, so maybe I'll at some point share those as I kind of go along. Okay.

So again, the background and kind of Stephanie explained that very nicely in her introduction of how this paper relates to what's been done before is that a lot of the existing prior work was focused on measurement of R-star, which as we have seen both in the first presentation and also in Stephanie's plot, is in and of itself a challenging endeavor. You see there's significant uncertainty about what exactly R-star is going to be. And now this paper tries to go another step further, which is kind of going to make the question even harder, it tries to learn about the causal effects of changes in that natural rate on various macroeconomic outcomes of interest and notably aggregate output and inflation and that in both the short run and in the long run. So what I'm going to focus on in my discussion today, I mean essentially going to be my comparative advantage, I think a lot about identification. So I'll here in particular think about what can I structurally think of what is a natural rate shock and what econometric/macroeconomic tools do we have access to, to actually get at these causal innovations to the natural rate, given that the challenge is we don't even observe the natural rate. We in the first step need to be estimating it. And then at the end I'll just briefly share a couple of additional thoughts on what I think this all means for structural modeling, and this is going to relate to also what Stephanie said at the end about do we see those same causal effects inside and outside of ZLB regions.

So let me begin with maybe a bit of a provocative question, but it is a question that I was thinking when I was reading the intro, as I think to myself, okay, what is an actual rate shock, right? What is the actual object of interest in that paper? As a kind of a quote here from the intro of this paper is that is the effect of natural rate shocks on output and inflation. So I know what shocks are. I know output inflation as outcomes of interest, but what really is a shock to the natural rate? Kind of the way I tend to approach things, like everything to me is a vector moving average, so I think of structural shocks, so the epsilon is hitting the macro economy. That maps into the various outcomes of interest, in particular output and inflation. Okay. And we can assign meanings to lots of these structural shocks. We think of innovations to investment and demand. We think of technology shocks. We think of changes in the fiscal stance, changes in the monetary stance. So what then is an innovation to the natural rate? Well one thing in this framework you can always do is you can say the natural rate and in particular these long-run movements of the like low frequency movements of your rate. That is one of the set of outcomes in this vector y. And let's suppose this is driven by a subset of primitive shocks, some of which Stephanie actually alluded to in her talk. So the canonical vanilla one in the small-scale New Keynesian model is just an innovation to the patience of households. I mean I don't know exactly what that's supposed to be capturing, right, but that is the usual shock we're thinking about. A little bit more structure we can think of demographic trends, as Stephanie touched upon, or we could think of changes in idiosyncratic and aggregate earnings risk or the ability of households to insure against that, how sophisticated financial markets are. All of this is going to map into what equilibrium real rates look like. Okay. So these could be more primitive, more causally interpretable structural shocks, what I call here the ISS, that would be moving around the long run real rate R-star. And so by extension and also we'll be moving the various other outcomes

we're interested in, output inflation, and so on. Okay so then kind of with this bit of abstraction in the background, what is it that kind of in my interpretation this paper is doing? So what you kind of saw with, kind of when Stephanie presented the main model setup that they're estimating, what this at the end of the day gets towards is it gives us an estimated series of the real rate. So here R-star in that notation XTR, and then essentially what the estimated model is doing is projecting on innovations in that smooth series of R-star, okay. So in particular what I'm saying here is so you have your outcome of interest y, and these can be impulse responses that Stephanie showed on what inflation, what output are doing are essentially just projections of future outcome variable of interest, say output, on innovations in the estimated real rate, okay. That's sort of indicating that as innovation [inaudible] controlling for everything else in the past, and that's the dynamic causal effect you're trying to tease out. This is the econometric estimate of what this paper is doing or trying to recover. So then if you combine that with what I had before, and it's just a couple of lines of algebra to see what is it that this paper is kind of giving us through the lens of the setup I've written down here, it's in general going to be some linear combination of the various primitive shocks that are actually driving the natural real rate R-star, okay.

