

# HOW COMPLICATED DOES THE MODEL HAVE TO BE?

PAUL KRUGMAN  
*Princeton University*

*Simple macroeconomic models based on IS-LM have become unfashionable because of their lack of micro-foundations, and are in danger of being effectively forgotten by the profession. Yet while thinking about micro-foundations is a productive enterprise, complex models based on such foundations are not necessarily more accurate than simple, ad-hoc models. Three decades of attempts to base aggregate supply on rational behaviour have not displaced the Phillips curve; inter-temporal models of consumption do not offer reliable predictions about aggregate demand. Meanwhile, the ease of use of small models makes them superior for many practical applications. So we should not allow them to be driven out of circulation.*

## I. A VANISHING ART

Two years before writing this piece I was assigned by my then department to teach Macroeconomics I for graduate students. Ordinarily this course is taught by someone who specializes in macroeconomics; and whatever topics my popular writings may cover, my professional specialities are international trade and finance, not general macroeconomic theory. However, MIT had a temporary staffing problem, which is itself revealing of the current state of macro, and I was called in to fill the gap.

The problem was this: MIT's first macro segment is a half-semester course, which is supposed to cover the 'workhorse' models of the field—the standard approaches that everyone is supposed to know, the models that underlie discussion at, say, the Fed, Treasury, and the IMF. In particular, it is supposed to provide an overview of such items as the IS-LM model of monetary and fiscal policy, the AS-AD approach to short-run versus long-run analysis, and so on. By the standards of modern macro theory, this is crude and simplistic stuff, so you might think that any trained macroeconomist could teach it. But it turns out that that isn't true.

You see, younger macroeconomists—say, those under 40 or so—by and large *don't know this stuff*. Their teachers regarded such constructs as the IS-LM model as too *ad hoc*, too simplistic, even to be worth teaching—after all, they could not serve as the basis for a dissertation. Now MIT's younger macro people are certainly very smart, and could learn the material in order to teach it—but they would find it strange, even repugnant. So in order to teach this course MIT has relied, for as long as I can remember, on economists who learned old-fashioned macro before it came to be regarded with contempt. For a variety of reasons, however, MIT couldn't turn to the usual suspects that year, and I had to fill the gap.

Now you might say, if this stuff is so out of fashion, shouldn't it be dropped from the curriculum? But the funny thing is that while old-fashioned macro has increasingly been pushed out of graduate programmes—it takes up only a few pages in either the Blanchard–Fischer (1989) or Romer (1996) textbooks that I assigned, and none at all in many other tracts—out there in the real world it continues to be the main basis for serious discussion. After 25 years of rational expectations, equilibrium business cycles, growth and new growth, and so on, when the talk turns to the next move by the Fed, the European Central Bank, or the Bank of Japan, when one tries to see a way out of Argentina's dilemma, or ask why Brazil's devaluation turned out relatively well, one almost inevitably turns to the sort of old-fashioned, small-model macro that I taught that spring.

Why does the old-fashioned stuff persist in this way? I don't think the answer is intellectual conservatism. Economists, in fact, are in general neophiles, always looking for something radical and different. Anyway, I have seen over and over again how young economists, trained to regard IS-LM and all that with contempt if they even know what it is, find themselves turning to it after a few years in Washington or New York. There's something about primeval macro that pulls us back to it; if Hicks hadn't invented IS-LM in 1937, we would end up inventing it all over again.

But what is it that makes old-fashioned macro so compelling? The answer, I would argue, is that we

need small, *ad-hoc* models as part of our intellectual tool-box.

Since the 1970s, macroeconomic theory has been driven in large part by an attempt to get rid of the ad-hockery. The most prominent and divisive aspect of that drive has been the effort to provide micro-foundations for aggregate supply. But efforts to provide micro-foundations for aggregate demand—by grounding individual decisions in intertemporal optimization—have been almost equally determined. The result has been a turning away from simple, *ad-hoc* models of the IS-LM genre.

What I would argue is that this tendency has gone too far. Of course we should do the more complicated models; of course we should strive for a synthesis that puts macroeconomics on a firmer micro-foundation. But for now, and for the foreseeable future, the little models retain a vital place in the discipline.

