Oxford Review of Economic Policy, Volume 34, Numbers 1–2, 2018, pp. 156–168

Good enough for government work? Macroeconomics since the crisis

Paul Krugman*

Abstract: This paper argues that when the financial crisis came policy-makers relied on some version of the Hicksian sticky-price IS-LM as their default model; these models were 'good enough for government work'. While there have been many incremental changes suggested to the DSGE model, there has been no single 'big new idea' because the even simpler IS-LM type models were what worked well. In particular, the policy responses based on IS-LM were appropriate. Specifically, these models generated the insights that large budget deficits would not drive up interest rates and, while the economy remained at the zero lower bound, that very large increases in monetary base wouldn't be inflationary, and that the multiplier on government spending was greater than 1. The one big exception to this satisfactory understanding was in price behaviour. A large output gap was expected to lead to a large fall in inflation, but did not. If new research is necessary, it is on pricing behaviour. While there was a failure to forecast the crisis, it did not come down to a lack of understanding of possible mechanisms, or of a lack of data, but rather through a lack of attention to the right data.

Keywords: macroeconomic models, pricing behaviour, measurement and data

JEL classification: E10, E30

I. Introduction

It's somewhat startling, at least for those of us who bloviate about economics for a living, to realize just how much time has passed since the 2008 financial crisis. Indeed, the crisis and aftermath are starting to take on the status of an iconic historical episode, like the stagflation of the 1970s or the Great Depression itself, rather than that of freshly remembered experience. Younger colleagues sometimes ask me what it was like during the golden age of economics blogging, mainly concerned with macroeconomic debates, which they think of as an era that ended years ago.

Yet there is an odd, interesting difference, both among economists and with a wider audience, between the intellectual legacies of those previous episodes and what seems to be the state of macroeconomics now.

Each of those previous episodes of crisis was followed both by a major rethinking of macroeconomics and, eventually, by a clear victor in some of the fundamental debates. Thus, the Great Depression brought on Keynesian economics, which became

* City University of New York, e-mail: pkrugman@gc.cuny.edu
doi:10.1093/oxrep/grx052

the subject of fierce dispute—and everyone knew how those disputes turned out: Keynes, or Keynes as interpreted by and filtered through Hicks and Samuelson, won the argument.

In somewhat the same way, stagflation brought on the Friedman–Phelps natural rate hypothesis—yes, both men wrote their seminal papers before the 1970s, but the bad news brought their work to the top of the agenda. And everyone knew, up to a point anyway, how the debate over that hypothesis ended up: basically everyone accepted the natural rate idea, abandoning the notion of a long-run trade-off between inflation and unemployment. True, the profession then split into freshwater and saltwater camps over the effectiveness or lack thereof of short-run stabilization policies, a development that I think presaged some of what has happened since 2008. But I'll get back to that.

For now, let me instead just focus on how different the economics-profession response to the post-2008 crisis has been from the responses to depression and stagflation. For this time there hasn't been a big new idea, let alone one that has taken the profession by storm. Yes, there are lots of proclamations about things researchers should or must do differently, many of them represented in this issue of the *Oxford Review*. We need to put finance into the heart of the models! We need to incorporate heterogeneous agents! We need to incorporate more behavioural economics! And so on.

But while many of these ideas are very interesting, none of them seems to have emerged as *the* idea we need to grapple with. The intellectual impact of the crisis just seems far more muted than the scale of crisis might have led one to expect. Why?

Well, I'm going to offer what I suspect will be a controversial answer: namely, macroeconomics hasn't changed that much because it was, in two senses, what my father's generation used to call 'good enough for government work'. On one side, the basic models used by macroeconomists who either practise or comment frequently on policy have actually worked quite well, indeed remarkably well. On the other, the policy response to the crisis, while severely lacking in many ways, was sufficient to avert utter disaster, which in turn allowed the more inflexible members of our profession to ignore events in a way they couldn't in past episodes.

In what follows I start with the lessons of the financial crisis and Great Recession, which economists obviously failed to predict. I then move on to the aftermath, the era of fiscal austerity and unorthodox monetary policy, in which I'll argue that basic macroeconomics, at least in one version, performed extremely well. I follow up with some puzzles that remain. Finally, I turn to the policy response—and its implications for the economics profession.

