What Have We Learned From The Crisis?

Paul Krugman

September 2016

According to most chronologies, the global financial crisis began in July 2007, when BNP Paribas closed withdrawals from two of its funds, the modern equivalent of a bank shutting its doors. By early 2008 the financial panic had translated into a global recession; in September 2008 the failure of Lehman turned it into a free fall. And the aftershocks are still very much with us: although the free fall ended in mid-2009, growth rates thereafter were generally lower than growth pre-crisis, so the world economy has never made up the lost ground.

At this point, then, we're talking about an 8- or 9-year and counting episode, which is longer than the famous era of stagflation in the 1970s and early 1980s. The costs of the crisis and post-crisis slump were also much larger than those of the stagflation era, with steeper and more prolonged drops in income, more unemployment, more social and political disruption.

But here's a funny thing, striking to those of us of a certain age – that is, old enough to have already been studying or doing economics in the 70s. Stagflation had a huge impact on economic thinking, both at the level of academic research and on conventional wisdom among policymakers. The global financial crisis and the recession/stagnation that followed seem to have had much less impact. To a remarkable extent, economists and economic policymakers are still saying the same things in 2016 that they were saying in 2007. For some reason, there doesn't seem to be a clear consensus about what, if any, lessons we should draw from years of terrible economic performance.

Yet I would submit that there are some very important lessons for those willing to see them, and those lessons are what I want to talk about in this lecture.

I was tempted, when I began writing up my thoughts here, simply to present a checklist of things we have learned or should have learned since 2007. It seems to me, however, that it's helpful to put some more structure on the discussion, and I ended up with three main categories of things we should have figured out by now given the past 9 years' events.

First, we've seen a lot of vindication for old, unfashionable ideas – oldies but goodies that got deemphasized, and in some cases effectively blackballed, in the decades following the 1970s, but have turned out to be remarkably useful practical guides to policy and its effects in the post-crisis world.

Second, there have been some revelations about financial markets, especially the role of liquidity and the failure of arbitrage when you need it most, that have definitely changed how I see the world, and have important policy implications.

Third, we've made some important and uncomfortable discoveries about the politics and sociology of economics itself – about the resistance of both the economics profession and public officials to changing their views in the face of contrary information.

As you might expect, I will end this lecture with a plea for doing better. But let that wait; right now, I want to get into the substance of what went down, how that compared with what we should have expected, and what we should learn from the difference.

Imaginary James Tobin

There's a widespread impression out there that the crisis and aftermath have been devastating for the credibility of economists, proving that they know nothing. But my personal experience has been almost the opposite.

Here's a confession: Until the crisis struck, and especially given what happened in the next few years, I was always a bit unsure about my own bona fides. Obviously I'd been a professional success, but why? Was it truly because I'd been making a real contribution to our understanding of how the world works, or was I simply good at playing an academic game? I wasn't trying to fake it, but where was the clear, unambiguous demonstration that the models I deployed to interpret the world actually added value?

Then came the crisis and policy response, and there were several immediate questions in which popular intuitions and simple macroeconomic models were very much at odds. Would budget deficits cause interest rates to soar? Practical men said yes; economists, at least those of us with certain tools in our boxes, said no. Would huge increases in the monetary base cause runaway inflation? Yes, said practical men, politicians, and a few economists; no, said I and others of like mind. Would fiscal austerity depress output and employment? No, said many important people; on the contrary, it would be expansionary, because it would raise confidence. Yes, a lot, said Keynesian-minded economists.

And my team won three out of three. Gooooaaal!

OK, seriously: there was a fairly old-fashioned framework for macroeconomic thinking that some of us, at least, carried into the post-crisis landscape with quite a lot of success. How old-fashioned? Well, I think I've said in the past that economists from 1970 or so – that is, from before the intellectual revolution brought on by stagflation – might well have done a better job responding to the crisis than the economists we actually had on hand. So let me enlarge on that point by introducing a character I sometimes think of as "imaginary James Tobin," a pre-stagflation Keynesian who thought deeply about macroeconomic policy and financial markets. What would IJT have said about the post-2008 environment?

Tobin these days is mainly remembered for his magnificent work on financial markets and the determination of aggregate demand, of which more shortly. But I want to start with an argument he was widely perceived as having lost – the argument about aggregate supply, specifically about unemployment-inflation tradeoffs.

