Published Quarterly by the Federal Reserve Bank of Minneapolis

September 2015

The Region

Amy Finkelstein

A Step in the Wrong Direction

On the Ethics of Redistribution

Research Digest

Volatility and Foreign Assets

Sovereign Default = Tight Credit?

Technology & Inequality

Volume 29 Number 3
September 2015
ISSN 1045-3369

The Region

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

2 Ninth District Notes
Raising Rates Now Would Be a Mistake
Narayana Kocherlakota

4 On the Ethics of Redistribution
Policy analyses often ignore world's poorest when advocating redistribution
V. V. Chari and Christopher Phelan

8 Interview with Amy Finkelstein
Douglas Clement

26 Research Digest
Volatility and foreign assets
Increased saving due to uncertainty helps explain international debt
Douglas Clement

Risky business
How the threat of government default makes banks leery of lending
Phil Davies

33 Student Essay Contest Winner
An Analysis of the Impact of Technology on Income Inequality
Solomon Polansky

36 Virtual Fed
Collegecast
Joe Mahon

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

Email: letters@mpls.frb.org
Web: minneapolisfed.org
Subscriptions: minneapolisfed.org/forms/publications-address-change

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

Raising Rates Now Would Be a Mistake

Narayana Kocherlakota

President
Federal Reserve Bank of Minneapolis

Editor's note: The following opinion piece was published in the Wall Street Journal, *Aug. 19, 2015.*

I participate in the meetings of the Federal Open Market Committee, the monetary policy-making arm of the Federal Reserve. In that capacity, I'm often asked by members of the public about the biggest danger facing the economy. My answer is that monetary policy itself poses the biggest danger.

Many observers have called for the FOMC to tighten monetary policy by raising interest rates in the near term. But such a course would create profound economic risks for the U.S. economy.

Why would a near-term tightening of monetary policy be so problematic? Because given the prevailing economic conditions, higher interest rates would push the economy away from the FOMC's economic goals, not toward them.

Congress has mandated that the Fed promote price stability and maximum employment. The FOMC has translated its price-stability mandate into a target 2% inflation rate, as measured by the personal consumption expenditures price index. Inflation has run consistently below that objective for more than three years and is currently at 0.3%.

The outlook is for more of the same. Most private forecasters do not see inflation reaching 2% for the next two years. Government bond yields are consistent with that same subdued inflation outlook. In June the FOMC's own staff forecast was that inflation would remain below the committee's 2% target until the 2020s.

The U.S. inflation outlook thus provides no justification for policy tightening at this juncture. Given that outlook, the FOMC should ease, not tighten, monetary policy by, for example, buying more long-term assets or by reducing the interest rate that it pays on excess reserves held by banks. Along these lines, the board of directors of the Minneapolis Fed has for the past few months been recommending a reduction in the interest rate that the Federal Reserve charges banks for discount window loans.

With inflation so low, higher rates will push the economy away from the Fed's price and employment goals.

Many observers have called for the FOMC to tighten monetary policy by raising interest rates in the near term. But such a course would create profound economic risks for the U.S. economy.

Now, this is not to say that increasing the federal-funds rate by a mere quarter of one percentage point, as many advise, would in and of itself have a huge direct impact on the U.S. economy. But even a small change toward tighter policy would send a strong message to financial markets.

If the Fed raises interest rates when inflation is so far below target, market participants and other members of the public could well conclude that the FOMC has implicitly lowered its inflation goal. That, in turn, poses two serious risks to the economy.

The first risk is near-term. Financial decisions depend on real—that is, net-of-inflation—interest rates. If the public believes the Fed has lowered its inflation goal, all real interest rates in the U.S. will be higher.

This will discourage people from borrowing money for homes and autos. It would likely raise the dollar's value relative to foreign currencies, making U.S. goods and services less attractive to households and firms here and abroad. Prices of homes and other assets would also feel downward pressure under higher interest rates. All of these changes would likely discourage spending and create a drag on U.S. economic activity and employment growth.

The second risk is longer-term. In late 2014 the FOMC ended its asset purchase program, even as the outlook for inflation was sliding further below the 2% goal. Financial markets logically interpreted this step as meaning that the FOMC had tacitly lowered its longer-term inflation objectives from the 2% goal established in January 2012. If the committee were to tighten monetary policy again by raising the federal-funds rate in 2015, when the inflation outlook has changed little since late 2014, markets would likely respond similarly.

Moreover, by again setting policy in a direction opposite its stated goals, the Fed would diminish the credibility of those goals. As we have seen in Japan over the past two decades, when the public comes to doubt a central bank's commitment to its goals, the economy can land in a permanent low-interest-rate trap. The central bank is then much less able to fight recessions effectively. Unfortunately, as we have also seen in Japan, such traps are extremely difficult to escape.

I am confident that the time will come when economic conditions will be appropriate for the FOMC to raise the federal-funds rate from its current low level. But that time is not now. Tightening monetary policy when inflation is projected to be so low is a step in the wrong direction. R

On the Ethics of Redistribution

Redistribution policy analyses violate their own behind-the-veil-of-ignorance criterion when they ignore poor country impact

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

Christopher Phelan
University of Minnesota
Federal Reserve Bank of Minneapolis

Executive summary

Analysts of optimal policy often advocate for redistributive policies within developed economies using a *behind-the-veil-of-ignorance* criterion. Such analyses almost invariably ignore the effects of these policies on the well-being of people in poor countries. We argue that this approach is fundamentally misguided because it violates the criterion itself.

Introduction

When evaluating economic inequality, economists frequently employ the ethical principle referred to as *behind-the-veil-of-ignorance.* Originated by Nobel Laureate John Harsanyi and philosopher John Rawls, this criterion imagines the social contract that would be developed by a society of risk-averse people who don't yet know where each of them will end up in that society's distribution of income.[1] So, for example, income transfers from those with high innate skills (and therefore high incomes) to those with low innate skills (and low incomes) are justified as a type of insurance. Such transfers represent the outcome of an insurance contract the mythical behind-the-veil individual would have been eager to sign if only he had been given the chance to do so.

But such insurance schemes also have incentive effects. For instance, policy mechanisms that transfer income from highly skilled people to those with low innate skills frequently require progressive income taxes. Such policies affect incentives regarding the acquisition of skills through effort and education.

If high incomes are highly taxed, high-innate-skills individuals may have less incentive to get, say, a medical degree. Economic arrangements seen as best using the behind-the-veil criterion typically trade off such output losses against the "insurance" or welfare gains associated with transfers.

From behind the veil of ignorance, no individual could know into which country (or economic class) he or she will be born. Behind-the-veil, risk-averse people would therefore want to ensure that people born in rich countries do not adopt policies that hurt people born in poor countries. Nevertheless, analysts almost invariably ignore the effects of domestic tax policy on those in other nations. But consistent use of the behind-the-veil criterion would mean that analysts cannot treat people who live in rich, developed economies differently than they treat people who live in poor, less-developed economies.

