

Jonathan A. Parker

Editor's Note: *This is an abbreviated version of EF's conversation with Jonathan Parker. For additional content, go to our website: www.richmondfed.org/publications*

Economists are sometimes pegged as either theorists or empiricists. But often this dichotomy is overstated. Many economists bring together theory and empirical analysis to study a broad range of questions. For Jonathan Parker, this approach is perhaps the defining characteristic of his work.

Parker, the Robert C. Merton (1970) Professor of Finance at the Massachusetts Institute of Technology's Sloan School of Management, uses data in novel ways to better understand a host of economic issues and the theories that underpin them. For instance, the economic stimulus program of 2008 offered the potential to examine the way households respond to an influx of liquidity — and with it, whether people smooth their consumption, as theory would predict. But to realize that potential required developing some investigational tools — in Parker's case, designing surveys for households belonging to the Nielsen Consumer Panel to better understand what they did with the payments they received and why.

Parker has also looked at such issues as whether people can hold incorrect but nonetheless utility-optimizing beliefs; which segments of the income distribution are most affected by economic shocks and how that has changed over time; and whether households respond to good economic news in a proportionate manner to bad economic news. As he says, he's an applied microeconomist, an asset pricer, a macroeconomist, a public finance economist, and a behavioral economist. Which one depends on the question at hand and the methods required to answer it.

Prior to joining the MIT faculty, where he is also the co-director of the Golub Center for Finance and Policy, Parker taught at Northwestern University, Princeton University, and the University of Wisconsin, and he was a research fellow at the University of Michigan. He edits the National Bureau of Economic Research's *Macroeconomics Annual*, serves on the board of editors of the *American Economic Review*, and is a member of the Congressional Budget Office's Panel of Economic Advisers. Aaron Steelman interviewed Parker at his office at MIT in December 2016.



EF: Among your work on economic stimulus programs is a recent paper with Daniel Green, Brian Melzer, and Arcenis Rojas on the Car Allowance Rebate System (CARS) of 2009, popularly known as “Cash for Clunkers.” Could you discuss the empirical findings of that paper as well as potential implications for structuring stimulus programs given what we know from participation in CARS?

Parker: One of the interesting things we saw about that program was that it was massively oversubscribed. The government originally allocated \$1 billion to a three-month program and exhausted that \$1 billion in about a week. It then reauthorized the program for another \$2 billion and still ran out of funds two months early. The other notable thing was that it was a program that provided liquidity. It paid households \$3,500 or \$4,500 to trade in and scrap an old vehicle. And that means it provided liquidity — and really enough liquidity for a down payment. So we wanted to know: Can we link these two, the provision of liquidity and the high take-up rate? Also, there was interesting existing research that had been done on the program, specifically work by Atif Mian and Amir Sufi, which produced a nice aggregate impact measure of the program but nothing at the micro level of how individual households were responding. And I wondered about the reversal of the impact, which is one of their main findings: The program generated sales, but within six to nine months afterward there was no cumulative difference in purchases for people eligible for the program and people who weren't.

We got access to the Bureau of Labor Statistics' Consumer Expenditure Survey data and made a precise measure of eligibility of vehicles based on fuel efficiency and used car value by make, model, and year. We then mapped the program responses to eligibility and the economic subsidy associated with any given car. If you owned a car, the economic subsidy was the program payment minus the value you could get for your car on the used car market. So a car worth \$4,500 on the used car market would get, in effect, no subsidy from the program, but a car worth \$1,000 would get a \$3,500 subsidy. We mapped from car value to the program response to see if people with really junky old cars used the program much more strongly. And indeed we found that to be the case. Typically, about \$1,000 of used car value reduced your probability of participating in the program by about half a percentage point. That suggests the government could have gotten as large a response with slightly smaller subsidies because the program ran out of funds and there was a lot of response from people with moderate-valued vehicles.

EF: But how can you know people's sensitivity to the subsidy in advance?

Parker: Exactly the right question. Because the program ran out of money, it's not a program for which we observe an unconstrained, equilibrium response. Instead, it was a response constrained by the funding amount. So we don't absolutely know; what we do know is that the subsidies were more generous than they needed to be to generate that many sales. And what we also know is, had the subsidy been lower, it probably would have been the people with the lousiest cars who would have traded them in and that would have resulted in less destruction of more valuable used cars. That's all easy to look at and say after the fact. But there was this massive underestimation of the response to the program, and we think that's because of liquidity.