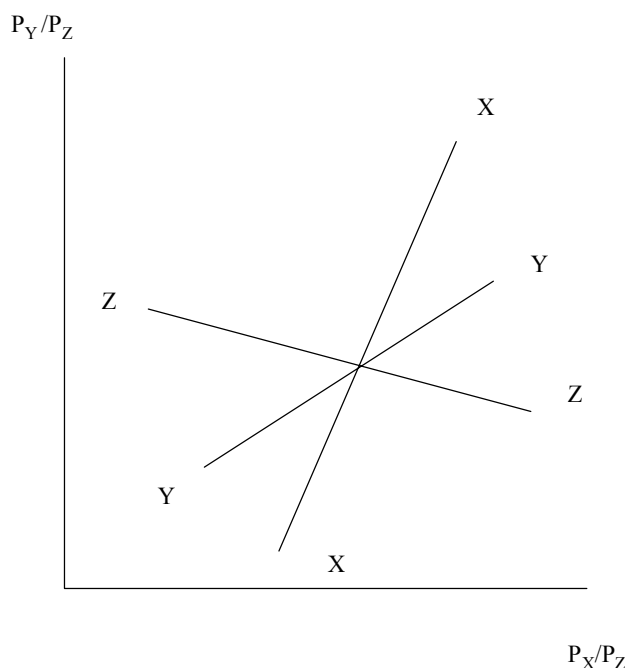
To make this point, let me first revisit where the little models come from in the first place.

## II. HOW THE HOC GOT ADDED

Aficionados know that much of what we now think of as Keynesian economics actually comes from John Hicks, whose 'Mr Keynes and the Classics' (Hicks, 1937) introduced the IS-LM model, a concise statement of an argument that may or may not have been what Keynes meant to say, but has certainly ended up defining what the world *thinks* he said. But how did Hicks come up with that concise statement? To answer that question we need only look at the book he himself was writing at the time, *Value and Capital*, which has in a low-key way been as influential as Keynes's *General Theory*.

*Value and Capital* may be thought of as an extended answer to the question, 'How do we think coherently about the *interrelationships* among markets—about the impact of the price of hogs on that of corn and vice versa? How does the whole system fit together?' Economists had long understood how to think about a single market in

Figure 1



isolation—that’s what supply-and-demand is all about. And in some areas—notably international trade—they had thought through how things fitted together in an economy producing two goods. But what about economies with three or more goods, where some pairs of goods might be substitutes, others complements, and so on?

This is not the place to go at length into the way that Hicks (and others working at the same time) put the story of general equilibrium together. But to understand where IS-LM came from—and why it continues to reappear—it helps to think about the simplest case in which something more than supply and demand curves becomes necessary: a three-good economy. Let us simply call the goods  $X$ ,  $Y$ , and  $Z$ —and let  $Z$  be the numeraire.

Now equilibrium in a three-good model can be represented by drawing curves that indicate combinations of prices for which each of the three markets is in equilibrium. Thus in Figure 1 the prices of  $X$  and  $Y$ , both in terms of  $Z$ , are shown on the axes. The line labelled  $X$  shows price combinations for which demand and supply of  $X$  are equal; similarly with  $Y$  and  $Z$ . Although there are three curves, Walras’s Law tells us that they have a common intersection, which defines equilibrium prices for the economy as a whole. The slopes of the curves are drawn on the assumption that own-price effects are

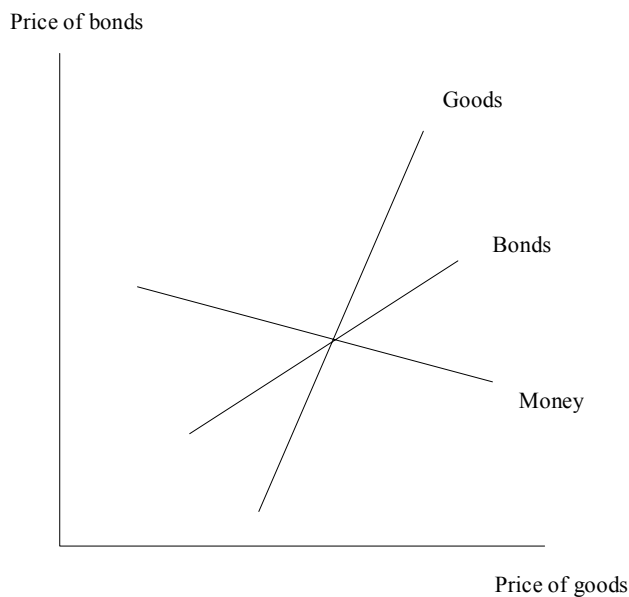
negative, cross-price effects positive—thus an increase in the price of  $X$  increases demand for  $Y$ , driving the price of  $Y$  up, and vice versa; it is also, of course, possible to introduce complementarity into such a framework, which was one of its main points.