Now the question is, okay, so we have this innovation. We know it comes from a bunch of potentially heterogeneous underlying, primitively interpretable structural things. So what is it that we're estimating? Well we're estimating a weighted average of various underlying, potentially heterogeneous treatment effects, okay. And I do think that heterogeneity is potentially important. Like if I think of these structural determinants, I think that changes in household earnings risk versus changes in demographic composition versus say changes in credit spreads and so financial sophistication could and probably in general do have different dynamic causal

effects on household labor supply, on consumer spending over the life cycle, and so in particular on output and inflation. So one part that I kind of would like to see made at the beginning, and it's part of that is actually already in the paper, and I'm going to say soon where it is, it's like what I would like to see is one, like an argument along one of the following two lines. One, and I think this is where the paper is kind of going with the demographic discussion that Stephanie had at the end that maybe R-star is indeed driven by a single and structurally interpretable kind of innovations. These super cycles that don't align with what demographic transitions look like, so maybe the shock even though in principle, demographic transitions are driving some of the variation in R-star, this is not the main shock whose causal effects we're picking up. Okay, so this is one argument. The other argument you could be making is that the various main shocks that drive R-star actually do end up having homogeneous effects in everything you're interested in, then I don't care that I'm computing a weighted average. I don't think that's a particularly promising argument. I think the first one is the more promising one, but I will also get back to one reason why I actually do think there's some heterogeneity that's kind of going on under the hood and actually some of the results Stephanie had showing up pointing me in that direction. Okay. So that's just kind of the bigger conceptual background on how I think, like how I think of what the estimate is, how we should think of natural rate shocks. So let's get a little bit into kind of the nitty-gritty detail of what methods do we have available to us to identify those shocks and the causal effects on the outcomes of interest.

So here I'll just very quickly repeat what Stephanie already wrote down. She explained it very well, so I don't have to say too much, and she's going to be splitting things into here evolution equations for the various cyclical components of three variables of interest, output, inflation, and interest rates, and then those are going to be appended by a set of measurement

equations. So here, first of all, the evolution of these cyclical components, simple AR1 there, and then you're going to have your five innovations. The blue ones are the one you really care about. They are going to be the ones with long-run permanent effects. So the technology shock XT, the real rate shock XTR, and the monetary shock XTM, okay. Those are going to be in particular the one we're going to be, we're going to be disentangling. Key to that disentangling are identifying assumptions that are implicitly embedded in the observation equations. So what we have here is the three observables on the left-hand side, so not the cyclical components but actual realizations of output inflation interest rates. Then you have your cyclical components. You have a measurement error at the end, and then you have all of these permanent shocks in there, and what's key to identification in this paper is which ones of those permanent shocks show up in which equation and with what coefficients, okay. Those are at the heart the exclusion restrictions that make everything in the identification tick. So in particular, the fact that the real shock, the TFP shock, only shows up in output. That's the only thing that's moving in the long run. The monetary shock, of course, moving along on inflation but one to one also moving nominal rates, but that is an assumption that's later on in the paper being relaxed, as Stephanie mentioned. And then we have our real rate shock XTR. It's moving nominal rates not inflation, so it can move around the long-run real rate. And then it also shows up in the output equation with an additional coefficient delta, and that delta is the main object of interest plus then filtering it through with the matrices B and C to get your overall impulse responses. So let me try to dig a little bit deeper into where exactly kind of those joint restrictions and the rest of the structure of the model are giving us identification. So, the first part, as I've already mentioned now is we have these three non-stationary shocks. They are disentangled via co-integrating relationships on which variable appears in which of the measurement equations. That part I've already mentioned. So from this

we can kind of arrive at a pretty clean view of where exactly identification is coming from, and in the hope that this will help some people maybe in the room, I kind of think of everything in terms of VAR identification, so I'm going to relate this to VAR identification. Everything to me at the end of the day kind of looks like a VAR, so I'll try to do this. Kind of what's going on here are kind of two things at the end of the day. Number one is you're saying at least long run movements of these three observables, they are only being driven by three shocks. So even though you at the end of the day have five shocks with three variables, the long-run movements, those you can collapse to as many shock as observables, this is going to be reminiscent of but not identical to the invertibility assumption that's usually being imposed in conventional VAR-type analyses. And then on top of that you need exclusion restrictions and those are the ones we've been talking about, which ones of the three shocks is entering where. Okay.