__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.

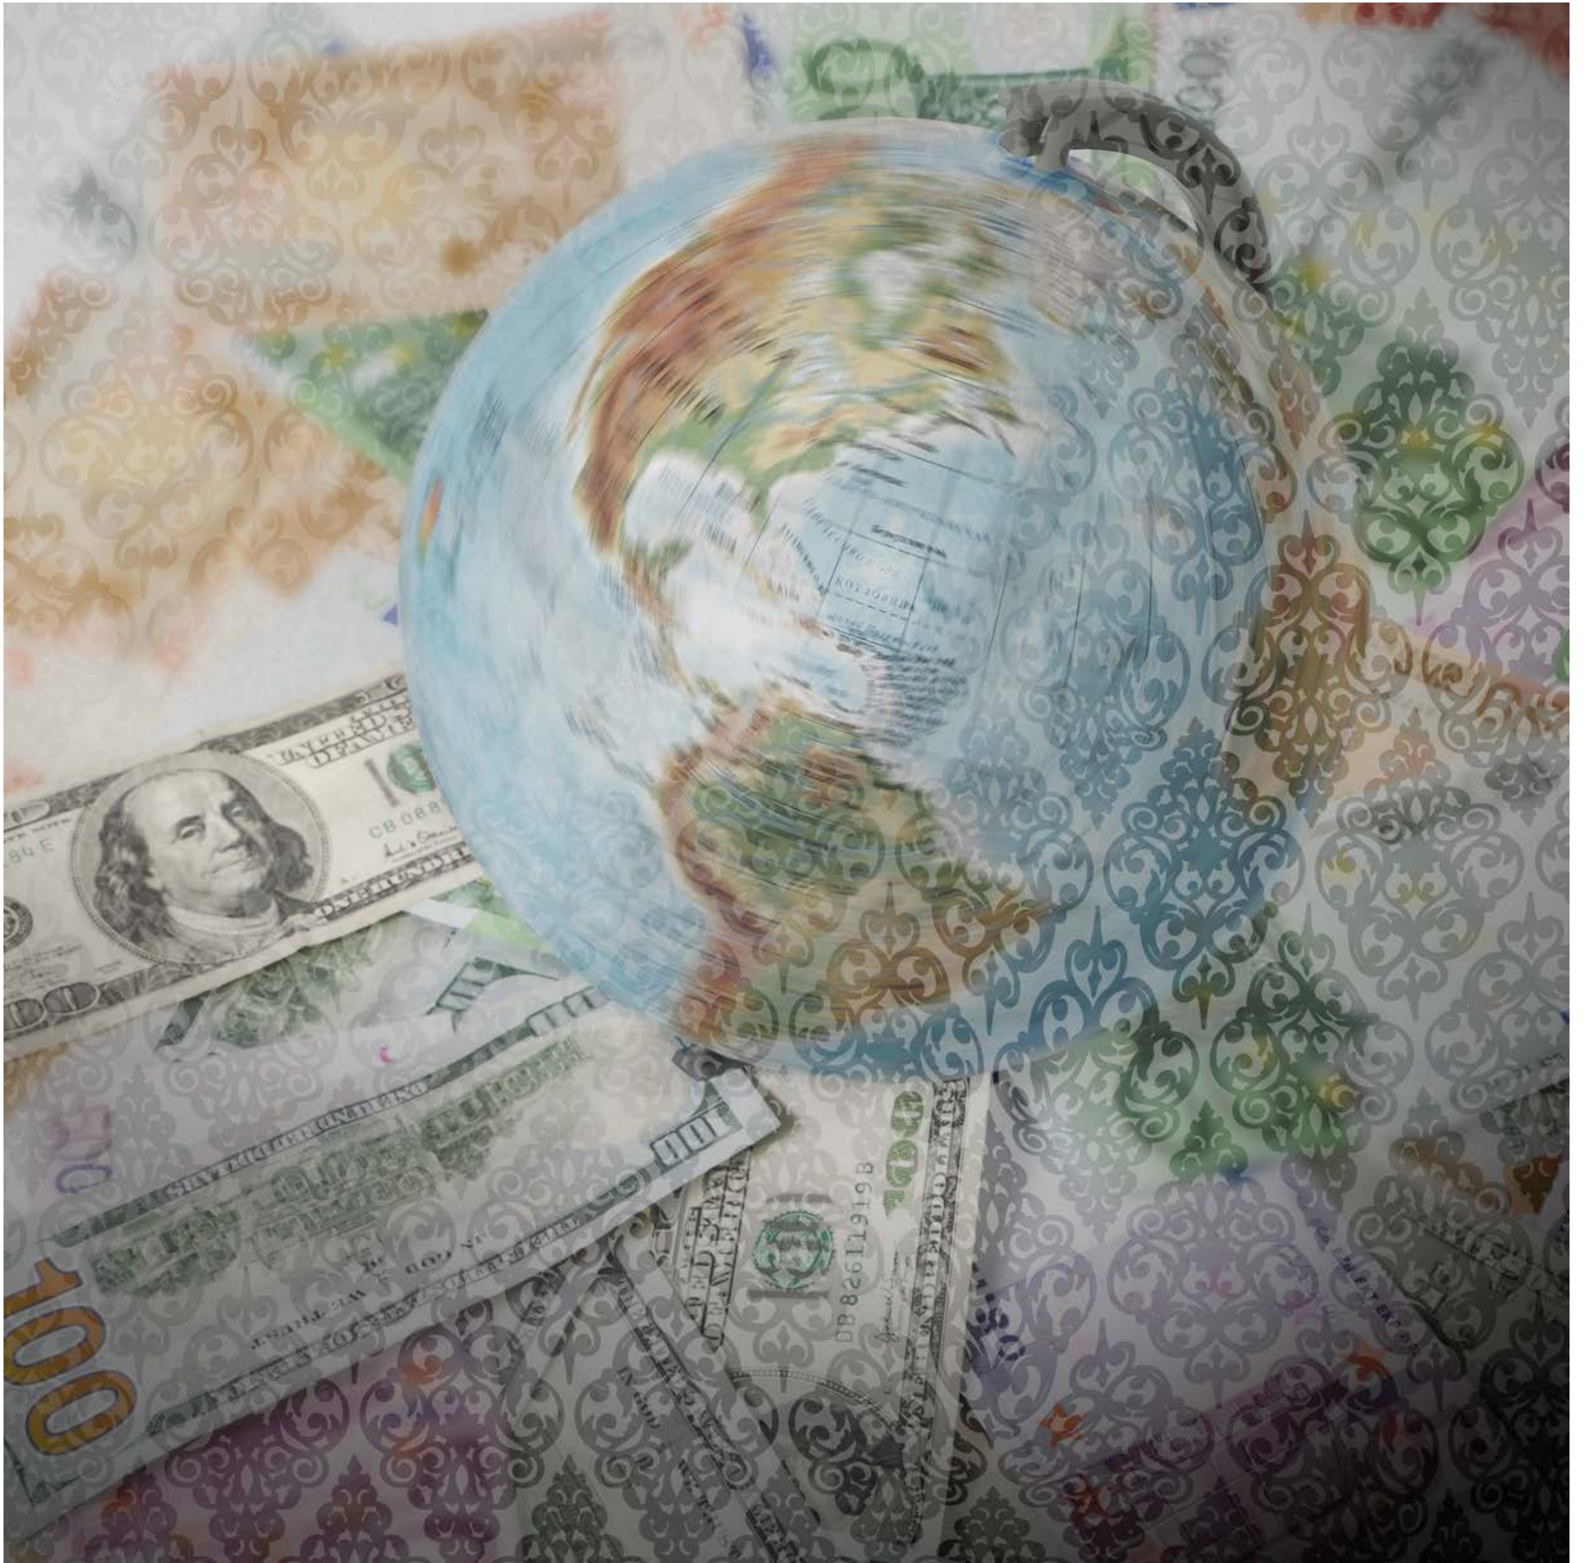

Are policies that purport to help the comparatively well-off (those at, say, the poverty line in developed countries) at the expense of the superlatively well-off (the rich in developed countries) desirable from the behind-the-veil perspective assuming that that perspective is global?

In particular, taking this perspective means analysts should care about *global* versus within-country inequality. It further implies that those considered poor in developed countries are, in absolute terms, quite wealthy compared to the vast majority of the world's population. A typical American in the lowest 5 percent of income (for America) has a higher income than 95 percent of Indians, 80 percent of Chinese and 50 percent of Brazilians.[2] In the United States, 99 percent of households have indoor plumbing (a toilet with a sewer connection).[3] In India, it's 12 percent.[4] For Americans below the poverty line, nearly three-quarters have a car (and 31 percent have two or more)[5] and 97 percent have air conditioning.[6] In India, only 5 percent of all households have cars and 2 percent of all households have air conditioning.[7]

This then begs the following question: Are policies that purport to help the comparatively well-off (those at, say, the poverty line in developed countries) at the expense of the superlatively well-off (the rich in developed countries) desirable from the behind-the-veil perspective assuming that that perspective is global?

Increasing world trade is an example of the tension between policies that help those in developing countries versus those that help those lower in the income distribution in developed countries. According to a World Bank Study, in the three decades between 1981 and 2010, the rate of extreme poverty in the developing world (subsisting on less than $1.25 per day) has gone down from more than one of every two citizens to roughly one of every five, all while the population of the developing world increased by 59 percent.[8] This reduction in extreme poverty represents the single greatest decrease in material human deprivation in history.

But this decrease in extreme poverty in the developing world has coincided with a marked increase in income inequality in the developed world, and the latter has received much more attention, at least from policy analysts in these richer nations.

One possible cause of both trends has been the increase in international trade, which lessens the market value of less-skilled labor in developed countries while increasing its value in developing countries.[9] If one uses a behind-the-veil criterion focused only on developed countries, then the increase in trade has made things worse. If instead one considers the entire world, then the trade increase has made the world phenomenally better.

But trade is not the only way in which policies in developed countries affect those in developing countries. Nobel Laureate Robert Lucas has developed an economic model in which countries "take off" by removing impediments to becoming a growing economy (China's abandonment of central planning, for example) but vary regarding when such impediments are removed.[10] Two key implications of his model, both of which match the data, are that the later historically a country takes off economically, the further its per capita output will be at that point from the rest of the world, and the faster its subsequent growth. (The country will, in a sense, go through a stage of "catch-up growth.") These implications can arise either from economic growth through diffusion of ideas or from diffusion of practical knowledge through education.

Consider the following highly stylized example: In a world with just two countries, one developed and the other poor, output is produced in each by a combination of skilled workers and unskilled workers. When they're young, unskilled workers have the opportunity to become skilled by working with older, skilled workers.

But imagine that young, unskilled workers can work with older, skilled workers from *either* country. In particular, assume that skilled, older workers (such as plant managers) from developed countries can train young, unskilled workers from developing countries. (Alternatively, imagine that young, unskilled workers from developing countries travel to developed countries to become educated and then return home as skilled workers.) When these young workers age, they in turn train future generations of young workers at home.

Suppose further that in each country only some young workers are born with an innate ability to acquire skills, while others are born without that ability. Suppose also that young workers who have this ability must exert effort to acquire skills and therefore must be provided with appropriate incentives to do so.

A rich-country policy to tax high incomes will redistribute income (within that country) from those with high innate abilities (and, by assumption, with the ability to become highly skilled) to those with lower innate abilities. In so doing, that policy will reduce inequality within the rich country, but it will also create disincentives there to becoming highly skilled and thereby reduce the global supply of skilled workers. This reduced supply of skilled workers from the developed country then reduces opportunities for young workers in the poor country to become skilled.

Applying the Harsanyi-Rawls behind-the-veil-of-ignorance criterion but considering *only* people in the developed country would appear to make this a beneficial policy because it helps the poor of that rich country. But, in our example, it hurts the poorest of the poor in the world, those in the developing nation. A *proper* application of the behind-the-veil-of-ignorance criterion—one that takes all people in all countries into consideration—can thus lead to the implication that such a policy is extremely undesirable. At the very least, a proper application of the criterion says that redistribution within rich countries imposes costs on people in other countries which need to be taken into account.

We conclude that using the behind-the-veil-of-ignorance criterion to advocate for redistributive policies within developed countries while ignoring the effect of these policies on people in poor countries violates the criterion itself and is therefore fundamentally misguided.

Many economic analysts use social welfare functions in which, implicitly, only the well-being of domestic residents matters. This type of analysis is acceptable as long as the analyst acknowledges that such a social welfare function is not developed from deeper ethical considerations. A giant literature in public finance justifies such social welfare functions by appealing to the veil of ignorance. Our point simply is that those who use this criterion should weight the welfare of poor people in Chad, the world's poorest nation, very heavily. To our knowledge, very little if any of the relevant research does so. R

Endnotes

[1] In *A Theory of Justice,* Rawls (1999, p.118) writes, "[N]o one knows his place in society, his class position or social status; nor does he know his fortune in the distribution of natural assets and abilities, his intelligence and strength, and the like."

[2] See Milanovic (2011, p. 116).

[3] U.S. Census Bureau, 2009-2013 5-Year American Community Survey.

[4] Population Reference Bureau. prb.org/Publications/Articles/2012/india-2011-census.aspx

[5] *Backgrounder.* heritage.org/Research/Welfare/bg2064.cfm

[6] U.S. Energy Information Administration. eia.gov/consumption/residential/data/2009/#undefined

[7] Population Reference Bureau. prb.org/Publications/Articles/2008/howindianslive.aspx

[8] World Bank. worldbank.org/content/dam/Worldbank/document/State_of_the_poor_paper_April17.pdf

[9] While it is clear that extreme poverty in the developing world has decreased, it is unclear whether *inequality* in the developing world has increased or decreased. See Goldberg and Pavcnik (2007).

[10] See Lucas (2000).

References

Goldberg, Pinelopi Koujianou, and Nina Pavcnik. 2007. "Distributional Effects of Globalization in Developing Countries." *Journal of Economic Literature* XLV: 39-82.

Lucas, Robert E. 2000. "Some Macroeconomics for the 21st Century." *Journal of Economic Perspectives* 14 (1): 159-68.

Milanovic, Branko. 2011. *The Haves and the Have-Nots.* New York: Basic Books.

Rawls, John. 1999. *A Theory of Justice.* Cambridge, Mass.: Harvard University Press.

PHOTOS BY PETER TENZER

Amy Finkelstein

Standard economic theory long predicted that people with more insurance coverage will make more insurance claims.* Standard theory also said that markets for both annuities and long-term care insurance should be large and robust. Conventional wisdom held as well that geographic variation in U.S. health care spending was due mostly to supply-side factors—doctors' practices, technology, hospital management—not patient demand.

Not one of these "truths" is valid. But only after MIT economist Amy Finkelstein analyzed their empirical realities and theoretical flaws did economists understand why. Her gift for combining data and theory has revealed subtleties of economic behavior that long eluded the profession. And she's applied this talent to improve understanding and policy in health insurance—one of the most complex, expensive and contentious areas of public discourse.

Her contributions have been widely recognized. Calling her "the leading scholar in Health Economics and one of the most accomplished applied micro-economists of her generation," the American Economic Association honored her in 2012 with the John Bates Clark Medal, given to the leading American economist under 40. She received a similar award last year from the American Society of Health Economists. In 2008, she received the Elaine Bennett Prize for outstanding economic research by a woman at the beginning of her career. She's been honored was well by the Sloan Foundation, the Econometric Society and the American Academy of Arts and Sciences.

While insurance, especially health insurance, has been her nearly exclusive research target, Finkelstein's energy is now turning to health care *delivery*—its efficiency, organization and design. "How do we design health care systems to efficiently deliver the care we think should be delivered?" she asks in the following interview. "There's a lot we don't yet know about how to best design these systems, [making] it an extremely fun and exciting area." That passionate curiosity is at the core of her research, powering an intellect that promises new truths for economics, health care, policy and the public.

* Moral hazard and adverse selection are the explanations. Moral hazard is the economist's term for people taking fewer precautions when they are insulated from loss by insurance. Adverse selection is when high-risk individuals—with dangerous jobs, lifestyles or health conditions—purchase more insurance than low-risk people.

MULTIPLE DIMENSIONS OF INFORMATION

Region: Standard theory says that because of adverse selection and moral hazard, people with more insurance coverage will make more claims. Yet the data don't always support that prediction. People with more auto insurance, for example, don't necessarily have higher claim rates.

In a 2006 paper with Kathleen McGarry, you developed an explanation having to do with multiple dimensions of information—in particular, differences among people in risk *aversion* as well as risk *type*. How does that distinction help resolve the ambiguous empirical findings?

Finkelstein: What I love about that whole body of work, which our paper is just a part of, is that it's a really, really nice interplay between theory and empirics. Seminal and influential theoretical work on this dates back to the 1970s, such as Rothschild and Stiglitz (1976). Their workhorse theory assumed that individuals differ only in privately known risk type. Their model generates the famous result that the private market may generate too little insurance coverage, and that there are potential welfare gains from government intervention. Models such as theirs and Akerlof's (1970) lemons model have been extremely influential in both academic research and public policy.

The empirical prediction of their model is that individuals with more insurance will be higher risk (that is, more likely to experience the insured event or "accident"). Somewhat amazingly, the wave of empirical work investigating the predictions of this influential theory really only took off in the late 1990s and early 2000s. People like Pierre-André Chiappori and Bernard Salanié (2000) started actually looking at the predictions empirically. And some papers started to find that there were markets in which those with more insurance *weren't* actually higher risk.

