

Christopher Carroll

Editor's Note: *This is an abbreviated version of EF's conversation with Christopher Carroll. For the full interview, go to our website: www.richmondfed.org/publications*

When the housing market took its precipitous negative turn in 2006, policymakers were plagued by a single nagging question: How much would a collapse in housing wealth drag consumer spending down with it?

There were two schools of thought. One was based on the notion of wealth effects, that wealth makes people feel richer, such that a dollar change in wealth pushes spending in the same direction by a few cents. The more ominous school of thought said that the unprecedented growth in consumption during the housing boom years was not due to wealth alone, but also a relaxation of credit constraints that gave people an increased ability to use their housing wealth for consumption. Take that cash cow away, research suggested, and consumption was likely to fall by two or three times as much as suggested by the wealth effect alone. A threat, indeed, for an economy comprised two-thirds by consumer spending.

Christopher Carroll, professor of economics at Johns Hopkins University, was one voice behind the more pessimistic estimates, and he says the evidence from the Great Recession has proved that view correct. Carroll is a long-time scholar of saving and consumption dynamics at the individual and aggregate level, studying questions that range from housing wealth effects to the consumption response of households to uncertainty, and from national saving patterns to the surprisingly modest spending of the wealthy. Much of Carroll's work came to the forefront of current events nearly simultaneously, leading to a second stint at the President's Council of Economic Advisers (CEA) that spanned the implementation of the historic 2009 American Reinvestment and Recovery Act — also known as the fiscal stimulus — an experience that Carroll describes as changing how he views public policy.

Carroll joined the faculty of Johns Hopkins University in 1995. In addition to serving twice on the CEA, he began his economics career as a staff economist at the Federal Reserve Board of Governors. Renee Haltom interviewed Carroll at his home in Columbia, Md., in February 2013.



EF: How well did existing theories of the wealth effect hold up during the housing boom and crash? Did economists learn anything new?

Carroll: The theory was never particularly clear about how large wealth effects should be and what would be the channels. There was empirical evidence that when the value of some set of assets goes up, whether it's house values or stocks or total wealth, then there's subsequently a growth in consumption spending. You could interpret the change in spending as a consequence of wealth changing. An alternative interpretation is that everybody got more optimistic: They saw that the economy was improving, and that's why the stock market boomed — it was anticipating the movement in consumption. That's not really a causal story. So that was always an issue.

There was a substantial literature showing that subsequent movements in consumption after house price changes were bigger than those associated with the stock market. But it was never clear from that literature whether people spend more when their house value goes up because they feel richer, or whether a collateral constraint has been reduced. That is, when your house is worth more, you have a greater ability to get a home equity loan or a second mortgage or

refinance, and use your house as an ATM. Or you buy a new house and sell the old one, which bids up house prices further. There's a lot of research using local geographical data, especially by Atif Mian at Princeton and Amir Sufi at the University of Chicago, finding that in places where house prices went up a lot, debt and consumption spending went up.

So one thing that has become much clearer in the last couple of years is that a lot of the relationship between house price changes and subsequent spending really was the result of collateral constraints being loosened. That now seems to be a consensus while it was a speculative idea five years ago. And, of course, that's not sustainable, especially if house prices peak and then start going back down. Then you've got a real hangover afterward.

EF: Before the crisis, many economists were lamenting the long-term decline of the household saving rate, which went from about 10 percent of disposable income in the early 1980s to a low of 1.3 percent in 2005. You've noted that economists are hard to please: They worry about the current economy when consumers spend too little, and they worry about our long-term welfare when consumers spend too much. Is there a middle ground that would keep economists happy?

Carroll: What the saving rate is ultimately about is the aggregate capital stock and aggregate national wealth. You're not going to put much of a dent in that with two or three years of a low saving rate. But if a country's saving rate is low for 20 or 30 years, then you end up a lot poorer.