So if you take that perspective and those two keys to identification, there are a couple of concerns that will kind of jump out at you for, yeah, kind of potential threats to the econometric design. The first one is let's for a second fully buy the exclusion restrictions. That's where we, like we think these shocks have those properties, where they're showing up. You are still relying reasonably heavily on the first part. So what I call here is a generalization of invertibility, I'm going to refer to as recoverability as kind of a term in some of recent econometric work people have used to describe this, recoverable because you look into the past and future, but anyway, we don't need to get into this. But what this is imposing here is there are exactly three shocks, and you know it's exactly three shocks that are driving these long-run fluctuations in these three variables. If your exclusion restrictions are right, then what you in general are going to be estimating as your shock of interest, the XTR is going to be a weighted average of those various primitive underlying real rate shocks, okay. Note here you -- what

matters for this is actually not that there's exactly one technology shock and exactly one monetary shock, you could have more of those, then you couldn't identify them but you don't care as long as you only care about the XTR shock, okay. This is what's going to be my kind of offhand comment on the side now on the Pi-star part of the extended version of the paper. There it does matter that you're committing to saying there's only one monetary shock that drives long run Pi-star. It's kind of putting a lot of time series structure by knowing that it's one shock rather than multiple ones. So we can talk about this more later if you're interested. Anyway, so this is the first part. If the exclusion restrictions are correct, you're going to be recovering that weighted average. If the exclusion restrictions are violated, of course, then all bets are off. You could be getting a weighted average of all kinds of shocks hitting the economy, and actually like thinking through what drives to me R-star and when I kind of then thought a little bit more structuring kind of the hanger OLG-type models that I tend to think of more in my own work, then I would think, okay, even something like the technology or real shock, if you think that, and if you have persistently worse growth prospects, but say the ability of households to self-insure hasn't changed, because the overall amount of liquidity hasn't changed, the severity of financial frictions hasn't changed, then still you would think that kind of shock would directly map into movements in the real rate on top of, by construction, directly moving output. That is just one kind of additional point on interpretation and what the exclusion restrictions are buying you but also how you need to think this through a little bit more. Then I have just a couple of additional quick points. How am I doing on time? Five, oh terrific. That's very good. Okay.

So just a couple of like little points on -- just details on the implementation. One is just, and everyone who's done a kind of long run identification where you're really relying a lot on movements at low frequencies, they have much longer time series than what the usual long-run

identification VAR paper has, but still, it's hard to estimate covariances, auto covariances at very long runs. That's why some econometric workers kind of try to go in the direction of let's try to mimic those longer identifying restrictions by just kind of imposing quantity balance at medium frequencies. That's just a little bit easier and more robust to estimate. So that was just one thought I had. The second thing this is something that Stephanie mentioned in her talk is there is a second mode for delta. Most mass is on the -- getting the sign right -- the positive coefficient for delta, there was a non-trivial mass on the negative coefficient. There's actually the impulse responses that you saw would have been truncated conditional on positive delta. So there is a lot of mass, so it is informative to report. But the negative mode, maybe this is just reflecting the fact that there are different types of natural rate shocks that you're picking up. One of them is the dominant one. That's why most of the posterior mass is there, and that's why that's an interesting thing to report, but there is another -- there's something else going on, okay. This is like one thing that I thought was maybe consistent with this multi-shock interpretation that I was having. And then final point just what I kind of initially thought when I just read like the very beginning of the paper, one thing you could be doing, of course, they are going much beyond that, but what you could be doing is let me take R-star from other work that has estimated that, let me make my own focus to just be this dynamic causal effect propagation. So you could take those series and just project on innovations of this simple LPIV internal IVVAR. Not saying this is in any way more credible what they're doing. I would just -- I would find the comparison interesting. Is it materially different, the estimation part of the R-star series versus the shock propagation, where are the differences coming from, really. So I would find that useful to see. And then just a couple of idle thoughts at the end on kind of implications for structural modeling.