MULTIPLE DIMENSIONS OF INFORMATION

We found that individuals with long-term care insurance were not more likely to go into a nursing home than those without it, as standard adverse selection theory would predict. In fact, they often looked less likely to go into a nursing home.

Region: Chiappori and Salanié found that lack of positive correlation in the French auto insurance market. There's no correlation between coverage comprehensiveness and frequency of accidents.

Finkelstein: What Kathleen and I realized is that it could be because the simple Rothschild/Stiglitz model, for tractability purposes, had people differing only on their risk type,[1] but, in fact, people differ also on their preferences: how risk averse they are, how worried they are, how cautious they are, et cetera.

Region: So the interplay between them—risk type and risk preferences—could be at the root of it.

Finkelstein: Exactly. Suppose you have people—in health insurance we often refer to them as the "worried well"—who are healthy, so a low-risk type for an insurer, but also risk averse: They're worried that if something happens, they want coverage.

Region: They will take out as much insurance as they can.

Finkelstein: Right. As a result, people who are low risk, but risk averse, will also demand insurance, just as high-risk people will. And it's not obvious whether, on net, those with insurance will be higher risk than those without.

So, you can have private information of the Rothschild/Stiglitz type—an individual purchasing insurance would know their risk type, but the insurance company wouldn't—and it can impair the functioning of an insurance market, but it wouldn't be detected by the standard test of comparing "accident rates" for people with and without insurance.

Our paper gave a proof by example.

We looked at long-term care insurance—which covers nursing homes—and rates of nursing home use. We found that individuals with long-term care insurance were not more likely to go into a nursing home than those without it, as standard adverse selection theory would predict. In fact, they often looked less likely to go into a nursing home. These results held even after controlling for what the insurance company likely knew about the individual, and priced insurance on.

The standard interpretation of this result would be that there wasn't private information in the long-term care insurance market. But our data gave us a way to detect private information: people's self-reported beliefs about their chance of going into a nursing home. And we showed that people who think they have a higher chance of going into a nursing home are both more likely to buy long-term care insurance and more likely to go into a nursing home.

And, again, these results held even after controlling for what the insurance company would have predicted. In other words, we found direct evidence that individuals have private information about their risk of nursing home use and that people who thought they were higher risk than the insurance company thought were more likely to purchase long-term care insurance. That certainly sounds like the standard adverse selection models!

But if you just look at the cross-section data, you don't confirm the theo-

retical prediction that those with more insurance are more likely to go into a nursing home.

So we realized, as a basic econometric decomposition, that there must be some other characteristic of individuals that was positively correlated with long-term care insurance purchases, but *negatively* correlated with nursing home use. That was the only way to reconcile the facts.

Then we found some examples in the data that we broadly interpreted as proxies for preferences such as risk aversion, and we found that individuals who report being more likely to, for example, get flu shots, or more likely to wear seatbelts, were both more likely to buy long-term care insurance and less likely to subsequently go into a nursing home.

Region: This research has been very influential in the field; it's one of your most widely cited papers, I believe. And it's a great example of what you highlighted earlier, a productive interplay between theory and data. Could you talk a little bit about the value of conducting research that way?

Finkelstein: I'm an empiricist mostly, although I'm very motivated by theory, and so the empirical work that I've done has gone in the direction of, "How do we design empirical tests that are robust to the fact that the real world is more complicated than the simple theory suggested, and how do we think about welfare in that context?"

I'm excited to see other researchers taking up the challenge of expanding and enriching the theory itself to model private information about risk type when there are multiple dimensions of heterogeneity. This type of interplay and conversation between theory and empirics is, I think, ideally how the field progresses.

This point was really underscored for me in an amusing experience I had in 2008, when I was asked to speak at a *festschrift* for Michael Rothschild. Of course, I was honored and thrilled to do so, and I talked about how influential the early work by Rothschild and Stiglitz had been, and how it had motivated a subsequent empirical literature.

Mike [Rothschild] spoke after I did, and his comment really stuck with me. After listening to my description of the empirical work trying to test his theory, he said something like, "Wow, honestly, when we wrote down that model, we never thought anyone would take it literally! It's a *model*."

ANNUITIES AND ADVERSE SELECTION

When you also have preference heterogeneity, the optimal contract that different people would choose may differ. Then, rather than the obvious theoretical solution—the mandate—you're in the world that I love to think about, which is one of a potential empirical trade-off. … Is the welfare lost from adverse selection in the unmandated market bigger or smaller than the welfare lost when imposing a uniform policy through a mandate?

ANNUITIES AND ADVERSE SELECTION

Region: You've done a lot of work on annuities and adverse selection, including some of your early work with Jim Poterba[2] and more recent work both with him and with others. The initial work provided important clues as to why the market for long-term annuities is so small when theory predicts that many people would benefit from purchasing them. What's your explanation? What did you find in that research?

Finkelstein: The work that Jim and I did showed that adverse selection exists in annuity markets. An annuity is a survival-contingent income product. So individuals who think they are likely to be long-lived are "high risk" from the insurance company's perspective.

We found that individuals who are longer-lived are more likely to buy an annuity and to buy annuities whose payments are more backloaded, meaning that adverse selection distorts not just the share of individuals with annuities, but also the annuity contract allocation. We found that because individuals who have private information that they are likely to live a long time are more likely to buy annuities, annuities are priced higher than they would be if annuitant mortality was typical of the general population.

But Jim has done other work with Jeff Brown, Olivia Mitchell and Mark Warshawsky (1999) in which he shows that even with this price markup (or "load"), most risk-averse individuals would still be willing to purchase annuities. This suggests that while adverse selection exists, it is unlikely to be the primary cause of the fact that so few people voluntarily purchase annuities.

Liran Einav and I, with Paul Schrimpf (2015), have done some related work estimating the welfare costs of adverse selection in the semi-compulsory UK annuities market and, using very different methods, we are finding results consistent with Jim's earlier work.

Region: That paper with Einav and Schrimpf indicated that the standard solution to adverse selection, a government mandate, actually might lower welfare.

Finkelstein: Yes. And it's exactly related to what we were just talking about in the work with Kathleen McGarry on preference heterogeneity. When you have private information about risk type, you get allocative distortions. Lower-risk people who would be willing to buy insurance at the actuarially fair price for them don't have that option and may end up inefficiently uninsured. A standard solution suggested by theory, and widely used and discussed in policy, is mandated insurance coverage.

But when you also have preference heterogeneity, the optimal contract that different people would choose if you got rid of asymmetric information about risk type, and each person faced prices that were actuarially fair for their risk type, may differ.

Then, rather than the obvious theoretical solution—the mandate—you're in the world that I love to think about, which is one of a potential empirical trade-off. On the one hand, market unraveling and allocative distortions due to adverse selection suggest that a mandate may be welfare-improving by counteracting the underinsurance that adverse selection generates. On the other hand, imposing a mandate, a one-size-fits-all policy, when some people would optimally choose different policies, may introduce its own allocative distortions.

So now we face an empirical question: Is the welfare lost from adverse selection in the unmandated market bigger or smaller than the welfare lost when imposing a uniform policy through a mandate? And if you are going to impose a uniform policy through a mandate, *which* policy should you mandate?

GEOGRAPHIC VARIATION IN HEALTH CARE SPENDING

Region: In a recent paper with Matthew Gentzkow and Heidi Williams, you analyze the source of large geographic differences in health care spending across the United States shown by the Dartmouth research—the fact that the average Medicare enrollee in Miami spends 80 percent more than his or her demographic counterpart in Minneapolis, for instance.[3]

Your goal was to understand whether those differences were driven by "supply" factors that might be amenable to policy interventions, like doctors' incentives or beliefs that could lead them to order excessive treatments, versus "demand" factors, such as patients in the higher-spending areas being less healthy or preferring more intensive care.

Would you describe this research a bit—your methods and findings?

Finkelstein: This is a very exciting new area of research for me. The work on insurance we've just been talking about is an area where—while there's obviously a lot more important work to be done—I'm starting to feel like I've gotten my head around the portion of it that I bit off to chew on 15 years ago.

But questions relating to the determinants of health care spending and practice are something I'm really just starting to think about. So it's a fun new area for me, and an extremely exciting research collaboration.

While it's new to me, the literature on the subject is, of course, quite rich already. There's a very well-known and influential body of work coming out of Dartmouth on geographic variations in spending, as you noted.

Region: It's gotten a lot of publicity.

Finkelstein: Yes, especially in the debates over the Affordable Care Act. A lot of the debate and much of the research and academic discussion have been: The fact is that high-spending places don't get better health outcomes. Does that mean you could cut spending? There's been a lot of good work on that.

Heidi, Matt and I came at this question from a different perspective, which is instead of asking what the *consequences* are of the geographic variation in spending, we tried to ask: What are its *causes?* We did this both because it's interesting in its own right and because different causes have potentially different implications for (a) whether we think the variation in spending is inefficient or not, and (b) if so, what policies would change things?

Matt has a previous paper in the *American Economic Review* with Bart Bronnenberg and Jean-Pierre Dubé (2012) that tries to understand differences in preferences for consumer brands by looking at how brand preferences change when people move across geographic areas with different consumer brand shares.

Coffee preferences, for example. I'm going to get the exact details wrong (I myself am not a coffee drinker), but in,

GEOGRAPHIC VARIATION IN HEALTH CARE SPENDING

We estimate that about half of the geographic variation in health care utilization reflects something "fixed" about the patient that stays with them when they move, and about half reflects something about the place. … One of our next steps is to get inside that and ask: What is it about the place? Is it doctors' beliefs? Is it doctors' past experience? Is it the number of MRI machines? We're going to investigate this by now looking at how doctors' practice styles change when the doctors move!

say, Miami, people tend to drink Folgers and in Minneapolis they drink Maxwell House. The point is there are large and persistent geographic differences in brand market shares for consumer products.

Sounds very similar to what the Dartmouth Atlas was showing for health care—large and persistent differences in practice patterns. And Matt and his co-authors have this *really* beautiful paper in which they try to look at the role of habit formation in explaining geographic variation in brand preferences: Is it that somehow what I used as a kid is what I stick with?

Well, if you think about what's going on in health care, the possibility that I stick with the style of treatment I get used to early in life has profound implications. It says that in dealing with rising health care spending, we're going to have a hard time changing anyone's current behavior; we have to change only new people's.

We started with the same idea as in Matt's previous paper: to look at people who moved geographically across areas with different patterns of health care utilization (i.e., high-utilization versus low-utilization areas) and whether their health care utilization changed. Originally, we were very focused on this issue of habit formation, which would suggest a very specific conceptual model and econometric specification.

But, as often happens with my projects, they don't go the way I expect. We found very clear patterns in the data on what happens when individuals move across areas that look nothing like what you'd expect in the type of habit formation model Matt and his co-authors had found for consumer brands.

With habit formation, what I did in the past affects me currently, although over time the importance of the past diminishes (depreciates, like capital does). In a model where habit formation is important, you would think if you moved from a high-spending place to a low-spending place, you'd be used to spending a lot on health care, so initially you would continue to do that. But over time, you might gradually reduce your health care spending as you adjust to the equilibrium in the new place.

Region: So if you went from Miami to Minneapolis, say, you'd reduce your spending eventually, but it would take some time.

Finkelstein: Exactly. They spend a lot per patient in Miami, but not in Minneapolis. So you would expect, in a model with habit formation, that maybe initially there wouldn't be much change in your health care utilization. But over time—whether it's because doctors would be urging you to do less or the people around you were like, "Why go to the doctor when you have a minor pain?"—you would gradually change your behavior toward the new norm.

But that's just not what we see at all. We have about 11 years of data on Medicare beneficiaries and about 500,000 of them who move across geographic areas. When they do, we see a clear, on-impact change: When you move from a high-spending to a low-spending place, or vice versa, you jump about 50 percent of the way to the spending patterns of the new place. But then your behavior doesn't change any further.

This is what I love about empirical research: I go into it with an idea—a question and an idea about the answer. But if I knew the answer, it wouldn't be fun to do it. And it certainly wouldn't be important if all we ever did was confirm our hypotheses. I have to have some idea to start, of course, but I often find myself radically rethinking it because it turns out just not to be right.

