Published Quarterly by the Federal Reserve Bank of Minneapolis

December 2014

The Region

Raj Chetty

Clarifying "Price Stability"

Illusion of School Choice

Too Correlated to Fail

Research Digest: *Recession and Revival*

Are Economists Worth $1 Trillion?

Volume 28 Number 4
December 2014
ISSN 1045-3369

The Region

Executive Editor: Kei-Mu Yi
Senior Editor: David Fettig
Editor: Douglas Clement
Managing Editor: Jenni C. Schoppers
Senior Writer: Phil Davies
Designers: John Biasi, Lori Korte, Mark Shafer

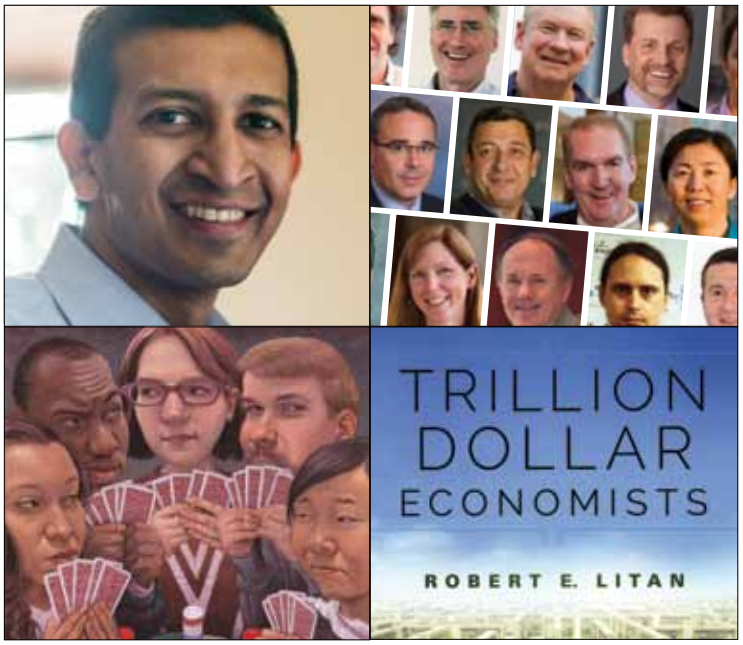

The Region
Federal Reserve Bank of Minneapolis
P.O. Box 291
Minneapolis, MN 55480-0291

Email: letters@mpls.frb.org
Web: minneapolisfed.org

The Region *is published by the Federal Reserve Bank of Minneapolis. The views expressed here are not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Articles may be reprinted if the source is credited and the Public Affairs Department of the Minneapolis Fed is provided with copies. Permission to photocopy is unrestricted.*

Cover photograph by Peter Tenzer

NINTH DISTRICT
NOTES

Clarifying the Meaning of Price Stability

Narayana Kocherlakota

President
Federal Reserve Bank of Minneapolis

Over the past few years, the Federal Open Market Committee has made great progress in formulating, and communicating, the objectives of monetary policy to the public. In this column, I'll discuss some of that progress and also describe two ideas about how the Committee can make further improvements along these lines. My discussion will build on a speech, "Clarifying the Objectives of Monetary Policy,"[1] that I gave several times in the fall, as well as on information in the recently released minutes from the Oct. 28-29 FOMC meeting.

The framework statement

The natural starting point for any discussion of monetary policy goals is the Federal Reserve Act, the law in which Congress created the Fed and defined its purposes. Through the Federal Reserve Act, Congress requires the Federal Reserve to make monetary policy so as to promote effectively the goals of maximum employment, stable prices and moderate long-term interest rates. Most economists believe that if the Fed achieved the first two mandates (maximum employment and stable prices), it would automatically achieve the third (moderate long-term interest rates). Hence, monetary policymakers in the United States are usually described as having a dual mandate: to promote price stability and maximum employment.

Congress' short overarching description of Federal Reserve objectives is the foundation for current monetary policymaking, but it does not address many specifics. In January 2012, in a key milestone in the evolution of the Fed's communications, the FOMC adopted a longer and more precise description of its long-run goals. I'll call this short but pathbreaking document the "framework statement." It contains a number of important ideas, and indeed I encourage all Americans to read the entire statement.[2]

In this column, I'll stress only what I see as the most important aspect of the statement: It translates the words "price stability" into a longer-run goal of a 2 percent annual inflation rate. Here, the term "inflation rate" refers specifically to the personal consumption expenditures (or PCE) inflation rate. This is a measure of the rate of increase in the prices of all goods and services, including those related to food and energy. The adoption of this explicit 2 percent target means that the American public need

I've suggested that the FOMC clarify that its inflation target is symmetric and that the Committee typically seeks to achieve that target within a two-year horizon. Let me emphasize that these two suggestions represent clarifications, not alterations. The framework statement, as written, is completely consistent with the formulations of price stability that I've proposed. However, a shortcoming with the current statement is that it is also consistent with other interpretations of price stability.

guess no longer about the Federal Reserve's inflation intentions—either on the upside or on the downside: 2 percent is our goal.

Possible clarifications in the formulation of the price stability mandate

The framework statement was adopted by the FOMC in January 2012. It has been reaffirmed, with only minor wording changes, in January 2013 and again in January 2014. However, the minutes for the January 2014 meeting note that FOMC participants saw the coming year as an appropriate time to consider whether the statement could be improved in any way. I concur: The time is right to consider sharpening the FOMC's statement of its objectives in several ways. In what follows, I'd like to explain, and express support for, two particular clarifications related to the FOMC's formulation of the price stability mandate.[3]

First, I believe the FOMC should be clear that its inflation objective is *symmetric*. Many observers emphasize the need to keep inflation from rising above 2 percent. But in my view, inflation *below* 2 percent is just as much of a problem as inflation *above* 2 percent. The central bank of Canada also has a 2 percent inflation target. Its language about symmetry is pretty clear, at least as central banking communications go: "the Bank is equally concerned about inflation rising above or falling below the target and will act … in order to bring inflation down, or to push it back up, to 2 per cent."[4] In my view, the FOMC should use similar language to characterize its inflation objective.

Why do I see symmetry as important? Without symmetry, inflation might spend considerably more time below 2 percent than above 2 percent. Inflation persistently below the 2 percent target could create doubts in households and businesses about whether the FOMC is truly aiming for 2 percent inflation, or some lower number. This kind of unmooring of inflation expectations would reduce the effectiveness of monetary policy as a mitigant against adverse macroeconomic shocks.

Second, I believe that the FOMC should consider articulating a benchmark two-year time horizon for returning inflation to the 2 percent goal. (Two years is a good choice for a benchmark because monetary policy is generally thought to affect inflation with about a two-year lag.) Right now, although the FOMC has a 2 percent inflation objective over the long run, it has not specified any time frame for achieving that objective. This lack of specificity suggests that appropriate monetary policy might engender inflation that is far from the 2 percent target for years at a time and thereby creates undue inflation (and related employment) uncertainty. Relatedly, the lack of a public timeline for a goal can sometimes lead to a lack of urgency in the pursuit of that goal. I believe that, if the FOMC publicly articulated a reasonable time benchmark for achieving the inflation goal, the Committee would be led to pursue its inflation target with even more alacrity.

Some might argue that this kind of time horizon is impractical. In fact, many central banks incorporate a similar timing benchmark. For example, the Bank of Canada typically makes its monetary policy choices so that the inflation rate is projected to return to 2 percent within two years.[5] I say "typically"—there are certainly situations in which the Bank of Canada chooses policy so that inflation is projected to return to target more slowly (sometimes taking as long as three years) or more rapidly (sometimes as quickly as 18 months). But it continues to treat two years as a benchmark, in the sense that it feels compelled to explain *why* it is choosing a different time horizon.

To sum up: I've suggested that the FOMC clarify that its inflation target is symmetric and that the Committee typically seeks to achieve that target within a two-year horizon. Let me emphasize that these two suggestions represent clarifications, not alterations. The framework statement, as written, is completely consistent with the formulations of price stability that I've proposed. However, a shortcoming with the current statement is that it is also consistent with other interpretations of price stability (such as a 10-year horizon for returning inflation to the desired target).

Discussion at the Oct. 28-29 FOMC meeting

As the minutes from the Oct. 28-29 FOMC meeting reveal, both of these possible clarifications to the framework statement were discussed at that meeting. With reference to the former clarification (symmetry), the minutes from the meeting say that "there was widespread agreement that inflation moderately above the Committee's 2 percent goal and inflation the same amount below that level were equally costly."[6] I am glad to see this kind of FOMC consensus on this important issue.

The minutes go on to say that "many participants thought that this view was largely shared by the public."[7] I am not one of these participants. The public's main reference on the FOMC's monetary policy objectives is the FOMC's framework statement. The statement makes no reference to symmetry. Without such a reference, we cannot expect the public to know that the FOMC views deviations from its inflation objective in a symmetric fashion. I would support including the above clause from the minutes in the framework statement itself. R

Endnotes

[1] Narayana Kocherlakota, "Clarifying the Objectives of Monetary Policy," speech at the St. Paul Rotary, St. Paul, Minn., Nov. 18, 2014, available at minneapolisfed.org/news_events/pres/speech_display.cfm?id=5426.

[2] "Statement on Longer-Run Goals and Monetary Policy Strategy," Board of Governors of the Federal Reserve System, as amended effective Jan. 28, 2014, available at federalreserve.gov/monetarypolicy/files/FOMC_LongerRunGoals.pdf.

[3] In May, I gave a speech at the Economic Club of Minnesota that discussed, without endorsing, the idea that the FOMC should use price level targeting as opposed to inflation targeting. I continue to think that the Committee should consider this change in approach. However, it's a topic that deserves more space than what is available in this column.

[4] "Monetary Policy," Bank of Canada, May 29, 2012, available at bankofcanada.ca/wp-content/uploads/2010/11/monetary_policy.pdf.

[5] "Monetary Policy," Bank of Canada, May 29, 2012, available at bankofcanada.ca/wp-content/uploads/2010/11/monetary_policy.pdf.

[6] "Minutes of the Federal Open Market Committee," Oct. 28-29, 2014, available at federalreserve.gov/monetarypolicy/fomcminutes20141029.htm.

[7] "Minutes of the Federal Open Market Committee," Oct. 28-29, 2014, available at federalreserve.gov/monetarypolicy/fomcminutes20141029.htm.

Too Correlated to Fail

Anticipation of bailouts encourages banks to invest alike, making bailouts more probable and crises more severe

V. V. Chari
University of Minnesota
Federal Reserve Bank of Minneapolis

Christopher Phelan
University of Minnesota
Federal Reserve Bank of Minneapolis

Economic Policy Papers *are based on policy-oriented research produced by Minneapolis Fed staff and consultants. The papers are an occasional series for a general audience. The views expressed here are those of the authors, not necessarily those of others in the Federal Reserve System.*

Executive summary

In this paper, we argue that the anticipation of bailouts creates incentives for banks to herd in the sense of making similar investments. This herding behavior makes bailouts more likely and potential crises more severe. Analyses of bailouts and moral hazard problems that focus exclusively on bank size are therefore misguided in our view, and the policy conclusion that limits on bank size can effectively solve moral hazard problems is unwarranted.

Introduction

Misery loves company.[1]

Since Kareken and Wallace (1978), it has been well understood that deposit insurance creates incentives for banks to take on excessive risk. Protected from losses by deposit insurance, bank depositors will rationally pay little or no attention to the riskiness of their bank's portfolio; consequently, the interest rate a bank needs to offer to attract deposits will not be sensitive to the risk characteristics of its portfolio—undermining the usual risk/return trade-off faced by investors. Banks that seek to maximize shareholder value therefore have an incentive to take on more risk than they would if their deposits were uninsured. Indeed, banks that trade equity on public markets have strong incentives to take on as much risk as regulators allow. This phenomenon of one party taking excessive risks because another party bears all or some of the cost of failure is often referred to, in banking and other spheres, as "moral hazard."

Stern and Feldman (2004) argue that when a large financial institution is confronted with the possibility of failure, policymakers concerned about broader systemic fallout from that failure have strong incentives to intervene. Even uninsured debtholders may be bailed out to prevent failure, and expectations of such bailouts induce them to be relatively unconcerned about the level of risk of their financial institutions. Just as with the *explicit* protection of deposit insurance, the lack of concern generated by *implicit* guarantees of government bailouts encourages banks to take on excessive risk.

Stern and Feldman's argument has been interpreted (or misinterpreted) to mean that policymakers should be concerned about potential failure of *large* financial institutions *only*. This interpretation suggests that a simple method of curing this moral hazard problem is to set regulatory limits to ensure that no individual financial institution is "too" big.

This policy conclusion is mistaken, we argue in this paper. Policymakers do not intervene when big banks

are threatened simply because those *banks* are too big. Rather, they intervene because the potential *systemic costs* resulting from bank failure are considered too big.

Bank size is not the issue

Consider two scenarios, one without regulatory limits on bank size and the other with such limits. Suppose that when regulations do limit bank size, small banks—all below the size limit—choose scaled-down versions of the large bank's portfolio. That is, each small bank's portfolio has holdings in the exact proportion, but smaller size, of that large bank's portfolio. (For simplicity, assume all depositors are identical so that the characteristics of depositors in all banks—large and small—are identical.)

The very prospect of government bailouts creates an incentive for banks—regardless of size—to take on highly correlated risks, which, in turn, raises the likelihood of financial crisis.

If the aggregate economy is hit with a shock that adversely affects investment portfolios, and the survival of both large and small banks is in doubt, would the aggregate costs of banking system failure differ under the two scenarios? Clearly not. Since the collective financial assets, liabilities and risk profiles are identical whether bank size is limited or not, the systemic costs of not bailing out banks are exactly the same. Therefore, policymaker incentives to undertake bailouts are unaffected by bank size limits, if collections of smaller banks assume the same or similar portfolio risk as would one big bank.