So what Stephanie mentioned is, well, she mentioned, she didn't show it, but it's very prominent and well explained, well derived in the paper, that these natural rate shocks don't just cause a recession and shift trend output over the entire sample but even on non-ZLB subsamples, okay. While in theory what we would expect is if monetary policy is not constrained, at least our vanilla notion of a natural rate shock, like it's in patients one monetary policy, it's just one-to-one offset. Nothing happens to output. Nothing happens to inflation. So then that's kind of a takeaway and inconsistent with standard models of secular stagnation. Just two quick things I wanted to mention on this. Number one is this does require on the monetary reaction being the correct one, right. It's given how difficult -- three minutes, very good -- given how difficult R-star is to estimate in real time, maybe it's difficult for the Fed, even though secular stagnation theory is completely correct at least in the short run to be offsetting things very well. This does not address our long-run findings, but at least on the short one that you saw an immediate bust, that's in principle something that where you say maybe the Fed just didn't lower interest rates enough. And one thing that I was kind of thinking what you could be doing there is ask given that they also have an identified monetary shock in the exact same system, how much of an additional innovation to nominal rates would you need to stabilize output or stabilize inflation? This could also tell you does this natural rate shock move us along an NKPC that you would have -- you could estimate an NKPC slope from the monetary shock, and you could see does the natural rate shock give us that same slope, that would also maybe buttress Stephanie's demand shock interpretation of the natural rate shock. And then the other point is, again, just the misidentification point from before, if what you're capturing is partially a long run, a real rate shock with long-run effects on output that just -- sorry not a real rate shock, another shock that moves output and also then tends to move R-star, then maybe the secular stagnation stuff is all

going on perfectly for the actual real rate shocks that are part of your shock, those are being offset, but the residual part is not being offset, and that's what we're seeing in the figure. That is just one potential interpretation. Okay.

So let me conclude, and I think this is a substantial move forward of the empirical natural rate literature, moving the goal from pure measurement to learning about dynamic causal effects. This is a very ambitious step, and so it's always good for discussion if a paper does something ambitious, because then you certainly have something to talk about. And then kind of at these two main points of just conceptually what do I think of is the object of interest, what is a real rate shock and then what's kind of precisely the set of identifying assumptions that allow us to talk about this. And kind of what I would find useful in the paper is a little bit more of a discussion of here's a class of structural data generating processes such that the estimate of the specific econometric method we're proposing is actually correctly that interpretable dynamic causal effect of a particular rear rate innovation that again I can make sense of through a class of models. And, yeah, that's all I have, so I'm really excited to also see future versions, see where this paper goes, and it was a pleasure discussing it. Thank you.

[Applause]

ANNETTE VISSING-JØRGENSEN. Stephanie will give a few replies, then we'll open it up for questions.

STEPHANIE SCHMITT-GROHÉ. So, first of all, thank you very much, Christian, for a very generous discussion. All these points are very well taken. I wish I could say we have a clear answer to this. First, I might give a little bit of a lame of an answer. Many people in this room have played with the NK model and in the New Keynesian model, the standard shock is not a permanent -- it's a natural rate shock. It's usually treated as a transitory shock, so there are

precedents to doing this, what you comment, this audacious step of having just a reduced form natural rate shock. They are typically not of the permanent type, but they are of the transitory type, but I totally agree with you. Why we came short on that dimension is that, as I said, we are looking for structural models that we can compare to our data, and as you said, we are struggling a little bit with how to interpret the bimodality. So is the bimodality telling you that there are actually two shocks that have this feature that we're imposing that they have no effect on inflation in the long run. They do have an effect on the real interest rate in the long run, and they do have an effect on output in the long run, and you know, how many shocks of that type are there, and at least from theory, we know that there could be various. And so maybe the way to, we thought identifying five shocks was already very careful, but I agree with you, what is a natural rate shock. Yeah, it's again, a lame answer, but as you mentioned, the preference shock is hard to say. So I think if we could come up with one, and I wanted to mention that makes sense to me, and I can understand as simple as the Ben Bernanke global saving glut shock. So there was a change in global financial markets, and there was suddenly a big supply, much more supply of assets and international financial market, and so market clearing in a free capital mobility world reduces that shock. So that's one way to think about it, and that makes sense. Yeah. So we, you know, I'm going to ask you for your slides, and I'm going to, you know, try to learn from that and see if we can improve the paper in that dimension. But thank you very much for your discussion.