For Tobin was one of the last prominent holdouts against the Friedman-Phelps natural rate hypothesis, which said that there is no long-run tradeoff between unemployment and inflation.

Friedman, Phelps, and their followers argued that any attempt to hold unemployment persistently below the natural rate would lead to ever-accelerating inflation; and their models implied, although this is rarely stressed, that an unemployment rate persistently above the natural rate would lead to ever-declining inflation and eventually accelerating deflation.

Tobin was, however, skeptical. In his 1972 presidential lecture to the American Economic Association he took on the natural-rate hypothesis, arguing that the prediction of accelerating deflation from high unemployment was contradicted by evidence from the 1930s. He suggested that there is a basic reluctance on the part of firms and workers to cut nominal wages, and that given the inevitable churn and disequilibrium of labor markets, at any given time there are likely to be some workers whose equilibrium wages – but not, perhaps, their actual wages -- are declining in nominal terms. How frequent such episodes are will depend on the overall rate of inflation, and the result will be Phillips tradeoffs that persist in the long run, at least at low inflation.

I was a graduate student in the 1970s, and I remember the attitude toward Tobin's views on this issue in the age of rational expectations: it was basically dismissive, even among those who honored his contributions elsewhere. But we've now had multiple years of unemployment clearly above any notion of the natural rate in most advanced countries, and while inflation has fallen, it hasn't turned into runaway deflation anywhere. So Tobin was right about that. And he also seems to have been right about downward nominal wage rigidity: if you look at the distribution of nominal wage changes in depressed economies, like the case of Portugal shown in Figure 1 (chart courtesy of Olivier Blanchard), you see a large spike at zero.

Figure 1. Distribution of nominal wage changes in Portugal, in 1984 and in 2012

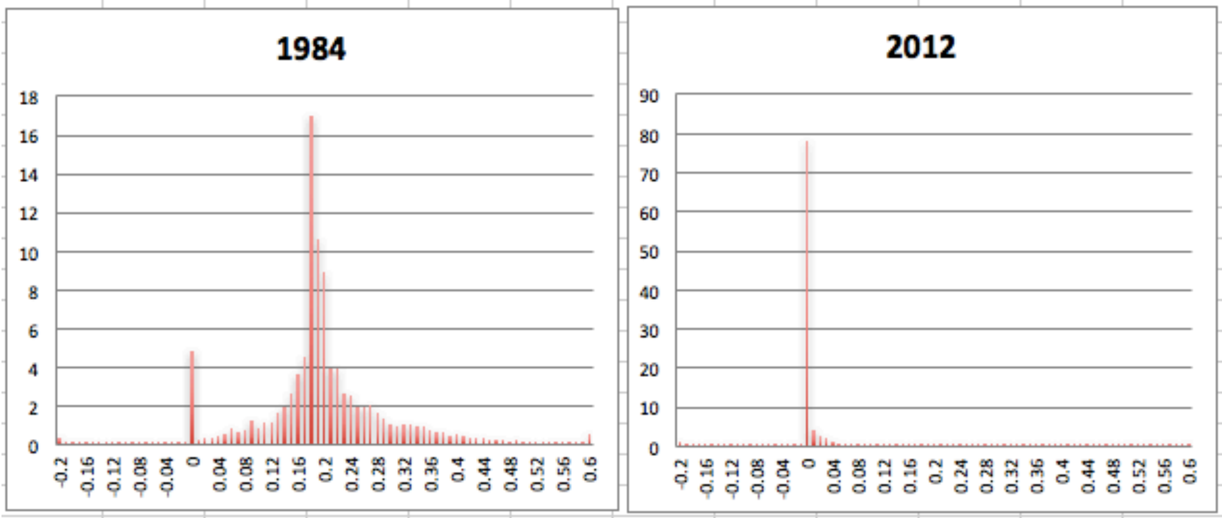

For reasons not completely persuasive to me, the standard response of macroeconomists to the failure of deflation to materialize seems to be to preserve the Friedman-Phelps type accelerationist Phillips curve, but then assert that expected inflation is "anchored", so that it ends up being an old-fashioned Phillips curve in practice. We can debate why, exactly, we're going this way. But what I don't think you can deny is that when it comes to inflation and aggregate supply, Tobin's 1972 last stand against the natural rate turns out to be a better guide to the post-2008 landscape than just about anything written in the 35 years that followed.