II. The Queen's question

When all hell broke loose in financial markets, Queen Elizabeth II famously asked why nobody saw it coming. This was a good question—but maybe not as devastating as many still seem to think.

Obviously, very few economists predicted the crisis of 2008–9; those who did, with few exceptions I can think of, also predicted multiple other crises that didn't happen. And this failure to see what was coming can't be brushed aside as inconsequential.

There are, however, two different ways a forecasting failure of this magnitude can happen, which have very different intellectual implications. Consider an example from a different field, meteorology. In 1987 the Met Office dismissed warnings that a severe hurricane might strike Britain; shortly afterwards, the Great Storm of 1987 arrived, wreaking widespread destruction. Meteorologists could have drawn the lesson that their fundamental understanding of weather was fatally flawed—which they would presumably have done if their models had insisted that no such storm was even possible. Instead, they concluded that while the models needed refinement, the problem mainly involved data collection—that the network of weather stations, buoys, etc. had been inadequate, leaving them unaware of just how bad things were looking.

How does the global financial crisis compare in this respect? To be fair, the DSGE models that occupied a lot of shelf space in journals really had no room for anything like this crisis. But macroeconomists focused on international experience—one of the hats I personally wear—were very aware that crises triggered by loss of financial confidence do happen, and can be very severe. The Asian financial crisis of 1997–9, in particular, inspired not just a realization that severe 1930s-type downturns remain possible in the modern world, but a substantial amount of modelling of how such things can happen.

So the coming of the crisis didn't reveal a fundamental conceptual gap. Did it reveal serious gaps in data collection? My answer would be, sort of, in the following sense: crucial data weren't so much lacking as overlooked.

This was most obvious on the financial side. The panic and disruption of financial markets that began in 2007 and peaked after the fall of Lehman came as a huge surprise, but one can hardly accuse economists of having been unaware of the possibility of bank runs. If most of us considered such runs unlikely or impossible in modern advanced economies, the problem was not conceptual but empirical: failure to take on board the extent to which institutional changes had made conventional monetary data inadequate.

This is clearly true for the United States, where data on shadow banking—on the repo market, asset-backed commercial paper, etc.—were available but mostly ignored. In a less obvious way, European economists failed to pay sufficient intention to the growth of interbank lending as a source of finance. In both cases the institutional changes undermined the existing financial safety net, especially deposit insurance. But this wasn't a deep conceptual issue: when the crisis struck, I'm sure I wasn't the only economist whose reaction was not 'How can this be happening?' but rather to yell at oneself, 'Diamond–Dybvig, you idiot!'

In a more subtle way, economists were also under-informed about the surge in housing prices that we now know represented a huge bubble, whose bursting was at the heart of the Great Recession. In this case, rising home prices were an unmistakable story. But most economists who looked at these prices focused on broad aggregates—say, national average home prices in the United States. And these aggregates, while up substantially, were still in a range that could seemingly be rationalized by appealing to factors like low interest rates. The trouble, it turned out, was that these aggregates masked the reality, because they averaged home prices in locations with elastic housing supply (say, Houston or Atlanta) with those in which supply was inelastic (Florida—or Spain); looking at the latter clearly showed increases that could not be easily rationalized.

Let me add a third form of data that were available but largely ignored: it's fairly remarkable that more wasn't made of the sharp rise in household debt, which should have suggested something unsustainable about the growth of the 2001–7 era. And in the aftermath of the crisis macroeconomists, myself included (Eggertsson and Krugman, 2012) began taking private-sector leverage seriously in a way they should arguably have been doing before.

So did economists ignore warning signs they should have heeded? Yes. One way to summarize their (our) failure is that they ignored evidence that the private sector was engaged in financial overreach on multiple fronts, with financial institutions too vulnerable, housing prices in a bubble, and household debt unsustainable. But did this failure of observation indicate the need for a fundamental revision of how we do macroeconomics? That's much less clear.