This diagram is simply standard, uncontroversial microeconomics. What does it have to do with macro?

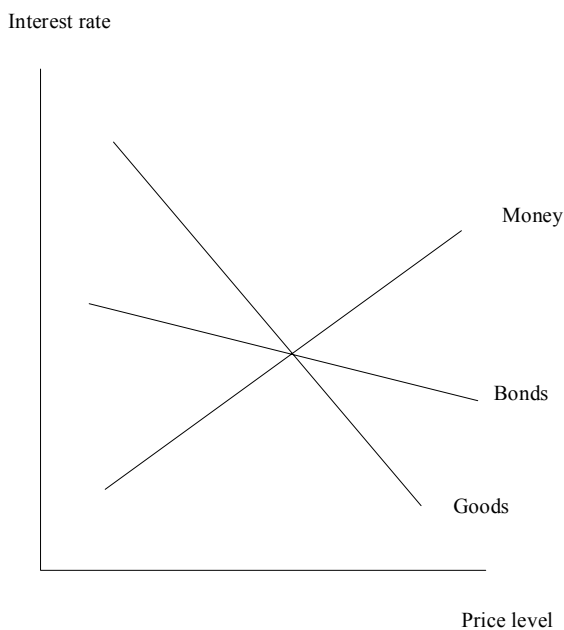
Well, suppose you wanted a first-pass framework for thinking coherently about macro-type issues, such as the interest rate and the price level. At minimum such a framework would require consideration of the supply and demand for *goods*, so that it could be used to discuss the price level; the supply and demand for *bonds*, so that it could be used to discuss the interest rate; and, of course, the supply and demand for money.

What, then, could be more natural than to think of goods in general, bonds, and money as if they were the three goods of Figure 1? Put the price of goods—aka the general price level—on one axis, and the price of bonds on the other; and you have something like Figure 2—or, more conventionally putting the interest rate instead of the price of bonds on the vertical axis, something like Figure 3. And already we have a picture that is essentially Patinkin’s flexible-price version of IS-LM.

**Figure 2**



**Figure 3**



If you try to read pre-Keynesian monetary theory, or for that matter talk about such matters either with modern laymen or with modern graduate students who haven't seen this sort of thing, you quickly realize that this seemingly trivial formulation is actually a powerful tool for clarifying thought, precisely because it is a general-equilibrium framework that takes the interactions of markets into account. (I have heard the story of a famed general-equilib-

rium theorist who, late in a very distinguished career, saw an IS-LM diagram and asked who came up with that brilliant idea.) Here are some of the things it suddenly makes clear.

**(i) What Determines Interest Rates?**

Before Keynes-Hicks—and even to some extent after—there has seemed to be a conflict between

the idea that the interest rate adjusts to make savings and investment equal, and that it is determined by the choice between bonds and money. Which is it? The answer, of course—but it is only ‘of course’ once you’ve approached the issue the right way—is both: we’re talking general equilibrium here, and the interest rate and price level are jointly determined in both markets.

### **(ii) How Can an Investment Boom Cause Inflation (and an Investment Slump Cause Deflation)?**

Before Keynes this was a subject of vast confusion, with all sorts of murky stuff about ‘lengthening periods of production’, ‘forced saving’, and so on. But once you are thinking three-good general equilibrium, it becomes a simple matter. When investment (or consumer) demand is high—when people are eager to borrow to buy real goods—they are in effect trying to shift from bonds to goods. So both the bond-market and goods-market equilibrium schedules, but not the money-market schedule, shift; and the result is both inflation and a rise in the interest rate.

### **(iii) How Can We Distinguish between Monetary and Fiscal Policy?**

Well, in a fiscal expansion the government sells bonds and buys goods—producing the same shifts in schedules as an investment boom. In a monetary expansion it buys bonds and ‘sells’ newly printed money, shifting the bonds and money (but not goods) schedules.