But let me move on to even bigger triumphs for the old-fashioned macroeconomics of imaginary James Tobin, which have come on the demand side.

The U.S. Federal funds rate hit zero in late 2008, with the economy still in a nosedive. The Fed responded with the first round of quantitative easing; later rounds would eventually lead to an almost 400 percent increase in the monetary base. Meanwhile, the budget deficit soared to heights never before seen in peacetime, mainly because of plunging revenues and increased mean-tested spending, but also to some extent because of deliberate fiscal stimulus. So what effect would these radically unusual policies have?

The answer from quite a few public figures was to predict soaring inflation and interest rates. And I'm not just talking about the goldbugs who infest TV business channels. Monetary economists like Allan Meltzer and Martin Feldstein warned about the coming inflation, joined by a Who's Who of the Republican establishment. Academics like Niall Ferguson and John Cochrane warned about massive crowding out of private investment.

But old-fashioned macro, with something like IS-LM at its base, offered startlingly contrary predictions at the zero lower bound. My imaginary James Tobin – OK, also me personally, not trying to emulate IJT -- predicted no inflationary impact of monetary expansion, in fact not much impact at all, because at the ZLB money and bonds are near-perfect substitutes; no crowding out, because the zero lower bound is also a regime of excess desired savings. And sure enough, inflation stayed low, as did interest rates.

IJT-style macro also made a prediction about the output effects of fiscal policy – namely, that it would have a substantial multiplier at the zero lower bound.

There was a lot of dispute about that proposition, of several kinds. Chicago's Cochrane insisted that the old-fashioned macro behind it had been "proved wrong." Robert Lucas denounced Christina Romer's use of multiplier analysis as "shlock economics," basing his argument on a garbled version of Ricardian equivalence. And, from a different perspective, Jean-Claude Trichet sunnily declared that warnings about the contractionary impact of austerity were "incorrect," because budget discipline would improve confidence.

A few years on, and the old-fashioned Keynesian analysis looks pretty good. There was a sort of natural experiment in the euro area, in which some countries were forced into severe austerity over the period 2009-2013, while others were not; Figure 2 shows what the outcome looks like.