We think that an economic subsidy should generate intertemporal substitution; it's a temporary price subsidy to a durable good. In this case, we figured out how the liquidity dimension could actually be measured separately a little bit from the economic subsidy. The economic subsidy is not the same for everyone, but for most people it is the same as the liquidity provided by the program. But some people have loans on their program-eligible vehicles. If a vehicle is securing a loan, then when it's brought into the dealer and scrapped as part of the program, the household has to pay off that loan, and so they lose some of the liquidity benefit of the program. In our study, we estimate this liquidity effect, separate from the intertemporal substitution effect of the economic subsidy, and we find that the effect of the program was much smaller on vehicles that were securing loans. In fact, it's almost nonexistent. So we find the impact of liquidity to be very strong — it was an accelerant for the economic subsidy in the target population.

We also find very weak evidence consistent with the reversal effect that Mian and Sufi first discovered, which feeds

into the question: Is this a worthwhile sort of program to do? It was a program that caused, at an annual rate, a \$44 billion increase in personal consumption expenditures on durable goods in the third quarter of 2009, which was the quarter in which the recession ended and in which GDP grew by about \$44 billion. And in the previous quarter GDP declined by about \$44 billion. So it looks pivotal. On the other hand, half of the content of the vehicles purchased under the program was imported, so that means that one has to take the number of new purchases and divide by two to get an estimate of the partial-equilibrium impact on demand. So really it wasn't pivotal at moving us from no growth to growth, and also the program seems to have been reversed over six to nine months because there's no cumulative impact in sales. On the other hand, it generated all that spending for a relatively small fiscal cost of only \$3 billion (\$12 billion at an annual rate). But these are all accounting, partial-equilibrium calculations.

For this to be optimal from a stabilization perspective, you need to believe that the government multiplier is much larger in the quarter in which CARS is run than six months later. And this is a period when we are having a slow recovery. So the net benefit of the program is ultimately a general equilibrium question that other people would need to answer, but the hurdle is significant given that one has to see such a significant swing in the size of the multiplier between those two periods. If one wants to do a similar program again, and similar programs have now been run in countries all over the world, our results generally emphasize that the liquidity was a crucial part of the program — not just people substituting over time due to a temporary price subsidy — and as such, our findings relate to the literature on investment tax credits for firms where liquidity also seems to be important.

EF: I would like to go back to some of your earlier work on household financial decisionmaking — in particular, your 2002 *Econometrica* paper with Pierre-Olivier Gourinchas. It seems consistent with the standard life cycle theory of saving and consumption. Would you say that's a fair characterization of that paper?

Parker: From the perspective of today, I think the contribution of that paper is more methodological in some sense. We worked out a framework for taking cohort-level analysis of microdata that had been used nicely before by Angus Deaton and Christina Paxson, Orazio Attanasio, and others, and combined it with a structural model of an income fluctuation problem so as to estimate the parameters governing the behavior of households using a simulated method of moments estimator. That said, as you noted, the model fits the life cycle profiles of consumption and saving with a model in which households differ solely based on their history of income shocks and their age. So age is a major determinant of the propensity to spend. Since then, the research has expanded in many ways to endogenize the choices we made exogenous in that paper or assumed away, including portfolio choice, labor supply, illiquid retirement saving, government

programs, housing, and some very nice work by Mariacristina De Nardi and Eric French and co-authors on retirement. People are also considering the liquidity of different investments now in structural models and stochastic credit constraints, all of which we pushed away, but the method remains a very useful one for evaluating these models.

EF: You've revisited some questions fundamental to life cycle theory in your recent paper, "Why Don't Households Smooth Consumption?"

Parker: In that paper, I use Nielsen Consumer Panel data to design and run my own survey on households to measure the effect of what was then the second of these large randomized experiments run by the U.S. government, the economic stimulus program of 2008. The key feature of that program was that the timing of the distribution of payments was determined by the last two digits of the Social Security number of the taxpayer, numbers that are essentially randomly assigned. So the government effectively ran a \$100 billion natural experiment in 2008, distributing money randomly across time to people, and this policy provides a way to measure quite cleanly how people respond to infusions of liquidity.