Proponents of bank size limits as a solution to the moral hazard problem induced by bailouts implicitly assume that the combined portfolio of a collection of smaller banks will be less risky than the portfolio of a large bank of equivalent size. This assumption is unwarranted, we contend. In fact, the very prospect of government bailouts creates an incentive for banks—regardless of size—to take on highly correlated risks, which, in turn, raises the likelihood of financial crisis.

Policymakers will intervene when the aggregate assets of threatened financial institutions are sufficiently large to represent a substantial risk to the broader economy should those institutions fail. The following example illustrates the manner by which this policy motivation creates an incentive for banks to take on correlated risks. (We provide a numerical version of this example in the second section of the paper.)

Consider an extreme case where U.S. banks can invest in mortgages to residents of just two states, either Florida or New York (both have basically the same size population). We'll further assume that just one of these states will have a high default rate, but that banks don't know which state that is until after the mortgages are sold. In a well-functioning market without regulator bailouts of failing banks, banks will invest roughly half their assets in each state, since default rates are not known in advance—thereby providing themselves, through diversification, with the highest possible level of protection from loss.

But suppose that, for some reason, all banks invest in Florida mortgages only. If Florida turns out to have the high default rate, then all banks are threatened with failure, and policymakers have a strong incentive to bail them out. From the perspective of an individual bank considering whether to buy Florida or New York mortgages, it is therefore rational to buy only Florida mortgages. Each bank knows that if Florida mortgages default, it is assured a government bailout precisely because *all* banks are threatened, and the government will therefore intervene to prevent broad systemic failure.

If, again, virtually all banks invest in Florida mortgages, but instead it is *New York* mortgages that have a high default rate, a particular bank that bucked the herd and invested in New York mortgages would not receive a bailout, since the system as a whole is not threatened. Thus, the existence of a bailout policy encourages all banks, regardless of size, to invest similarly and thereby correlate their risk portfolios.

How would banks go about correlating their risks in a more realistic world? One way to do this is through securitization, a practice that has become extremely prevalent in recent decades. Bank loans are securitized by selling claims to a pool of those loans. Securitization of this form allows banks to diversify their portfolios and ensures that their profits are not unduly dependent on the idiosyncratic risk of the loans that they have originated. But because securitized loans are usually held by other banks, the practice ensures that all banks end up holding very similar portfolios and thus have highly correlated risk.

This paper argues that limits on bank size miss the point. What truly matters to the well-being of the broad economy is not the risk profile of any given bank portfolio, large or small, but the risk profile of the *entire banking system*. Regulators therefore need to understand what kinds of events are likely to threaten a significant fraction of the aggregate assets of the entire banking system, rather than concentrate (as current policies do) on a limited number of large banks. In particular, they must focus on how the portfolio of the entire banking system is exposed to such events. Regulation of a given bank then should deal with whether that particular bank's behavior is mitigating or aggravating the risk exposure of the entire system. In brief, we need stress tests of the entire banking system, not just of individual banks.

This paper argues that limits on bank size miss the point. What truly matters to the well-being of the broad economy is not the risk profile of any given bank portfolio, large or small, but the risk profile of the entire banking system.

A numerical illustration

In this section, we provide a numerical (though still extreme and hypothetical) example of the New York/Florida scenario.

A large number of banks have access to investment funds, and they can invest only in New York or Florida mortgages. Each bank separately chooses what fraction of its funds to put into New York mortgages versus Florida mortgages. In each state, mortgages face both idiosyncratic risk (meaning a risk situation particular to that mortgage) and aggregate risk (experienced by the entire state). Every bank makes an individual decision about the fraction of its funds it will invest in each state, and all banks make their investments before anyone knows what the future risk scenario will be. Banks are aware of investment decisions made by other banks.

We'll assume that, after investment decisions are made, the aggregate economy can be in one of three situations:

- With 5 percent probability, a housing collapse occurs in New York, but not in Florida.
- With 5 percent probability, a housing collapse occurs in Florida, but not in New York.
- With 90 percent probability, no housing collapse occurs anywhere in the United States.

Let's also assume (generously) that a $1 mortgage investment returns $1.50, but only if the borrower doesn't default. In addition, we'll say that 30 percent of borrowers will default if their state suffers a housing collapse, but 10 percent will default if there is no collapse.

Bank investments in New York work as follows (and symmetrically for Florida): If New York experiences a housing collapse, each dollar invested there has a 70 percent chance of yielding $1.50, but a 30 percent chance of yielding nothing. If Florida experiences the housing collapse or the nation as a whole is free of a housing crisis, a dollar invested in New York has a 90 percent chance of a $1.50 return and a 10 percent chance of total loss. In sum, New York mortgage investments (logically) are more likely to yield nothing if a housing market collapses there than if it collapses in Florida or not at all.

Similarly, a dollar invested in Florida mortgages has three times the likelihood of returning nothing if Florida housing collapses relative to the chance of a total loss if the New York market fails or there is no housing crisis anywhere in the United States.

The point of this setup is to present a situation where the banking system's exposure to aggregate risk is determined by the choices of many small actors—in this case, small banks. Here's how it would work, depending on where those small actors invest:

If all banks invest in New York, there's a 95 percent chance that each will get $1.50 back for 90 percent of total dollars invested (given the 10 percent default rate) and a 5 percent chance that each will get a full return on just 70 percent of the bank's total investment (since 30 percent of mortgages will default). That works out to a mean return of $1.335 =1.5*(.95*.9 +.05*.7) per dollar invested.

And if all banks split their investments 50-50 between Florida and New York, or half the banks invest totally in Florida and half just in New York? Here, too, the mean return is $1.335 per dollar invested.[2]

Given this situation, where aggregate risk is determined by many small banks (not simply those considered "too big"), what role does government policy play? More particularly, how does the presence or absence of policy intervention through bailouts affect bank decisions and aggregate risk?

But if banks diversify over states, either by each bank diversifying between New York and Florida or by half the banks investing in New York and half in Florida, the mean total return remains the same, but the variance is lower and the portfolio's worst case scenario is better. (If all banks invest in one state, the worst-case scenario is 30 percent of loans fail. If banks diversify over states, the worst case scenario is 20 percent of loans fail.)

Given this situation, where aggregate risk is determined by many small banks (not simply those considered "too big"), what role does government policy play? More particularly, how does the presence or absence of policy intervention through bailouts affect bank decisions and aggregate risk?

In a world without government intervention, if banks are at all risk-averse, they will each invest half their funds in a large number of New York mortgages and the other half in a large number of Florida mortgages. This ensures that each bank makes a return of $1.35 (= 90 percent of $1.50) per dollar invested if there is no housing collapse and $1.20 (= 80 percent of $1.50) per investment dollar if either Florida or New York suffers a housing collapse. Investing in any other proportion is a "mean preserving spread," something that risk-averse entities, by definition, avoid if costless to do so, as is the case here.

Regardless of the fraction invested in each state, if there is no collapse, a bank's return is $1.35 (= 90 percent of $1.50) per dollar invested. And if there is a collapse—in either New York or Florida—investing half in each market ensures $1.20 (= 80 percent of $1.50 return per dollar). Investing any other proportion introduces further risk to the bank because then its mean return stays constant, but its return when either New York's or Florida's housing market collapses depends on which occurs.

Now introduce government bailouts. In particular, assume that if 25 percent or fewer mortgages fail, this is considered by government to be within the range of "nonemergency" states of the economy, and thus policymakers do nothing. But if more than 25 percent of mortgages fail, the government declares a financial crisis, triggering a bailout of all failed mortgages, possibly using lump-sum taxes on banks to fund these bailouts.

Again we ask, what will banks do? Unlike the situation without bailouts, now what makes sense for a particular bank to do depends on *what other banks do.*

First, suppose all (or almost all) banks invest half their funds in each state, as is the case without the possibility of bailouts. In this case, at most 20 percent of mortgages will fail; thus, a government bailout will never occur. Given no possibility of bailouts, any particular bank should invest 50-50 as well. Thus, all banks investing half in each state is a set of mutually reinforcing behaviors—an equilibrium.

But with a bailout policy in place, there are two other equilibria as well: one in which all banks invest only in New York and one in which all banks invest only in Florida. To see this, suppose a bank sees all (or almost all) other banks investing all their funds in New York. Does that bank profit from investing all in New York as well? If so, then all banks investing only in New York is a set of mutually reinforcing behaviors.

And it indeed does make sense for each bank to invest all in New York if all the other banks are doing so. To see this, consider what happens, scenario by scenario, to a bank that "goes along with the herd" and invests all in New York when all other banks are doing so versus a bank that doesn't go along with herd (and invests 50-50 in each state), again when all other banks invest only in New York.

If no housing collapse happens in either state, it makes no difference whether this bank goes along with the herd or not. It gets a return of $1.50 on 90 percent of its mortgages regardless of where they are.

Next, if the New York housing market collapses, 30 percent of all mortgages will fail, triggering, by assumption, a government bailout of all mortgages. Thus, in the "New York collapse" scenario, it also makes no difference whether this particular bank goes along with the herd or not. Its profits are $1.50

per dollar invested (since all failing mortgages are paid off by the government) less a bailout tax, again, regardless of which states the mortgages are in.

Finally, if a housing collapse occurs in Florida, the 50-50 strategy returns a lower amount than investing all in New York, since 20 percent of the bank's mortgages fail versus 10 percent if the bank had invested only in New York.

Thus, in two scenarios (no collapse and a New York collapse), it makes no difference whether a bank goes with the herd or not, and in the remaining scenario (a collapse in Florida), a bank is strictly better off having gone with the herd. Since banks must choose how to invest before they know which scenario occurs, it makes financial sense for each bank to invest only in New York if all other banks do so as well. (Symmetrically, there is also an equilibrium where all banks invest only in Florida.)

Note here that these two "extra" equilibria—all banks investing only in New York mortgages and all banks investing only in Florida mortgages—exist only because of the anticipation of bailouts. The anticipation of bailouts causes a financial fragility due to the coordinated behavior of small banks that would not exist otherwise.

Conclusion

In this paper, we have argued that the anticipation of bailouts creates incentives for banks to herd. This herding behavior makes bailouts more likely and crises more severe. Analyses of bailouts and moral hazard problems that focus exclusively on size are therefore misguided, in our view, and the policy conclusion that limits on bank size can effectively solve moral hazard problems is unwarranted. R

Endnotes

[1] Attributed to John Ray, English naturalist and botanist. Poet and dramatist Christopher Marlowe is also cited as a source through his use of a similar Latin phrase, Solamen miseris socios habuisse doloris. Doctor Faustus, Sc. 5.

[2] This calculation is the sum of a 90 percent chance of no collapse (and thus a 90 percent repayment rate) plus the 10 percent probability of an 80 percent repayment rate, where the 80 percent repayment rate is the result of averaging a 20 percent default rate over the total investment since both New York and Florida face 30 percent default rates if their market collapses but just 10 percent default if their state market remains healthy. That is, $1.335 = 1.5* (.9*.9+.1*.8).

References

Kareken, John H., and Neil Wallace. 1978. "Deposit Insurance and Bank Regulation: A Partial-Equilibrium Exposition." *Journal of Business* 51 (3): 413-38.

Stern, Gary H., and Ron J. Feldman. 2004. *Too Big to Fail: The Hazards of Bank Bailouts.* Washington, D.C.: Brookings Institution Press.

Raj Chetty

Raj Chetty wrote an essay in high school that questioned assumptions and conclusions in *Time on the Cross* by Robert Fogel, a Nobel laureate in economics. As a Harvard freshman, he sent it to the eminent economist Martin Feldstein, asking to be his research assistant. Impressed, and though he rarely picked freshmen, Feldstein gave Chetty the position.

Good choice. Chetty excelled in economics, graduating *summa cum laude* in three years and completing his Ph.D. in another three. He taught at Berkeley from 2003 to 2009; then he returned to Harvard as one of the youngest tenured professors in the university's history.

Chetty's research is characterized by uncommon insight, powerful analysis and a refusal to accept conventional theories at face value. His specialty is public economics, and his work "has transformed the field," observed Feldstein in honoring Chetty as the 2013 John Bates Clark medalist—at 33, one of the youngest recipients ever for the award, given to the American economist under 40 judged to have made the most significant contribution to economic thought and knowledge.

Chetty has focused primarily on social insurance and taxation, and more recently education and income mobility, but he's made important contributions in risk aversion, interest rates and corporate investment, and a variety of methodological issues.

Honors (and they are many) haven't distracted Chetty from a deep and rigorous research agenda. He picks crucial questions, collaborates generously, improves theory, uses novel methods (often with massive databases) and ultimately distills his findings clearly for very distinct audiences: fellow economists, policymakers and the general public.

In the following *Region* conversation, Chetty explores work on income mobility, education, labor supply, taxation and a range of other topics, providing solid evidence for the Clark award statement referring to him as "arguably the best applied microeconomist of his generation."

PHOTOS BY PETER TENZER

TEACHER QUALITY

Region: I'd like to begin with your research on teacher quality that found such substantial long-term impact on student outcomes. Can you tell us a bit about your research approach, your findings and perhaps your testimony in the California court case?

Chetty: Certainly. With John Friedman and Jonah Rockoff, I've looked at the long-term impacts of teachers on student achievement and students' long-term success. We studied that question by taking advantage of incredible new data sets—and that approach is basically part of the larger theme of my recent work, which brings "big data" to bear on public policy questions. In much the same way that Google and Amazon use very large data sets to improve the quality of the products they offer, we are trying to use large data sets to improve public policy decisions.

In the context of teacher quality, we were focused on one very important issue in education that's receiving a lot of attention in the current policy debate: How can we measure and improve, possibly, the quality of teachers in public schools in America? We tackled that question by getting data from one of the biggest urban school districts in the United States, on 2½ million children over a 20-year period, during which they wrote 18 million tests.