Figure 2

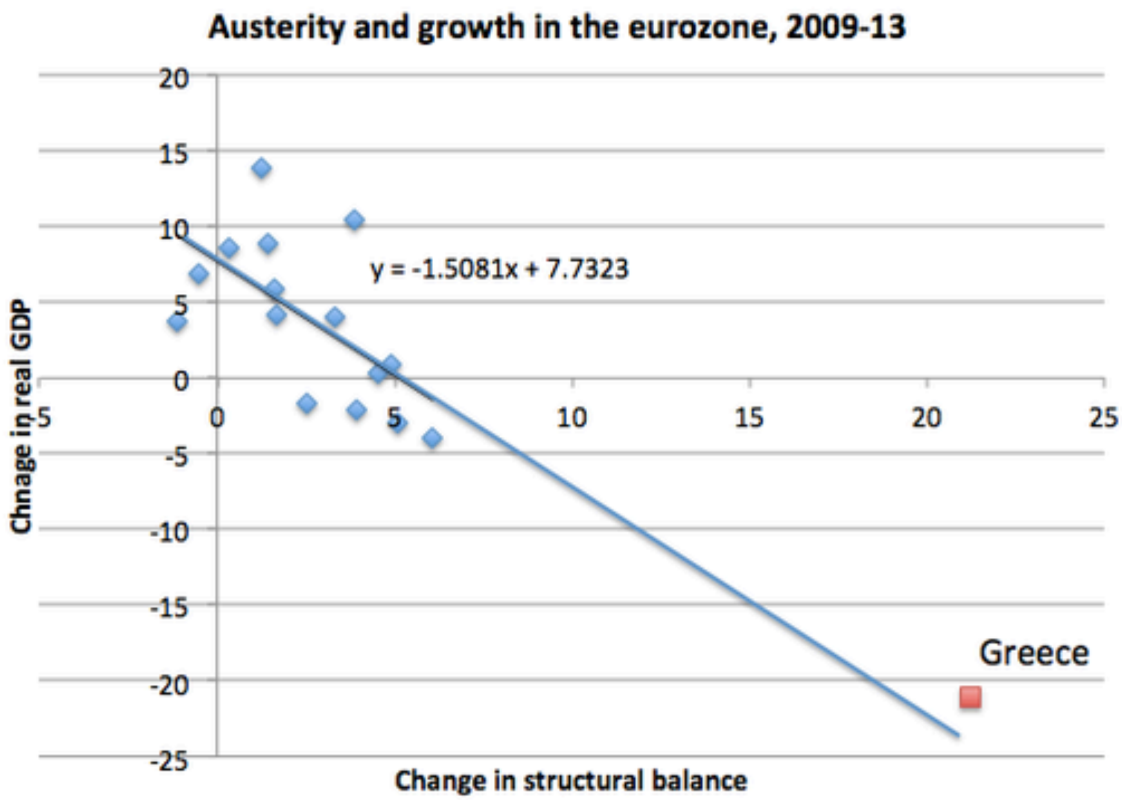

It sort of looks like a multiplier around 1.5, doesn't it? Which just happens to be the multiplier Christy Romer was assuming in her stimulus analysis.

Now, this wasn't a perfect natural experiment in the sense that there may have been common factors driving both austerity and economic contraction. But more sophisticated estimates of the multiplier, like Blanchard-Leigh or Nakamura-Steinsson, also seem to converge on a number around 1.5, which is pretty much what IJT analysis of an economy with large automatic stabilizers from a big public sector would have suggested.

But wait, we're not quite done. One aspect of the post-2008 story that apparently surprised many people, even smart economists like Martin Feldstein, was that huge increases in the monetary base didn't seem to produce much rise in broader monetary aggregates, leading to claims that something strange was going on – that maybe it was all because the Fed was paying interest on excess reserves. But the same thing happened in Japan in the early 2000s, without any special interest payments. And it was, in fact, completely predictable if you were aware of Tobin's 1960s work with William Brainard on his "general equilibrium" approach to monetary theory.

What was this approach? Basically, it was an application to asset markets, including the role of intermediates, of an approach similar to IS-LM: general equilibrium, yes, but with an ad hoc if plausible treatment of aggregate behavior. This approach told you right away that the volume of bank deposits, which are the non-cash component of broader money aggregates, was determined not by some mechanical multiplier but by incentives – and that in a liquidity trap just swapping monetary base for zero-interest securities would have no effect on these incentives, and hence no effect on deposits.

Oh, and this wasn't an ex post rationale: it's what those of us who knew the golden oldies were saying in advance. Me in 2009: "Central banks don't control the money supply, they only control the monetary base. Broad aggregates like M2 may well be unaffected by what the central bank does: increase the monetary base, and all that happens is an offsetting fall in the money multiplier."

The bottom line is that the crisis and its aftermath have actually provided a powerful vindication of macroeconomic models. Unfortunately for many economists, the models it vindicates are more or less vintage 1970. It's far from clear that anything later added to our ability to make sense of events, and developments in macro over the course of the 80s and after may even have subtracted value.

Whatever it takes

Ask macroeconomic theorists what we learned from the crisis, what their models have been missing, and you don't often hear that we need to relearn 1970-vintage macro. Yet it's pretty hard for DSGE theorists, even of the New Keynesian brand, to say with a straight face that their models worked well. So what you usually hear is that macro needs to incorporate the financial sector in a way it hasn't. Is this the right response?

I don't think so. It's true that banks and their role were underemphasized in the formal models, and that this should be better handled. But I don't see this as a key failing. As I'll explain in a minute, I would argue that (a) we had a fairly reasonable understanding of the logic of banking crises – our failure was more empirical than conceptual; and (b) the financial sector ended up being less central to the story than it might have seemed in 2009. The real conceptual surprises have come elsewhere.

So, let's talk briefly about banking crises in theory and practice. If I had to summarize the way I talked to myself during 2007, it went something like this: "This housing bust is going to be nasty, but banks are protected by deposit insurance. Wait – more than half the system is shadow banking? Yikes! Diamond-Dybvig!"

The point is that we had a pretty good story about how bank runs, even contagious bank runs, can happen, formalized in the famous Diamond-Dybvig paper. True, it wasn't integrated with the DSGE models that had come to dominate journal articles, but real-world oriented economists knew about it. If they didn't make allowances for a modern version of the early-30s banking crises, it was because they imagined that regulation and safety nets had contained that threat. What was missing was institutional understanding, the realization that new forms of finance had recreated the old risks by repackaging the functions of banking in forms that weren't regulated and protected.