The goal of "Why Don't Households Smooth ..." is to provide evidence of the structural model underlying the observed importance of liquidity on household spending behavior. And in theory, while the buffer-stock model might correctly match the behavior, it also might be that people spend expected income gains only when they arrive because of problems stemming from self-control, inattention, inability to plan, some sort of rule of thumb or mental accounting behavior, or the like. So I designed a bunch of questions trying to get at these alternative behaviors. I should clarify that they are not really alternatives, in the sense that they all interact with liquidity constraints.

The first thing I found out is that illiquidity is still a tremendous predictor of who spends more when a predictable payment arrives. But it's not only liquidity. People with low income have a very high propensity to spend, and not just people who have low income today, as would be associated with the standard buffer-stock model. You can imagine a situation where you've had a bad income shock, you happen to have low liquidity, and you spend a lot. But illiquidity one or even two years prior to the payment is just as strongly

Jonathan A. Parker

► Present Position

Robert C. Merton (1970) Professor of Finance, Massachusetts Institute of Technology

► Previous Faculty Positions

Northwestern University (2007-2013); Princeton University (1999-2007); University of Wisconsin (1997-1999)

► Education

Ph.D. (1996), Massachusetts Institute of Technology; B.A. (1988), Yale University

► Selected Publications

"Why Don't Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment," *American Economic Journal: Macroeconomics*, forthcoming; "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013 (with co-authors); "Who Bears Aggregate Fluctuations and How?" *American Economic Review*, 2009 (with Annette Vissing-Jorgensen); "Optimal Expectations," *American Economic Review*, 2005 (with Markus Brunnermeier); "Consumption Over the Life Cycle," *Econometrica*, 2002 (with Pierre-Olivier Gourinchas)

associated with a propensity to spend out of liquidity, as illiquidity at the time of the payment. This same set of people who have persistently high propensities to consume are also the people who characterize themselves as the type of people who spend for today rather than save for tomorrow when I asked them specifically about their type, not their situation. They are also the people who report that they have not sat down and made financial plans.

What you end up with is that a high propensity to consume correlates with low liquidity, which is useful for theorizing but also presents a little bit of a chicken-and-egg problem. Is it different preferences, objectives, or behavioral constraints that are causing both the low liquidity and the propensity to spend, or is it the low liquidity that is causing the lack of planning and high spending responses? So for many purposes, what I take my findings to mean is that the buffer-stock model is a quite reasonable model with one critical ingredient. The critical difference relative to the way I modeled households

in the 2002 paper with Gourinchas is that I think there's much more heterogeneity in preferences across households. While in that paper we looked at differences in preferences across occupation and industry, I think there's just much more persistence in heterogeneity in behavior, consistent in the buffer-stock model with differences in impatience. Partly I say this because I do not find a big relationship between age and propensity to spend in a number of studies, and partly from the persistence of the high-spending propensities I find in this recent paper. But it's also visible in some sense in even older data. Low liquidity, or low financial wealth, is a very persistent state across households, suggesting the propensity to spend is not purely situational. A lot of it is closer to an individual-specific permanent effect than something transient due to temporary income shocks.

EF: Did people generally understand the magnitude of the 2008 stimulus program prior to receiving payments? And if they didn't, did that show up in consumption patterns?

Parker: In my study, one of the questions I asked people was: So you got this economic stimulus payment, did you expect it? Was it more than you expected? Was it less than you expected? Was it a surprise? First of all, about 80 percent of households got basically what they expected. That means

you're never going to explain the spending response by people being surprised, as say in some versions of an inattention model. That is a nonstarter. Expectations about the program were reasonably accurate, with the important caveat that people may not be answering the survey truthfully. Interestingly, there is a slightly higher propensity to spend, though not statistically significantly so, among those who were surprised and received more than they expected. But there is also exactly the same response among those people who got less than they expected. So it looks more like the people who weren't expecting the right thing are worse at consumption smoothing.

EF: How do you define the distinction between “optimal expectations” and “rational expectations”? What are the differences in the ways agents with each set of expectations tend to behave? And if agents with optimal expectations may make “poorer” decisions, in some sense, how may that ultimately be advantageous or desirable?