We take that data, which tells us how students did in math and English, what teachers they had, which classrooms they were assigned to and so forth, and link that to administrative records from tax returns and social security databases on students' earnings, college attendance outcomes and various other markers of success later in life. So, essentially, the type of question we are able to ask is, how did the third-grade teacher that you had affect your success 25 years later?

We're ultimately interested in evaluating the long-term impacts of teacher quality, but the first step in that analysis is to define a way of measuring teacher quality. One measure that has received a lot of attention recently is what are called "value-added measures." The basic concept of measuring a teacher's value-added is quite simple, although there are various technical issues to be worked out. The idea is to use changes in test scores as a measure of teacher quality. For instance, if you are a fourth-grade teacher, we take your students' test scores at the end of fourth grade and subtract their test scores at the beginning of fourth grade. The average change is essentially what we call the teacher's value-added.

Region: And across the entire database, that's using a *standardized* test, not each teacher's pop quizzes or exams.

Chetty: Right, using standardized tests administered at the city or state level so that everybody is measured on the same scale. There's been a very controversial debate about the use of these measures for two main reasons. First, quite naturally, people are concerned that test scores might not be a very good measure of teacher quality. Maybe some teachers are really great teachers who inspire their students to succeed in the long run, but that doesn't show up on a standardized math test.

Another important concern is that these measures may not be picking up the causal effects of teachers. Rather, they may be picking up something about which types of students a teacher is assigned. This gets to the idea that value-added estimates may be statistically biased. This is potentially quite important because if you are a teacher who is assigned students who are doing really well and gets rated as a high value-added teacher, and I get a worse draw on the students and I am rated a low value-added teacher, we might be equally good teachers, but I might end up losing my job or not getting a promotion just because of the mix of students I happen to get. So it's very important to figure out how much bias there is.

In light of these concerns, we set out to answer two questions with our data. First, how much bias is there in value-added estimates? And second, do they really pick up something on a teacher's long-term impacts, or are they just picking up who is good at teaching to the test and who's not?

In a nutshell, we basically conclude, first, that value-added measures largely capture the causal effect of teachers rather than differences in the types of students they get. That is, a child who is randomly assigned to a teacher who is high value-added rather than low value-added will end up having higher test scores at the end of the school year. Of course, this result only establishes that some teachers are able to raise test scores more effectively than others; it is not clear whether this is driven by teaching to the test or "deep learning" that has persistent benefits.

So we then move on to our second question: If you're assigned a high value-added teacher in third grade—that is, the teacher who is systematically improving test scores—and I happen to get a low value-added teacher, does that impact last? Are you, in fact, doing better many years later, or are we both doing as well as each other?

The prior literature in education would lead us to think that these impacts are *not* that long lasting. Many studies have shown that test score gains tend to "fade out" over time. What that means is that if a child is assigned to a better teacher in third grade, we see her doing better on third grade tests, but a lot of that gain shrinks by the end of fourth grade and virtually disappears by fifth or sixth grade. Based on that evidence, you might have thought, well, by the time we're looking at people's earnings years later, so many other things have happened in their lives, and we're not really going to find a meaningful effect of these teachers.

Region: That dissipation has even been found in the early childhood research, I think.

Chetty: Exactly, that's a generic pattern found in studies of early childhood interventions: the Perry Preschool study, Head Start, et cetera. And so going into this work, our prior assumption was we might find something, but more likely we might not find any lasting impact, which would also be useful to know. So we were very curious to look at the data.

Much to our surprise, it immediately became evident that students who were assigned to high value-added teachers showed *substantially* larger gains in terms of earnings, college attendance rates, significantly lower teenage birth rates; they lived in better neighborhoods as adults; they had higher levels of retirement savings. Across a broad spectrum of outcomes, there were *quite* substantial and meaningful impacts on children's long-term success, despite seeing the same fade-out pattern for test scores.

Region: No wonder this research has received so much public attention and criticism.

Chetty: Yes. The study received quite a bit of attention in the media and in the policy debate and ultimately in the legal realm, where a lot of these issues are currently being contested.

Region: As in *Vergara versus California*, where you testified. What were the issues there?

Chetty: Yes, one of the places where this played out was in a lawsuit in California called *Vergara versus California*, in which I was an expert witness for the plaintiff. That case was partly motivated by the findings in this study, but focused on a slightly different issue: on whether teachers should be granted tenure, and in what manner and how long they should be given for evaluation before tenure was granted. There was a complex set of issues at play in that legal decision beyond the particular issues surrounding value-added.

But the fact that emerges from this study—that we are able to measure teacher quality at a relatively early stage using test score data and are able to identify teachers who have long-lasting impacts on students' achievement and later outcomes—is important for the Vergara decision. It shows that teachers matter and that teachers vary in effectiveness, so implementing policies that keep better teachers in school districts might actually have a meaningful impact on students' outcomes, which was a core argument in the lawsuit.

ON TEACHER QUALITY

If you're assigned a high value-added teacher in third grade—that is, the teacher who is systematically improving test scores ... does that impact last?

Much to our surprise, it immediately became evident that students who were assigned to high value-added teachers showed *substantially* larger gains in terms of earnings, college attendance rates, significantly lower teenage birth rates; they lived in better neighborhoods as adults; they had higher levels of retirement savings.

U.S. INCOME MOBILITY

Region: Another body of your very recent work that's received a lot of attention is that on U.S. income mobility. You found both that U.S. intergenerational income mobility hasn't changed very much over the past 40 years or so, but that it does differ substantially across the U.S., evidently due to impact of several factors: residential segregation, social capital, income inequality, primary school quality and family stability. I hope that's a fair synopsis. Would you elaborate on that work?

Chetty: A lot of my current research, the education work being one example, is focused on understanding how we can improve outcomes for disadvantaged youth. In my view, the bigger-picture question here is how all these factors can contribute to intergenerational (or social) mobility. One of the core ideals, I think, of American society—and in some ways, the reason my own parents came to the U.S. like many other immigrants, in search of the American dream—is the idea that no matter what your background, you have a great chance of succeeding in America and of moving up in the income distribution relative to where you started. Our education research tries to approach that from one particular angle, as one factor that might matter.

More recently, we've been studying the level of social mobility in the United States from a broader lens. How has intergenerational mobility changed over time in America, and how does it vary

ON U.S. INCOME MOBILITY

While we've identified some potential factors that are good predictors of differences in mobility ... what we need to do to improve upward mobility in the United States is much less clear.

One of the intriguing preliminary findings from this work is that … every extra year you spend in a better environment, your own outcomes improve and converge to the outcomes of the prior residents.

across places within the U.S.? There's a popular conception that the U.S. once was a great land of opportunity and that that's no longer true today. Unfortunately, we've had relatively little data to actually be able to study the degree of social mobility systematically in the United States, so it is has been hard to know whether this conception is accurate or not.

When we actually looked at the data over the past 30 to 40 years or so—a period for which we have good information from de-identified tax returns on children's parents' income as well as their own income—we find that, much to our surprise, there isn't that much of a difference in social mobility in the United States today relative to kids who were entering the labor force in, say, the 1970s or 1980s. That is, children's odds of moving up or down in the income distribution relative to their parents have not changed a whole lot in the past few decades.

We find that where there is much more variation is across space rather than over time. So the big story is that it's not that things are changing over time necessarily, but rather that some places have, and have always had, much higher levels of social mobility than others in the United States.

To take one example, let's focus on a simple statistic: the odds of moving from the bottom fifth of the U.S. income distribution to the top fifth, so kind of a Horatio Alger story of leaping from the bottom to the top. In the U.S. as a whole, your odds of moving from the bottom fifth to the top fifth are 7½ percent. That compares with about 11 percent in Denmark and 13 percent in Canada.

All of those numbers might seem pretty small at first glance, but you have to remember that you can't have more than 20 percent of people in the top 20 percent. And so the fact that Canada is at 13 percent means that Canada actually has quite a high level of social mobility relative to the 7½ percent in the U.S. It says that a child's odds of achieving the "American dream," in some sense, are twice as high if she is growing up in Canada rather than in the U.S.

Those cross-country comparisons draw a lot of interest, but they are difficult to interpret because there are many differences across countries, starting from the fact that the income distribution is much more compressed in Canada and Denmark than in the U.S. (making it easier to climb from the bottom fifth to the top fifth there than in the U.S.). What's more striking and informative, in my view, is that there is actually even more variation in your odds of moving from the bottom to the top, *within* the United States than among countries.

For example, for children growing up in places like Salt Lake City, Utah, or San Jose, California, the odds of moving from the bottom fifth of the national income distribution to the top fifth are more than 12 percent or even 14 percent in some cases, *more* than virtually any other developed country for which we have data.

In contrast, in cities like Charlotte, North Carolina, Atlanta, Georgia, or Indianapolis, Indiana, a child's odds of moving from the bottom fifth to the top fifth are less than 5 percent—less than *any* developed country for which we currently have data.

Within the United States, there's this incredible spectrum of variation in social mobility, which means that we shouldn't really think of social mobility purely at the national level. Is the U.S. the land of opportunity or not? That question doesn't really have a clear answer. Rather, we need to think about it at a much more local level and try to understand why some places have much more mobility than others and what we can do about it

.

Region: You looked at many possible correlates with levels of mobility and identified those five that I mentioned earlier. What really are the implications of *correlates*, per se, in this instance? And where are you and your colleagues headed with this research now? What are the key questions that still need to be addressed?

Chetty: Right. We've identified a set of correlates which you mentioned, things like school quality, social capital and so forth, that are correlated with differences in mobility across areas. But I want to stress that that does not mean that those are the *causal* determinants of differences in mobility.

To take one example, consider social capital, the idea that the social cohesiveness in a community matters. That idea was popularized by my colleague, Bob Putnam, who wrote a famous book called *Bowling Alone*. The title comes from one of the ways in which Putnam measures social capital: the number of bowling alleys in an area.

I was actually amazed to find in our own data that the number of bowling alleys is strongly correlated with differences in upward mobility across areas. But that, I think, nicely highlights the point that these are correlations rather than causal effects, because I'm pretty sure the policy lesson here is not that we should be building more bowling alleys to try to improve social mobility in the U.S.

The point is that while we've identified some potential factors that are good predictors of differences in mobility, what that means in terms of what we need to do to improve upward mobility in the United States is much less clear.

That's exactly where I think this research needs to go and where my colleagues and I are now working. One set of studies currently under way is looking at families that move across areas. We're studying 20 million families that moved with their kids between metro areas of the United States. We ask if you move, say, as a 5-year-old, from Atlanta to Salt Lake City, do your outcomes improve? Do you look more like the kids who grew up in Salt Lake City and did really well? And secondly, how does that play out, depending upon *when* you moved? If you moved when you were 10 years old or 15 years old, rather than as a 5-year-old, do you get less of the benefit?

One of the intriguing preliminary findings from this work is that there's a

linear "exposure effect." Every extra year you spend in a better environment, your own outcomes improve and converge to the outcomes of the prior residents. This type of evidence strongly suggests that the differences in upward mobility across places are actually a causal effect of growing up in, say, Salt Lake City rather than Atlanta, as opposed to just differences in the types of people who live in Salt Lake City versus Atlanta.

It's *that* type of work that we think will help us move toward characterizing the causal effects of each place and, ultimately, toward understanding what one might actually be able to change in a city like Atlanta to improve mobility.

Region: Have you looked at how this might relate to the phenomenon of brain drain? You used "commuting zones" as your geographic unit in this research, and I wonder if you'd find that within some of those zones there's a lot of brain drain in the sense of the "best and brightest" moving from small towns to larger cities to find better-paying jobs. In Minnesota, for example, St. Cloud has high income mobility, and it's common for St. Cloud natives to migrate to the Twin Cities. Do you think this might play a role in intergenerational mobility patterns? (See the October 2014 *fedgazette* online at minneapolisfed.org.)

Chetty: Brain drain does appear to be an important factor, particularly in rural areas. One of the striking patterns in the data is that some rural areas exhibit very high levels of upward mobility—namely, the Great Plains, places like Iowa and rural Minnesota and so forth.

What's particularly remarkable about these places is that they suffer from, as you put it, a brain drain effect, where the talented kids who are doing really well end up leaving those areas and moving to Minneapolis, Chicago or New York, where they're earning high incomes and they're very successful. But what that means is the talent pool of the people who are left in that area is reduced, yet it seems like cohort after cohort, these places continue to produce very good outcomes.

That again suggests to me that it's something about the institutions, the structure of those places, that's leading to these excellent outcomes, and not merely the types of people who live there.

SALIENCE OF TAXES

Region: Let's shift entirely to some research that has not gotten the same level of public attention as the two we've discussed—though I believe it's your most cited article in professional literature—on salience of taxes. You've showed that people's awareness of sales tax significantly affects their purchasing behavior; you used two empirical experiments or data sets, one on beer and the other about grocery store sales. Can you tell us a bit about both findings and what implications they have for public finance?

Chetty: Sure. One of my long-standing interests has been in bringing behavioral economics—insights from psychology and economics—to bear on public policy issues. One of the central assumptions when I was in graduate school that we made in all of our courses and most of the papers we read is that people were fully aware of and optimized perfectly with respect to very complicated tax schedules, welfare policies and so forth. From introspection and from talking with friends and family, my instinct was that that doesn't really seem like an accurate description of behavior and attention for most people outside the economics department.

In order to test the assumption that people optimize perfectly with respect to taxes and quantify how large any deviations from it might be, we set about doing some *very* simple experiments. We first worked with a grocery store, where we decided to test the idea that people are paying attention to sales taxes when they purchase products. For example, if you go to a grocery store and you buy something like a hair brush (a product that's

ON SALIENCE OF TAXES

This finding suggests that the *salient* tax, the one that's included in the price that people focus on, has a much larger effect than the nonsalient tax that people may not have in mind.

These results show that even for simple sales taxes, people don't really seem to be paying attention. This suggests that for much more complicated taxes, such as income taxes ... these issues are likely to be all the more severe.

not edible and hence is generally subject to sales tax), the price you see quoted on the shelf in the United States doesn't include that tax of, say, 5 to 9 percent that you might pay at the register when you actually pay your bill. How does that affect purchasing behavior?