ANNETTE VISSING-JØRGENSEN. All right, let's open it up for questions. John you get [inaudible]. Can we get a microphone up to John here in the front row?

JOHN WILLIAMS. That's a, you know, very interesting paper, great discussion, great presentation, and discussion. I guess I have two comments, maybe a little bit reaction. Stephanie,

your point about the Euler equation. Yeah, so if you look at our original paper with my original paper, Thomas, we do use the Euler equation to kind of motivate the notion that there's a relationship between the equilibrium real interest rate or R-star and the growth rate of consumption, but we also have a footnote in here that, you know, says a lot of models with a lot of kind of different frameworks would have this feature. We used a Solow growth model, has a similar relationship, and I think that, you know, Christian's comment that if you go to an OLG framework or HANK kind of framework, you're going to find different ways to find this, and maybe my reaction to your comment is that probably we shouldn't be so focused on the word permanent and maybe I too. It's merely about medium over, you know, maybe over a decade or something if the, you know, I liked your example of, Christian's, you know, if people are pessimistic about income growth for many years, that's going to affect the equilibrium in the economy. So that's how I've thought about it, not so much of, you know, is it literally permanent changes in growth rate? Is it literally permanent changes of R-star, but maybe a more medium-term relationship, which I do think comes out a lot of models. The second question comes out of the discussion, which I thought was very, very helpful to understand this, and I think this identification issue is fundamentally hard. And I'll give you the example of some, you know, work that Thomas Mertens and I have done over recent years, just a very theoretical model, that, you know, yes, I'm going to bring up the lower bound, but in a world where there's a lower bound that binds, you know, a part of the time, there is a correlation between the level of the natural rate of interest and the inflation rate. Basically, it's the effect of the ELB binding policy part of the time that affects the mean inflation rate for a given inflation target. But more generally, I get to this issue of, you know, if you're thinking about the fundamental shocks or the kind of the deeper shocks, it might be affecting, you know, trend growth rates and inflation rates

and interest rates, to me, it's hard to really feel confident that we're able to say, well, there's a monetary shock, there's a real shock, there's a this shock. So I think more kind of robustness or work on make, you know, understanding that would be helpful, and I think that's really the discussion kind of solidified that view. But again, this is a great research program and paper. Thanks. Thank you.

ANNETTE VISSING-JØRGENSEN. Benoît.

BENOÎT MOJON. Great paper. Benoît Mojon from the BIS. So on the interpretation of the shock to R-star, you just mentioned that you were thinking of it as possibly the savings glut. Can you elaborate on how if it's a savings glut shock you would have an effect on GDP growth. So I could see it like in a closed economy that if I have an increased preference for savings and for savings in safe assets rather than stocks that then you're going to increase the risk premium on stocks, but if it's coming from abroad, then it is less clear to me how it would impact your growth, your domestic growth in this case.

[Inaudible]

STEPHANIE SCHMITT-GROHÉ. So I'm going to be shameless and answer to Benoît because I know what to answer, but to you it's more difficult so let me first answer to Benoît. So okay, so let's be clear what I said. He says what is an R-star shock, and I totally agree what is an R-star shock. In [inaudible] it's a shock to the discount rate of private agents, and that is like answering a question with another thing that is un-understandable. So I said one of the things I could interpret as a true change in the real interest rate and in the permanent component of the real interest rate if the world changed and many more people come participate in in the global financial market, in particular, they choose to supply tons of assets to the global financial market since they joined the WTO. That is what I said. The second thing is do I have a theory in which

if there's an exogenous decline, a permanent decline in the real interest rate, that that will permanently lower the trend path of output, my answer is I come up empty. And so what can I go to? We're actually trying to work on the theory of that, and maybe here in this room some people know some papers we are not aware of, but we don't think they exist on a coherent theory about that. Maybe they are, you know maybe you think all the people are riding on previous episodes of low real interest rates can get, I can't get a recession, but I need to get the permanent component to decline, right. So and that's a little bit to John's argument, this is not like transitory or persistent effects, but it's like what we saw in the recovery from 2008. It was basically the path of GDP was pushed down, and then we kept going in a parallel fashion. To come up with a model like that is what we are working on and trying, but I do not have a good answer.