Parker: In some sense, the starting point for my work with Markus Brunnermeier came from a number of observations in the social psychology literature that people just tend to be optimistic or overconfident, the type of behavioral biases that lead people to believe they're better drivers than average — that sort of optimism. Looking at the objective functions that we usually consider, the simplest way to maximize the expected present discounted value of anything is to put more probability on better outcomes — simply to be more optimistic. You can see how that can be a source of happiness today. If we think about how good we are at many different things, it's nice to have confidence and believe you're maybe better looking or smarter than you actually are. On the other hand, to the extent that you actually allow yourself these sorts of enjoyable biases, you're likely to make slightly worse decisions. You might leave insufficient time to complete a project, for instance, which would make you worse off.

So the basic idea of that optimal expectations paper is to think of the optimal trade-off between those two — the idea that you will get more expected future utility today by expecting better outcomes, but on the other hand you're going to make some decision errors because of that expectation. It turns out that this sort of a simple trade-off has many interesting implications. The first is basically that you're always somewhat optimistic. The reason is that moving a small amount of probability from, say, the worst state out in the future to the best leads to a first-order gain in expected present discounted value of utility flows of consumption. But a small change in probability causes a small change in behavior, and a small change in behavior from the optimal has very small — second-order — welfare costs. So, overall, the benefits outweigh the costs.

There are also some interesting implications that come from the fact that optimism is situational. For example, when considering investing, one way to be optimistic is to think the stock market's going to go up more than everybody else believes, and to go longer into it. But you can also be

optimistic by shorting the market and believing it's going to crash when everybody else thinks it's going to go up. So when do you short? It turns out that there are conditions under which you will actually invest in an unfair bet if it's positively skewed enough. That gives you a theory that looks like people buying lottery tickets, which are unfair gambles with a very small probability of a very high positive payout. They provide a very nice future state to believe in at a pretty low-dollar cost today. So the observed unfair gambles, lottery tickets, are exactly the sort of unfair gambles that our theory predicts people should prefer. This type of behavior looks like the big short. That's a theme that runs through several of our results: People with optimal expectations want something that has a very high positive payoff to dream about that at the same time isn't very costly to invest in.

There is also a natural nonconvexity in the model, which we didn't expect. When I am more optimistic about a certain outcome in the future, that means I want to buy more consumption in that state of the world. When I buy more consumption in that state, that means I want to be more optimistic about it, which in turn means I want to buy more consumption there. And this natural nonconvexity means that people are going to do something like hold a reasonably well-diversified portfolio and then invest excessively in a particular asset, such as one or two individual stocks. We didn't expect that sort of behavior to pop out, but that's what the model taught us. This leads to our work with Christian Gollier that looked at the conditions under which the model generated disagreement and could raise the return on negatively skewed assets.

EF: How would you describe the changes we have seen in the way high-income and high-consumption households have become exposed to aggregate economic fluctuations over the last 30 years roughly?

Parker: Due to some difficult data issues, I have not really been able to track the consumption of high-consumption households, but in work with Annette Vissing-Jorgensen we have looked at how the labor income of high-income households has changed significantly. What we zoomed in on is that high-income households used to live a relatively quiet life in the sense that the top 1 percent would earn a relatively stable income, more stable than the average income. When the average income dropped by 1 percent, the incomes of the top 1 percent would drop by about only six-tenths of a percent. In the early 1980s that switched, so that in a recession if aggregate income dropped by 1 percent, the incomes of the top 1 percent dropped more like 2.5 percent — quadrupling the previous cyclical volatility. So now they're much more exposed to aggregate fluctuations than the typical income. We also show that decade by decade, as the top income share increased, so did its exposure to the business cycle in the 1980s, 1990s, and 2000s. And as you go further and further up the income distribution, that top share — not just the top 1 percent, but the top 10th of a percent, and the top 100th of

a percent — there's also been a bigger increase in inequality and a bigger increase in the exposure to the business cycle.

EF: What's the story for that?