The classic economic model assumes that everyone is taking into account the sales tax when deciding what to buy. To test that assumption, we took 1,000 products in the grocery store, such as cosmetics and hair care accessories and so forth. You might ask, why pick that unusual set of products? The reason is that when we approached the grocery store manager and said, "We want to do an experiment that we think will show that people aren't taking taxes into account," he said, "I absolutely believe your hypothesis that this is going to reduce sales and there's no way I'm going to let you do it with *all* the taxable products in the store." So that's why we concentrated on this particular subset of products.

Our experimental intervention was to post the *tax-inclusive* price of the good. So if, for example, a lipstick was selling for $5.99, we added a tag saying $5.99 *plus* California Sales Tax = $6.73. The standard economic model would say that this intervention should have had no impact. Under the standard model, consumers already knew they were paying this tax, and we were not giving them any new information.

In contrast with this prediction, when we looked at the data, we found that there were clear reductions in sales of the products we tagged with the tax information relative to other products and other stores where we had not done this intervention. That's one piece of evidence suggesting that providing salient information on tax rates does, in fact, seem to affect the behavior.

A concern with the experimental strategy is that we introduced something artificial in the environment. Any experiment like this has the concern that we might be seeing an effect not because of the information we provided, but because we're doing something really unusual. Consumers come into this aisle of the store and see a thousand tags that they've never seen before. Maybe they just think, "This is kind of confusing, I don't understand what's going on and I'm not going to buy any of these products this week." In this case, we are going to see a reduction in demand, but it might not be because of the tax information effect that we're really after.

To complement that experiment in the grocery store, we conducted a second study that doesn't suffer from that problem. We used existing data, and we basically compared the effect of *tax* changes to the effect of *price* changes. We do this is by focusing on one particular good—beer—because alcohol has a very useful property: It's subject to two taxes. One is an excise tax, and it's included in the price that you see on the shelf or on a restaurant menu. And then there's the sales tax, which is added later at the register.

The standard economic theory tells us that raising the sales tax or the excise tax by an equivalent amount should have the same effect on alcohol consumption because it doesn't matter to the consumer—either way you're paying the same amount.

We test that prediction by using changes that states have implemented over time in excise and sales tax rates, and we look at how such changes affect beer sales. We find that the tax that's *included* in the price, the excise tax, has much larger effects, five or 10 times as big, as a sales tax change of an equivalent amount. Once again, this finding suggests that the *salient* tax, the one that's included in the price that people focus on, has a much larger effect than the nonsalient tax that people may not have in mind.

These results show that even for simple sales taxes, people don't really seem to be paying attention. This suggests that for much more complicated taxes, such as income taxes, capital gains taxes, estate taxes—things that people might not even fully understand—these issues are likely to be all the more severe.

The policy implications of this result, which have since been developed in many subsequent papers by various others, are basically that when we think about designing income tax systems, we shouldn't make the assumption that everybody's paying attention to every provision of the income tax code. We need to take into account the fact that lots of people might completely ignore some of the incentives. This has important implications in a variety of domains.

For instance, trying to increase the amount that people save for retirement is a common policy goal. We spend about $100 billion a year in the U.S. to encourage saving by providing subsidies for retirement savings accounts. If people don't understand those incentives and don't pay attention to them, we basically are spending money without any bang for our buck. If we are trying to change behavior, it is important to use salient incentives that people actually see and understand rather than just focusing on the dollars and cents. The way in which policies are framed might be just as important as the amount we're spending on them.

LABOR SUPPLY AND UNEMPLOYMENT INSURANCE

Region: I'll come back to your research on retirement subsidies, but first I'd like to discuss earlier work on labor supply and unemployment insurance, published in a 2008 *Journal of Political Economy* article. You looked at the extent to which the inverse relationship between labor supplied by workers and the level of unemployment insurance is due to so-called moral hazard—that is, the disincentive to work effort created by insurance. You found that liquidity constraints also play a major role. Would you tell us more about that research?

Chetty: A well-established fact in the literature on unemployment insurance is that when you provide people higher levels of unemployment benefits, they take more time out of work and you drive up

ON LABOR SUPPLY AND UNEMPLOYMENT INSURANCE

I found that the effects of unemployment benefits on unemployment durations were *much* larger for liquidity-constrained individuals than non-liquidity-constrained individuals, suggesting that liquidity effects are quite important.

Having a relatively generous unemployment benefit system, somewhere along the lines of what we have in the U.S. today, might actually be desirable. The moral hazard costs are not as large as economists previously thought.

unemployment rates. Traditionally, that's thought to be driven by a moral hazard effect, as you say, meaning that when I get higher unemployment benefits, my effective wage from returning to work is lower because I lose those higher levels of benefits when I find a new job. So, my incentive to find a job is essentially reduced, creating so-called moral hazard.

In this study, I thought about a different effect that might drive the relationship between unemployment benefits and labor supply, which I call a liquidity effect. It's just the idea that if you have more cash on hand while you're unemployed, you can take longer to find the job that suits your skills best, for example. Or from a different perspective, if you have a very low level of unemployment benefits and very little cash in your savings account, you might need to take the first job you can get in order to put food on the table and feed the family.

This effect suggests that we might see a relationship between the level of benefits and how quickly people find jobs *not* because people are thinking that the incentive to find a job has changed, a price effect, but rather because they cannot "afford" to search for the right job.

To evaluate the relative importance of liquidity effects vs. moral hazard effects, I looked at variation across people in terms of the amount of money they have in the bank when they lose their job. Some people happen to lose their job at times when they have essentially no assets so they're really liquidity-constrained. Other people have a few thousand dollars of savings when they lose their job and so they're not quite as pressed to find a job immediately.

If the relationship between unemployment benefits and unemployment rates is driven *purely* by moral hazard, we would expect to see that moral hazard effect, both for the people who have significant assets in their bank accounts and for the people who don't, because everyone's incentives are being distorted by insurance. If, in contrast, the liquidity effects are very important, we would expect to see unemployment benefits having a bigger effect on the liquidity-constrained individuals relative to the people who have assets when they lost their job.

In the data, I found that the effects of unemployment benefits on unemployment durations were *much* larger for liquidity-constrained individuals than non-liquidity-constrained individuals, suggesting that liquidity effects are quite important. Based on this and related analysis, I end up concluding that something like two-thirds of the relationship between unemployment benefits and unemployment rates, is actually due to a liquidity effect, rather than a distortionary moral hazard effect.

That result has implications for how you want to set the level of unemployment benefits. If the moral hazard effects are extremely large, then we are hurting the economy by having a high level of unemployment benefits and one would want to scale them back. If the liquidity effects are important, then we're providing a benefit while people are out of work and so having benefits is actually useful. I end up concluding that having a relatively generous unemployment benefit system, somewhere along the lines of what we have in the U.S. today, might actually be desirable. The moral hazard costs are not as large as economists previously thought.

LABOR SUPPLY ELASTICITY

Region: Let me ask another question related to labor supply, about measurement of elasticity, or responsiveness of workers to changes in wage rates. There's been a long-standing dispute, in a sense, between micro- and macroeconomists over the actual level of elasticity. Some macroeconomic models of business cycle fluctuations depend on an elasticity level that microeconomic evidence can't support.

You've made some progress toward reconciliation here, but still it seems there's a large gap when it comes to estimates of intertemporal or Frisch elas-

ON LABOR SUPPLY ELASTICITY

What creates this big difference between micro and macro estimates of elasticities?

One important factor we think matters is that *micro* estimates of elasticities often are based on short-run changes in policies.

Macro evidence in standard models still requires *much* larger elasticities than micro evidence suggests.

A potential resolution is so-called labor wedge models or search-theoretic models of the labor market ... which say that something is wrong in the market, in a sense, that's making it difficult for people to find jobs.

ticity on the *extensive* margin, people looking for jobs (as opposed to workers varying hours of labor supplied, the intensive margin). Can you elaborate on your findings? You seem to find some support for the idea of labor wedges, which I think some Minneapolis Fed economists were glad to hear. Can you explain that a bit too?

Chetty: The concept of labor supply elasticity is fundamental in many parts of economics. It's relevant in macroeconomics for understanding business cycles. The idea here is that if wage rates are higher during booms relative to recessions, people might have less of an incentive to work in recessions relative to booms, which would affect the number of people participating in the labor force. That would create fluctuations in unemployment rates and labor force participation rates if these elasticities are large.

Macroeconomists have traditionally thought that these elasticities have to be quite large in order to match the patterns we see over the business cycle in standard business cycle theories—many of which were pioneered, in fact, in Minneapolis.

The problem is that if you go to the microeconomic level and try to directly estimate this elasticity, you find much smaller estimates. There have been hundreds of studies over the past few decades that essentially ask, "If I change a person's wage rate by changing tax rates or their wage rates, how much do they actually change the amount they work?" The uniform finding of those studies is you get quite small elasticities, around 0.1 or 0.2. That means that a 10 percent increase in the wage changes the amount that people work by something like 2 percent, *far* too small to explain macroeconomic fluctuations in standard models of the business cycle that do not allow for what are called "labor wedges," market imperfections that make the economy deviate from equilibrium.

In some recent work I've done with various coauthors, we've tried to understand what creates this big difference between micro and macro estimates of elasticities. And we've made some progress, although as you correctly noted, we haven't fully explained the gap in understanding the difference.

One important factor that we think matters is that *micro* estimates of elasticities often are based on short-run changes in policies. For instance, I might change tax rates by 10 percent next year. Is that going to affect the amount that you work substantially? Well, it might not because you might have to find a different job or go negotiate with your employer for a different pay package in order to do so. All of that might take quite a bit of time. That is, there are a lot of adjustment costs involved. Micro estimates of elasticities may get attenuated by such adjustment costs. This leads us to think that the estimates might actually be a little bit bigger than suggested by the micro evidence.

But on the other side, we think that some macro estimates, which suggest elasticities well above 1, are likely to be overstated because they're driven by other omitted factors that are varying at the same time as the wage rate. For instance, consider differences across countries. Countries with higher tax rates have lower labor supply, not just because of the direct effect of the tax, but because there are many other things that are different across these countries. There are different labor structures, they have different social welfare programs and so forth.

Region: And some of the macro estimates are considerably higher than 1. Ed Prescott's research, for example, suggests a Frisch elasticity of about 3.

Chetty: Yes, I think that microeconomic evidence strongly suggests that it's very hard to believe that the actual elasticities are as large as 3, or even above 1. The elasticities are likely more like 0.5, and may be much smaller at business cycle frequencies because workers tend not to adjust their behavior rapidly in the short run, as I mentioned earlier.

That still leaves you with a substantial gap between what we're finding in the micro data and in the macro data. The gap arises in particular on the extensive margin, that is, the fraction of people who choose to work when wage rates change. There's a pretty good alignment between micro and macro evidence in terms of hours people work, conditional on working.

Region: So, there *is* general agreement on elasticity at the intensive margin, but not on the extensive.

Chetty: Exactly. Alignment on the intensive margin, but in terms of how many people choose to work—the extensive margin—macro evidence in standard models still requires *much* larger elasticities than micro evidence suggests.

A potential resolution is so-called labor wedge models or search-theoretic models of the labor market, which do *not* make the assumption that we're at a market-clearing equilibrium in a recession. They're basically disequilibrium models, which say that something is wrong in the market, in a sense, that's making it difficult for people to find jobs. Stated differently, it's not purely a choice of workers not to work when there's a recession. It fits with the intuitive idea that there's involuntary unemployment: Lots of people are looking for jobs and want to work, but something in the economy is not working right, and there is potentially room for policy intervention.

JOB MARKETS

Region: That leads us to the next question I wanted to ask, about the health of current job markets. The very slow rebound of labor markets from the Great Recession has weighed heavily on workers in both the United States and Europe. Policymakers have had little success in efforts to address that stagnation. Given the work that you've done on social insurance programs, labor supply elastic-

ON JOB MARKETS

Social insurance programs and tax incentives are *not* the primary reason that we are not seeing a recovery. ... Other factors, like a shortfall in aggregate demand and the financial crisis, are more important in delaying a full recovery.

It's also very important to keep sight of the important long-term factors that are affecting the U.S. economy … related to education and the skill level of the U.S. economy in an increasingly globalized environment.

ity and related areas, what's your general sense of the major factors that are involved in the slow recovery?

Chetty: It's, of course, difficult to gauge exactly what's driving the slow recovery at the macro level, but I think there are a host of complex factors involved. To begin, as we know from work by my colleagues Carmen Reinhart and Ken Rogoff, recovering from any financial crisis historically has been a slow process, and it takes quite a bit of time for financial markets to recover and for the economy to start functioning more normally. From this historical perspective, what we are seeing is not all that abnormal given what happened.

Importantly, I think factors like social insurance programs and tax incentives are *not* the primary reason that we are not seeing a recovery. My best sense of the evidence, looking at the impacts of unemployment benefits and the numerous studies that have been done on how they affect workers' job search behavior, it's not because we have—or *had*—fairly generous unemployment benefits that the recession has been prolonged. Other factors, like a shortfall in aggregate demand and the financial crisis, are more important in delaying a full recovery.

I think it's also very important to keep sight of the important long-term factors that are affecting the U.S. economy. This comes back to some of the issues we started out with related to education and the skill level of the U.S. economy in an increasingly globalized environment. Are U.S. workers adequately trained to be able to get jobs, not just coming out of this recession, but going forward more broadly? Are we going to see wage growth in the middle of the income distribution, which has really stagnated over the past few decades? In the U.S., I think a lot of the answers there have to do with long-term human capital investment and not just changes in incentives to work, which is what people have focused on, to some extent.