[Inaudible]

AUDIENCE MEMBER. So thanks very much. This is very fascinating and important, of course. So think about it from a policy perspective. I mean the conclusion you had what the inflation today is more permanently higher. I mean what are the intuitions behind why that's the case, and is that driven by the identifying restrictions? Could that be also because of technology being weak, meaning the technology shock driving inflation up? Could it also be because of monetary policy? So and maybe there's also a question for Christian. I mean how do both of you think about the permanence of inflation in this current environment?

STEPHANIE SCHMITT-GROHÉ. Okay. So I'm an academic. I have not followed the milieu things, but here the issue is if you feed that empirical model with data where there were five or six bouts of very rapid and very large changes in inflation, and I mean it's mechanical, so it's not maybe a very smart answer, but if you had in '18, in '23, and there were another bunch of bouts of very rapid either spikes up or spikes down of inflation, right, and you have, in that

model, you have five shocks competing for the attention. And the model says I saw a lot of sudden spikes, then it's, I think, intuitive that if you now see the sudden spike of inflation of going up to, I don't know, 8 percent, I forgot what the peak was here, that then this algorithm says, yeah, no, there are a lot of shocks. They could be persistent, some of them, but it is not necessarily to put a policy word on it and on anchoring, right. That is what the algorithm tells you, saying if you look at history, there were -- even though monetary policy might have been very different, we try some ways to control that that's the interpretation, I mean.

CHRISTIAN WOLF. Yeah, and I guess the only other thing I would add to elaborate on this is just kind of the mechanical part of the answer Stephanie gave, I think, very nicely illustrates just how much you're getting from the restriction of saying there is a certain number of shocks. I know there is like usually as many observables as shocks here, as many shocks, long-run shocks as observables, and because then you have that, you're forcing the model to say everything is equally persistent, that the underlying drivers are going to be heterogeneous across different cyclical episodes. So that's always kind of -- I mean this is just kind of my personal pet peeve with the literature, but I always felt like it was more credible to identify individual shock episodes, trace out the dynamic effects without imposing the restriction of everything needs to be following the same AR1 with the same coefficient or the same ARMA1-1, which is kind of what's implicit in a lot of say the business cycle DSG estimation and work.

[Inaudible]

AUDIENCE MEMBER. I wanted to come back to this question of what are these natural rate shocks? It looked to me just from the time series that the natural -- the big natural rate shocks were the Great Depression and the financial crisis, and that suggests to me that what these shocks really are in some underlying sense is a financial crisis. And so big financial crises,

we think we know cause big declines in output, persistent effects on output, persistent effects on inflation from the work of Reinhart and Rogoff, for example. And so I wonder if the way to think about these natural rate shocks is really as financial crisis shocks, which would be hard to capture in a model, but still for a policy maker, that may feel like that's what's really driving things.

STEPHANIE SCHMITT-GROHÉ. I mean we also think that they are financial crisis shocks, right. And so, all I then -- but I would have to come up with a financial crisis shock that is not tied to the zero lower bound, right. That's one challenge. And the other one, so we do not know a model in which maybe you know one where I can get that, that there is, what is the financial crisis -- what is the nature of the financial crisis and why does it give me the two things at the same time, that the real interest rate is depressed for a very long period of time and that the trend level of output is sort of pushed down, so but that's just a challenge, but we, I think you are naturally led to look at that. Maybe you are led to look at zombie banking, you know. So I think it -- [inaudible] mediation or something. Yeah. So I mean you -- you are pushed in that direction, in the direction of misallocation of credit, but we don't have a good answer at this point.

ANNETTE VISSING-JØRGENSEN. Anyone else? All right. Right on time. Let's wrap it up. Thank you very much, guys, for doing a great job and thank you [inaudible].

[Applause]