Region: And, in this case, you find that there's essentially 50 percent brand loyalty but a 50 percent shift toward the new location preference pretty much as soon as the person arrives.

Finkelstein: Yes. We estimate that about half of the geographic variation in health care utilization reflects something "fixed" about the patient that stays with them when they move, such as their health or their preferences for medical care. And about half of the geographic variation in health care utilization reflects something about the place, such as the beliefs and styles of the doctors there, or the availability of various medical technologies.

This gives you a very different perspective on how to think about the geographic variation in health care spending than the prior conventional wisdom that most of the geographic variation in the health care system was due to the supply side—that is, something about the *place* rather than the *patient*.

If we think the geographic variation is all due to supply side differences—they just practice differently in Miami than in Minneapolis—then you might start to think about policies designed to make high-spending Miami more like low-spending Minneapolis in order to reduce health care costs.

But if half of the geographic variation reflects the fact that people in Miami are sicker or have preferences for more intensive health care treatments than people in Minneapolis, you might think about such policies differently.

Region: So *some* part of it is amenable to policy that addresses the health care system, but perhaps less than previously thought.

Finkelstein: Sure. The glass half full is that about 50 percent of the geographic variation in health care spending is due to the supply side. And, relatedly, the fact that we don't find evidence of habit formation suggests that if you can figure out what policies can affect the provider side, those should have a relatively quick effect.

I mentioned that this research is what I hope is the beginning of a long and fruitful collaboration with Heidi and Matt. One of our next steps is to look at the 50 percent of the geographic variation that we've found is due to "something

about the place" and try to get inside that black box and ask: What is it about the place? Is it doctors' beliefs? Is it doctors' past experience? Is it the number of MRI machines? These are all things we want to look into. And, because we're a one-trick pony, we're going to investigate this by now looking at how *doctors'* practice styles change when the doctors move!

DYNAMIC INCENTIVES AND MORAL HAZARD

Region: You've been looking lately at how consumers respond to pricing for health care. Your January 2015 paper with Einav and Williams on marginal pricing responses—in breast cancer treatment specifically—is one example. But I'd like to ask you now about another recent piece that looks at dynamic incentives and moral hazard, investigating whether people consider future prices as well as current prices when making decisions about health care. Would you describe that work, including how you were able to find data on that current/future price distinction? [4]

Finkelstein: Liran Einav and I, together with several different co-authors, have now, I think, three related papers on this topic.[5] We're looking at the fact that health insurance contracts don't create *a* price for medical care; they create a nonlinear budget set.

Typically, you start off in a deductible range in which you pay dollar for dollar for your medical care. After you've spent a certain amount, you move into some cost-sharing range where maybe you pay 20 cents on the dollar for your medical care. And then, after you've spent enough, you hit some catastrophic, out-of-pocket maximum, at which point you pay essentially nothing for further medical care.

Now consider the classic health economics question: "How does consumer health care spending respond to the price of health care?" Well, *which* price?

DYNAMIC INCENTIVES AND MORAL HAZARD

This is something we do too little of in economics[:] a replication study within the original paper. We have the same basic design but two very different contexts and, in both, we find that people are forward-looking; i.e., they take the future price of care into account in making current medical decisions.

The first question in thinking about that is: What do consumers do when they're making health care decisions? Do they say, "Oh, today I'm in the deductible range and, gosh, if I go get my headache treated, I'm going to have to pay every dollar for that"? Or do they think, "Well, it's January and, yes, I'm in the deductible region, but I have chronic diabetes and I easily spend way past the deductible every year and end up in the cost-sharing arm at 20 cents on the dollar. So, really, the marginal price of my going to get my headache checked out in January is not dollar for dollar, it's 20 cents on the dollar." So, which way do they think?

Region: In essence, do they look at just the current price, or do they think about future costs as well in making a decision about what to do now?

Finkelstein: Exactly.

Region: Would you describe how you managed to tease out the data on that—finding a way to distinguish between decisions on *just* current price and those on current *plus* future prices? Your method was ingenious.

Finkelstein: That was really challenging for us, and a lot of fun to work on. To understand whether consumers look at only the current price of care or also take into account future costs, the ideal would be to randomly vary the future price of care (or the expected end-of-the-year price of care because contracts are annual), holding the current price—the spot price—constant. That's hard because most of the things you think of that would change the future price usually also change the spot price.

But we realized that most insurance contracts are specified annually; i.e., you return to the beginning (the deductible) part of the contract each January. Yet people can sometimes join a contract at different points in the year. That generates people in the same contract with the same initial price, but facing different future prices of care because they have different durations in the contract.

We found two different institutional settings where we could look at this. One was employer-provided health insurance. Plans are always specified annually, so the deductible is an annual deductible, and that deductible always resets January 1. But, obviously, people join firms throughout the year. So, what happens when you join in, say, September? Well, your annual deductible is going to reset in four months rather than in 12.

Now imagine someone who joins a firm in February, as opposed to September. They face the same initial or spot price of care. They both have a deductible, but they face a very different end-of-year price for care because one of them has much less time to go past the deductible.

Similarly, in Medicare Part D, which provides prescription drug coverage for the elderly, you can't join until you're 65. But people turn 65 in different months of the year and, again, it's an annual contract.

Region: So, in this one paper you look at both settings—employer-provided health insurance and Medicare Part D—and results from the second confirm those from the first.

Finkelstein: Yes, exactly. We were really excited. I think this is something we do too little of in economics. In some sense, we have a replication study within the original paper. We have the same basic design but two very different contexts and, in both, we find that people are forward-looking; i.e., they take the future price of care into account in making current medical decisions.

We then said, "OK, we tested the hypothesis that people are forward-looking. Now let's try to quantify it." Is it important? They could be forward-looking, but not very much. Or it might not matter because most health shocks occur once in the year or something. We looked into this in a separate paper with Paul Schrimpf.[6]

In the Part D context, we look at how people's drug purchases respond to the famed "donut hole"—that region of health care spending in which insurance suddenly becomes less generous on the margin and individuals have to pay dollar for dollar for their prescription drugs for a while. We examined what the effects will be of "filling the donut hole" in Part D—i.e., getting rid of that region where individuals face the full costs of their purchases—which is going to happen under the Affordable Care Act in 2020.

We see that a lot of the response is actually anticipatory. It's not just that people who end up in the donut hole spend more when you fill the donut hole and provide coverage in that region. Also, people who are *worried* about ending up in the donut hole and were therefore cutting back their spending earlier in the year to try to avoid reaching the donut hole are affected. In other words, when ACA covers the donut hole, we may find that is going to increase spending not just among people who end up in the donut hole, but also those who anticipate they will.

THE OREGON HEALTH INSURANCE EXPERIMENT

There's been a lot of conjecture that ... because Medicaid reimbursement rates to providers are so low, providers wouldn't want to treat Medicaid patients. ... Our findings reject this view. We find compelling evidence from a randomized evaluation that relative to being uninsured, Medicaid does increase use of health care.

THE OREGON HEALTH INSURANCE EXPERIMENT

Region: You've done a lot of important work recently on interactions between public policy and health, health care and health insurance. I'd like to ask in particular about your work on the Oregon Health Insurance Experiment.[7] What have you discovered about the impact of Medicaid funding on those who receive it? You looked at everything from emergency room use to employment. And could you begin with some background?

Finkelstein: The Oregon Health Insurance Experiment is a randomized evaluation of the impact of covering low-income uninsured adults with Medicaid. In 2008, the state of Oregon realized it had enough money to cover some but not all individuals with its Medicaid expansion program—a program that covers low-income uninsured adults who are not categorically eligible for Medicaid. That is, they are not in a specific eligibility category such as receiving disability insurance or cash welfare. Think of them as low-income but "able-bodied" adults.

So the state had to decide the fairest way to allocate a limited number of health insurance spots. State policymakers felt that first-come-first-served actually wasn't fair because it privileged people who had their act together, those who were more in the know, better informed. They decided that the fairest thing to do was to run a lottery.

We realized that this created an unprecedented opportunity for a randomized evaluation of the impact of Medicaid. The "we" is important here—this was a huge team effort. My co-principal investigator, Kate Baicker, at the Harvard School of Public Health, and I worked with a large team of researchers, including other academics as well as individuals in the state of Oregon.

We looked initially at the three major domains where you think health insurance might have an effect: health care use, financial security and well-being, and health. We looked at the impact of Medicaid in the first one to two years of coverage.

For health care use, we found across the board that Medicaid increases health care use: Hospitalizations, doctor visits, prescription drugs and emergency room use all increased. On the one hand, this is economics 101. Demand curves slope down: When you make something less expensive, people buy more of it. And what health insurance does, by design, is lower the price of health care for the patient.

On the other hand, there were ways in which these results were surprising. For Medicaid, in particular, there's been a lot of conjecture that while in general, health insurance would increase use of health care, that because Medicaid reimbursement rates to providers are so low, providers wouldn't want to treat Medicaid patients. There have been claims in the *Wall Street Journal* and other places that "Medicaid is worthless or worse than worthless." I read an article, I think it was in the *New York Times*, where someone said, "I have Medicaid, but it's a useless piece of plastic. I can't get in to

see a doctor."[8] Our findings reject this view. We find compelling evidence from a randomized evaluation that relative to being uninsured, Medicaid *does* increase use of health care.

Another result that some found surprising was on use of the emergency room. There had been claims in policy circles that covering the uninsured with Medicaid might get them out of the emergency room …

Region: Because people would have greater access to preventive care that might lower the need for ER visits.

Finkelstein: Right. And we *do* find that Medicaid increases doctor visits. And it increases preventive care. For example, we find that Medicaid increases mammogram rates by 60 percent. But when we look at the emergency room, we don't find that Medicaid decreases ER use. In fact, we find evidence of the opposite: We found that Medicaid increases emergency room use by 40 percent. That's a really big effect. And it occurs across the board: Whether you looked by patient demographics or by type of care—on-hour care, off-hour care, people who'd had a lot of previous ER visits versus people who hadn't—in every subgroup, we find that Medicaid increases ER use.

How do we understand these results? The point is, Medicaid doesn't just make the doctor free, it makes the emergency room free. And when something is cheaper, we expect people to use more of it. So that's one reason ER use should go up. The hope that ER use would go down comes from the belief that doctor visits are substitutes for the ER, so when the doctor also becomes free, you go to the doctor instead of the emergency room. Maybe this is the case (or maybe it isn't), but on net, our results show any substitution for the doctor that may exist is just not outweighed by the direct effect of making the emergency room free. On net, Medicaid increases use of the emergency room, at least in the first one to two years of coverage we are able to look at.

THE OREGON HEALTH INSURANCE EXPERIMENT

> First and foremost, health insurance is designed to provide financial security. Like fire insurance; fire insurance doesn't prevent your house from burning down. But if it does, insurance provides you with money to either rebuild your house or move to another house.
>
> It's the same with health insurance. It's nice if it improves your health, but its first purpose is to smooth your consumption so that when you have these big medical bills, you don't have to forgo food, housing, utilities, et cetera. … Our results show that the low-income uninsured do face out-of-pocket costs for medical care and that Medicaid substantially reduces this financial risk exposure.

The second set of results—which to me are the most important in the sense that they get too little attention in public policy discussion—is the basic economic or financial security aspect of health insurance. First and foremost, health insurance is a financial product. What it's designed to do is provide financial security. Like fire insurance; fire insurance doesn't prevent your house from burning down. But if it does, insurance provides you with money in exactly the state of the world in which you need resources to either rebuild your house or move to another house.

It's the same with health insurance. It's nice if it improves your health, but its first purpose is to smooth your consumption so that when you have these big medical bills, you don't have to forgo valuable food, housing, utilities, et cetera.

There was a question, though, with a very low-income population like ours, of whether there is really any financial risk of medical events, even when they are nominally uninsured. Maybe, in fact, all their care is paid for by ex ante or ex post charity care—that is, charity pays for their care before or after they get sick. However, our results show that the low-income uninsured do face out-of-pocket costs for medical care and that Medicaid substantially reduces this financial risk exposure. For example, we find that Medicaid virtually eliminates catastrophic, out-of-pocket medical spending. So it definitely has this financial security element.