RETIREMENT SAVINGS PROGRAMS

Region: Earlier you touched on the work that you've done on retirement savings programs, and I'd like to come back to that now. You mentioned that in the United States, the government spends a great deal to encourage workers to save money through their retirement savings programs, but that it's possible these subsidies may not be very effective and that some people might simply shift their savings from a taxable to a nontaxable account without increasing total savings—the so-called crowding out effect.

When we spoke with Richard Thaler last year (see the September 2013 *Region* online at minneapolisfed.org), he mentioned your work on this question and that you'd used Danish data sets to better understand it. Would you tell us about that study?

Chetty: Sure. What are some ways to encourage workers to save for retirement? As I noted above, one approach is to subsidize retirement savings, which we do in the U.S., and many other developed countries do, in the form of tax-subsidized retirement savings accounts—IRAs or 401(k)s. Basically, the idea is to make it cheaper to save for retirement and thereby try to encourage people to save more. That's the way economists have traditionally thought about such problems: If you want to encourage more of an activity, reduce the price of that behavior.

A completely different approach, which is motivated by behavioral economics and pioneered by people like Dick Thaler, David Laibson, Brigette Madrian and many others, is to exploit the fact that people don't actually seem to pay much attention to things like saving for retirement. The idea is to use defaults or automatic enrollment to encourage people to save more for retirement. The way this might work is your employer might say, "We're going to have an 'opt-out' rather than 'opt-in' retirement sav-

ON RETIREMENT SAVINGS PROGRAMS

What are some ways to encourage workers to save for retirement? One approach is to subsidize retirement savings, in the form of tax-subsidized retirement savings accounts.

A completely different approach ... is to use defaults or automatic enrollment to encourage people to save more for retirement.

[The defaults] work on the 85 percent of the population who are *not* paying attention to the tax subsidies, the passive savers. So the default, in my view, is better than the tax subsidy for the goal of trying to raise savings.

ings program. When you come to work at our firm, the default option is going to be that we take 3 percent of your paycheck and put it in your 401(k) account. You can opt out of that if you want."

Well, what we find empirically is 80 percent of people do not opt out of that. They just go along with the default.

Region: So that's the nudge, as Dick Thaler would call it.

Chetty: Right, that's the idea of using a behavioral nudge to try to influence behavior. Now, the big question in this area has been, to what extent are we just shifting the account in which people save? To what extent are these increases in savings in some accounts crowded out by reduced savings in other accounts, and to what extent are we actually raising total savings by implementing defaults in 401(k) accounts?

For example, when I default you into saving more in your 401(k) account, you might say, "OK, I don't need to worry about saving as much, so I'm going to just run down my bank account balance a little bit more." And that leaves us (as policymakers trying to encourage retirement savings) in exactly the same place overall in terms of total savings, and we wouldn't have really accomplished anything. The same issue could arise with tax incentives.

To tackle this question, in a recent study, we used excellent data from Denmark where we have information on the portfolios of the entire Danish population. We look at a series of changes in Danish policy that changed tax incentives for retirement savings. We also look at changes in defaults that firms implement for their workers and see how they affect workers' savings, both in the retirement accounts and in other accounts.

To summarize the findings, suppose we cap the 401(k) tax subsidy at a maximum rate of, say, 25 percent, the type of policy reform currently being discussed in the U.S. What effect would this reduction in the 401(k) subsidy have? It turns out that 85 percent of people, whom I'm going to call "passive" savers, totally ignore this reform and don't respond at all. Fifteen percent of people sharply reduce the amount they save in the 401(k) when the subsidy is reduced. But, critically, they take the money that they were saving there and just shift 95 percent of it to another account, thus leaving total savings almost unchanged.

In light of this evidence, we think that 401(k) and IRA subsidies are just inducing a small number of active, tax-savvy savers to shift the money they would have saved elsewhere into tax-preferred retirement savings accounts.

Region: About 15 percent of people do this.

Chetty: Yes, 15 percent of the population. In contrast, if you look at the defaults, they work on the 85 percent of the population who are *not* paying attention to the tax subsidies, the passive savers. What's interesting about the default is, not only does it make you save more in the retirement account; it actually looks like people are not saving less in any other account. We can default people to save more in their employer pension and that just leads to roughly a one-for-one increase in total savings. There's no evidence of crowd-out in other accounts. So the default, in my view, is better than the tax subsidy for the goal of trying to raise savings.

And that's for three reasons. *First*, the default doesn't cost tax revenue. That is, we don't actually have to spend tax revenue to implement the default, unlike the 401(k) subsidy. *Second*, the tax subsidy induces some people to save more in retirement accounts, but most of that just comes from shifting; whereas, the default is actually inducing new saving. *Third*, if you think about whose savings you most want to increase, it's the passive savers who are not paying attention to retirement and are going to end up retiring without having enough assets to sustain their retirement. The active savers who are financially savvy and paying attention to these tax incentives already have these retirement savings portfolios.

In the end, this really strikes me as a case where the insights from behavioral economics—the types of issues one thinks about once one allows that people may not always be optimizing perfectly—really point in quite a different direction in terms of policy.

ON A RETURN TO INDIA

I hope that the type of work that I'm doing here on education, on human behavior, on tax policy and so on has implications just as much for countries like India as it does the United States.

RETURN TO INDIA?

Region: Our earlier discussion of brain drain and your parents moving here when you were about 9 from Delhi made me realize that you and your family are an example of the classic country-to-country brain drain.

Chetty: Yes.

Region: Do you see yourself returning at some point to a research or a policymaking role back in India, similar to what your father did as an economic adviser? Do you ever consider working there?

Chetty: Yes, certainly. I think that some of the most important challenges that the world faces are in developing countries like India and in Africa where the problems are very important, and I've certainly thought about working on

those issues and also possibly trying to play a more direct role in that context.

I hope that the type of work that I'm doing here on education, on human behavior, on tax policy and so on has implications just as much for countries like India as it does the United States. One of the advantages of working in the developed country context is that this is where we have great data and are able to make progress in obtaining empirical insights that will hopefully apply more generally.

At the moment, I find myself most excited about staying focused on doing research, as opposed to becoming directly involved in the implementation of policy, which I think involves many complexities beyond the pure research findings, naturally. I'm happy that a lot of our research is playing an active role in policy debates, that people are citing it and making use of it to make more informed and hopefully better decisions. I am hopeful that our research group will be able to continue to produce research findings that are highly relevant to the policy debate in the coming years.

—Douglas Clement
Sept. 24, 2014

See video excerpts from this interview at minneapolisfed.org

More About Raj Chetty

Current Positions

William Henry Bloomberg Professor of Economics, Harvard University, since July 2013; Affiliate, Department of Statistics, since January 2013; Director, Lab for Economic Applications and Policy, since 2009

Co-director, Public Economics Program, National Bureau of Economic Research, since 2008

Previous Positions

Professor, Department of Economics, Harvard University, 2009-13

Professor, Department of Economics, University of California, Berkeley, 2008-09; Associate Professor with tenure, 2007-08; Assistant Professor, 2003-07

National Fellow, Hoover Institution, Stanford University, 2007-08

Professional Affiliations

Member, Congressional Budget Office Panel of Economic Advisers, since 2011

Editor, *Journal of Public Economics*, since 2009; Co-editor, 2007-08

Faculty Research Fellow and Research Associate, National Bureau of Economic Research, since 2003

Co-director, State Capabilities Group, International Growth Centre, 2008-09

Member, Board of Editors, *Journal of Economic Literature*, 2007-10

Honors and Awards

Richard T. Ely Lecturer, American Economic Association, 2015

Fellow, American Academy of Arts and Sciences, 2014

Calvó-Armengol International Prize in Economics, 2013

John Bates Clark Medal, American Economic Association, 2013

Fellow, Econometric Society, 2012

MacArthur Foundation Fellowship, 2012

Economic Policy Best Paper Prize, *American Economic Journal*, 2012

Mahalanobis Memorial Medal, Indian Econometric Society, 2012

Young Labor Economist Award, IZA, 2010

Distinguished Research Affiliate Award, CESIfo, 2008

Alfred P. Sloan Research Fellowship, 2008

American Young Economist Award, 2008

Publications

Dozens of research papers with a focus on effective government policies, including extensive work on tax policy, unemployment and education

Education

Harvard University, Ph.D., economics, 2003

Harvard College, B.A., *summa cum laude* in economics, 2000

University School of Milwaukee, valedictorian, 1997

For further background, visit
rajchetty.com/

VIRTUAL FED

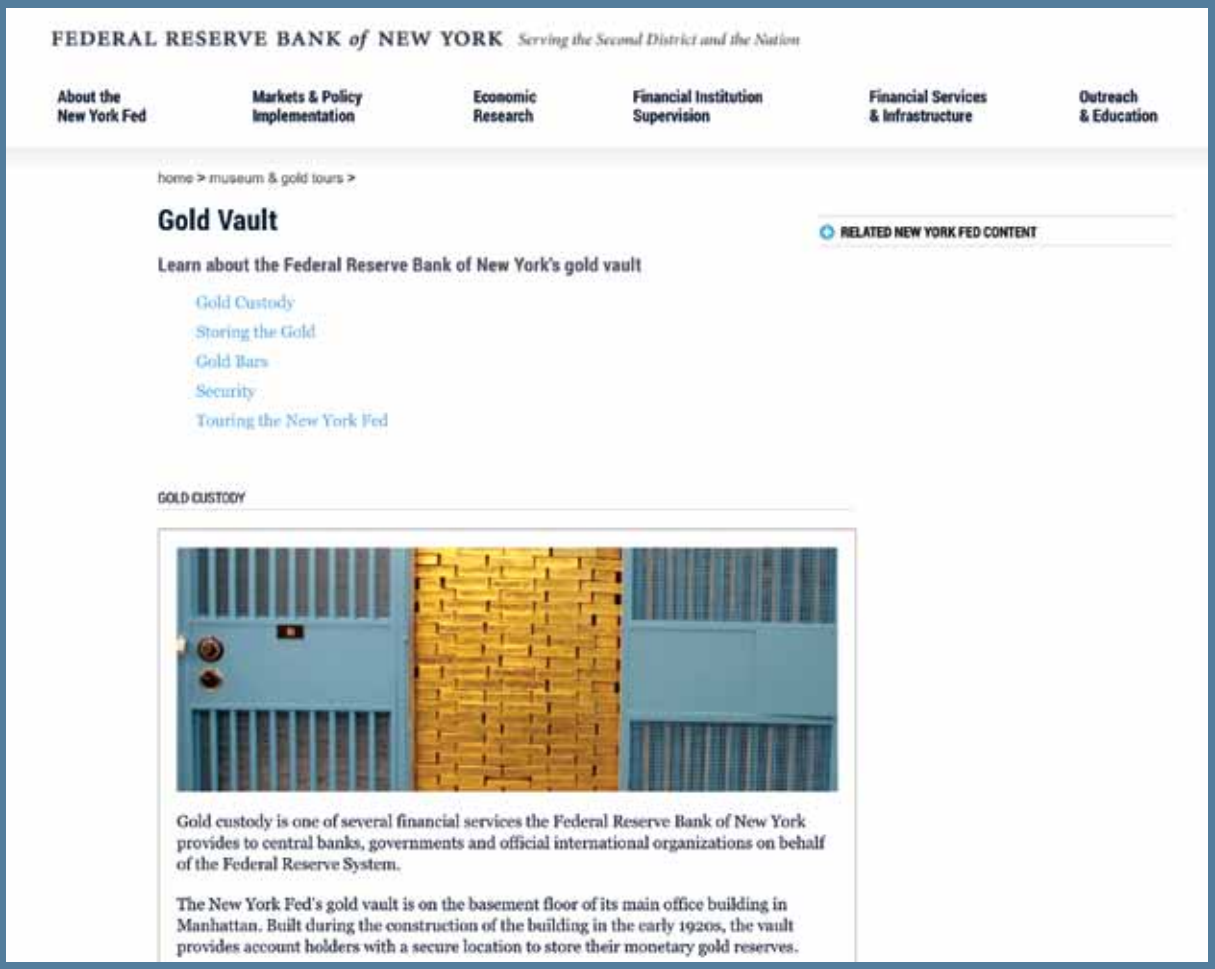

Guardian and custodian

While references to the Federal Reserve might conjure visions of an enormous vault full of gold coins and bars, the Fed doesn't actually own gold bullion. But resting on Manhattan bedrock 80 feet below Wall Street, the Federal Reserve Bank of New York does indeed have a huge vault holding about 530,000 bars of bullion. The gold belongs to other central banks, foreign governments, the U.S. Treasury and official international bodies. The New York Fed maintains it as a financial service to those institutions, keeping the gold secure and carefully accounting for every ounce. The Fed even transfers bars physically among the vault's 122 compartments when one owner transfers holdings to others.

Despite high security, the New York Fed does offer tours of the gold vault: newyorkfed.org/aboutthefed/visiting.html. A virtual tour is also available—if slightly less impressive—with photos of the vault and background information. Look, but don't touch, at newyorkfed.org/aboutthefed/goldvault.html.

—Joe Mahon

Gaming the (School) System

In a widely used school choice mechanism, parents avoid picking their favorite schools, opting instead for the low-risk choice of neighborhood schools

Phil Davies
Senior Writer

Every spring, parents face the trial of choosing the school that their child will attend in the fall—the K-12 institution that they hope will provide a nurturing learning environment and pave the way for success in college and the working world.

Many public school systems allow families to take their pick of schools in the broad community, rather than assigning students to attend a school in their immediate neighborhood. The goal is to match children to schools that best meet their academic and social needs.

Because some schools are in higher demand than others, school districts have devised ways to distribute seats among students—school choice, or student assignment mechanisms. As in any method of allocating scarce resources, there are winners and losers: Some children are assigned to their preferred school, others are relegated to a school ranked lower on their choice list and some are rejected by all their desired schools.