I do think that before the crisis our saving rate was lower than is wise or sustainable. There's an emerging consensus that the decline mostly reflected the fact that it was getting easier to borrow during that time. So I think that something did have to ultimately put an end to it, but exactly when was unclear. The problem from a macroeconomic point of view is when you try to reverse that all of a sudden. If we could have gradually inched the saving rate up 1 percentage point a year for 10 years, that would have been a healthy way to deal with the problem. But having it go up by 5 percentage points in the course of a year is a huge economic shock.

It would be very hard to come to an agreement about what an "equilibrium" saving rate should be. What's clearer is that when there is a really dramatic change in the saving rate, either an increase that we saw in the Great Recession or the drop that we saw in the mid-2000s, that ought to be a danger signal for policymakers. The economy really can't efficiently handle rapid changes in aggregate demand.

One way of saying a little bit more about that is to look at a longer history for countries that have been in a reasonably

When there is a really dramatic change in the saving rate, either an increase that we saw in the Great Recession or the drop that we saw in the mid-2000s, that ought to be a danger signal for policymakers.

stable developed equilibrium for a long time. Most such countries tend to have personal saving rates somewhere in the 5 percent to 8 percent range. I think when our saving rate gets below that range for a sustained period of time, that's something that one ought to worry about. And it had for quite a while been below that range.

EF: What are the major unresolved puzzles in consumption theory? Are there areas where theory doesn't quite match up with reality?

Carroll: One is the research on default retirement contribution rates. There's an impressive body of new research that finds that people's retirement saving decisions are very much influenced by the default choices in their retirement saving plan. I recently discussed the latest paper in this literature at the National Bureau of Economic Research's Economic Fluctuations and Growth meeting in San Francisco. The authors had data that basically covered the entire population of Denmark; 45 million data points, and they could see people for 15 years. They found that if an employer has a default 401(k) contribution rate of 6 percent, 85 percent of people will just go with 6 percent, rather than changing the contribution rate or opting out. If the default is 10 percent, then 85 percent of people will go with 10 percent. I think the evidence for default contributions is just overwhelmingly persuasive.

That is a really big challenge to the economists' standard modeling approach, which is to say that people rationally figure out how much they need to have when they retire and they figure out a rational plan to get there. The problem is, now that we have discovered serious flaws in the rational optimizing model for how people make those decisions, we're kind of a bit at sea at being able to say, "Suppose we changed the tax rates on 401(k)s, or suppose we do this policy or that policy. What consequence would it have?" given that we don't know why people are making those decisions in the first place.

The explanation I proposed at the conference was to say that, within some range, people trust that their employer has figured this out for them. The job of the human resources department is to figure out what my default contribution ought to be, and it would be too hard to solve this problem myself, so I'm just going to trust that somebody else has done it. It's not different from when you take an airplane and you trust that the FAA has made sure that it's safe, or when you go to the doctor and you trust that the advice makes sense and is not going to poison you. Maybe people trust that the default option is going to be a reasonable choice for them.

That makes a little bit of progress in the sense that you

could think though under what circumstances one would expect people to trust that decision has been made well. Are people who are not very trusting less likely to go with the default decision? What are the forces that reinforce people's trust in the employers to make a good decision? What are the circumstances that encourage employers to make a decision that deserves to be trusted? Maybe the employer needs to have some fiduciary responsibility to have made a good decision. If people are going to trust the employer to make a good decision, we ought to make some effort to give the employer the incentives to actually make that good decision.

EF: What about puzzles at the macro level?

Carroll: I think there's a really big one for which the profession has not reached a consensus or even come close. That is the very strong relationship across countries going from high growth to high saving. The theory in every textbook says that if you know you're going to be richer in the future because you're a fast-growing country, why in the world would you save now, when you're poor, making your future rich self better off? It makes much more sense to borrow now since it'll be easy for you to pay off that debt in the future when you're richer.