And once that realization struck home, it took no time at all to recognize what we were seeing. I don't recall seeing anyone agonizing over the events of 2008, wondering how such things were possible given our models. On the contrary, it was more or less immediately obvious that we were seeing a new version of an old, fairly well-understood story.

Furthermore, while the financial sector played a central role in the hottest part of the crisis, over the medium term it's far from obvious that banking is the most important thing to focus on. For financial disruption was a big issue for a relatively short time, while economic troubles have gone on and on.

To see what I mean, look at any of the widely used measures of financial stress, like the St. Louis Fed stress index (Figure 3):

Figure 3

What you see in all cases is a severe but brief spike in 2008-9, then a quick return to normal conditions; yet recovery took a very long time in the US, and Europe went on to have a whole new set of problems. Maybe banking wasn't that central after all?

But if it wasn't banking, what was it? Well, we did have a hell of a housing bust: residential investment in the U.S. as a percentage of potential GDP fell by more than four points. Add in the effect of the lost wealth from plunging housing prices on consumer demand, and maybe an additional squeeze from deleveraging when the housing bust led to a reevaluation of what levels of indebtedness are acceptable. Given all this, it really isn't clear that banking is all that central to the story of the past 8 years.

Still, something important did happen during that period of financial stress, and again in Europe in 2011-2012. In both cases we saw prices of important classes of assets drop sharply in a short period of time, then recover with almost equal speed despite little obvious change in the fundamentals. I know these episodes aren't usually treated as closely related, but I think they are, and they both have big implications for policy.

Let's start with 2007-9. What people (myself included) often focus on when looking at that period is the TED spread – the difference between the rates at which banks lend to each other and the rate at which the U.S. government can borrow. This spread spiked as the extent of the subprime mess became apparent. And this is often taken as evidence that banks really doubted each others' solvency.

But here's the thing: lots of spreads, many of them not obviously tied to banks, spiked during the same period.

Figure 4 shows the difference between yields on 10-year U.S. government inflation-protected securities (TIPS) and ordinary nominal securities of the same maturity. That spread is normally negative, reflecting expected positive inflation, but it shot up to zero in the height of the crisis. Why?