The widespread adoption of student assignment mechanisms in the United States and other countries has raised concerns about their efficiency and fairness. Economists who study these mechanisms observe that they're imperfect instruments; some designs are prone to manipulation and bias. To improve their chances of getting their children into an acceptable school, many parents hide their true school preferences—subverting the intent of school choice. And unsophisticated players—those with little education and low income, for example—may be at a disadvantage in the school choice game.

Much economic research over the past decade on what's known as the school choice problem has relied on theory or laboratory experiments to try to understand the strategies of households and the outcomes of different, commonly used student assignment mechanisms. But these approaches can go only so far in revealing the flaws of current school choice mechanisms and suggesting improvements.

Recent research by Maia Güell, an economics professor at the University of Edinburgh, and Caterina Calsamiglia, an associate professor at the Autonomous University of Barcelona, breaks new ground by using real-world data to analyze the issue of school choice.

In "The Illusion of School Choice: Empirical Evidence from Barcelona" (Minneapolis Fed Working Paper 712, online at minneapolisfed.org), Güell—who in July finished a one-year residency at the Minneapolis Fed—and Calsamiglia exploit a change in that city's school administrative districts to gain fresh insight into student assignment mechanisms and their ramifications.

The results of their analysis upend the presumption that most parents act on their preferences in picking schools. In the widely used school choice mechanism they study, school assignment is large-

ILLUSTRATION BY C.F. PAYNE

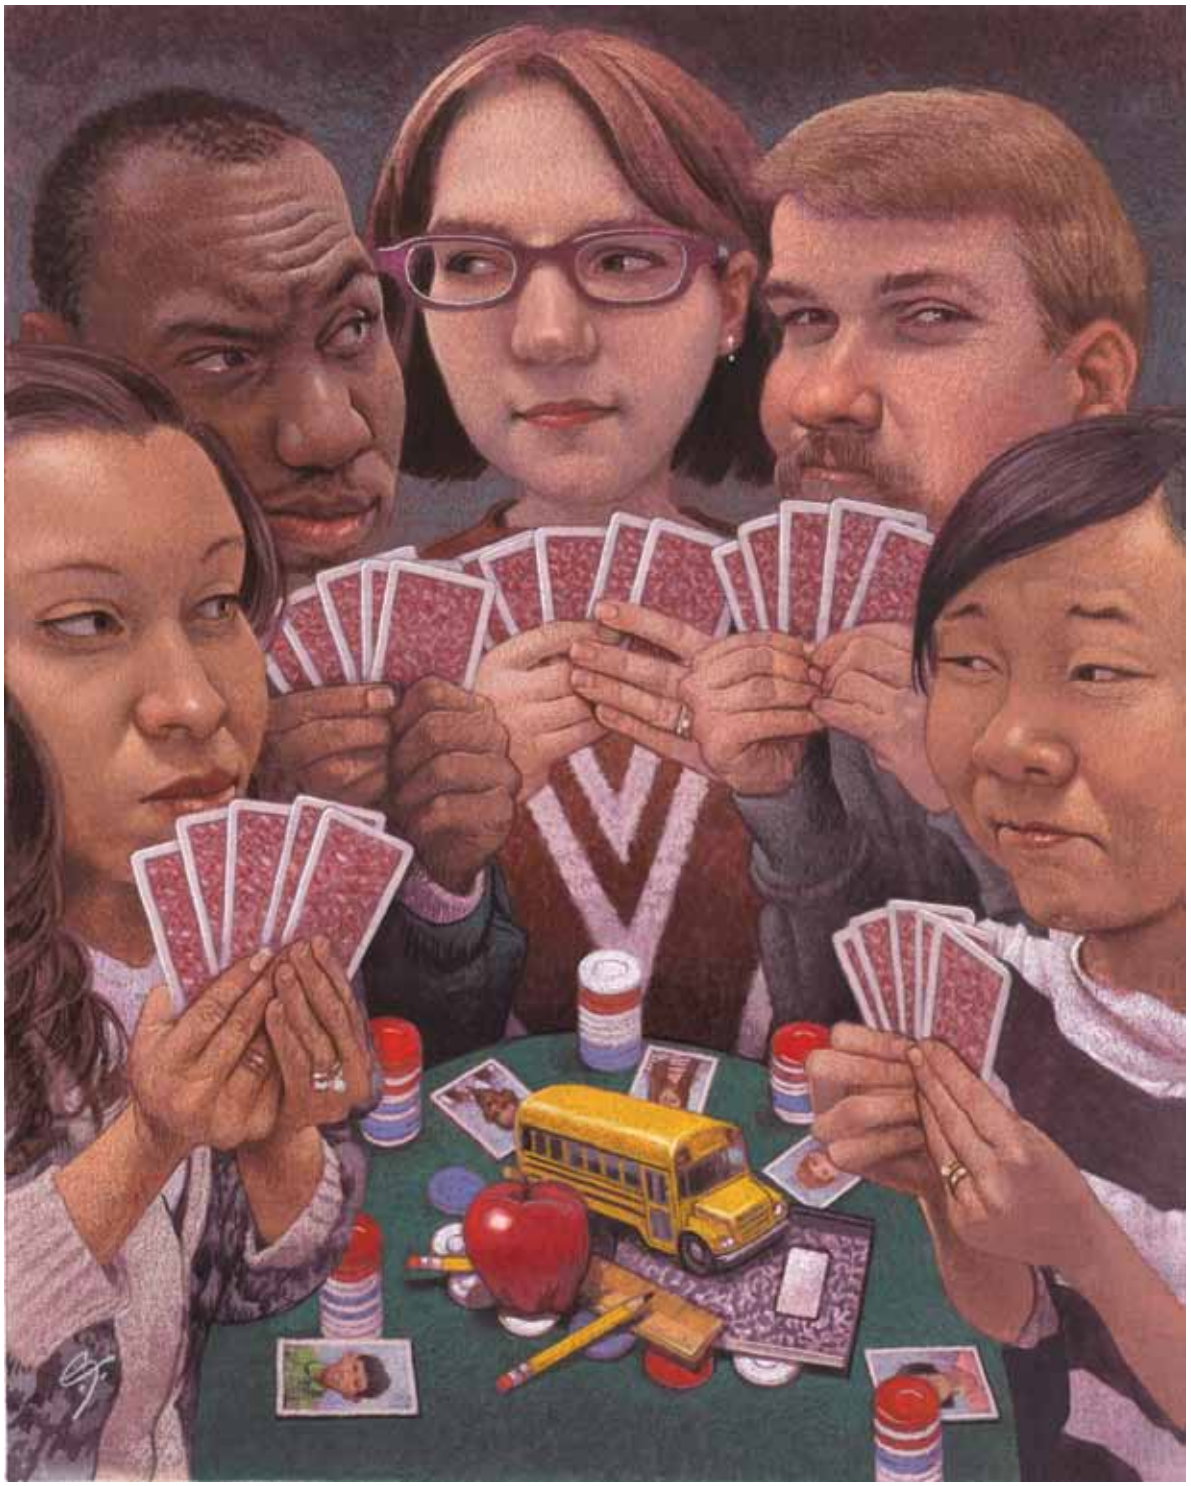

ly determined by school district rules that induce parents to avoid the perceived risk of not enrolling their children in a school that is even minimally acceptable to them. To improve their odds of being allocated to an acceptable school, families systematically apply for schools in their neighborhood because the system's rules give them priority in the local school. They make a "choice" as if they had no choice.

Just as college applicants include one school they're confident will accept them, along with applying to less likely but more desirable colleges, parents of elementary school children go for a "safe" option. "The fact that you've implemented a choice system doesn't mean that people are going to actually choose," Güell said in an interview.

This implies that in communities with such a system, the benefits of offering school choice are limited, because not all families exercise genuine choice.

The economists also discover a new, subtle form of inequity in the student assignment method used in Barcelona and many other cities: Not only does the system harm less-educated families unversed in the rules of the school choice game; it also benefits some better-educated—and likely richer—parents because they can take greater risks, thereby gaining increased access to the best public schools.

Take your choice

Traditionally, public schools have assigned children to neighborhood schools close to where they live. But over the past 30 years, many school districts in the United States and other developed countries have adopted school choice systems that expand access beyond the neighborhood to other schools in the community or even outside it. (In 1987, Minnesota became the first U.S. state to authorize interdistrict school choice.)

By achieving a better fit between pupil and school in learning goals and teaching methods, for example, school choice is thought to contribute to academic success. Proponents say it also raises educational standards by fostering competition among schools. Educators consider school choice particularly important for low-income households in areas with subpar schools. "There are inequalities in our society," Güell said, "so we want the kid who was born in a poor neighborhood to be able to go to a school outside that neighborhood."

Capacity in each school is limited, requiring school districts to devise a means of fairly allocating seats in oversubscribed schools. In the typical student assignment mechanism, families submit a list of schools ranked according to preference. Then the school system applies a set of rules to those picks to determine how available seats are allotted to students.

The most common school choice mechanism in the United States is the "Boston mechanism," named for the city that developed it after a 1974 court ruling enforced desegregation in Boston's public schools. Variants of this system have been used in Seattle, Denver, Minneapolis and other communities across the country.

Under the Boston mechanism, students who list a school as their first choice are assigned to that school, with priority given to students who meet certain criteria, such as having a sibling in the school or living nearby. Priority points are awarded based on these criteria, and ties in points are broken through a random lottery. If no seats are left in that school, rejected students are considered for the school they ranked second, but only after children who ranked it as their first choice are assigned. As in the first round, students with priority go to the front of the line. The same procedure plays out in subsequent rounds until all students are placed in a school.

Other matching methods used by school districts include the "student deferred acceptance" mechanism, which makes tentative school assignments and reconsiders them at each step in the selection process based on students' priority, and the "top trading cycles mechanism," which allows students with priority for a school to trade places with other students with equivalent priority.

Flawed mechanisms

These and other proposed assignment mechanisms have generated lively debate among researchers and educators and have fostered a new area of research in what economists refer to as mechanism design theory. As explained by the Nobel committee that awarded its 2007 prize to pioneers in the field, mechanism design is "the art of produc-

ing institutions that align individual incentives with overall social goals." Put otherwise, it's about structuring the rules of a game to achieve the best outcome. A branch of game theory, mechanism design has been applied to real-life problems such as allocating broadcast spectrum, filling intern slots at hospitals and allotting kidneys to waiting transplant recipients.

School choice presents a similar challenge: identifying flaws in current assignment methods and finding ways to make them fairer, more efficient and less confusing for students and their parents.

There is no foolproof method; all systems currently used in schools have shortcomings, according to the tenets of mechanism design. But the Boston mechanism has elicited the severest criticism by researchers.

In 2003, economists Atila Abdulkadiroğlu and Tayfun Sönmez published a seminal paper that revealed inherent flaws in the mechanism. Their theoretical analysis shows that it encourages families to lie about their true school preferences and apply to schools that are less popular though acceptable, out of fear of being assigned to a truly bad school (in their view) if they're rebuffed by their first-ranked school. This is inefficient, because students forgo attending their preferred schools. And it puts families who state their true preferences—heedless of priority points—at a disadvantage.

After an article on the research appeared in the *Boston Globe*, the Boston school system reevaluated its student assignment mechanism, ultimately replacing it with the student deferred acceptance method. Subsequent economic research has confirmed these faults in the Boston mechanism and found it inferior to other methods.

Much of this research involves lab experiments that simulate various school mechanisms. For example, a 2006 paper co-authored by Sönmez found that only about 20 percent of subjects in a controlled experiment stated their true school preferences under the Boston mechanism. A 2010 experimental study by Calsamiglia and other researchers showed that imposing a limit on the number of schools that participants can list makes them more likely to try to manipulate the system by misstating their preferences.

Calsamiglia and Güell advance this line of research into the empirical arena by using real-world data to analyze choices made by parents in Barcelona. Drawing upon detailed application, admission and enrollment data, the economists focus on the role of priority points in shaping school decisions under the Boston mechanism.

Maia Güell

There goes the neighborhood

In communities that use the Boston mechanism, over three-quarters of households pick a school in their neighborhood as their number one choice. But do families pick nearby schools because they genuinely prefer them, or for some other reason, such as a perception that they stand a better chance of entering those schools? Or, for families that can move to the neighborhood of their preferred school, are those motives intertwined?

Teasing out the truth empirically is no easy task. Simply observing the rankings families submit to the school district doesn't reveal the motivations of families navigating the Boston mechanism; the strategies behind those selections are opaque. But changes in the way neighborhoods are defined in

the Barcelona public school system afforded an opportunity for Calsamiglia and Güell to lift the veil on the decision-making process.

The city of Barcelona has used the Boston mechanism to place children in public schools since the mid-1990s. In the spring of 2007, school system administrators redrew the city map, abolishing old neighborhoods and establishing smaller ones based on the distance to schools from residents' homes.

The researchers were aware of the changes; Calsamiglia lives in Barcelona and Güell was an associate professor at Pompeu Fabra University in Barcelona before leaving to join the faculty at Edinburgh. Both immediately realized the implications for investigating school choice. As a result of the change in neighborhood definition, the set of schools for each family that were "safer"—that is, they conferred more priority points than other schools, increasing the odds of acceptance—also changed.

So "what was safe yesterday is not necessarily safe today, and what was not safe yesterday now becomes safe," Güell said. Because the change occurred right before the school application deadline, it also separated residential and school decisions; parents could not move to the neighborhood of their preferred school.

By observing demand for different sets of schools before and after neighborhoods were redrawn, the economists hoped to determine which impulse—genuine preference for a certain school or a desire for safety—is strongest under the rules of the Boston mechanism. "The key insight for us," Güell said, "is that if parents choose according to their preferences, the [change in neighborhood definition] should not change their preferences—they still pick their top-ranked school."

However, if families seek to avoid risk, the changeup in neighborhoods should lead them to alter their behavior, opting for safer schools in their new neighborhoods.

To prove the matter, Calsamiglia and Güell needed school application and assignment data for Barcelona children who entered elementary school between 2005 and 2010. Home addresses of these children were also required, to chart the effect of shifting neighborhood boundaries on family decisions. Working with the school district administration to allay any privacy concerns, the researchers obtained a rich data set chronicling school choices made on behalf of over 77,000 children in the Barcelona area.