The third set of results are the impacts on health. Here our findings on the impacts of Medicaid are more mixed. On mental health, we find substantial effects. We find that Medicaid lowers the risk of probability of screening positive for depression by 9 percentage points, or 30 percent off baseline. We also find that Medicaid improves self-reported health.

However, we did not detect statistically significant effects on the physical health measures we studied: blood sugar, cholesterol and blood pressure. Now, on the one hand, we picked those measures because they are things that clinical trials have shown are responsive to medical treatment within a short time frame unlike, say, weight loss, which is very hard to move around. So you might have expected them to have an effect. On the other hand, we're only looking one to two years out. Long-run effects could be different.

Another issue is that for some of the health measures, like blood sugar

(a marker of diabetes), our results just lack statistical precision. We found that Medicaid increases diabetes medication use. If you look in the clinical trial literature at what reduction in blood sugar you would expect given the increase in medication we saw, we can't rule out that that reduction in blood sugar may have occurred. We simply lack the statistical power to reach a conclusion here.

But for others of our health measures, the "null" findings are informative. For example, our results for blood pressure. There was earlier quasi-experimental work on Medicaid in the 1980s suggesting that Medicaid reduces hypertension. The confidence intervals on our estimate of the impact of Medicaid on blood pressure allow us to rule out the magnitudes found in the previous quasi-experimental literature. So I think here we update negatively on the likelihood that Medicaid will reduce hypertension.

Region: You also looked for impact on labor activity and found none.

Finkelstein: Yes, we found no impact on labor market activity on employment or earnings, and there we can rule out reasonably sized effects.

MEASURING THE WELFARE IMPACT

Region: In your recent paper with Luttmer and Hendren, you took an overall look at the welfare impact of Medicaid based on the Oregon Health Insurance Experiment results and found, I think, that the range was about 20 to 40 cents on a dollar of government expenditure in terms of direct benefit to a recipient.

That indicates that Medicaid may not be a worthwhile program, in a sense. But then you found a substantial *indirect* effect. Would you explain what you found there?

Finkelstein: This paper has been a long time in the works, and it's been very fun working with Nathan [Nathaniel Hendren] and Erzo [Erzo F. P. Luttmer] on it. In the papers I've written with the Oregon Health Insurance Experiment team of researchers, we deliberately steered away from trying to do any welfare analysis. The experimental results themselves are straightforward experimental analysis—clear, easy to explain, and (I think) very compelling.

Welfare analysis is much trickier and requires the researcher to make a number of assumptions. For example, how much do you value a statistical life? How risk averse are people? You can do a better or worse job on that—and Nathan and Erzo and I certainly tried our best!—but by necessity welfare analysis adds a layer of complexity and assumptions to the clear-cut empirical results. So we wanted to keep those distinct.

MEASURING THE WELFARE IMPACT

The nominally "uninsured" are not really completely uninsured. They have substantial implicit insurance. … A lot of spending on Medicaid is going to a set of people who, for want of a better term, we refer to as "external parties."… So, in terms of the total welfare impact of Medicaid, you have to grapple with the transfers Medicaid delivers to these providers of implicit insurance. Is the ultimate incidence to Medicaid recipients and their families? Does it accrue to hospital CEOs?

But then having been careful not to do any sort of casual, armchair welfare analysis in presenting the experimental results, it was very striking that the public didn't shy away from doing so. The media and the public policy world were eager to jump to welfare conclusions—often wildly different ones, depending on which results they focused on. Conclusions in the media based on the Oregon results ranged, for example, from "Medicaid makes a big difference" to "Medicaid doesn't actually help the poor."

So Erzo and Nathan and I asked: Can we say something more systematic and objective? And the first answer we came to is: It's hard because this is not a good that's traded in the market; it's a publicly provided good. Economists' standard way of doing welfare analysis is to look at demand.

Region: Right: "willingness to pay."

Finkelstein: Yes, but what's demand for [government-funded health insurance]? This isn't a traded good where individuals face prices for Medicaid and you can observe demand, or willingness to pay. So we take a variety of approaches and, in each one, we do a bunch of sensitivity analysis to the inevitable assumptions.

Region: And this is why you present a range of welfare estimates?

Finkelstein: Yes, and our central estimate is that the value of Medicaid to a recipient is about 20 to 40 cents per dollar of government expenditures. A priori, you might have thought it would be much larger than a dollar because there's a value to insurance, or it could have been smaller because of issues such as moral hazard.

The *other* key finding is that the nominally "uninsured" are not really completely uninsured. We find that, on average, the uninsured pay only about 20 cents on the dollar for their medical care. This has two important implications. First, it's a huge force working di-

rectly to lower the value of Medicaid to recipients; they already have substantial implicit insurance. This gives me a lot of confidence that our central welfare estimates of a value of Medicaid to recipients of about 20 to 40 cents per dollar of government spending are "real"—that is, they are not just driven by our inevitable assumptions, but are coming pretty directly from the data.

Second and, crucially, the fact that the uninsured have a large amount of implicit insurance is also a force saying that a lot of spending on Medicaid is not going directly to the recipients; it's going to a set of people who, for want of a better term, we refer to as "external parties." They're whoever was paying for that other 80 cents on the dollar.

Region: So, a relative, or the health care system itself.

Finkelstein: Right. And, in fact, there's a paper being presented here[9] tomorrow by Craig Garthwaite, Tal Gross and Matt Notowidigdo (2015) that suggests that a lot of the incidence of Medicaid is actually on uncompensated care by hospitals, so it's actually hospitals that serve the poor that benefit [from Medicaid].

Region: They write, "Each additional uninsured person costs local hospitals $900 each year in uncompensated care." That's a lot.

Finkelstein: Right, and I think that work is very complementary to ours. Matt is a co-author of mine in other work, and we have joked that it's good we wrote these two papers separately, because they complement each other so well. If we had written them together, we would have been accused of colluding!

The fact that so much of the health care costs of the "uninsured" are borne by people other than them is incredibly important for thinking about our welfare results. Welfare benefits to Medicaid *recipients* are only 20 to 40 cents per dollar of government spending, but whoever was providing the implicit insurance to the previously "uninsured" are also getting large benefits.

So, in terms of the *total* welfare impact of Medicaid, you have to grapple with the question of the ultimate economic incidence of the transfers Medicaid delivers to these providers of implicit insurance for the uninsured. Is the ultimate incidence to Medicaid recipients and their families? Does it accrue to hospital CEOs?

In some sense, our paper raises as many questions as it answers. The clear next step is to think about the ultimate economic incidence of these transfers to external parties. How much of it is accruing to low-income, sick individuals or their families? How much is it accruing higher up the income distribution?

Thinking about our own work and other related work over the last year, my view of what it means to be "uninsured" has changed. The "uninsured" are not as uninsured as we might have thought.

Now, that doesn't mean there aren't benefits to insurance. Some people respond to our results by saying, "This insurance isn't as valuable as real insurance. You might wait to go to the doctor," et cetera, et cetera. The results from the Oregon Health Insurance Experiment allow us to quantify these potential benefits. What are the health benefits of substituting this implicit insurance for formal insurance? What are the financial benefits?

EXPANDING RANDOMIZED CONTROLLED TRIALS

Region: One thing that makes the Oregon Health Insurance Experiment so valuable is that it is a randomized controlled trial (RCT). Would you discuss that aspect in particular?

Finkelstein: There have been literally hundreds of studies on the impact of Medicaid. I think the reason the Oregon Health Insurance Experiment gets a lot of attention in the media and in public policy speaks to the power and credibility of randomized controlled trials, not just in academia, but the broader public, which really understands and values it.

But the truth is, what we did in Oregon was not rocket science. And in my mind, that's a feature, not a bug. Unfortunately, one reason the Oregon experiment gets so much attention is that randomized trials on important domestic health policy questions are too rare.

Region: But you and your colleagues are addressing that—trying to expand their use in the United States.

Finkelstein: Yes. When I saw the attention the Oregon Health Insurance Experiment was getting, I realized that some of it is because it's a very exciting experiment and we hopefully did a good job analyzing it. But a lot of it, as I said, is because it's rare to have these randomized trials domestically on questions of how health care services are delivered.

And then I just looked down the hall at MIT, and my colleagues are running dozens of experiments around the world, through J-PAL, which is the Abdul Latif Jameel Poverty Action Lab. It was founded at MIT back in 2003, and J-PAL has been promoting randomized trials on a wide range of antipoverty programs.[10] They've had an enormous influence on changing the norms in the field of international development to doing more randomized evaluations and helping policymakers understand—and act on—the results.

They've been working in a host of countries for years, with regional offices around the world: J-PAL Africa, J-PAL Southeast Asia, J-PAL South Asia, J-PAL Latin America, J-PAL Europe. And it's like, "Gee, which continent is missing there? Not Antarctica, but North America."

So, two years ago, together with Larry Katz at Harvard, I founded J-PAL North America. It's J-PAL's newest regional

center, also based at MIT, and is designed to support, encourage and promote randomized trials on important domestic policy issues. Over the past two years, we've expanded J-PAL's network to include many of the leading academics who have been pioneering the use of randomized trials in the United States, across a wide range of sectors, like education, energy, housing or employment.

Region: You and Sarah Taubman (2015) just wrote a paper that makes a strong case for broadening the use of randomized controlled trials in U.S. health care delivery and suggests a number of ways to design RCTs to overcome cost and ethical issues that sometimes stand as obstacles. Could you tell us about that?

Finkelstein: When we looked at the data, we discovered that 80 percent of intervention studies on medicine in the United States are randomized.

Region: Is this just drug trials?

Finkelstein: No, not exclusively. And even if you leave out drugs, about two-thirds of medical intervention studies were randomized. This includes intervention studies on medical devices, surgical procedures, et cetera. Whereas, if you look at *health care* interventions, it's less than 20 percent.

Now, a lot of dollars and efforts are going into health care policy and issues of how we deliver health care, not just the medical side of it. So it seemed to us unfortunate that it's so rare.

Region: But you also discuss reasons for that scarcity, that there are objections to carrying them out having to do with ethics and cost. And you propose potential solutions to both of those problems. Would you elaborate on that?

Finkelstein: Sure. On the ethics side—that actually relates to what we were just talking about in Oregon—it's unethical to do a randomized trial when you know one policy or intervention is better than another, *and* you have the resources to give it to everyone.

Often in health care policy, there is equipoise; we don't actually know which form of health care delivery is better. But more to the point, even when we have a sense that Medicaid or something else helps people (even if we don't know exactly how or how much), resources are often very limited. The Oregon Health Insurance Experiment, as I said, came about for fairness reasons. Usually, policymakers running programs are constrained, for logistical reasons and often for financial reasons, so they're effectively rationing care, or rationing insurance.

EXPANDING RANDOMIZED CONTROLLED TRIALS

Unfortunately ... randomized trials on important domestic health policy questions are too rare. … Historically, randomized controlled trials on health care delivery have been conducted the way medical trials are done, which is extremely expensive, in terms of both time and money. …

To address the cost obstacle to RCTs in health care, another proposal was to realize the vast and largely untapped potential of administrative data, which allows for essentially costless follow-up on a census of individuals with extremely rich, detailed data.

Region: Oregon's policymakers had a limited budget and wanted to spend it wisely, and fairly.

Finkelstein: Right. They had to decide the fairest way to allocate a limited number of Medicaid spots.

Region: Various pundits have mocked such plans, referring to them as "gambling for health" or "health care lotteries."

Finkelstein: The truth behind that joke is that if you had the funding to cover everyone and you withheld it from half the people simply to run a research experiment, that clearly would be unethical. But if you're going to be allocating scarce spots in an ad hoc manner, why not make it *systematically* ad hoc?

Region: In addition to ethical concerns, there are often cost concerns about randomized controlled trials—not just the cost of intervention, but of the *research* itself.

Finkelstein: Historically, randomized controlled trials on health care delivery have been conducted the way medical trials are done, which is extremely expensive, in terms of both time and money. You individually recruit people, get their consent and then follow up through primary data collection of additional surveys. The follow-up is not only extremely expensive, but runs into methodological issues since you can't always find the people on follow-up, and it may be easier to find the people who are in the treatment arm of the experiment because you've been having more contact with them.