Figure 4

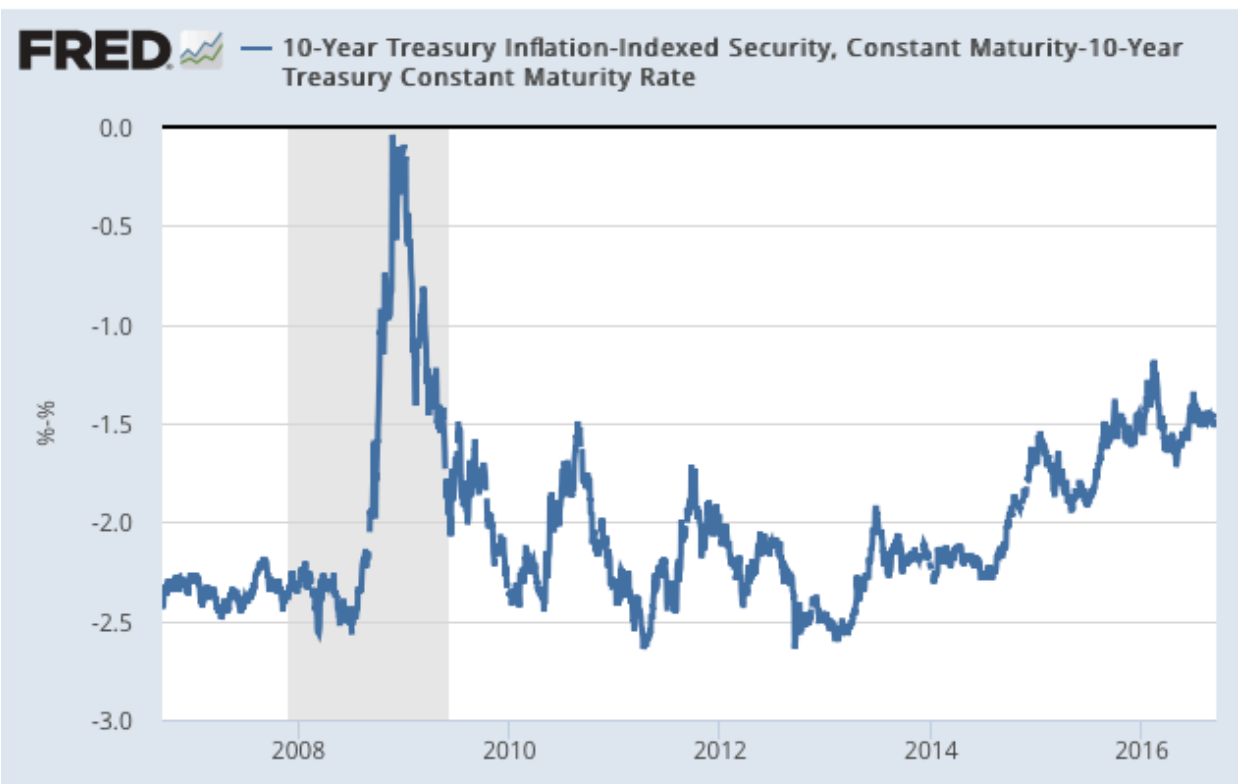

It wasn't about solvency, since these were two kinds of debt instruments issued by the same, ultra-solvent borrower. Was it a collapse of inflation expectations? Probably not: survey-based measures of expected inflation, like that collected by the Survey of Professional Forecasters, don't show anything like the apparent plunge. By all accounts, what happened was a drying up of liquidity for all but the most heavily traded assets – basically, everything except plain-vanilla Treasury securities ended up being sold at a huge discount.

But if that's what it was, why weren't investors rushing in to buy the underpriced assets? The answer, I think, is to ask, which investors?

Again, we have an existing model we can pull off the shelf to understand the issue, but one less familiar to macroeconomists than Diamond-Dybvig. Back in 1997 Shleifer and Vishny published an insightful paper, "The limits of arbitrage," which pointed out that the investors we often assume will step in to profit from obvious underpricing of assets are usually a small group of specialists with limited capital and substantial leverage. More than that: they are heavily invested in the very assets that are underpriced.

Their argument was, in part, that this reality means that really severe cases of mispricing can be self-fulfilling. Think of a fairly thinly traded security – for example, debt of a smallish developing country. The main investors in that debt are likely to be specialized intermediaries for whom it is a large part of their portfolio. If some event, or even a rumor, causes the price of that debt to plunge, those investors will also be plunged into financial distress, and therefore be unable to buy into the profit opportunity.

What I'm suggesting is that what happened between the fall of Lehman and the late spring of 2009 was a giant example of this phenomenon. The prices of just about everything except the most widely held, plain-vanilla securities – basically nominal U.S. government securities – plunged, creating widespread distress among leveraged investors of all kinds, and therefore preventing arbitrage that would correct this mispricing.

The virtue of this story, as I see it, is that it explains both how things got so bad and why financial conditions improved so rapidly in mid-2009. At the time, many people – myself included – were very skeptical about the Treasury's "stress tests" and all that. But by giving banks a relatively clean bill of health, even as the Fed made it clear that it would provide liquidity as needed, Treasury ended up breaking the vicious circle: prices of distressed securities began to rise, freeing leveraged investors to buy more of these securities, causing them to rise further.

And the same logic helps us to understand one of the most dramatic episodes of the ongoing crisis: the moment when Paul DeGrauwe saved the euro. OK, he had a bit of help from Mario Draghi.

In 2011-2012 it seemed very plausible that the whole euro system would be blown apart by speculation against southern European debt. Greece was a small player; but Italy and Spain combined account for around a third of the euro area economy. And in that scary period their soaring interest spreads against Germany seemed like the harbingers of imminent doom: markets seemed to have decided that they were fundamentally insolvent.

Yet DeGrauwe, in an analysis that deserves to be considered an instant classic, noted that Spain's finances looked no worse than those of Britain, which was paying very low rates. Why? He argued that it was really about liquidity: Britain, which retained its own currency, faced no risk of a cash crunch, while Spain, having joined the euro, was very much at risk of running out of money.