They were ready to investigate how parents, constrained by the rules of the Boston mechanism, play the school choice game.

A rush to "safety"

The change in neighborhoods reshuffled the deck for Barcelona families, forcing them to rethink their school choices. Over 80 percent of households saw a change in the group of schools for which they had priority. For families living in the center of an old neighborhood, a new smaller neighborhood meant that many previously safer schools were no longer so. And some families living near the edge of an old neighborhood had new relatively safe schools to consider, because the new neighborhood included schools that previously lay outside it.

Calsamiglia and Güell map pupils' addresses to tell which families applied to a school in their neighborhood—thus qualifying for priority points—both before and after 2007. (Because children are generally assigned to a school only once, the analysis necessarily compares the choices of different households over time.) The researchers examine the top choice of families—the school ranked first on their application form.* If preferences are the main driver of school choice, demand for various schools in the system should not change with new neighborhood boundaries; in 2007 and in later years, parents should still apply to their favorite schools, regardless of location.

In fact, the neighborhood reshuffle prompted a marked shift between 2006 and 2007 in the proportion of families applying to different types of schools. Demand decreased for former neighborhood schools that now lay outside the new neighborhood and increased for schools once situated outside the neighborhood that now fell within the new neighborhood. From 2006 to 2007, the share of households that ranked first former nonneighborhood schools that were redefined as neighborhood schools rose from 9 percent to 17 percent. Also, for

*To focus on the role of priorities based on neighborhoods, not family ties, Calsamiglia and Güell exclude from their study applicants with older siblings in that school—almost half of the children in their database.

families living in the center of the old neighborhood, demand shifted from former neighborhood schools to schools that remained in the new neighborhood.

These findings point to a rush to safe choice schools in 2007; families shunned schools that no longer gave them priority, choosing instead schools that gave neighborhood residents a better chance of admission. This suggests that student assignment under the Boston mechanism with priority points is to a large degree determined by these points, not by parents' predilection for certain schools. "To sum up, we find evidence that families' preferences play a limited role in school choice because a change in the definition of neighborhood makes families change their choices," Calsamiglia and Güell write.

Being relatively certain of getting into an acceptable school, though not their most preferred, is so attractive to families that the allocation of school seats in Barcelona is quite similar to the distribution that would result if children were simply assigned to schools in their neighborhoods. Previous theoretical studies or lab experiments overlooked the strong influence of priorities on parents' choices. "The way priority points shape all these mechanisms was largely ignored, but they matter a lot," Güell said.

For educators as well as economists, the implication is inescapable: If school choice is illusory under the Boston mechanism, so for the most part are the benefits believed to stem from permitting families to apply to the schools of their choice.

Not so naïve, after all

In an extension of their analysis, Calsamiglia and Güell shed light on another important question related to school choice—why some parents take big risks under the Boston mechanism by boldly stating their preferences for oversubscribed schools outside their neighborhood.

For each family in Barcelona, the chance of being assigned to a public school depends on the submitted choices, the priority points conferred by those schools and total demand for each school in the system. By calculating the odds of children being assigned to their first-ranked school—and if they are rejected, to their second-ranked school—the economists identify those applicants who could be considered "naïve"; that is, they apply to schools (usually outside their neighborhood) to which they have no chance of being assigned. About 23 percent of applicant families behave this way.

But further analysis shows that some "naïve" families are not so naïve, after all. Calsamiglia and Güell also obtained information on enrollment—the schools that children actually attended over the study period—from the Barcelona school district. The authors use these data to analyze the outcomes of risky behavior under the Boston mechanism.

It turns out that high risk takers comprise two types of households: families that seem truly unable to grasp the consequences of their choices and those that take a calculated risk because they have an alternative, an outside option in the parlance of game theory.

Like most school systems, Barcelona's maintains a waiting list for oversubscribed schools, giving children who are rejected by those schools a second chance. However, even with waiting lists, families with fairly low priority for popular schools court disappointment: Almost one-third of such families are shut out of their top-ranked school. And, of that group, the naïve—families with less than the requisite points for their first-ranked school—fare the worst. Compared with households that incur less risk by applying to neighborhood schools, a smaller share of naïve families gain entry to any of their ranked schools, and a larger share fail to enter any selected school.

But about 14 percent of "naïve" families who miss out on their preferred school enroll their children in a school outside the Barcelona school system, either a public school in another city or a private school. These are households with an outside option.

Previous research had established the naïvete of some applicants under the Boston mechanism. But the role of the outside option in school decisions had not been explored. Calsamiglia and Güell's analysis demonstrates that a significant share of outwardly naïve families are in fact sophisticated players who shrewdly weigh the odds.

Because parents have the option of enrolling their child elsewhere—most often in a private school—they can take the risk of applying to a preferred school outside their neighborhood. If this strategy works, the child is assigned to his or her first-ranked school from the waiting list (the re-

searchers assume, based on the mathematical odds, that at least as many families win entry to their top choices as are rejected). If the risky bid fails, the family plays its get-out-of-jail-free card.

"Many of these so-called naïve parents know exactly what they're doing," Güell said. "If it doesn't work, they go to a private school." Thus, under the Boston mechanism, parents willing to roll the dice have a better chance of enrolling their children in the best public schools than families without an outside option.

Who are these savvy high rollers? Census data on the socioeconomic characteristics of Barcelona school applicants reveal that among families that exhibit naïve behavior, those that take the outside option are more highly educated than those that do not. These families are likely to have higher incomes, too. "Our empirical evidence suggests a new and important source of inequality that the [Boston mechanism] induces," write Calsamiglia and Güell.

Not only are less-educated families more likely to choose unwisely and end up in lower-quality schools; those with higher education benefit the most from openly stating their preferences.

Priorities reconsidered

By divining the motivations of Barcelona families from the school system data, the economists show the extent to which a yearning for safety dominates the strategy of parents under the Boston mechanism. The pursuit of priority points trumps school preference, virtually eliminating school choice for all but a few. "The risk involved in stating preferences is not worth taking, leading most of the applicants to apply for one of the neighborhood schools," the authors write.

Some families stay true to their school, but only better-educated applicants playing the outside option exercise informed choice; children in less-educated households that rank a nonneighborhood school number one are likely to wind up in the least-desirable schools.

Calsamiglia and Güell's work has important implications for the design of student assignment mechanisms. Given the power of priorities in influencing school choice under the Boston mechanism, their findings suggest that children would benefit from other methods of allocating school seats that are fair to all families and encourage them to sincerely state their preferences.

In a follow-up paper, Güell teams with Calsamiglia and Chao Fu, an economist at the University of Wisconsin, to assess the potential gains of switching from the Boston mechanism to either the deferred acceptance or the top trading cycles mechanism, a method first proposed by Abdulkadiroğlu and Sönmez.

Once again tapping the Barcelona school system data, the researchers show that adopting top trading cycles—a method that allows families to pick their favorite schools without fear of losing priority for lower-ranked schools—increases the share of families whose children are assigned to schools outside their neighborhood.

However, neither study is likely to be the last word in the debate about the efficacy and fairness of school choice mechanisms. Some researchers, including Abdulkadiroğlu, have come to the defense of the Boston mechanism in recent years, comparing it favorably with other mechanisms. That may be one reason why many U.S. school districts have stuck with this method to govern the annual ritual of assigning students to their places. R

References

Abdulkadiroğlu, A., and T. Sönmez. 2003. School choice: A mechanism design approach. *American Economic Review* 93 (3), 729-47.

Calsamiglia, C ., C. Fu and M. Güell. 2014. Structural estimation of a model of school choices: The Boston mechanism versus its alternatives. Mimeo, University of Wisconsin.

Calsamiglia, C., and M. Güell. 2014. The illusion of school choice: Empirical evidence from Barcelona. Working Paper 712, Federal Reserve Bank of Minneapolis.

Calsamiglia, C., G. Haeringer and F. Klijn. 2010. Constrained school choice: An experimental study. *American Economic Review* 100 (4), 1860-74.

Chen, Y., and T. Sönmez. 2006. School choice: An experimental study. *Journal of Economic Theory* 127 (1), 202-31.

Research Digest

Great Recession

RESEARCH REVIVAL

Just as the Great Depression led to the revitalization of economic theory and empirical methodology, the Great Recession has sparked a renaissance of research. Since the start of the financial crisis in 2007, economists have launched (or redoubled) efforts to understand how such crises lead to recessions, why this recession endured so long, and how policymakers might both revive the economy and prevent a repetition.

At the Minneapolis Fed, economists have pursued several avenues toward the same end. In the first several months of 2014 alone, over a dozen working papers and staff reports were published on related topics. Some of this new work has already been featured in previous *Region* issues through digests, articles or policy papers. While space limitations preclude in-depth reviews of all remaining 2014 research, brief synopses are offered here.

Research Digest

Debt default and European bailout

Several papers have dealt with international dimensions of the crisis and recession. Fed senior research economist Cristina Arellano, with Yan Bai at the University of Rochester, examined international default contagion in sovereign debt markets (SR 491; see June 2014 *Region*) and in a related paper (SR 495) analyzed optimal renegotiation policies for sovereign defaults.

In February, Minneapolis Fed consultant Tim Kehoe, with Stony Brook's Juan Carlos Conesa, published a staff report (SR 497) asking whether the success experienced in the mid-1990s in preventing Mexico's debt crisis from leading to sovereign default through massive bailout assistance could be repeated in the Eurozone.

In short, maybe not.


Cristina Arellano


Yan Bai


Tim Kehoe


Juan Carlos Conesa

Kehoe and Conesa's analysis suggests that a similar bailout by the European Commission, European Central Bank and International Monetary Fund might not succeed. Debt levels are so high among the affected nations (primarily Greece, Ireland, Italy, Portugal and Spain) that they could choose to default on their debt rather than reducing it through loan paybacks as Mexico did by early 1997, thereby regaining access to international credit markets.

A primer on federal funds

The efforts of central banks worldwide to address the financial crisis, deep recession and slow recovery continue to be scrutinized; in that vein, several Minneapolis Fed papers focused on various aspects of monetary theory and practice.

A series of three working papers (WP 708, 710 and 711) released in March and April 2014 by Gara Afonso of the New York Fed and Ricardo Lagos of New York University and the Minneapolis Fed examines the federal funds market through which financial institutions trade their "reserves"—dollar balances that are held at the Fed to meet legal requirements, earn interest or clear transactions.

As Afonso and Lagos observe, this market is important to banks for managing their reserves and offsetting liquidity and payment shocks. It is also the "epicenter of monetary policy implementation," they note, since the Fed uses it—through selling and buying bonds—to influence the amount of liquidity in the national economy. A thorough understanding of this market, therefore, "is of first-order importance to economists interested in monetary theory and policy," they write.

The first of the series is an empirical study of the market's trade dynamics. The second develops a model of the market to answer key questions—what determines the fed funds rate, the rate at which banks borrow and lend to one another overnight? How does this market reallocate funds among banks, and can it achieve efficient reallocations?—and to analyze the effectiveness of central bank policies that use the interest rate paid on reserves to manage the fed funds rate. The third is a "primer" that pulls together major elements of the first two while extending investigation into market structure changes and central bank tools such as open market and discount window operations as well as the interest rate on bank reserves.


Gara Afonso


Ricardo Lagos

Research Digest

Analyzing the "unconventional"


S. Boragan Aruoba

In August, S. Boragan Aruoba of the University of Maryland and the Minneapolis Fed issued a staff report (SR 502) that studies the effects of the Fed's recent "unconventional" monetary policy in which it sought to stimulate the economy by purchasing assets such as mortgage-backed securities, moving beyond the Fed's traditional reliance on the fed funds rate (since the nominal rate was essentially at zero interest per year and could be reduced no further in efforts to boost investment and spending). Such policies were largely untested, and economists expressed "wildly different views" about their impact, notes Aruoba—some confident of success, others predicting ineffectiveness and still others foreseeing that they would trigger damaging inflation.

Aruoba approaches this research by looking at inflation expectations over various time horizons measured in various surveys and combines those with nominal interest rates over the same period, thereby creating a "term structure of real interest rates"—basically a picture of actual interest rates on different maturities, adjusted for inflation, from 1992 to the present. He gauges how this picture responds after the Fed's various policy actions: the initial and second rounds of quantitative easing, the maturity extension program (aka Operation Twist) and the announcement of an explicit inflation target.

Ultimately, Aruoba concludes that the Fed's unconventional policies (along with its ability to sustain zero nominal interest on short-term assets) "kept long-run inflation expectations anchored." They also provided a large level of monetary stimulus, he observes, as indicated by "real interest rates on all horizons … about 3.5% lower than their pre-crisis averages."

So much money, so few loans


Javier Bianchi


Saki Bigio

Another paper that examines the effectiveness of Fed policies was issued in September (SR 503) by Javier Bianchi of the University of Wisconsin and the Minneapolis Fed, and Saki Bigio of Columbia University. Bianchi and Bigio seek to understand why, despite unprecedented policy interventions by central banks, including the Fed, to reduce long-term interest rates and also provide large amounts of liquidity to financial institutions, bank lending seems not to have been stimulated much at all. In the argot of monetary theorists, something has interfered with the "transmission of monetary policy." Banks have plenty of resources to increase loans, thanks to Fed policy, but they seem unwilling to do so.

The economists develop a model to understand how monetary policy operates through a national banking system—two realms that economists have traditionally analyzed separately (monetary economics and financial theory). The model's focus is the liquidity management problem that banks face because they use low-return demand deposits to finance higher-return loans.

The well-known problem is liquidity mismatch: Deposits must be immediately accessible to depositors, but the loans are longer-term assets. Banks hold their reserves at the central bank and use those reserves to settle transfers of deposits with other banks. Therefore, central bank policy actions can alter interbank dynamics by affecting the rates at which banks borrow and lend from one another.