So to address the cost obstacle to RCTs in health care, another proposal that Sarah Taubman and I made in that *Science* piece, drawing on our experiences in Oregon, was to realize the vast and largely untapped potential of administrative data, which allows for essentially costless follow-up on a census

of individuals with extremely rich, detailed data.

I think there's real potential for more RCTs in U.S. health care delivery. J-PAL North America is working to help realize that potential in a number of ways. For one thing, we have some very generous funders who have given us money to allocate to researchers in our network who want to do RCTs to improve the efficiency of U.S. health care delivery. In addition, we do a lot of matchmaking. J-PAL North America staff have many conversations with practitioners who are trying to improve health care delivery—be they a health care system, a state government, an employer or an insurer—and learn which problems they want to solve. J-PAL staff then connect those practitioners with researchers who want to study these questions.

Beyond this, we provide support for researchers and practitioners so they don't have to reinvent the wheel for each study. We also create and share research "public goods," such as tips on how to design a study's recruitment and consent, examples of data use agreements and help with many of the other small hurdles that may otherwise delay or derail a promising research opportunity.

FUTURE WORK

The Affordable Care Act tried to slow the growth of health care spending. That's a much harder problem. There's both a lot of overuse of unnecessary procedures and a lot of underuse of low-cost, effective things. How do we design health care systems to efficiently deliver the care we think should be delivered?

METHODOLOGY

Region: You're somewhat unusual among economists in that you use both experiments and semistructural econometric techniques in your research, with great success. What guides your choice of one approach versus the other, and what are their relative merits?

Finkelstein: One of the fun things about research is getting to learn and use different techniques as appropriate. In terms of my use of both "reduced form" experimental techniques and "structural" techniques, the most important thing to emphasize is that I view these techniques as complements rather than substitutes.

So-called reduced form methods—be they literal experiments or quasi-experiments—are invaluable for providing transparent and compelling estimates of causal effects. But often the use of economic models and modeling techniques is important in translating the experimental "treatment effects" into economic objects that can be used out of sample.

I'll give you an example from the paper we were talking about earlier on prescription drug purchase decisions and health insurance contract design. Liran, Paul and I focused on Medicare Part D, the program for prescription drug insurance for the elderly, and especially on the famous "donut hole," where insurance suddenly becomes less generous on the margin, with people jumping from paying about 30 cents on the dollar to about 90 cents.

We show very clear visual evidence of a response to this increase in price: A graph of the distribution of annual drug spending shows that a lot of people "bunch" right at the donut hole—that is, they stop buying drugs once they enter the donut hole, where drugs suddenly become a lot more expensive for them. This is pretty compelling evidence that there *is* a behavioral response to insurance: When consumer cost-sharing goes up, people buy fewer drugs.

So, this reduced form evidence of "bunching" is useful in rejecting the null of no behavioral response to insurance in a simple and clear way. That's statistics jargon. When I say "rejecting the null of no response," I mean being able to reject a hypothesis that nothing occurs and conclude, rather, that there is a response to price change.

But we wanted to go beyond this "rejecting the null" to actually quantify the spending response. For example, we wanted to try to forecast how drug spending would respond to contracts we don't see in the data, such as the requirement under the Affordable Care Act that the donut hole be "filled in"—that is, that cost sharing *not* increase—by 2020. Well, to do that you need a model—both an economic model of behavior and a set of additional econometric assumptions to estimate it—that allows you to take the reduced form evidence of a behavioral response and use it to make predictions.

These approaches are complements, not substitutes. Without the bunching evidence, I wouldn't be confident that there is an underlying behavioral response. But without the additional modeling assumptions and estimation, I wouldn't know how to "use" that bunching in an economic sense.

FUTURE WORK

Region: Let's jump to the future. You're very active here at the National Bureau of Economic Research in health care and also with Raj Chetty on public economics.[11] What do you see as some of the most pressing issues in those two arenas and some of the promising research avenues?

Finkelstein: Well, there are many important and active areas of research in both public economics and health care. I won't pretend to cover them all. But I can

mention where I see my own research heading—which, by revealed preference, I presumably view as some of the most pressing and promising avenues!

As I mentioned, a lot of my work has been focused on insurance, particularly health insurance. There is naturally a lot more to learn here. But for myself, I feel like I'm starting to hit diminishing returns in that area. I feel myself pivoting—and it may be a subtle pivot to anyone except me—from health insurance to health care delivery: thinking about issues related to the efficiency of health care delivery, different organizational forms of health care delivery, different ways of designing health care systems.

I watched as an outsider—I was not involved in the policy process at all—the discussions around the Affordable Care Act. The act was intended to do two things. One is cover the uninsured, which we kind of know how to do. There are more or less efficient ways of doing it, and we know a lot about that now, thanks to a number of health economists who have done a lot of work on that question.

The other thing the Affordable Care Act tried to do is slow the growth of health care spending. That's a much harder problem. We think that there's both a lot of overuse of unnecessary procedures and a lot of underuse of low-cost, effective things. How do we design health care systems to efficiently deliver the care we think should be delivered and reduce use of the care we think is unnecessary?

At the micro level, I'm eager to start a bunch of randomized controlled trials to look at specific interventions that try to improve the efficiency of health care delivery. For example, we're working with Dr. Jeff Brenner and the Camden Healthcare Coalition to see if we can reduce hospital readmissions among super-utilizers of the health care system. We're doing another RCT with Mt. Sinai Healthcare System in New York City looking at whether clinical decision support software can help reduce overscanning. I'd like to do more studies like this!

In addition to studying the impact of particular interventions at the micro level, I'm also interested in thinking about questions that are more systemwide: How do we design public insurance and different types of incentive structures to try to get more efficient health care delivery? In other words, to try to get the market to adopt the most effective interventions and designs. These are hard questions! But hopefully we can make some progress.

What excites me about this whole set of questions on health care delivery is that it's an area that, to me, is at that sweet spot for research of being both an incredibly important set of issues and ones where we don't already know the answers.

There are areas of economics that are incredibly important and the policy world has not caught up, but where the *economists* are mostly in agreement on what the optimal solution is. But what's exciting to me about this work on health care delivery is, well, if you made me king of the world, I wouldn't actually know what we should do.

The constraints in health care delivery aren't just constraints of the political process; there are a lot of real intellectual constraints. There's a lot we don't yet know about how best to design these systems, and that makes it an extremely fun and exciting area to work in and to advise students in.

—*Douglas Clement*
July 23, 2015

More About Amy Finkelstein

Current Positions

Ford Professor of Economics, Massachusetts Institute of Technology, since 2012; Professor of Economics, 2008-12; Associate Professor of Economics, 2007-08; Assistant Professor of Economics, 2005-07

Co-Scientific Director, J-PAL North America, since 2013

Co-Director, Public Economics Program, National Bureau of Economic Research, since 2008; Research Associate, since 2007; Faculty Research Fellow, 2001-07; Visiting Scholar, Demography of Aging, 2001-02

Previous Positions

Visiting Professor of Economics, Booth School of Business, University of Chicago, 2010-11

Junior Fellow, Harvard Society of Fellows, 2002-05

Staff Economist, Council of Economic Advisers, Washington, D.C., 1997-98

Professional Affiliations

Associate Editor, *Journal of Economic Perspectives,* since 2014

Member, Executive Committee, American Economic Association, since 2013

Member, Panel of Health Advisers, Congressional Budget Office, since 2013

Study Section Member, Social Sciences and Population Studies, National Institutes of Health, since 2010

Fellow, TIAA-CREF Institute, since 2009

Honors and Awards

ASHEcon Medal, 2014

Arrow Award for Best Paper in Health Economics, iHEA, 2013

Fellow, Econometric Society, 2012

John Bates Clark Medal, 2012

American Academy of Arts and Sciences, elected 2012

Institute of Medicine, elected 2009

Presidential Early Career Award for Scientists and Engineers, 2009

Elaine Bennett Research Prize, 2008

Alfred P. Sloan Research Fellowship, 2007-09

Publications

More than three dozen research articles and working papers, particularly focused on health care spending, health insurance markets and policy analysis

Education

Massachusetts Institute of Technology, Ph.D., economics, 2001

Oxford University, M.Phil., economics, 1997

Harvard University, A.B., summa cum laude, government, 1995

For further background, visit
economics.mit.edu/faculty/afink

ENDNOTES

[1] "We make the bald assumption that individuals know their accident probabilities, while companies do not. Since insurance purchasers are identical in all respects save their propensity to have accidents, the force of this assumption is that companies cannot discriminate among their potential customers on the basis of their characteristics" (Rothschild/Stiglitz 1976, p. 623).

[2] See interview with Poterba, June 2008, *The Region*, online at minneapolisfed.org/publications/the-region/interview-with-james-poterba.

[3] See the Dartmouth Atlas of Health Care at dartmouthatlas.org/.

[4] See Aron-Dine, Cullen, Einav and Finkelstein (Forthcoming).

[5] See Aron-Dine, Einav and Finkelstein (2013); Aron-Dine, Cullen, Einav and Finkelstein (Forthcoming); Einav, Finkelstein and Schrimpf (2015).

[6] See Einav, Finkelstein and Schrimpf (2015).

[7] See nber.org/oregon and povertyactionlab.org/publication/insuring-uninsured.

[8] See Pear (2011).

[9] Interview was held at the NBER Summer Institute in Cambridge, Mass.

[10] See interview with Duflo, December 2011, *The Region*, online at minneapolisfed.org/publications/the-region/interview-with-esther-duflo.

[11] See interview with Chetty, December 2014, *The Region*, online at minneapolisfed.org/publications/the-region/interview-with-raj-chetty.

REFERENCES

Akerlof, George A. 1970. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84 (3): 488-500.

Aron-Dine, Aviva, Mark Cullen, Liran Einav and Amy Finkelstein. Forthcoming. Moral Hazard in Health Insurance: Do Dynamic Incentives Matter? *Review of Economics and Statistics.*

Aron-Dine, Aviva, Liran Einav and Amy Finkelstein. 2013. The RAND Health Insurance Experiment, Three Decades Later. *Journal of Economic Perspectives* 27 (1): 197-222.

Bronnenberg, Bart J., Jean-Pierre H. Dubé and Matthew Gentzkow. 2012. The Evolution of Brand Preferences: Evidence from Consumer Migration. *American Economic Review* 102 (6): 2472-2508.

Brown, Jeffrey R., Olivia S. Mitchell, James M. Poterba, and Mark J. Warshawsky. 1999. New Evidence on the Money's Worth of Individual Annuities. *American Economic Review* 89 (5): 1299-1318.

Chiappori, Pierre-André, and Bernard Salanié. 2000. Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy* 108 (1): 56-78.

Einav, Liran, Amy Finkelstein and Paul Schrimpf. 2015. The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D. *Quarterly Journal of Economics* 130 (2): 841-99.

Finkelstein, Amy, and Sarah Taubman. 2015. Randomize Evaluations to Improve Health Care Delivery. *Science* 347 (6223): 720-22.

Garthwaite, Craig, Tal Gross and Matthew J. Notowidigdo. 2015. Hospitals as Insurers of Last Resort. Working Paper 21290. National Bureau of Economic Research.

Pear, Robert. 2011. "Cuts Leave Patients With Medicaid Cards, but No Specialist to See," *New York Times,* April 1.

Rothschild, Michael, and Joseph Stiglitz. 1976. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *Quarterly Journal of Economics* 90 (4): 629-49.

Research Digest

Volatility and foreign assets

Increased saving due to economic uncertainty helps explain international debt patterns

When people are uncertain about the future, they save more: Precautionary savings is the term economists use for this phenomenon, and it's been used to explain lifetime saving patterns and financial booms and busts. Might it also be a factor in how much countries owe one another? That is, could people's saving behavior, surging when times are tough and shrinking when paychecks seem secure, explain the net foreign asset position of a country—the extent to which it is a creditor or debtor nation?

Minneapolis Fed economists Alessandra Fogli and Fabrizio Perri examine the relationship in a recent staff report, "Macroeconomic Volatility and External Imbalances" (SR 512, online at minneapolisfed.org), and determine that economic volatility "is an important determinant of the medium/long run evolution of external imbalances in developed countries." The key mechanism leading from one to the other is caution about the future.