Of course, this is all still a theory of ‘money, interest, and prices’ (Patinkin’s title), not ‘employment, interest, and money’ (Keynes’s). To make the transition we must introduce some kind of price-stickiness, so that incipient deflation is at least partly translated into output decline; and then we must consider the multiplier impacts of that output decline, and so on. But the basic form of the analysis still comes from the idea of a three-good general-equilibrium model in which the three goods are ‘goods in general’, bonds, and money.

Sixty years on, the intellectual problems with doing macro this way are well known. First of all, the idea of treating money as an ordinary good begs many

questions: surely money plays a special sort of role in the economy. (For some reason, however, this objection has not played a big role in changing the face of macro.) Second, almost all the decisions that presumably underlie the schedules here involve choices over time: this is true of investment, consumption, even money demand. So there is something not quite right about pretending that prices and interest rates are determined by a static equilibrium problem. (Of course, Hicks knew about that, and was quite self-conscious about the limitations of his ‘temporary equilibrium’ method.) Finally, sticky prices play a crucial role in converting this into a theory of real economic fluctuations; while I regard the evidence for such stickiness as overwhelming, the assumption of at least temporarily rigid nominal prices is one of those things that works beautifully in practice but very badly in theory.

But step back from the controversies, and put yourself in the position of someone who must reach a judgement about the likely impact of a change in monetary policy, or an investment slump, or a fiscal expansion. It would be cumbersome to try, every time, to write out an intertemporal-maximization framework, with micro-foundations for money and price behaviour, and try to map that into the limited data available. Surely you will find yourself trying to keep track of as few things as possible, to devise a working model—a scratch-pad for your thoughts—that respects the essential adding-up constraints, that represents the motives and behaviour of individuals in a sensible way, yet has no superfluous moving parts. And that is what the quasi-static, goods–bonds–money model is—and that is why old-fashioned macro, which is basically about that model, remains so popular a tool for practical policy analysis.

Of course, if we knew—really knew—that one got much more reliable results by doing the things that *ad-hoc* macro does not, it would be a tool to be used only for the most preliminary examination of issues. But do we know that? Let me turn briefly to the two main areas of contention, the two main areas in which the search for micro-foundations has been most aggressively pursued: aggregate supply and intertemporal spending decisions. In each case I want to ask: how sure are we that the micro-founded version is really better?

What I would argue is that in each case the first efforts to derive some micro-foundation had a big pay-off. That is, in each case a preliminary, rough application of the idea that there was a deeper structure underneath the quasi-static model produced a result that was both compelling and empirically confirmed. But in each case, also, the subsequent work, the elaboration of the micro-foundations, has produced little if any gain in predictive power. *Thinking* about micro-foundations does you a lot of good; taking care to put them into the model all the time, it seems, does not.

### III. AGGREGATE SUPPLY

The story of the aggregate supply wars is familiar to most economists, though everyone tells it a bit differently. Let me give my own version.

As I see it, the effort to put micro-foundations beneath aggregate supply went through four stages.

#### (i) The Natural-rate Hypothesis

When macroeconomists first began thinking seriously about why nominal shocks really have real effects, they were led into a variety of models in which imperfect information might cause a confusion between real and nominal shocks. The famous Phelps volume (1970) included many of these models. The key implication of these models was some version of the principle that you can't fool all of the people all of the time, suggesting not just that prices would be flexible in the long run, but that persistent inflation would get built into expectations too. Thus was the natural-rate hypothesis born.

While there was some resistance to the natural-rate idea—and some sophisticated challenges have been posed in recent years—it was an idea that most economists found compelling. And its empirical predictions were soon seen to be roughly true. The Phillips curve did turn out to be unstable; for the USA, at least, a time series of inflation and unemployment seemed to go through clockwise spirals, just as you would have expected. (European data have never fitted very well, but that was easily explained away as the result of a secular upward trend in the natural rate.)

Better yet, the natural-rate hypothesis in effect predicted stagflation in advance, a victory that gave huge credibility to the whole enterprise of micro-founded macro. It therefore created a favourable environment for the second stage.

#### (ii) The Lucas Project

Few ideas in economics have been so influential, yet left so little lasting impact, as the idea that nominal shocks have real effects because of 'rational confusion'. When I was in graduate school the Lucas supply curve, with its lovely metaphors about islands and signal processing, but its alas very unlovely formal apparatus, was *the* central development in macroeconomics. But hardly anyone seems to think it a useful tool today.

