

# Raj Chetty

**R**AJ 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

#### ON TEACHER QUALITY

**I**f 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.

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.

#### 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 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](http://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 be-

cause 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 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 under-

stand 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

**S**ocial 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](http://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, Brigitte 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 sav-

#### 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.

ers 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.

#### 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](http://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/](http://rajchetty.com/)