"The intuition is simple," write Fogli and Perri. "In response to increases in domestic uncertainty agents increase their precautionary saving balances. Decreasing returns [and] increasing risk of domestic capital (arising from the increase in uncertainty) … imply that the bulk of the additional precautionary saving will go into foreign assets." Put briefly: Uncertain about their economic future, businesses and households spend less and save more; those higher savings flow, in part, into foreign assets.

The empirical picture

The report begins with a deep empirical analysis of the relationship between macro volatility and foreign asset positions in 20 OECD countries from 1970 to 2012. A cursory look at the 20 national graphs comparing net foreign assets and volatility hints at a strong positive link between the two. Japan and Belgium, for instance, have experienced substantial volatility over the past 40 years and become creditor nations as their net foreign assets (gross foreign assets minus gross foreign liabilities) have increased. Australia and the United Kingdom, in contrast, have experienced declining economic volatility and their net foreign assets have also declined. They are now international debtors.

The economists use regression analysis to study the relationship more carefully, accounting for factors other than uncertainty that have an effect on both volatil-

Research Digest

"The main take-away is that for OECD countries there is a robust, economically and statistically significant positive association, over the medium/long run, between changes in country specific volatility and changes in net foreign asset position."

PHOTOGRAPH BY STEVE NIEDORF

Fabrizio Perri and Alessandra Fogli

Research Digest

ity and foreign asset positions. "Country fixed effects" and "time fixed effects," for instance, measure national characteristics and global economic events, respectively, that might impact volatility and/or foreign asset positions. GDP growth is another factor they consider. If high growth periods coincide with low volatility periods, and quickly growing nations borrow internationally to finance further investment, that too would result in a correlation between volatility and net foreign assets. The economists include these and other factors in their regressions and find that "even after controlling for a very wide range of factors, the volatility of GDP growth is always significantly … associated with the net foreign asset position of a country."

Might the seeming relationship between the volatility and net foreign assets simply be an artifact of how the economists measure volatility or the specific sample they selected? Fogli and Perri investigate how robust their results are to a variety of volatility measures and several subsamples from their 20-country, four-decade data set. The association between volatility and net foreign assets remains strong. "The main take-away," conclude the economists, "is that for OECD countries there is a robust, economically and statistically significant positive association, over the medium/long run, between changes in country specific volatility and changes in net foreign asset position."

A model to test the mechanism

The next steps are to develop a model economy that includes the consumption, saving and investment behavior they hypothesize is the central mechanism linking volatility and foreign debt, and to then use that model to assess whether it generates quantitative results that are consistent with their real world data.

They use a standard one-good, two-country business cycle model. But, crucially, they extend it to allow for holding of foreign stocks and for business cycle volatility that varies over time. Thus modified, the model provides for a precautionary saving motive: As volatility shocks change, the desire to save also changes. "This naturally generates … external imbalances, with the more volatile country accumulating a net positive external position vis-à-vis the less volatile one," write Fogli and Perri. The model, therefore, "is a good laboratory to check whether precautionary saving motive can account for the observed association between volatility and imbalances."

Using their model, they focus first on how a volatility shock affects a country's net foreign asset position and find that, indeed, their model does replicate the hypothesized relation—that is to say, volatility does link to international imbalances via precautionary savings responses.

But can it do so faithfully? Can it generate the *quantitative* association they've documented empirically across four decades and 20 developed nations? And, perhaps more importantly, they use the model to assess the relative contribution of volatility shocks to net foreign asset positions compared to other applicable factors.

Their analysis shows, on the one hand, that the model with the precautionary savings motive as a key element generates results that are comparable to those seen in data. But "the coefficients in the model are lower than the ones estimated in the data, suggesting that some of the association between NFA [net foreign assets] and volatility in the data might be driven by factors not captured in our simple model." On the other hand, if volatility shocks are shut down in the model, the association seen in the data virtually disappears. "The main conclusion from this is that country specific shocks to volatility/uncertainty are quantitatively important determinants of the evolution of global imbalances among developed countries."

Summing up and extending

The analysis convinces Fogli and Perri that their central notion is valid: Volatility affects net foreign

Research Digest

assets via the precautionary savings motive. "More macro risk translates into more saving and more saving leads to accumulation of foreign assets," they write. "Macro uncertainty, as well as features shaping the precautionary motive, should be a major factor to consider when discussing the causes, the sustainability and desirability of observed global imbalances."

The authors offer several directions for future research. First, they suggest exploring the causes of changes in aggregate uncertainty; one "leading candidate," they note, is uncertainty about policy. Another direction: Modify the model so that uncertainty has a detrimental effect on growth as much literature suggests; such a change may well improve the model's explanatory power, they suspect. A third avenue of future research: Extend the empirical analysis from OECD countries to emerging markets, where high uncertainty may explain the now classic conundrum of low capital inflows despite high returns. The final extension they mention is consideration of "idiosyncratic" risk—that is, risk faced by an economy's individual actors—in addition to economywide, or aggregate, risk. "In the presence of large idiosyncratic risk," they observe, "even small increases in aggregate risk can have a large impact."

—Douglas Clement

Luigi Bocola

Risky business

How the threat of government default makes banks leery of lending, now and in the future

When governments roll the dice with debt, a shudder goes through financial markets—and the broader economy. In 2010-12, on the heels of the Great Recession, several eurozone countries with large fiscal deficits and/or debt suffered sovereign debt crises characterized by soaring interest rates on government debt, higher borrowing costs for households and firms, falling stock indexes and other asset prices, and faltering economic activity. (In Greece, a long-running debt crisis culminated last summer in a missed payment to the International Monetary Fund and nationwide bank closures.)

A standard explanation for these financial disruptions is that exposure to devalued government bonds

Research Digest

increases the cost of raising funds for banks; they pass along these higher costs to firms, raising interest rates and discouraging capital investment. This view was a motive for massive lending to banks by the European Central Bank (ECB) during the sovereign debt crisis.

But what if constrained liquidity isn't the only factor that leads banks to tighten credit when government debt gets out of hand? Recent research by Luigi Bocola, a former research economist at the Federal Reserve Bank of Minneapolis and an assistant professor at Northwestern University, suggests that heightened lending risk in the shadow of a potential sovereign default also plays a role.

In "The Pass-Through of Sovereign Risk" (Minneapolis Fed Working Paper 722, online at minneapolisfed.org), Bocola describes this additional source of credit tightening and quantifies its impact on lending behavior and economic output in Italy during the European debt crisis of 2010-11. Constructing a model economy in which bankers are mindful of the possibility of default, Bocola finds that a large share of the increase in borrowing costs during the debt crisis can be attributed to banks' belief that lending to firms has become riskier.

The economist also uses his model to evaluate the effectiveness of the ECB's efforts to maintain a cushion of liquidity for cash-

Constructing a model economy in which bankers are mindful of the possibility of default, Bocola finds that a large share of the increase in borrowing costs during the debt crisis can be attributed to banks' belief that lending to firms has become riskier.

strapped European banks. It turns out that the ECB's loan lifeline did little to improve credit conditions or increase economic output in Italy during the debt crisis.

Sand in the works

The bankers in Bocola's model conduct business as usual, collecting savings from households and using these funds, along with their own accumulated equity or capital, to invest in long-term government bonds and to lend to firms. But in the model, a microcosm of conditions in Italy during the sovereign debt crisis, there is a risk of the government defaulting on its debt. The mere prospect of this event throws sand into the works of financial intermediation.

Mathematical rules drawn from his model, and estimated from data on the exposure of Italian banks to sovereign debt risk during the crisis, govern the financial frictions in the model—the constraints imposed on banks when the risk of default increases.

In the model, a sharp rise in the likelihood of a government default (Bocola doesn't dwell on the reasons for this increased default risk) sets off a series of events that restricts access to bank credit. One reason for this credit squeeze is heightened expectations of default, which slash the value of government bonds. Banks saddled with low-value government securities face a decline in net worth. This raises the cost of borrowing from investors and households; as a result, banks raise interest rates and lending to firms drops.

Another source of credit tightening in the model—one that has received scant attention from other researchers—is the perceived increased risk of commercial lending as the threat of sovereign default looms. Even when banks have sufficient liquidity, they anticipate funding constraints in the future and worry that firms hurt by recession triggered by default may not repay their loans. Thus these loans are seen as riskier "because bankers now attach a higher likelihood to a state of the world (a sovereign default) in which [loans] pay out little precisely when bankers are in most need of wealth," Bocola writes.

As a precaution, banks charge firms higher interest as compensation for holding risky assets and sell

Research Digest

Sovereign risk increases firms' borrowing cost

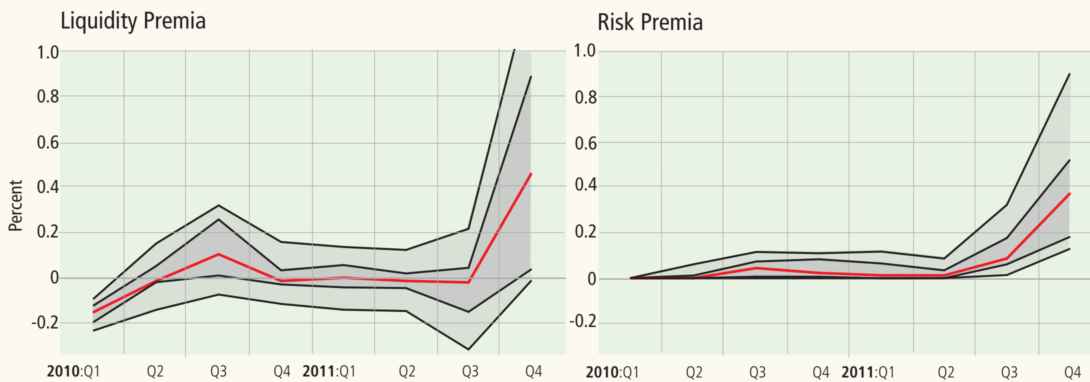

Shaded area = Probability intervals

Source: Federal Reserve Bank of Minneapolis calculations

some of their loans to reduce their exposure. Market demand for loan assets falls, and firms find credit harder to obtain.

By restricting access to credit, both mechanisms discourage capital spending by firms, with negative consequences for the Italian economy in the model.

The price of risk

Model in hand, Bocola uses it to conduct an experiment, comparing a sovereign-risk scenario with one in which there is no risk of government default. The results quantify the impact of the sovereign debt crisis on the borrowing costs of firms in Italy and on economic activity in the country. They also allow Bocola to isolate the effect of banks' precautionary motives on credit conditions, as distinct from current constraints on their liquidity.

In the model, the higher probability of an Italian default in the second half of 2011 leads to a marked increase in interest rates paid on loans by firms (see chart). The experiment also shows that bankers' concerns about the riskiness of firms contributed significantly to credit tightening in Italy. In the last quarter of 2011, this "risk channel" accounted for 45 percent of the increase in borrowing costs due to the sovereign debt crisis.

By reducing firms' ability to invest in production, higher interest rates sapped Italy's economic growth; the model predicts that in the first quarter of 2012, output would have been 1.4 percent higher without the financial drag induced by the debt crisis. Thus "a mere increase in the probability of a sovereign default depresses real economic activity," Bocola writes.

In response to the sovereign debt crisis, the ECB in early 2012 launched a lending program in Italy and in several other eurozone countries meant to prevent lending activity from seizing up. Bocola uses his model to assess the effects of these longer-term refinancing operations (LTROs), in which European banks borrowed hundreds of billions of euros from the ECB.

The LTROs were expected to lower interest rates and to spur increased economic output. Instead, in the model—reflecting the state

Research Digest

of the Italian economy at the time—these measures improved very little. Because banks' reluctance to lend was in large part based on aversion to risk, shoring up their liquidity was largely ineffective in easing credit conditions.

By highlighting the importance of firm risk in raising borrowing costs, Bocola's research offers lessons for policymakers weighing the best course of action to restore credit flows during debt crises. "If this mechanism is quantitatively important, policies that address the heightened liquidity problems of banks but do not reduce the increased riskiness of firms may prove ineffective in encouraging bank lending," he writes.

Ongoing research by Bocola, including a joint paper with Cristina Arellano, a senior research economist at the Minneapolis Fed, examines other aspects of sovereign debt crises, such as the role of self-fulfilling expectations and the impact of the European debt crisis on the private sector.