With this mechanism, Bianchi and Bigio analyze how monetary policy steps transmit through the banking system and conclude that while an early interbank market freeze was probably important at the onset of the recession, "a persistent decline in demand [for credit] seems the most plausible explanation" for increased central bank reserves along with decreased lending since 2008.

Research Digest

Credit contractions, not rigid prices

When interest rates hit the "zero lower bound" in late 2008, the Fed used two unconventional tools to get the economy on track. It engaged in large-scale purchases of long-term government-backed assets ("quantitative easing"), seeking to drive down long-term interest rates. It also used "forward guidance," specifying how long and under what conditions the Fed's traditional tool, the federal funds rate, would remain at zero.

Both tools could be justified by the New Keynesian approach to monetary policy, which emphasizes the importance of price rigidities. These rigidities hindered market adjustment, according to New Keynesians, and thus deepened and prolonged the recession. (On the flip side, price rigidities provide a role for monetary policy in the revival of economic health.)

But what if that was the wrong diagnosis? What if prices were actually quite flexible, and the root cause of recession was a credit crunch due to a tightening of collateral constraints? Then optimal policy might be very different.

In a September working paper (WP 714), Francisco Buera and Juan Pablo Nicolini of the Chicago Fed and the Minneapolis Fed, respectively, build a model that has credit and collateral constraints at its heart. This model, *with* flexible prices, replicates many of the recession's key features (such as the collapse in investment and low inflation despite liquidity injected by the Fed's asset purchases), but has very different policy implications than the New Keynesian model. "On the contrary, the model we study stresses a different and novel trade-off between ameliorating the initial recession and delaying the recovery."


Francisco Buera


Juan Pablo Nicolini

Maintaining the economy at the "zero bound" for nominal interest rates—as the Fed sought to do in order to stimulate investment—or avoiding it by targeting a somewhat higher interest rate "implies non-trivial trade-offs," write Buera and Nicolini, in particular, the choice between a less severe recession and a shorter one. These trade-off decisions are even more difficult when policy impacts on various economic actors (workers, entrepreneurs and savers, for example) are taken into account. The economists' "heterogeneous agent" model allows them to analyze those trade-offs as well as those for the aggregate economy.

Sorting out the costs

Several other papers have looked at the impact of the crisis and recession nationally, generationally and on various industry sectors. Ralph Koijen of the London Business School and Motohiro Yogo of the Minneapolis Fed look at the insurance industry in a November 2014 report (SR 500). Traditional theory about insurance markets assumes fair pricing, with efficient capital markets and supply policies.

But according to Koijen and Yogo, the financial crisis undermined that conventional wisdom. The economists document that life insurers reduced prices for long-term policies dramatically, (in contrast to standard theory that falling interest rates would lead them to *raise* policy prices). Prices for 30-year term annuities were marked down 19 percent relative to actuarial value for life annuities at age 60 and down 57 percent for universal life insurance at age 30. (In ordinary times, insurers earn a 6 percent to 10 percent markup.) They find larger price reductions for policies with looser statutory reserve requirements. This evidence suggests that even large insurance firms ran short of cash during the crisis and needed to raise money quickly by reducing prices to boost sales.


Ralph Koijen


Motohiro Yogo

Koijen and Yogo build a model to understand the forces behind such extraordinary pricing behavior and

conclude that "financial and regulatory frictions have a large and measurable impact on insurance markets." These findings emphasize how important *supply*-side frictions are in consumer financial markets, say the economists, supplementing the demand-side frictions (borrowing constraints, asymmetric information, moral hazard and bounded rationality) that most research has studied. Their study also provides microeconomic evidence for macro models based on financial frictions, "a leading explanation for the Great Recession," they observe.

Generation gaps, legacy costs and reassessments

Research by Andrew Glover, Jonathan Heathcote, Dirk Krueger and José-Víctor Ríos-Rull examines the recession's varying economic impact on the young and the old (SR 498); an earlier version of this paper was described in the September 2011 *Region*.

In a July staff report (SR 501), Tom Holmes and Lee Ohanian study the impact of legacy costs such as pay-with-promises compensation plans when cities (like Detroit) suffer economic shocks, as during the recession. This work also took the form of an economic policy paper (EPP 14-4; also in the June 2014 *Region*).

In another staff report (SR 494), described in the September 2014 *Region*, Ellen McGrattan and Ed Prescott reassess real business cycle theory to see if its central tenets were undercut by the recession, as critics suggest. They find to the contrary. The idea that business cycles are driven, in part, by fluctuations in factor productivity, as RBC argues, is not undermined by government data showing that labor productivity actually rose during 2008-2009, McGrattan and Prescott contend. If investment in intangible capital is accounted for, the link between factor productivity and business cycles persists.

In a similar reexamination of theory, Terry Fitzgerald and Juan Pablo Nicolini released a working paper (WP 713) in May that looks at the Phillips curve relationship between unemployment and future inflation, a link that many economists believe no longer holds. Fitzgerald and Nicolini find that because U.S. monetary policy over recent decades has sought to stabilize nationwide prices, data aggregated at the national level "is uninformative" about the relationship that may exist at smaller geographic levels such as cities or regions.

In fact, among U.S. metropolitan statistical areas from 1976 to 2010, "we find that a 1 percentage point increase in the unemployment rate is associated with a roughly 0.3 percentage point decline in inflation over the next year." They qualify this finding strongly, however, noting that it applies only if particular assumptions are made. "Our results do not prove Phillips curve skeptics wrong," they caution. (This research was also discussed in EPP 13-6, November 2013.)

By the time this *Region* is published, economists at the Minneapolis Fed and elsewhere will have released new research on other aspects of the Great Recession. It has inspired reconsideration of old theory and formulation of new—a silver lining to the cloud of severe economic downturn.

—Douglas Clement


Andrew Glover


Jonathan Heathcote


Dirk Krueger


José-Víctor Ríos-Rull


Tom Holmes


Lee Ohanian


Ellen McGrattan


Ed Prescott


Terry Fitzgerald


Juan Pablo Nicolini

BOOK REVIEW

What Good Are Economists, Anyway?

Darn good, according to *Trillion Dollar Economists*

Trillion Dollar Economists: How Economists and Their Ideas Have Transformed Business
By **Robert E. Litan***
Bloomberg
400 pages

Reviewed by **John Yanish**
Vice President and Deputy General Counsel
Federal Reserve Bank of Minneapolis

A few years back, a T-shirt worn by economics graduate students at the University of Minnesota read, "Sure, it works in practice, but how does it sound in theory?" In light of the recent financial crisis and ensuing years of economic challenge, that ironic academic poke arguably leads to a more serious question: "For all their theories, do economists add much value in the real world?"

Robert Litan, an economist, lawyer and senior fellow at the Brookings Institution (see bio below) answers that question with a resounding yes in his most recent book, *Trillion Dollar Economists*. Indeed, Litan portrays economists as unsung heroes who have created enormous value in business and a wide range of day-to-day activities, from investing to air travel, energy policy, online dating and beyond.

The book is unconventional in many respects. It's completely devoted to economics, but the author consciously writes for those who know little or nothing about it. An opening chapter is laid out as a primer on basic economic principles (e.g., micro versus macro) designed to give the otherwise uninformed reader a baseline of knowledge. Litan then leads readers, armed with those fundamentals, on a whirlwind journey through real-world examples. The journey's destination: Economists and economic theory have "created trillions of dollars of income and wealth for the United States and the rest of the world, hence the title of the book."

One surprise is the book's lack of focus on monetary policy. Although Litan has kind words at various points for Ben Bernanke, Paul Volcker and others who have shaped monetary policy over time, he gives little attention to Federal Reserve policies or other macroeconomic matters. Indeed, he laments that, in his view, most people are under a mistaken impression that economists focus their efforts on trying to determine the future trend of the economy, as the Fed does when setting policy. He states early on: "The truth is that relatively few economists actually engage in economic soothsaying or prediction. Yes, let me repeat that, or phrase it in a slightly different way: The popular perception of what most economists do is wrong, and I intend to show that in [this book]." Litan's central point: The unrecognized extraordinary value of economists is their direct (or near-direct) contribution to daily business practices and other matters of common interest.

*Robert Litan is a nonresident senior fellow at the Brookings Institution, where he formerly was senior fellow, vice president and director of economic studies. A professional economist, practicing attorney and prolific writer, Litan has authored or co-authored over two dozen books, edited many more and written scores of articles, reports, commentaries and reviews for professional and popular journals. He gives frequent speeches; testifies regularly before Congress, often on issues related to regulatory reform and/or banking; and has worked at nonprofit, corporate and government organizations and agencies. He earned his J.D. at Yale Law School and his Ph.D. in economics at Yale University.

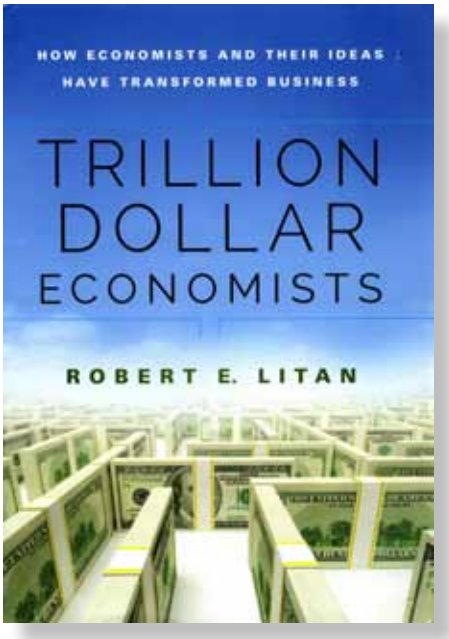

The bulk of *Trillion Dollar Economists* is devoted to examining what the author believes are examples of real-world successes driven by economists and economic theory. One such case: efficient air travel. The author reminds readers of years past, when airlines routinely "bumped" ticketed passengers from overbooked flights. Recognizing (finally!) the tremendous ill will this generated among paying customers, airlines began to leave seats empty (thereby losing revenue) to avoid having to bump. But doing so, of course, cost airlines the revenue those empty seats could have generated.

Enter economic theory. Litan relates how economist Julian Simon and others helped airlines avoid the issue by using an auction, the now familiar practice of offering vouchers for future travel to passengers willing to voluntarily forgo the overbooked flight in favor of a later flight. This low-cost practice eliminated ill will, allowed fully booked flights and, according to Litan, increased airline revenue by $100 billion over the past 30 years.

Building off that example, the book provides a range of other success stories brought about by economic theory. Index investing, the now-popular practice of investing in a very broad range of stocks rather than attempting to pick individual "winners" is a typical Litan case study. In the following excerpt, he traces index investing directly back to the influence of several economists on the mutual fund industry with informal, accessible storytelling style:

> Ideas can have real-world commercial impacts in many different ways. Sometimes, entrepreneurs and executives at established firms read a book or article with a clever idea and they proceed to make it operational. Others get commercial ideas from economists they hire as consultants. And frequently, as you will see in this book, entrepreneurs are motivated by an economic idea they learn while attending school.
>
> John "Jack" Bogle, the founder of the Vanguard family of mutual funds, is a prime example of the last way economists have had an impact. Bogle reports that he was heavily influenced by both [Burton] Malkiel and Paul Samuelson … two of the champions of indexing and critics of active money management, especially by individual investors, both in writing his senior thesis at Princeton on the idea of index funds and then actually implementing that idea at Vanguard.

According to Litan, when the remainder of the mutual fund industry followed suit, this new-found practice not only made investing simpler for the average individual, but yielded superior results and "transformed the investment world."

Litan relates similar positive contributions to energy policy (lower prices after economic theory established the ineffectiveness of mandatory price controls), sports (econometrics help team owners find value, as depicted in the book and movie *Moneyball*) and online dating (economic game theory yields better matches).

While making his case, Litan also provides dozens of profiles of individual economists, amounting to a type of Hall of Fame from the author's perspective. This feature, along with Litan's low-key, "this is not about me" presentation style, adds to the book's charm. It is clear that the author has genuine admiration and, in some cases, personal affection for the economists he profiles, many of whom were colleagues or had otherwise influenced his own career.

The book isn't without flaws. At points, such as his detailed description of the rise of telecommunications and the Internet, Litan seems to stretch too far. He details at great length contributions made by economists to important developments, such as providing economic analysis to support the legal case regarding the breakup of AT&T, which set the stage for greater competition in the field. While this and other efforts by economists had substantial impact, it is a far reach to extrapolate from these events to arrive at Litan's statement that "it is no overstatement to claim that the modern Internet-based business landscape of today owes much of its shape, if not existence, to the behind-the-scenes thoughts and research of numerous economists." An argument that, but for economists, we wouldn't have the Internet seems to go beyond the evidence provided.

This generous portrayal of economists contrasts with Litan's take on the Wall Street activities that sparked the recent financial crisis. He seems to suggest that economists had little if anything to do with shaping risky investment practices and strategies. While this might be valid, the reader is left with the sense that the landscape is being selectively reviewed, with an eye toward taking credit for the good and, as for the bad, giving economists the benefit of generous doubt.

That said, *Trillion Dollar Economists* makes an interesting contribution to a discussion of the relative merits of the economics profession. As a non-economist, and therefore a member of Litan's target audience, I found the book a challenging read, but worth the effort. The writing, while not always elegant, is ultimately digestible, even if concepts like regression analysis, performance optimization and congestion pricing remain a bit fuzzy. I finished the book largely convinced by his case for the value economists have added, in both business and other activities not associated with the field.

Ironically, economists are likely to be the audience that would most enjoy this book intended for non-economists. Those knowledgeable in the field would no doubt understand Litan's arguments better than lay readers, would be familiar with the economists he lovingly profiles and would certainly enjoy the accolades he showers on the profession. A potential problem, though, for the humorous T-shirt slogan business? "Sounds good in theory and works in practice" just doesn't have much punch to it. R