First, was the failure of prediction a consequence of failures in the economic framework that can be fixed by adopting a radically different framework? It's true that a significant wing of both macroeconomists and financial economists were in the thrall of the efficient markets hypothesis, believing that financial overreach simply cannot happen—or at any rate that it can only be discovered after the fact, because markets know what they are doing better than any observer. But many macroeconomists, especially in policy institutions, knew better than to trust markets to always get it right—especially those who had studied or been involved with the Asian crisis of the 1990s. Yet they (we) also missed some or all of the signs of overreach. Why?

My answer may seem unsatisfying, but I believe it to be true: for the most part what happened was a demonstration of the old line that predictions are hard, especially about the future. It's a complicated world out there, and one's ability to track potential threats is limited. Almost nobody saw the Asian crisis coming, either. For that matter, how many people worried about political disruption of oil supplies before 1973? And so on. At any given time there tends to be a set of conventional indicators everyone looks at, determined less by fundamental theory than by recent events, and big, surprise crises almost by definition happen due to factors not on that list. If you like, it's as if meteorologists with limited resources concentrated those resources in places that had helped track previous storms, leading to the occasional surprise when a storm comes from an unusual direction.

A different question is whether, now that we know whence the 2008 crisis came, it points to a need for deep changes in macroeconomic thinking. As I've already noted, bank runs have been fairly well understood for a long time; we just failed to note the changing definition of banks. The bursting of the housing bubble, with its effects on residential investment and wealth, was conceptually just a negative shock to aggregate demand.

The role of household leverage and forced deleveraging is a bigger break from conventional macroeconomics, even as done by saltwater economists who never bought into efficient markets and were aware of the risk of financial crises. That said, despite the impressive empirical work of Mian and Sufi (2011) and my own intellectual investment in the subject, I don't think we can consider incorporating debt and leverage a fundamental new idea, as opposed to a refinement at the margin.

It's true that introducing a role for household debt in spending behaviour makes the short-run equilibrium of the economy dependent on a stock variable, the level of debt. But this implicit role of stock variables in short-run outcomes isn't new: after all, nobody has ever questioned the notion that investment flows depend in part on the existing capital stock, and I'm not aware that many macroeconomists consider this a difficult conceptual issue.

And I'm not even fully convinced that household debt played that large a role in the crisis. Did household spending fall that much more than one would have expected from the simple wealth effects of the housing bust?

My bottom line is that the failure of nearly all macroeconomists, even of the saltwater camp, to predict the 2008 crisis was similar in type to the Met Office failure in 1987, a failure of observation rather than a fundamental failure of concept. Neither the financial crisis nor the Great Recession that followed required a rethinking of basic ideas.

III. Not believing in (confidence) fairies

Once the Great Recession had happened, the advanced world found itself in a situation not seen since the 1930s, except in Japan, with policy interest rates close to zero everywhere. This raised the practical question of how governments and central banks should and would respond, of which more later. For economists, it raised the question of what to expect as a result of those policy responses. And the predictions they made were, in a sense, out-of-sample tests of their theoretical framework: economists weren't trying to reproduce the historical time-series behaviour of aggregates given historical policy regimes, they were trying to predict the effects of policies that hadn't been applied in modern times in a situation that hadn't occurred in modern times.

In making these predictions, the deep divide in macroeconomics came into play, making a mockery of those who imagined that time had narrowed the gap between saltwater and freshwater schools. But let me put the freshwater school on one side, again pending later discussion, and talk about the performance of the macroeconomists—many of them trained at MIT or Harvard in the 1970s—who had never abandoned their belief that activist policy can be effective in dealing with short-run fluctuations. I would include in this group Ben Bernanke, Olivier Blanchard, Christina Romer, Mario Draghi, and Larry Summers, among those close to actual policy, and a variety of academics and commentators, such as Simon Wren-Lewis, Martin Wolf, and, of course, yours truly, in supporting roles.

I think it's fair to say that everyone in this group came into the crisis with some version of Hicksian sticky-price IS-LM as their default, back-of-the-envelope macroeconomic model. Many were at least somewhat willing to work with DSGE models, maybe even considering such models superior for many purposes. But when faced with what amounted to a regime change from normal conditions to an economy where policy interest rates couldn't fall, they took as their starting point what the Hicksian approach predicted about policy in a liquidity trap. That is, they did not rush to develop new theories, they pretty much stuck with their existing models.