What went wrong? Analytically, it all comes down to the case of the economic agent who knew too much. In the real world recessions last far too long for us to believe that people are voluntarily withholding labour because they believe they are facing idiosyncratic shocks. And in theoretical models, even if we somehow imagine that agents cannot read about the ongoing recession in the *Financial Times*, just about any sort of financial market will, in conjunction with individual experience, eliminate the confusion that Lucas-type models rely on.

Empirically, the main supporting evidence for the signal-processing aspect of the Lucas-type approach was the observation that the short-run aggregate supply curve appears steeper in countries with highly unstable inflation rates. But as Ball *et al.* (1988) pointed out, countries with very unstable inflation rates are also countries with *high* inflation rates, and one can think of many models in which high inflation would in effect reduce money illusion.

In short, the Lucas project—which tried to build an aggregate supply curve in which nominal shocks had real effects on genuine micro-foundations—failed. There were two responses to this.

#### (iii) The Great Schism

At the risk of a great oversimplification, in the 1980s macroeconomists reacted to the failure of the Lucas project in two ways. One reaction was to say that



because we had not managed to find a micro-foundation for non-neutrality of money, money must be neutral after all; it might look to you as if the Fed has power to affect the real economy, but that must be a statistical illusion. And that led in the direction of real business-cycle theory.

The other reaction was to say that we need some other explanation of apparent non-neutrality, resting in something like menu costs or bounded rationality. And that led in the direction of New Keynesian theory.

Each of these approaches had something to commend it, but surely at this point we can say that neither approach really succeeded.

I could go through the analytical and empirical objections to real business cycle theory at some length, but perhaps the crucial point to make is that the evidence for price stickiness has stubbornly refused to go away. My personal impression was that the tide really began to turn against equilibrium business-cycle models with the accumulation of data about co-movement of nominal and real exchange rates. As inflation rates in advanced countries fell during the 1980s, while nominal exchange rates continued to fluctuate wildly, it became more or less impossible to ignore the startling extent to which price levels remained stable even while large swings in the exchange rate took place. It is no accident that Obstfeld and Rogoff (1996)—which arguably is the defining book for the state of modern macro—basically uses exchange-rate evidence as the main basis for arguing that prices are, sure enough, sticky and hence that nominal shocks have real effects.

But if prices are sticky, why? New Keynesian macro offered an intellectually almost-satisfying answer. Suppose that there are distortions in labour and product markets, which give price-setters some monopoly or monopsony power. Then the loss to an individual price-setter from not reducing its price in the face of a fall in aggregate demand would be second-order in the size of the decline—but the cost to the economy as a whole of such price stickiness would be first-order. All that one needs, then, is some fairly trivial reason not to change prices—menu costs, bounded rationality, or something—and

one can justify the kind of price inflexibility that makes recessions possible.

It's a terrific insight; I remember being thrilled when I realized that we finally had a way other than confusion to reconcile sticky prices with more or less rational behaviour. But what does one do, exactly, with that insight? Fifteen years after the publication of the original menu-cost papers of Mankiw (1985) and Akerlof and Yellen (1985), it seems that this line of thought has produced not so much a micro-foundation for macro as a micro-excuse. That is, one can now explain how price stickiness *could* happen. But useful predictions about when it happens and when it does not, or models that build from menu costs to a realistic Phillips curve, just don't seem to be forthcoming.

#### (iv) Undeclared Peace

So where are we now? The passion seems to have gone out of the aggregate supply wars. The typical modelling trick, now used by both sides, is to suppose that prices must for some reason be set one period ahead. This assumption is somewhat problematic empirically, since there is considerable evidence for inflation inertia that goes beyond preset prices. But what I want to emphasize here is that the end result of the whole attempt to provide rational foundations for aggregate supply has ended up with something that is almost as *ad hoc* as the original assumption of sticky prices.

Notice the pattern. We got a big pay-off—the natural-rate hypothesis—from the first, crude, application of rational modelling. The subsequent elaboration, the stuff that makes the models so complicated, has added to our understanding of the intellectual issues, but not noticeably to our ability to match the phenomena.

## IV. AGGREGATE DEMAND

While there has been considerable theorizing and empirical work over other components of aggregate demand, consumption spending has clearly been the centre of research and debate. The history of this debate is almost, though not quite, as familiar as that

of the aggregate supply wars; let me again give my own version.

It began, of course, with the Keynesian consumption function, with consumption simply a function of disposable income. But this consumption function ran into empirical paradoxes and a forecasting débâcle almost as soon as it was formulated.