The latest example that's on everybody's minds is, of course, China, a country that has grown very fast for the last 20 years and has had a saving rate that just seems to get higher every year. If China were the only example, then it might be plausible to say that the phenomenon reflects some unique aspect of China's history or culture. There are some papers that argue the one child policy has something to do with it, or it's the end of communism and the transition to capitalism, or that it's Confucian values. But what China is doing right now actually looks virtually identical to Japan 30 years ago. Japan didn't have a particularly high saving rate in the 1950s, and by the 1970s it had the highest saving rate in the world, and that was a period of high growth in Japan. It's also true in South Korea. It grew at a very rapid rate starting from the early 1960s, and its saving rate went up and up. We also see this in Taiwan, Singapore, and Hong Kong. And it's not just East Asian countries; the same is true of Botswana and Mauritius. It's also true in the opposite direction for European countries, which were growing pretty fast after World War II. That fast growth came to an end in the early 1970s, and afterward the saving rate declined, just as it declined in Japan after Japan slowed down starting about 1990. So it seems to be a pretty pervasive, large effect that is really very much the opposite of what you'd expect from the standard off-the-shelf models.

I have a couple of papers proposing that habit formation has something to do with it. There are a lot of Chinese people whose idea of a good standard of living was formed back in the 1960s and 1970s, when China was much poorer. If you have this reference standard in your mind, you might

respond to rapid income growth by saving more because it's easier to save if you feel rich.

I have another paper that asks whether it's really about a precautionary motive. In that paper, a country makes a deal: In order to get the rapid growth, everybody is going to have to live with an economy that is constantly transforming itself, experiencing churn and creative destruction. All of the old ways of doing things have to be abandoned and everyone has to live through lots of disruptions. Then maybe the increases in saving reflect a precautionary motive.

In fact, what I really think is the right story is one that combines habit formation and a precautionary motive, such that they intensify each other. If I have these habits, then a good reason to resist spending when my income goes up is uncertainty over whether the factory that I'm working for will close down and I'll have to go back to my rural peasant roots. But in the academic publishing context, it's hard enough to introduce one novel thing in a paper.

EF: Milton Friedman's work in the 1950s on the "permanent income hypothesis," the idea that people smooth consumption over their lifetimes, was initially seen as a very important contribution. Yet many economists spent a lot of time in the 1970s and 1980s seemingly disproving his main predictions. What does that debate reveal about how economics is done?

Carroll: When Friedman wrote his famous book, the available mathematical tools were very primitive compared to what we know how to do today. So he used his gifts as a writer to lay out in good solid prose, of course supported by data and charts, his vision of how he thought things worked.

The book was very famous, so everybody wanted the prestige of being the one to formalize the model's main predictions. When you have a rigorous mathematical model, everyone can agree on what that model means. They might not agree on whether it's right as a description of how the world works, but they can all agree on what it says. So a big priority in the economics profession in the 25 years after Friedman wrote was coming up with the mathematical tools to analyze the optimal consumption choice problem that Friedman described informally. Friedman himself wrote a couple of papers trying to clarify his own views.

The first generation of those models had to make the radical simplifying assumption of perfect foresight: no uncertainty in the world, everyone knows what's going to happen for all of future history. There was a lot that those models said which was directly contradictory to things that Friedman said. For one thing, Friedman emphasizes the role of uncertainty and precautionary buffers, and he presents some data showing that people who face greater uncertainty tend to hold larger buffers. That, of course, is completely outside the cognizance of a perfect foresight model. Perfect foresight models also predicted that your spending out of a windfall shock to income — a 100 dollar bill on the

sidewalk — would be about one-tenth of the size that Friedman predicted. One reason is that Friedman defined “permanent income” to mean roughly what you would expect your income to be on average over a three-year period, whereas the perfect foresight model’s definition was your entire income stream from now to infinity. The marginal propensity to consume is so low in perfect foresight models because you’re spreading your windfall over all of history.

In the subsequent 25 years, we learned how to incorporate uncertainty seriously into the models, so we don’t have to have this silly perfect foresight assumption anymore. And we have learned how to incorporate financial constraints. In the perfect foresight models, if you know your income is going to be high in the future, you can borrow 100 percent of that future income to finance your spending today. The moment that you get admitted to medical school, your spending should triple because you’re going to have a high doctor’s salary. In the real world, maybe the bank is not willing to believe that you’re going to repay them if you go on a big spending spree right now. We now have the mathematical tools and technology to build in these kinds of constraints on people’s access to their future income.