These existing models made at least three strong predictions that were very much at odds with what many influential figures in the political and business worlds (backed by a few economists) were saying.

– First, Hicksian macroeconomics said that very large budget deficits, which one might normally have expected to drive interest rates sharply higher, would not have that effect near the zero lower bound.
– Second, the same approach predicted that even very large increases in the monetary base would not lead to high inflation, or even to corresponding increases in broader monetary aggregates.
– Third, this approach predicted a positive multiplier, almost surely greater than 1, on changes in government spending and taxation.

These were not common-sense propositions. Non-economists were quite sure that the huge budget deficits the US ran in 2009–10 would bring on an attack by the 'bond vigilantes'. Many financial commentators and political figures warned that the Fed's expansion of its balance sheet would 'debase the dollar' and cause high inflation. And many political and policy figures rejected the Keynesian proposition that spending more would expand the economy, spending less lead to contraction.

In fact, if you're looking for a post-2008 equivalent to the kinds of debate that raged in the 1930s and again in the 1970s—a conflict between old ideas based on pre-crisis thinking, and new ideas inspired by the crisis—your best candidate would be fiscal policy. The old guard clung to the traditional Keynesian notion of a government spending multiplier somewhat limited by automatic stabilizers, but still greater than 1. The new economic thinking that achieved actual real-world influence during the crisis and aftermath—as opposed, let's be honest, to the kind of thinking found in this issue—mostly involved rejecting the Keynesian multiplier in favour of the doctrine of expansionary austerity, the argument that cutting public spending would crowd in large amounts of private spending by increasing confidence (Alesina and Ardagna, 2010). (The claim that bad things happen when public debt crosses a critical threshold also played an important real-world role, but was less a doctrine than a claimed empirical observation.)

So here, at least, there was something like a classic crisis-inspired confrontation between tired old ideas and a radical new doctrine. Sad to say, however, as an empirical matter the old ideas were proved right, at least insofar as anything in economics can be settled by experience, while the new ideas crashed and burned. Interest rates stayed low despite huge deficits. Massive expansion in the monetary base did not lead to inflation. And the experience of austerity in the euro area, coupled with the natural experiments created by some of the interregional aspects of the Obama stimulus, ended up strongly supporting a conventional, Keynesian view of fiscal policy. Even the magnitude of the multiplier now looks to be around 1.5, which was the number conventional wisdom suggested in advance of the crisis.

So the crisis and aftermath did indeed produce a confrontation between innovative new ideas and traditional views largely rooted in the 1930s. But the movie failed to follow the Hollywood script: the stodgy old ideas led to broadly accurate predictions, were indeed validated to a remarkable degree, while the new ideas proved embarrassingly wrong. Macroeconomics didn't change radically in response to crisis because old-fashioned models, confronted with a new situation, did just fine.

IV. The case of the missing deflation

I've just argued that the lack of a major rethinking of macroeconomics in the aftermath of crisis was reasonable, given that conventional, off-the-shelf macroeconomics performed very well. But this optimistic assessment needs to be qualified in one important respect: while the demand side of economy did just about what economists trained at MIT in the 1970s thought it would, the supply side didn't.

As I said, the experience of stagflation effectively convinced the whole profession of the validity of the natural-rate hypothesis. Almost everyone agreed that there was no long-run inflation–unemployment trade-off. The great saltwater–freshwater divide was, instead, about whether there were usable short-run trade-offs.

But if the natural-rate hypothesis was correct, sustained high unemployment should have led not just to low inflation but to continually declining inflation, and eventually deflation. You can see a bit of this in some of the most severely depressed economies, notably Greece. But deflation fears generally failed to materialize.

Put slightly differently, even saltwater, activist-minded macroeconomists came into the crisis as 'accelerationists': they expected to see a downward-sloping relationship between unemployment and the *rate of change* of inflation. What we've seen instead is, at best, something like the 1960s version of the Phillips curve, a downward-sloping relationship between unemployment and the *level* of inflation—and even that relationship appears weak.