Figure 5

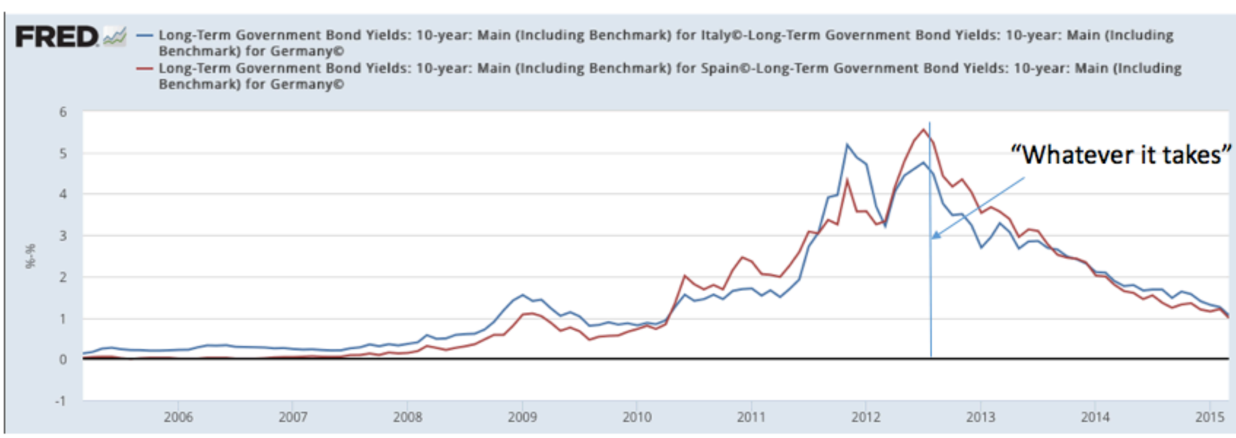

DeGrauwe suggested that Spain's situation gave rise to multiple equilibria – that investors were refusing to buy Spanish debt over fears that a cash crunch could push Spain into default, and that such a default might lead to repudiation of debt, which would end up justifying investors' fears. One might alternatively argue that a Shleifer-Vishny story was at work: the main buyers of Spanish and Italian debt were the countries' own banks, and plunging prices of that debt were pushing the banks into financial distress, forcing them to pull back even as the debt became greatly undervalued. (This was pretty much the gist of the "doom loop" diagnosis one heard at the time.)

Either way, the analysis suggested that relatively small actions by the ECB could salvage the situation – that all it would take was for the ECB to act as a lender of last resort, or even simply promise to do so. Sure enough, Mario Draghi said three words – "whatever it takes" – and all was well. OK, it wasn't quite that easy: Draghi had to back up those words with masterful political maneuvering, getting his board to accept in principle the idea of outright monetary transactions. But it was still a remarkable turnaround.

And my point is that it should be seen as fundamentally similar to the financial recovery that took place in the United States in spring 2009. In both cases we're talking about obvious asset mispricing that markets couldn't correct without an assist from central banks and other public institutions.

Which brings me back to the issue of how we do macroeconomics.

When macroeconomists talk about integrating the financial sector in DSGE models, my question is, what would that integration look like? Would it involve adding banks, but doing so in a way that preserves financial market efficiency except for the possibility of Diamond-Dybvig-type runs? If so, it wouldn't help much in making sense of how the crisis and aftermath actually played out.

On the other hand, making allowances for the possibility of liquidity crunches and failures of arbitrage would be really helpful. But can this even be done in the kind of macroeconomic frameworks we've been using for the last few decades? I really doubt it.

What looks useful is a sort of looser-jointed approach: ad hoc Hicks-Tobin-type models, with simple models of financial market failure on the side to be applied as seems necessary. For those seeking a definitive, integrated approach this will seem pitifully inadequate; and if I were a young academic seeking tenure I'd run away from all of this and either do empirical work or shun macro altogether. But models don't have to rigorously dot all i's and cross all t's – let alone satisfy the peculiar criteria that modern macro calls "microfoundations" --to be very useful in practice.

And let me further suggest that this loose-jointed framework suggests a broader rationale for policy activism than most macroeconomists – even self-proclaimed Keynesians – have generally offered in recent decades. Most of them – or, I guess I should say, most of us, since I was fairly comfortable with the Great Moderation policy consensus for a while – have seen the role for policy as pretty much limited to stabilizing aggregate demand. Correcting asset markets when they go wrong wasn't part of the mandate, because who were policymakers to claim that they were smarter than private investors?