The forecasting débâcle was the failure of consumption functions based on inter-war data to predict the post-war consumption boom. Indeed, consumption functions from the 1930s—with a marginal propensity to consume much less than the average propensity—seemed to lend plausibility to the ‘secular stagnation’ hypothesis, with its suggestion that it would be persistently difficult to get the public to spend enough to use the economy’s growing capacity. Happily, it didn’t work out that way.

More broadly, the stylized facts on consumption didn’t fit together if you tried to explain them with a consumption function of the form  $C = C(Y)$ . The way I teach undergrads is to describe three facts: the procyclicality of the savings rate; the fact that high-income households have higher savings rates than low-income households; and the constancy of the savings rate over the very long run. The first two facts seem to suggest a marginal propensity to consume substantially less than the average propensity. But the third fact seems to imply a marginal propensity to consume that is more or less the same as the average propensity.

Rational behaviour came to the rescue. The permanent-income / life-cycle approach pointed out that a rational individual bases his consumption on his wealth, including the present discounted value of expected future income, not just on current income. And all the stylized facts immediately fell into place: cyclical income gets saved because it is temporary, high-income families save more because they are also typically families with unusually high income for their own life cycle, but the long-term upward trend in income, because it is permanent, does not get reflected in a higher savings rate.

As with the natural-rate hypothesis, then, the first application of rational behaviour to the model produced a big pay-off. And then?

As originally used, the permanent-income/life-cycle approach was a bit loose; it was used to suggest variables that might proxy for wealth, rather than be taken literally. But from the 1970s on it became much more hard-edged: we were supposed to abandon *ad-hoc* consumption functions and base everything on intertemporal maximization.

What that meant, above all, was that the little macro models I described earlier were no longer allowable. Quite aside from the sticky-price issue, such models were now condemned as attempts to shoehorn an essentially dynamic, forward-looking problem into a quasi-static framework.

And at one level this is certainly right. In international macroeconomics, there are some issues, notably current account determination, where you really want to have an intertemporal view—otherwise it is just very hard to get any good bearings. But should we rule out the use of the simple, *ad-hoc*, quasi-static models altogether?

If fully fledged intertemporal consumption models were simply radically better at describing consumption behaviour, there would be a case for doing away with the *ad-hoc* models. But after a rousing initial start in Hall’s (1978) famous ‘random walk’ paper, the evidence seems to have become distinctly mixed. The basic point seems to be that predictable changes in income should not lead to changes in consumption—they should already be incorporated in the current level of consumption, just as predictable changes in profits are supposed to be reflected in the current stock price. But as the survey in Romer (1996) points out, there is considerable evidence that in fact such predictable changes in income have substantial effects, though not as much as a crude Keynesian consumption function would have predicted. And evidence for the strong implications of intertemporal consumption optimization, like Ricardian equivalence when it comes to budget deficits, is at best questionable.

Why the disappointing results? Arguments similar to those used in New Keynesian models of aggregate supply also apply to aggregate demand. A consumer who uses a sensible rule of thumb will typically suffer only a small loss in lifetime utility compared with an obsessively perfect optimizer. If there are



some costs to gathering and processing information, the rule of thumb may well be a better choice. Yet the implied consumption behaviour may be very different from that of a full optimizer. On the other hand, like new Keynesian price rigidity, this serves more as a micro-excuse than a micro-foundation for whatever we actually assume about spending.

In short, a careful intertemporal approach to aggregate demand does not turn out to be all that helpful. It does not, it turns out, offer a clear improvement in predictive power over that of simpler, *ad-hoc* approaches.

So the history of aggregate demand theorizing is in a fundamental sense very similar to that of aggregate supply. Introducing a whiff of rational behaviour into the model greatly improves its ability to reflect reality. Going the rest of the way, to use a fully specified model of rational choices over time, is at best an ambiguous improvement; it produces a model that is more satisfying intellectually, but not necessarily better in terms of accuracy.

## V. IN DEFENCE OF LITTLE MODELS

We now come to the moral of the sermon: the case for little, simple models. That means IS-LM or some variant for domestic macro, with an adaptive-expectations Phillips curve; it means a slightly souped up version of Mundell–Fleming for international macro, with some kind of regressive expectations on the real exchange rate to avoid the conclusion that all countries must have the same interest rate.