Obviously this empirical failure has not gone unnoticed. Broadly, those attempting to explain price behaviour since 2008 have gone in two directions. One side, e.g. Blanchard (2016), invokes 'anchored' inflation expectations: the claim that after a long period of low, stable inflation, price-setters throughout the economy became insensitive to recent inflation history, and continued to build 2 per cent or so inflation into their decisions even after a number of years of falling below that target. The other side, e.g. Daly and Hobijn (2014), harking back to Tobin (1972) and Akerlof *et al.* (1996), invokes downward nominal wage rigidity to argue that the natural rate hypothesis loses validity at low inflation rates.

In a deep sense, I'd argue that these two explanations have more in common than they may seem to at first sight. The anchored-expectations story may preserve the outward form of an accelerationist Phillips curve, but it assumes that the process of expectations formation changes, for reasons not fully explained, at low inflation rates. The nominal rigidity story assumes that there is a form of money illusion, opposition to outright nominal wage cuts, that is also not fully explained but becomes significant at low overall inflation rates.

Both stories also seem to suggest the need for aggressive expansionary policy when inflation is below target: otherwise there's the risk that expectations may become unanchored on the downward side, or simply that the economy will suffer persistent, unnecessary slack because the downward rigidity of wages is binding for too many workers.

Finally, I would argue that it is important to admit that both stories are *ex post* explanations of macroeconomic behaviour that was not widely predicted in advance of the post-2008 era. Pre-2008, the general view even on the saltwater side was that stable inflation was a sufficient indicator of an economy operating at potential output, that any persistent negative output gap would lead to steadily declining inflation and eventually outright deflation. This view was, in fact, a key part of the intellectual case for inflation targeting as the basis of monetary policy. If inflation will remain stable at, say, 1 per cent even in a persistently depressed economy, it's all too easy to see how policymakers might give themselves high marks even while in reality failing at their job.

But while this is a subjective impression—I haven't done a statistical analysis of recent literature—it does seem that surprisingly few calls for a major reconstruction of macroeconomics focus on the area in which old-fashioned macroeconomics did, in fact, perform badly post-crisis.

There have, for example, been many calls for making the financial sector and financial frictions much more integral to our models than they are, which is a reasonable thing to argue. But their absence from DSGE models wasn't the source of any major predictive failures. Has there been any comparable chorus of demands that we rethink the inflation process, and reconsider the natural rate hypothesis? Of course there have been some papers along those lines, but none that have really resonated with the profession.

Why not? As someone who came of academic age just as the saltwater–freshwater divide was opening up, I think I can offer a still-relevant insight: understanding wage- and price-setting is hard—basically just not amenable to the tools we as economists have in our kit. We start with rational behaviour and market equilibrium as a baseline, and try to get economic dysfunction by tweaking that baseline at the edges; this approach has generated big insights in many areas, but wages and prices isn't one of them.

Consider the paths followed by the two schools of macroeconomics.

Freshwater theory began with the assumption that wage- and price-setters were rational maximizers, but with imperfect information, and that this lack of information explained the apparent real effects of nominal shocks. But this approach became obviously untenable by the early 1980s, when inflation declined only gradually despite mass unemployment. Now what?

One possible route would have been to drop the assumption of fully rational behaviour, which was basically the New Keynesian response. For the most part, however, those who had bought into Lucas-type models chose to cling to the maximizing model, which was economics as they knew how to do it, despite attempts by the data to tell them it was wrong. Let me be blunt: real business cycle theory was always a faintly (or more than faintly) absurd enterprise, a desperate attempt to protect intellectual capital in the teeth of reality.

But the New Keynesian alternative, while far better, wasn't especially satisfactory either. Clever modellers pointed out that in the face of imperfect competition the aggregate costs of departures from perfectly rational price-setting could be much larger than the individual costs. As a result, small menu costs or a bit of bounded rationality could be consistent with widespread price and wage stickiness.

To be blunt again, however, in practice this insight served as an excuse rather than a basis for deep understanding. Sticky prices could be made respectable—just—allowing modellers to assume something like one-period-ahead price-setting, in turn letting models that were otherwise grounded in rationality and equilibrium produce something not too inconsistent with real-world observation. New Keynesian modelling thus acted as a kind of escape clause rather than a foundational building block.