Parker: First of all, we used to think the income cyclicality was exactly the reverse, because low-income workers would lose their jobs in recessions and high upper-income workers would not. And so while high-income households might get lower raises in recessions, they wouldn't actually go down to zero. Since job losses are concentrated among lower-earning workers, you have much greater cyclicality in overall incomes among low-wage workers. In another paper I did with Annette Vissing-Jorgensen, we looked at cross-country evidence in the recent decades of high inequality. The countries with the biggest earnings inequality were also the countries with the largest high-income cyclicality relative to the average. So what explains these sorts of findings? We thought there were two leading hypotheses.

First, starting around the end of the 1980s, we see the adoption of incentive-based pay for CEOs and other highly placed managers. Incentive compensation over this time rises, and it happens to be that the incentive compensation is not based on relative performance, which would therefore differ out what goes on in the macroeconomy, but instead is based on absolute performance. And in the U.S. case, that could partly be due to simply what counts legally as incentive-based compensation and so is not subject to corporate profits tax. Pay in the form of stock options, for example, counts as incentive-based compensation. Pure salary does not and so is taxed as corporate profits above \$1 million.

The other possibility is that it's purely technological. Something like incentive-based compensation may be a sideshow. The idea is that new information and communication technologies allow the best managers to manage more people, to run bigger companies, and therefore to earn more; the best investment managers to manage more money and to make more for themselves; the best entertainers and performers to reach more people and therefore earn a larger share of the spending on entertainment goods. High earners have become small businesses. While it is not universally true that such a shift to high-volume low-markup profits for the winners necessarily leads to greater cyclicality, it is true for some reasonable functional forms of production.

We do know that increased cyclicality in income among high earners can't come simply from the financial sector. That sector just isn't quantitatively big enough, and you see the increase in earnings share and in cyclicality across industries and occupations. It's not the case that just the top hedge fund managers have become the high earners and they're very cyclical; Oprah is also.

EF: What do you think are the most important unanswered or understudied questions in household behavior and household financial decisionmaking?

Parker: The big one is: Do we need a different model than the canonical stochastic life cycle model with credit constraints to understand consumer behavior? Do we need to introduce inattention or hyperbolic discounting, for instance, to make it richer? My sense is that for a lot of questions so far, the answer is still no, but we now have a few pieces of evidence that in a few places the answer is yes. As we get better data, and we think about questions like credit market equilibria and consumer financial regulation, we have the information to evaluate rich models of behavior and the need for models that are as complete as possible in describing behavior.

In my work, liquidity is first order, consistent with the buffer-stock model. But liquidity almost seems to explain too much. In the Nielsen study that we discussed earlier, people don't spend the money the week before it shows up — they spend it the week it shows up. And it seems like you're going to have a lot of difficulty quantitatively fitting that little foresight into a life cycle model unless people are often literally liquidity constrained, absolutely at their debt limits.

In the Cash for Clunkers program, liquidity mattered critically. One interpretation is that this importance is consistent with the canonical model in which some people lack the liquidity for a down payment. But there is an alternative interpretation. Again, our main finding is that people who have outstanding loans on their vehicles are much less likely to participate in the program, presumably because to buy a new car using the program, they would have to put some cash down along with the payment in order to make the down payment. Such people are much less likely to take advantage of the program than people who don't have loans on their vehicles but instead have unsecured debt, like on a credit card. Sounds like liquidity. But perhaps the people who have the secured debt could walk into the dealer and turn into that other person — that is, use their credit card to buy the car, so they leave the dealer with unsecured debt, just like the other person. In this case, liquidity matters, but maybe not according to strictly the life cycle model with liquidity constraints. Instead, such behavior sounds more like people using heuristics or mental accounts. The big question: In what combination do we need each ingredient — rationality and heuristics? And where do the heuristics come from?

The other question that I think research is really exploring is what equilibria look like for saving and borrowing. What equilibrium supports high-fee mutual funds, index funds, and so on, and how does that change the flow of funds between the corporate and household sector and the pricing of risk? How does the market for lending to households evolve as risk is repriced and interest rates move, and how does this feed back into spending? The interplay between borrowers and lenders in these markets is a very interesting and active area of research because we're getting a lot of the data on mortgages, credit cards, retail investment, and financial accounts. These data are allowing us to look at and understand the equilibria in those markets, which is really fun. **EF**