The combination of uncertainty and borrowing constraints pretty radically changes the implications of the mathematical models. And the thing that’s really striking is that what you get is something that corresponds remarkably well to the words that Friedman wrote in 1957. Arguably, he had a very good mathematical intuition. He didn’t know how to formalize that math, but he could see the contours of what optimal behavior looked like.

It’s an interesting story, not only because it makes you think, “Boy, that Friedman guy was pretty smart,” but also because now it’s very hard to get anything published until you have already worked out the fully specified rigorous mathematical formulation. You can’t just say, “Well, my intuition tells me something works like such and such, and it would be nice if somebody could work out the math for that in the future.” Friedman was able to get away with that before the profession got so hung up on rigorous mathematical proofs. Today, for example, we discussed that maybe the reason people go with their employer’s default retirement contribution is that they’re trusting the employer to have worked out the problem. I could never publish a paper making that claim. I would need to have the formal dynamic optimizing model of trust, and the formal set of beliefs that people have about the trustworthiness of their

Christopher Carroll

► Present Position

Professor of Economics,
Johns Hopkins University

► Other Positions

Senior Economist, President’s Council
of Economic Advisers (1997-1998 and
2009-2010); Staff Economist, Board of
Governors of the Federal Reserve
System (1990-1995)

► Education

Ph.D., Massachusetts Institute of
Technology (1990)

► Selected Publications

Author of numerous articles in such
journals as the *American Economic
Review*, *Quarterly Journal of Economics*,
Econometrica, *Journal of Monetary
Economics*, and *Journal of Economic
Perspectives*

employer, and the equilibrium determination of trustworthiness. What you can do is publish empirical papers that reject a rigid mathematical model as a test of that model, but then we’re left in the nihilistic position of saying, “We know that this benchmark model that everyone understands is wrong, but until the complete fully specified alternative is generated in someone’s brain, we can’t propose half-baked theories that may have a lot of truth to them like Friedman did in 1957.” I wish the profession would back off on that degree of rigidity. And maybe we have backed off a little bit.

That’s one of the reasons blogs are where some of the most interesting economics is being done these days. That is an outlet where you can say, “Here’s how I think this is working,” and people can criticize you and point out places where you’ve made

factual errors, but there’s not the counterproductively high barrier to having something to say that we have in formal academic publishing.

EF: So, would you say the permanent income hypothesis is back in favor (if it was ever really out)?

Carroll: There’s been a lot of evidence in the last 10 or 15 years confirming the basic dynamics that Friedman was talking about for how households make their year-to-year consumption saving choices. The term that is often used now for such models is “buffer stock saving” models, and I’ve written a number of papers on that topic. There are a lot of ways in which those models match our data reasonably well. So I suspect that a good description of the typical household’s behavior is that they figure that their employer has got the retirement saving thing figured out, and they just go with whatever the default is, and then they do this buffer stock saving thing with respect to whatever money is left over. A lot of the data that we use to test these models have been really focused on the buffer stock aspect of things and has ignored the retirement saving part of things.

I think people who work in this area would say that the buffer stock model is a pretty good description of everything except for the retirement saving part of people’s behavior. And the buffer stock saving model is essentially just providing the mathematical formalization of what Friedman was trying to say in 1957. So in that sense, I think the permanent income hypothesis has come into its own: We have a rigorous mathematical formulation of what Friedman was trying to say.

The terminology has changed somewhat. For a while, the

profession referred to the “permanent income hypothesis” as being the perfect foresight formulation that was developed after Friedman, but that I think is really inconsistent with what Friedman himself said. That’s why what I’m speaking of tends to be called the buffer stock model today. Although my name is associated with the buffer stock terminology because I wrote some of the early papers on it, my own interpretation of it is that Friedman got it right and we’ve finally just figured out the math.