But is that escape clause good enough to explain the failure of deflation to emerge despite multiple years of very high unemployment? Probably not. And yet we still lack a compelling alternative explanation, indeed any kind of big idea. At some level, wage and price behaviour in a depressed economy seems to be a subject for which our intellectual tools are badly fitted.

The good news is that if one simply assumed that prices and wages are sticky, appealing to the experience of the 1930s and Japan in the 1990s (which never experienced a true deflationary spiral), one did reasonably well on other fronts.

So my claim that basic macroeconomics worked very well after the crisis needs to be qualified by what looks like a big failure in our understanding of price dynamics—but this failure didn't do too much damage in giving rise to bad advice, and hasn't led to big new ideas because nobody seems to have good ideas to offer.

V. The system sort of worked

In 2009 Barry Eichengreen and Kevin O'Rourke made a splash with a data comparison between the global slump to date and the early stages of the Great Depression; they showed that at the time of writing the world economy was in fact tracking quite close to the implosion that motivated Keynes's famous essay 'The Great Slump of 1930' (Eichengreen and O'Rourke, 2009)

Subsequent updates, however, told a different story. Instead of continuing to plunge as it did in 1930, by the summer of 2009 the world economy first stabilized, then began to recover. Meanwhile, financial markets also began to normalize; by late 2009 many measures of financial stress were more or less back to pre-crisis levels.

So the world financial system and the world economy failed to implode. Why?

We shouldn't give policy-makers all of the credit here. Much of what went right, or at least failed to go wrong, reflected institutional changes since the 1930s. Shadow banking and wholesale funding markets were deeply stressed, but deposit insurance still protected a good part of the banking system from runs. There never was much discretionary fiscal stimulus, but the automatic stabilizers associated with large welfare states kicked in, well, automatically: spending was sustained by government transfers, while disposable income was buffered by falling tax receipts.

That said, policy responses were clearly much better than they were in the 1930s. Central bankers and fiscal authorities officials rushed to shore up the financial system through a combination of emergency lending and outright bailouts; international cooperation assured that there were no sudden failures brought on by shortages of key currencies. As a result, disruption of credit markets was limited in both scope and duration. Measures of financial stress were back to pre-Lehman levels by June 2009.

Meanwhile, although fiscal stimulus was modest, peaking at about 2 per cent of GDP in the United States, during 2008–9 governments at least refrained from drastic tightening of fiscal policy, allowing automatic stabilizers—which, as I said, were far stronger than they had been in the 1930s—to work.

Overall, then, policy did a relatively adequate job of containing the crisis during its most acute phase. As Daniel Drezner argues (2012), 'the system worked'—well enough, anyway, to avert collapse.

So far, so good. Unfortunately, once the risk of catastrophic collapse was averted, the story of policy becomes much less happy. After practising more or less Keynesian policies in the acute phase of the crisis, governments reverted to type: in much of the advanced world, fiscal policy became Hellenized, that is, every nation was warned that it could become Greece any day now unless it turned to fiscal austerity. Given the validation of Keynesian multiplier analysis, we can confidently assert that this turn to austerity contributed to the sluggishness of the recovery in the United States and the even more disappointing, stuttering pace of recovery in Europe.

Figure 1 sums up the story by comparing real GDP *per capita* during two episodes: Western Europe after 1929 and the EU as a whole since 2007. In the modern episode, Europe avoided the catastrophic declines of the early 1930s, but its recovery has been so slow and uneven that at this point it is tracking below its performance in the Great Depression.

Now, even as major economies turned to fiscal austerity, they turned to unconventional monetary expansion. How much did this help? The literature is confusing enough to let one believe pretty much whatever one wants to. Clearly Mario Draghi's 'whatever it takes' intervention (Draghi, 2012) had a dramatic effect on markets, heading off what might have been another acute crisis, but we never did get a clear test of how well outright monetary transactions would have worked in practice, and the evidence on the effectiveness of Fed policies is even less clear.