— Phil Davies

2014–2015 Student Essay Contest Economic Inequality

This spring the Minneapolis Fed held its 27th Annual Student Essay Contest, which is open to all high school students in the Ninth Federal Reserve District. The contest drew 269 essays from schools throughout the district. The winning essay is published here. Other top essays can be found at minneapolisfed.org under the Student Resources section of the Community & Education tab.

Thirty finalists each received $100. The third-place winner received an additional $200, and the second-place winner an additional $300. The first-place winner, Solomon Polansky of the Blake School in Minneapolis, received an additional $400 and was offered a paid summer internship at the Minneapolis Fed.

Economic inequality

Inequality takes many forms: racial, gender and political, to name a few. Among them, economic inequality—the unequal distribution of the national economic "pie" across different households—has gained a lot of attention in recent years. That is due in part to a wealth of new research on the topic that demonstrates that inequality has increased over the past generation.

Though other forms of inequality are important, contest entrants were asked to write specifically about economic inequality. And rather than debate the moral or political aspects of the question, they were asked to think like economists about the causes of economic inequality. Indeed, some essays focused on other forms of inequality as determinants of economic inequality.

Student Essay Contest Winner

An Analysis of the Impact of Technology on Income Inequality

Solomon Polansky
The Blake School
Minneapolis, Minn.

Luddite (n.): "broadly, one who is opposed to especially technological change."[1] Luddite finds its origin from a certain Ned Ludd, who smashed two knitting machines in early 19th century England to protest the developing frontier of technology and its effect on the workforce.[2] The Luddites' concerns are not without merit and remain relevant today in the United States. Over the past 30 years, U.S. productive output has soared while the number of labor hours has remained constant.[3] During this same time period, the top 1 percent of income earners doubled their percentage of income, while the bottom 90 percent fell from 70 percent to 60 percent.[4] Ongoing technological advances enable these productive strides, but also drive increasing income inequality by spawning two very distinct groups of winners and losers: those who benefit from technology, such as inventors of technology and workers whose productivity is enhanced by technological advance, and those who are negatively impacted through substitution of labor by technology.[5]

Inventors of new technology are the first to benefit from that new technology. In a free market, individuals are compensated based on the economic output of their factors of production. These factors of production include physical holdings (land, money) as well as intangibles (labor time, creativity). If an entrepreneur or inventor can successfully develop and market a desirable invention, the market will reward him/her by offering tremendous profits. Note that this unequal distribution of income is not necessarily a bad thing for the economy—in fact, the U.S. government openly supports new innovation by offering patents through the Patent and Trademark Office, thereby granting a (time-limited) legal monopoly (and the monopoly profits that follow).[6] But once an inventor earns these large incomes, the wealth inequality over others is unlikely to dissolve easily. There is a "snowballing effect on wealth distribution: top incomes are being saved at high rates, pushing wealth concentration [further] up," perpetuating the cycle of inequality.[7] While by no means will every inventor "strike gold" with his/her invention (in fact, most do not succeed), a skilled and lucky few will reap tremendous income; thus, propelling them into the highest echelon of income.[8] In short, "the people who benefit most are those with the expertise and creativity to use these advances."[9] And that drives both the incentive to invent and income inequality.

Skilled employees who use technology as a "tool" to increase their productivity also benefit. Consider highly skilled hedge-fund managers: These managers are already making a good income and would not be replaced with a computer (as of current technology) because they use human judgment to select investments. However, they become much more productive (and profitable for the firm) with the addition of computerized data and the skill to use it. Thus, their marginal revenue has increased, and the price the firm will be willing to pay, in salary, will also increase. These traders' incomes therefore increase with the addition of technology.[10] As technology is applied to skilled jobs (which are already high paying), the productivity of those workers increases and their income increases too, further extending the income inequality between skilled and unskilled laborers.[11]

However, not everyone benefits from advances in technology; laborers whose jobs can be substituted by technology are negatively affected. Businesses, by investing in capital such as new technology, will increase outputs while decreasing labor inputs (e.g., automation where purchasing a robot will replace a human worker). The Bureau of Labor Statistics reports that manufacturing employees' real output per hour increased from 51.2 units (which is proportional to dollars) per hour in 1990 to 110.3 in 2013; businesses produced 42 percent more output in 2013 than 1998.[12] However, the

total number of manufacturing workers actually decreased from 17.4 million in 1990 to 12.1 million in 2013.[13] A few skilled, knowledgeable employees are required to operate these advanced, high producing machines—in contrast to the hordes of unskilled laborers they replace.[14] In the early phases of technological development, it was largely simple manufacturing work being replaced by technology, as manufacturing firms sought to cut costs.[15] But now, with the advent of "big data" and analytical tools, even clerical work and professional services (both traditionally secure, white collar jobs) are being rendered obsolete by technology.[16] Technology leads companies to, inevitably, eliminate the workers whose labor has been replaced by a more efficient process in order to remain competitive in their markets. Thus, these workers' income has dropped to zero, forcing them into other lower-skill industries, such as food and restaurant services, that already have an ample supply of workers and thus driving wages downward.[17] Additionally, rapid globalization, enabled by advances in technology in transportation and communication, has opened up cheaper foreign labor markets for U.S. companies, further eroding the domestic manufacturing base.[18]

Applying technology to the economy thus creates both "winners" and "losers." It enables entrepreneurs and inventors, people with natural creativity and determination, to have the chance for great profits. It also increases the productivity (and therefore, income) of those whose "jobs are enhanced by machines"; these groups are the "winners."[19] However, technology eliminates the jobs of less-skilled (already lower-paid) workers by providing a more productive, albeit less "human," alternative and forcing workers into lower-paying service jobs; these workers are the "losers."[20] There is a clear schism widening between those benefiting and those being harmed by technology, and it is reflected in increasing income inequality. Ned Ludd was right to be concerned, and there is no easy answer to closing the gap. R

Endnotes

[1] Merriam-Webster.com. Accessed Nov. 23, 2014. merriamwebster.com/dictionary/luddite

[2] Porter, Eduardo. 2014. "Tech Leaps, Job Losses and Rising Inequality." *New York Times,* April 15. nytimes.com/2014/04/16/business/economy/tech-leaps-job-losses-and-rising-inequality.html?_r=0

[3] Sprague, Shawn. 2014. "What Can Labor Productivity Tell Us About the U.S. Economy?" *Beyond the Numbers* 3 (12). Bureau of Labor Statistics. bls.gov/opub/btn/volume-3/what-can-labor-productivity-tell-us-about-the-us-economy.htm

Workers in U.S. businesses worked the same number of labor hours (194 billion) in 2013 as in 1998, yet productive output increased 42 percent over that same time frame.

[4] Saez, Emmanuel, and Gabriel Zucman. 2014. "Wealth Inequality in the United States Since 1913: Evidence from Capitalized Income Tax Data." Working Paper 20625. National Bureau of Economic Research. gabriel-zucman.eu/files/SaezZucman2014.pdf

[5] This paper will address income inequality primarily. However, income inequality goes hand in hand with wealth inequality, as excess income allows one to invest in other capital, such as stocks and bonds, leading to the accumulation of wealth.

[6] U.S. Patent and Trademark Office. Accessed Nov. 21, 2014. uspto.gov/about/index.jsp

[7] Rotman, David. 2014. "Technology and Inequality." *MIT Technology Review.* Oct. 21. technologyreview.com/featuredstory/531726/technology-and-inequality/

[8] For a real world example of the potential for inventors, consider Bill Gates, founder of Microsoft. Rotman refers to these individuals as technology "superstars" who invent new technologies or generate new ideas for creative uses of technology.

[9] See Saez and Zucman.

[10] *The Economist.* 2012. "Who Exactly Are the 1%?" Jan. 21. economist.com/node/21543178

[11] Acemoglu, Daron. 2003. "Technology and Inequality." *NBER Reporter.* National Bureau of Economic Research. nber.org/reporter/winter03/technologyandinequality.html

[12] See Sprague.

[13] See Sprague.

[14] See Acemoglu.

[15] Katz, Richard, Robert Z. Lawrence and Michael Spence. 2011. "Manufacturing Globalization: The Real Sources of U.S. Inequality and Unemployment." *Foreign Affairs.* Council on Foreign Relations. November/December. foreignaffairs.org/articles/north-america/2011-11-01/manufacturing-globalization

[16] See Rotman.

[17] See Sprague. While the number of manufacturing jobs has decreased from 1990 to 2013, the number of food and restaurant service workers has increased from 6545.3 to 10487.1 (in thousands) during that same time period.

[18] Reich, Robert B. "How to Shrink Inequality." *The Nation.* May 26. thenation.com/article/179715/how-shrink-inequality

[19] See Porter.

[20] See Porter.

VIRTUAL FED

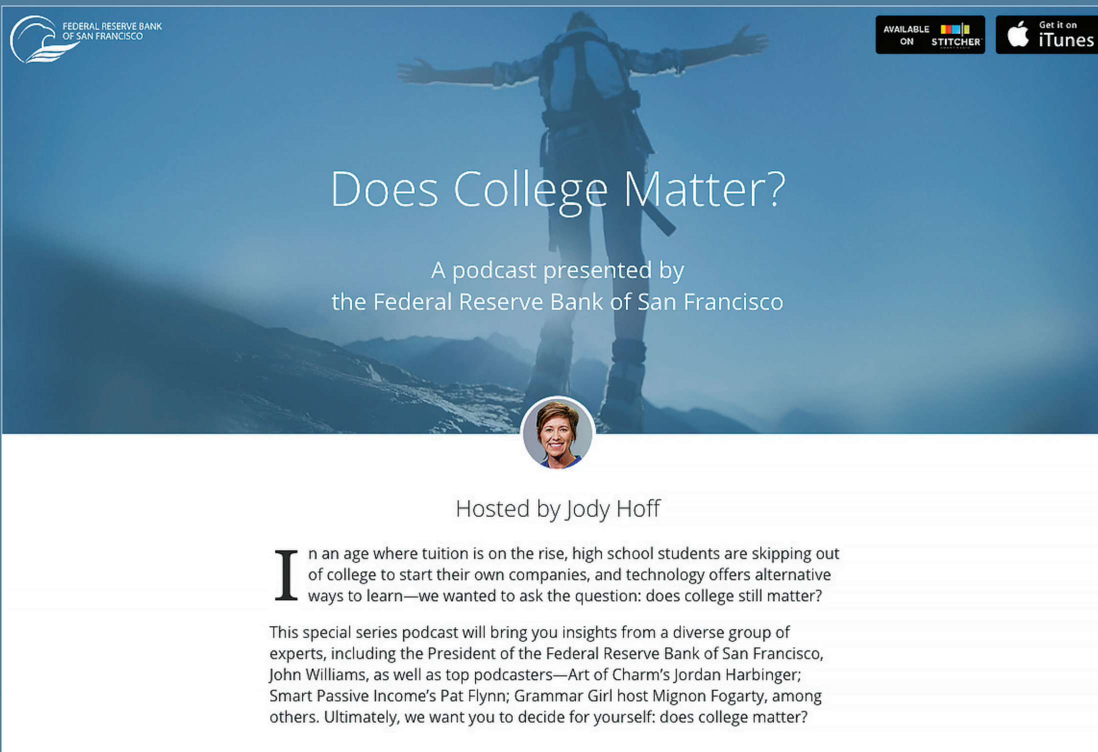

Collegecast

In recent years, the value of a college education has been re-examined by many (including *Region* writers). While the contribution of an advanced degree to earnings potential is clear, the rapid growth of tuition costs and the debt often required to cover them have made college less tempting for many potential students. Anecdotes about college dropouts earning millions at tech startups don't make the choice any easier.

Perhaps because of its location near the center of the information technology and social media universes, the Federal Reserve Bank of San Francisco has devoted a podcast series to examining this question. Episodes of "Does College Matter?" feature successful entrepreneurs—both those who dropped out of college and those who finished—and experts on the economics of education and labor markets from the San Francisco Fed. Rather than pushing a certain viewpoint, the series aims to present a variety of perspectives for listeners to consider as they decide for themselves.

Tune in at doescollegematter.org

—Joe Mahon