EF: You were a senior economist at the Council of Economic Advisers in 2009 and 2010. Was there a stark juxtaposition of views about the 2009 fiscal stimulus inside the CEA versus outside of it?

Carroll: I came on Aug. 1, 2009, so the stimulus had already been passed by the time I got there. A lot of what we were trying to do was monitor it, and figure out what effects it was having and how to explain those effects to the public. That was a difficult task. The public was not necessarily going to be persuaded by regression equations and statistical evidence. But it was a fascinating experience. When you’re working at a job like that, of course, you read everything that’s in the popular press and you see what’s on TV. Seeing things from the two perspectives of being inside and the outside was interesting.

There’s one particular point that I was struck by several times. The CEA tends to vet speeches that the president and sometimes other officials are going to make, and to help set the priorities for what’s going to be in the speeches. A number of times we would help to reshape the speech to make sure that key points were highlighted, and the arguments that we thought were the soundest economic arguments were made. And then the president would go out and give the speech, and I would later hear from economist friends, who would write to me complaining, “Why didn’t the president say this obvious point in the speech that he just made?” And that obvious point was the thing that the CEA had deliberately made sure was actually a highlight of the speech! But, of course, what your friend actually sees is the 15 seconds that gets excerpted on the news or some blogger’s two-paragraph reaction to the president’s speech.

So the narrowness of the communications channel is something that you get a very different perspective on from the inside. It has made me more circumspect in my own criticisms of the White House and the communications strategies they’ve pursued after I have come back to Johns Hopkins, because now I understand they might well agree with everything that I have to say on the subject and just not be able to get the message through. The president has a greater ability to express his point of view and get it heard than any other single person. But I think the extent to which even the president can’t penetrate through the fog of information and the vast number of sources of data that people pay attention to is underappreciated.

EF: You’ve been the Placement Director for new economics Ph.D.s at Johns Hopkins since 2002. Given what you’ve described as an overemphasis on math relative to concepts in the economics profession, what can Ph.D. programs do to better prepare students to become effective professional economists?

Carroll: It’s sort of an equilibrium problem. The profession demands a high level of mathematical expertise, and so nobody can responsibly back off of making sure that their students have that training. To do so would endanger their ability to get jobs.

I do think that the profession is much too insistent on the proposition that the only good economics is highly mathematical economics. For example, one of the most insightful things that I have read about the current crisis in Europe is not about the current crisis at all. It’s a book called *Lords of Finance* by Liaquat Ahamed, about Europe in the interwar period and the collapse of the gold standard. It’s a brilliant book. It includes all sorts of fascinating and compelling economics that I think really sheds light on the problems of the eurozone today, and there’s not a single equation in it.

The profession ought to be more eclectic, I think. We ought to recognize that a much better knowledge of history, the history of economic thought, and insights from evolutionary psychology and all sorts of other fields have a lot to contribute. At present we, as a profession, are not willing to tolerate that. Partly it’s an arms race problem in the sense that mathematical tools are easy to judge and rank people on. So we tend to focus on that.

I think most of my colleagues in the macro group at Hopkins would agree with most of what I have just said. What is a feasible choice for us in the current environment is to focus preferentially on real world policy questions. Of course, students need to have the ability to use the latest statistical techniques and to understand and to manipulate state-of-the-art models, but it’s a real talent to be able to take those mathematical tools and use them to illuminate practical policy questions that the International Monetary Fund or the central bank or a fiscal policymaker might face. A lot of macroeconomics doesn’t even try to address serious real world policy questions. Our department, for a variety of historical reasons, is full of people for whom I think those are the most interesting and important questions to study. That’s for us, I think, the sweet spot. They use the full range of techniques that are available, but they use them to a purpose and not as a goal in and of themselves, which is often what they seem to become in the hands of many academics.

So that has been the response of Johns Hopkins in partial equilibrium. One consequence is that the students that we train tend to be particularly attractive to policy institutions like the IMF and the Fed and the European Central Bank and places where you need some ability to grapple with the real world.

EF