The purpose of this paper is not, however, to evaluate the decisions of policy-makers, but rather to ask what lessons macroeconomists should and did take from events. And the main lesson from 2010 onwards was that policy-makers don't listen to us very much, except at moments of extreme stress.

This is clearest in the case of the turn to austerity, which was not at all grounded in conventional macroeconomic models. True, policy-makers were able to find some economists telling them what they wanted to hear, but the basic Hicksian approach that did pretty well over the whole period clearly said that depressed economies near the zero lower bound should not be engaging in fiscal contraction. Never mind, they did it anyway.

Figure 1: Slumps and recoveries in two crises, real GDP *per capita*, pre-crisis = 100

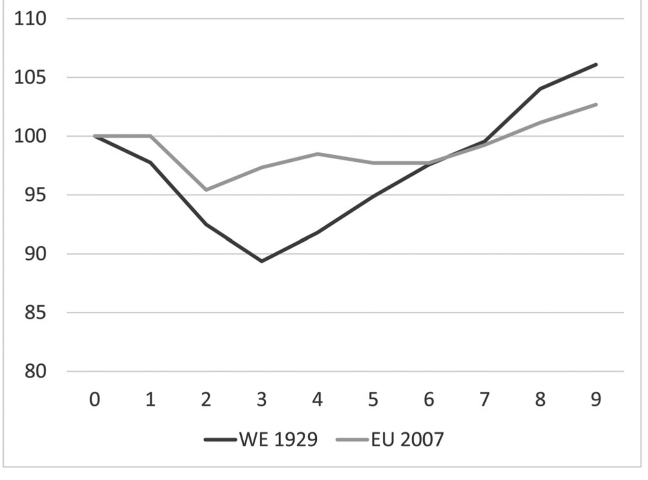

Note: The x-axis shows the number of years after each crisis.

Even on monetary policy, where economists ended up running central banks to a degree I believe was unprecedented, the influence of macroeconomic models was limited at best. A basic Hicksian approach suggests that monetary policy is more or less irrelevant in a liquidity trap. Refinements (Krugman, 1998; Eggertsson and Woodford, 2003) suggested that central banks might be able to gain traction by raising their inflation targets, but that never happened.

The point, then, is that policy failures after 2010 tell us relatively little about the state of macroeconomics or the ways it needs to change, other than that it would be nice if people with actual power paid more attention. Macroeconomists aren't, however, the only researchers with that problem; ask climate scientists how it's going in their world.

Meanwhile, however, what happened in 2008–9—or more precisely, what didn't happen, namely utter disaster—did have an important impact on macroeconomics. For by taking enough good advice from economists to avoid catastrophe, policy-makers in turn took off what might have been severe pressure on economists to change their own views.

VI. That 80s show

Why hasn't macroeconomics been transformed by (relatively) recent events in the way it was by events in the 1930s or the 1970s? Maybe the key point to remember is that such transformations are rare in economics, or indeed in any field. 'Science advances one funeral at a time,' quipped Max Planck: researchers rarely change their views much in the light of experience or evidence. The 1930s and the 1970s, in which senior economists changed their minds—e.g. Lionel Robbins converting to Keynesianism, were therefore exceptional.

What made them exceptional? Each case was marked by developments that were both clearly inconsistent with widely held views and sustained enough that they couldn't be written off as aberrations. Lionel Robbins published *The Great Depression*, a very classical/Austrian interpretation that prescribed a return to the gold standard, in 1934. Would he have become a Keynesian if the Depression had ended by the mid-1930s? The widespread acceptance of the natural-rate hypothesis came more easily, because it played into the neoclassical mindset, but still might not have happened as thoroughly if stagflation had been restricted to a few years in the early 1970s.

From an intellectual point of view, I'd argue, the Great Recession and aftermath bear much more resemblance to the 1979–82 Volcker double-dip recession and subsequent recovery in the United States than to either the 1930s or the 1970s. And here I can speak in part from personal recollection.

By the late 1970s the great division of macroeconomics into rival saltwater and freshwater schools had already happened, so the impact of the Volcker recession depended on which school you belonged to. But in both cases it changed remarkably few minds.