Once we admit that there can be big asset mispricing due to liquidity and collateral constraints, however, the case for intervention becomes much stronger. It's still easy to see how this could go overboard, with officials deciding that any market results they don't like are the result of misperceptions or manipulation by the gnomes of Zurich. (I know, that was always a silly phrase – everyone knows that the gnomes are actually in Basel.) But a fixed belief in financial market efficiency would have ruled out both the successful stabilization of U.S. markets in 2009 and the Draghi stabilization of 2012. There is more potential for and power in intervention than was dreamed of in efficient-market models.

Delusions of progress?

Let me now turn to the question of what economists have learned about themselves and about economic policymakers as a result of the crisis. It is not a happy story.

Earlier I pointed out that events suggest that we need to revisit the Friedman-Tobin debates of the 1960s. The way we all learned the story was that it was a straightforward battle of ideas, decided by evidence: stagflation proved that Friedman was right and Tobin was wrong, and both the profession and policymakers adjusted accordingly. Macro models emphasized monetary policy over fiscal, and assumed a vertical long-run Phillips curve; policy focused on stabilization, not full employment, and turned to independent central bankers and Taylor rules to prevent inflationary license.

But in the light of events since 2008, the first part of the story looks all wrong. Actually, Tobin was right both about the limits of monetary policy and the long-run Phillips curve under low inflation.

And what about the second part, in which economists did what they are supposed to do, and adjusted their views in the light of compelling evidence? If that were how we actually behave, the long slump after 2008 should have produced changes in theory and practice comparable to those of 1970s stagflation. But it hasn't.

To be fair, some economists have shifted their views; and there's a lot of excellent empirically-oriented work that implicitly or explicitly acknowledges the realities of large fiscal multipliers at the zero lower bound, apparent downward nominal wage rigidity, and more.

But equilibrium macro theorists have shown themselves utterly unwilling to admit that anything they had been saying proved wrong. What we get instead are elaborate excuses – which is how, for example, I see the so-called "neo-Fisherian" models (Fisher himself would surely have considered them ridiculous) that claim that cutting interest rates is deflationary rather than inflationary.

At the level of supposedly policy-relevant economics, we also see a remarkable pattern of learning nothing. People who published articles in 2009 declaring “inflation is coming” just kept on publishing articles declaring that “inflation is coming,” with no hint of being chastened by their earlier errors.

And while the ice may be breaking a bit among actual policymakers, we have yet to see any major central bank admit that the 2 percent inflation target is wrong, or many finance ministers openly call for bigger deficits.

What all this suggests to me is that we need to rethink our account of intellectual history. Maybe the evidence in favor of a natural rate, against old-fashioned Keynesianism, wasn’t all that compelling; maybe stagflation seemed to have such a profound effect on thinking mainly because it provided a plausible cover story for a shift that many economists wanted to make for other reasons. For there was a clear rightward drift of politics, a shift toward free-market ideology, already underway in the 1970s. So perhaps stagflation was not so much a rude shock as an excuse for going all in on that ideology, without much second-guessing.

And conversely, the crisis of 2008 and its aftermath have taken place in an environment in which conservative ideology retains a powerful position in real-world politics and the academy alike. So relatively few economists or policymakers have been willing to reconsider their views despite overwhelming empirical refutation.

Or to put it another way, one thing we seem to have learned from the crisis is that many of our colleagues are less engaged in something like science, an attempt to understand the world as it is, than we would like to think. Instead, when they invoke evidence it's the way a drunkard uses a lamppost: for support, not illumination.

The best excuse one can offer is that even hard scientists are often reluctant to change their views – "Science progresses one funeral at a time," said Max Planck. But what I'm pointing out here isn't just that too few economists were willing to learn from the Great Recession, but that there's a notable contrast with the way the profession seized on the troubles of the 1970s. This asymmetry is what's troubling, and suggests that politics and ideology have distorted our field.

OK, at this point you're going to ask me for a solution. And I don't really have one, except to urge everyone who does or talks about economics to be a bit self-aware. Nobody is pure; everyone is tempted to read evidence as supporting what he or she wants to believe. But some people fight it; they make a conscious effort to avoid seeing what they want to see, they always ask, "Is this the evidence talking, or my preconceptions?"

And you want to be one of those people. If your initial reaction to the incredible and terrible events of the past 9 years is that they just show that you were right all along, consider how unlikely that is, and challenge yourself. If there’s any offsetting benefit to economic crisis, it is that it can be a learning experience. Let’s not waste that opportunity.