The point is not that these models are accurate or complete, or that they should be the only models used. Clearly they are incomplete, quite inadequate to examining some questions, and remain as full of *hoc* as ever. But they are easy to use, particularly on real-world policy questions, and often seem to give more or less the right answer.

The alternative is to use models that are less *ad hoc*, that are careful to derive as much as possible from individual maximization. Nobody can seriously say that such models should be banned from discourse; where they confirm the insights of the simpler models they are reassuring, and where they contradict those results they serve as a warning sign. But

what we know pretty well, from decades of trying to give micro-foundations to macro, is that logical completeness and intellectual satisfaction are not necessarily indications that a model will actually do a better job of tracking what really happens. For many purposes the small, *ad-hoc* models are as good as or better than the carefully specified, maximizing intertemporal model.

Let me discuss an exception that proves the rule. One of the most influential macro models of the 1990s, and deservedly so, is the revisitation of Mundell–Fleming–Dornbusch by Obstfeld and Rogoff (1995). It's a beautiful piece of work, integrating a new Keynesian approach to price stickiness (albeit with the *ad-hoc* assumption that prices are set for only one period) with a full intertemporal approach to aggregate demand; it addresses the classic question of the effect of a monetary shock on output, interest rates, and the exchange rate. It's also pretty hard work: as restated in their 1996 book, the model and its analysis take 76 equations to describe, even though the authors restrict themselves to a special case, and to log-linearized analysis around a steady state.

The big pay-off to all this is a gorgeous welfare result, which says that the benefits of a monetary expansion accrue equally to both countries, wherever the monetary expansion takes place—a result that is immediately understandable in terms of the welfare economics of monopoly in general.

But ask two questions. First, is this model actually any better at predicting the impact of monetary expansions in the real world than a three- or four-equation, back-of-the-envelope Mundell–Fleming model? The authors give us no reason to think so.

Second, do you really believe that welfare result? Certainly not—indeed, it is driven to an important extent by the details of the model, and can quite easily be undone. The result offers a tremendous clarification of the issues; it's not at all clear that it offers a comparable insight into what really happens.

None of this should be taken as a critique of the 'new open-economy macroeconomics', which is wonderful stuff. The point is that while I can and will use this approach along with the old *ad-hoc* models, I am not going to give up those old models. And when

I write my next popular piece on international macro, it will probably be based on *ad-hoc* macro, not models with proper micro-foundations.

In the long run, what this says is that we haven't really got the right models. The small models haven't gotten any better over the past couple of decades; what has happened is that the bigger, more micro-founded models have not lived up to their promise. The core of my argument isn't that simple models are good, it's that complicated models aren't all they

were supposed to be. But pending the arrival, someday, of models that really do perform much better than anything we now have, the point is that the small, *ad-hoc* models are still useful.

What that means, in turn, is that we need to keep them alive. It would be a shame if IS-LM and all that vanish from the curriculum because they were thought to be insufficiently rigorous, replaced with models that are indeed more rigorous but not demonstrably better.

## REFERENCES

- Akerlof, G., and Yellen, J. (1985), 'A Near-rational Model of the Business Cycle, with Wage and Price Inertia', *Quarterly Journal of Economics*, **100**, 823–38.
- Ball, L., Mankiw, G., and Romer, D. (1988), 'The New Keynesian Economics and the Output–Inflation Tradeoff', *Brookings Papers on Economic Activity*, 1, 1–65.
- Blanchard, O., and Fischer, S. (1989), *Lectures on Macroeconomics*, Cambridge, MA, MIT Press.
- Hall, R. (1978), 'Stochastic Implications of the Life Cycle Permanent Income Hypothesis', *Journal of Political Economy*, **86**(December), 971–87.
- Hicks, J. R. (1937), 'Mr Keynes and the Classics', *Econometrica*, **5**(April), 147–59.
- Mankiw, N. G. (1985), 'Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly', *Quarterly Journal of Economics*, **100**, 529–39.
- Obstfeld, M., and Rogoff, K. (1995), 'Exchange Rate Dynamics Redux', *Journal of Political Economy*, **103**, 624–60.
- — (1996), *Foundations of International Macroeconomics*, Cambridge, MA, MIT Press.
- Phelps, E. (ed.) (1970), *Microeconomic Foundations of Employment and Inflation Theory*, Norton.
- Romer, D. (1996), *Advanced Macroeconomics*, New York, McGraw-Hill.