For saltwater macroeconomists, the recession and recovery came mainly as validation of their pre-existing beliefs. They believed that monetary policy has real effects, even if announced and anticipated; sure enough, monetary contraction was followed by a large real downturn. They believed that prices are sticky and inflation has a great deal of inertia, so that monetary tightening would produce a 'clockwise spiral' in unemployment and inflation: unemployment would eventually return to the NAIRU (non-accelerating inflation rate of unemployment) at a lower rate of inflation, but only after a transition period of high unemployment. And that's exactly what we saw.

Freshwater economists had a harder time: Lucas-type models said that monetary contraction could cause a recession only if unanticipated, and as long as economic agents couldn't distinguish between individual shocks and an aggregate fall in demand. None of this was a tenable description of 1979–82. But recovery came soon enough and fast enough that their worldview could, in effect, ride out the storm. (I was at one conference where a freshwater economist, questioned about current events, snapped 'I'm not interested in the latest residual.')

What I see in the response to 2008 and after is much the same dynamic. Half the macroeconomics profession feels mainly validated by events—correctly, I'd say, although as part of that faction I would say that, wouldn't I? The other half should be reconsidering its views—but they should have done that 30 years ago, and this crisis, like that one, was sufficiently well-handled by policy-makers that there was no irresistible pressure for change. (Just to be clear, I'm not saying that it was well-handled in an objective sense: in my view we suffered huge, unnecessary losses of output and employment because of the premature turn to austerity. But the world avoided descending into a full 1930s-style depression, which in effect left doctrinaire economists free to continue believing what they wanted to believe.)

If all this sounds highly cynical, well, I guess it is. There's a lot of very good research being done in macroeconomics now, much of it taking advantage of the wealth of new data provided by bad events. Our understanding of both fiscal policy and price dynamics are, I believe, greatly improved. And funerals will continue to feed intellectual progress: younger macroeconomists seem to me to be much more flexible and willing to listen to the data than their counterparts were, say, 20 years ago.

But the quick transformation of macroeconomics many hoped for almost surely isn't about to happen, because events haven't forced that kind of transformation. Many economists—myself included—are actually feeling pretty good about our basic understanding of macro. Many others, absent real-world catastrophe, feel free to take the blue pill and keep believing what they want to believe.

References

Akerlof, G. A., Dickens, W. T., and Perry, G. L. (1996), 'The Macroeconomics of Low Inflation', *Brookings Papers on Economic Activity*, **1**, 1–52.

Alesina, A., and Ardagna, S. (2010), 'Large Changes in Fiscal Policy: Taxes versus Spending', in J. Brown (ed.), *Tax Policy and the Economy*, **vol. 24**, NBER, University of Chicago Press.

Blanchard, O. (2016), 'The US Phillips Curve: Back to the 1960s?', PIIE Policy Brief 16-1.

Daly, M., and Hobijn, B. (2014), 'Downward Nominal Wage Rigidities Bend the Phillips Curve', Federal Reserve Bank of San Francisco, Working Paper 2013-08.

Draghi, M. (2012), Speech by Mario Draghi, President of the European Central Bank at the Global Investment Conference in London, 26 July, available at https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html

Drezner, D. (2012), 'The Irony of Global Governance: The System Worked', *Foreign Policy*, 1 October.

Eichengreen, B., and O'Rourke, K. (2009), 'A Tale of Two Depressions', Voxeu.com, 1 September.

Eggertsson, G., and Krugman, P. (2012), 'Debt, Deleveraging, and the Liquidity Trap: A Fisher–Minsky–Koo Approach', *Quarterly Journal of Economics*, **127**(3), 1469–513.

— Woodford, M. (2003), 'The Zero Bound on Interest Rates and Monetary Policy', *Brookings Papers on Economic Activity*, **1**.

Krugman, P. (1998), 'It's Baaack: Japan's Slump and the Return of the Liquidity Trap', *Brookings Papers on Economic Activity*, **2**.

Mian, A., and Sufi, A. (2011), 'House Prices, Home Equity-based Borrowing, and the US Household Leverage Crisis', *American Economic Review*, **101**(5), 2132–56.

Tobin, J. (1972), 'Inflation and Unemployment', *American Economic Review*, **62**(1), 1–18.