The swamps of DSGE despair

Noah Smith, March 29, 2013

David Andolfatto (of the St. Louis Fed and the blog MacroMania) points me to this interesting recent working paper by Braun, Korber, and Waki. The paper bears the somewhat unwieldy title of "Some unpleasant properties of log-linearized solutions when the nominal rate is zero." Basically, the authors of this paper take a New Keynesian model somewhat similar to the ones used by "Keynesian" macroeconomists - e.g. Paul Krugman - to help justify the use of fiscal stimulus in a depressed economy. They note that in most papers, the models actually used to measure the effect of government policy are "linearized" versions.

For the uninitiated: A DSGE model starts with the assumption of optimization by various economic agents such as households and firms, which spits out a system of nonlinear equations representing people's optimal choices. These nonlinear equations are then "log-linearized" around a "steady state", and the linearized forms of the equations, which are very easy to work with mathematically and computationally, are used to compute the "impulse responses" that tell you what the model says the effect of government policy will be. The linearization is equivalent to the assumption that the economy undergoes only small disturbances. That might not be a good assumption when it comes to major events like the recent crisis/depression, but it does make the models a LOT easier to work with. It also generally makes the equilibria *unique* - in other words, if you use the full, nonlinear version of a DSGE model, you are likely to come up with a bunch of different possible paths for the economy, and which path the economy takes will be determined purely by quantitative factors (like, whether variable X is more or less than 2.076, or something like that) - not the kind of thing that DSGE models are good at getting right. Since "multiple equilibria" generally means "we really don't know what's going to happen," macroeconomists tend to stick to the linearized versions of models, so that they can say "we *do* know what's going to happen."

Anyway, Braun et al. decide to venture into no-man's-land, and work with a non-linearized version of a New Keynesian model with a Zero Lower Bound. They find that, unsurprisingly, there are multiple equilibria. In some of these equilibria, the kind of special ZLB effects found by Eggertsson and Krugman - for example, the "paradox of toil" - are present, but small in size. In other equilibria, the effects go away entirely.

So can we conclude that the concerns of Keynesian economists about the ZLB are overblown, and that fiscal policy isn't the answer? Not so fast. Here is another working paper, by Fernandez-Villaverde, Gordon, Guerron-Quintana, and Rubio-Ramirez, which

conducts a similar exercise with a slightly different model. Fernandez-Villaverde et al.'s model is extremely hard to solve and their results come from picking some interesting-sounding cases and then doing numerical experiments (simulations) to see what happens in those cases. In the main case they consider interesting, the ZLB ends up being pretty important, and the fiscal policy multiplier is around 1.5 or 2.

(Update: A commenter points me to [this response by Christiano and Eichenbaum](#), two of the leading New Keynesian theorists. They show that most of the multiple equilibria found by Braun, et al. are not supported by a specific model of learning. Also, here is a [multiple-equilibrium DSGE paper by Mertens and Ravn](#) showing that in some equilibria, fiscal policy actually makes recessions worse. The Mertens and Ravn result also conflicts with the learning model of Christiano and Eichenbaum.)

So what do we learn from these sorts of exercises? In my opinion, we learn relatively little about the real economy, but that's OK, since we do learn some important things about DSGE models. Namely:

1. Almost every DSGE result you see is the result of linearization, If you drop linearization, very funky stuff happens. In particular, equilibria become non-unique, and DSGE models don't give you a good idea of what will happen to the economy, *even in the fictional world* where the DSGE model's assumptions are largely correct! As Braun et al. write:
There is no simple characterization of when the loglinearization works well. Breakdowns can occur in regions of the parameter space that are very close to ones where the loglinear solution works. In fact, it is hard to draw any conclusions about when one can safely rely on loglinearized solutions in this setting without also solving the nonlinear model.
So even putting aside the question of whether DSGE models accurately represent reality, we see that most of the DSGE models you see *don't even accurately represent themselves*.

2. In order to be usable, DSGE models have to have a LOT of simplification. These nonlinear New Keynesian models go so haywire that they often have to be simulated instead of solved. Furthermore, *neither of these models has capital or investment*. Since investment is the component of GDP that swings most in recessions, you'd think this would be an important omission. But putting in capital would make these already mostly intractable models into utterly hopelessly intractable models (And if you don't believe me, ask Miles Kimball, who has spent considerable time and effort working on the problem of putting capital into New Keynesian models). Never mind putting in other realistic stuff like agent heterogeneity!

Basically, every time you model a phenomenon, you face a tradeoff between realism and tractability - the more realistic stuff you include, the harder it is to actually use your model. But DSGE models face an *extremely unfavorable* realism/tractability tradeoff. Adding even a dash of simple realistic stuff makes them get very clunky very fast.

3. DSGE models are highly sensitive to their assumptions. Look at the difference in the results between the Braun et al. paper and the Fernandez-Villaverde et al. paper. Those are pretty similar models! And yet the small differences generate vastly different conclusions about the usefulness of fiscal policy. Now realize that every year, macroeconomists produce a vast number of different DSGE models. Which of this vast array are we to use? How are we to choose from the near-infinite menu of very similar models, when small changes in the (obviously unrealistic) assumptions of the models will probably lead to vastly different conclusions? Not to mention the fact that an honest use of the full nonlinear versions of these models (which seems only appropriate in a major economic upheaval)

wouldn't even *give* you definite conclusions, but instead would present you with a menu of multiple possible equilibria?

Imagine a huge supermarket isle a kilometer long, packed with a million different kinds of peanut butter. And imagine that all the peanut butter brands look very similar, with the differences relegated to the ingredients lists on the back, which are all things like "potassium benzoate". Now imagine that 85% of the peanut butter brands are actually poisonous, and that only a sophisticated understanding of the chemistry of things like potassium benzoate will allow you to tell which are good and which are poisonous.

This scenario, I think, gives a good general description of the problem facing any policymaker who wants to take DSGE models at face value and use them to inform government policy.

So what's my suggestion? First I'd suggest detailed studies of consumer behavior, detailed studies of firm behavior, lab experiments, etc. - basically, huge amounts of serious careful empirical work - to find out which set of microfoundations are approximately true, so that we can focus only on a very narrow class of models, instead of just building dozens and dozens of highly different DSGE models and saying "Well, maybe things work this way!" Second, I'd suggest incorporating these reliable microeconomic insights into large-scale simulations (like the ones meteorolgists use to forecast the weather); in fact, any DSGE model that incorporates all of the actual frictions we find is likely to be so complicated, and so full of multiple equilibria in the full nonlinear case, that it demands this kind of approach. Third, and in parallel to the weather-forecasting effort, I'd echo Bob Solow's call to use *simple* models when trying to explain ideas to other economists and to the public (explanation of ideas being what DSGE models are mainly used for, given their abysmal performance at actually predicting anything about the economy). Note that I don't have a ton of confidence in these alternatives; after all, it's a lot easier to find flaws in the dominant paradigm than it is to come up with a new paradigm.

But in any case, few people in the macroeconomics field seem to be particularly interested in that sort of alternative approach, or any other. And the scientific culture of macroeconomics doesn't seem to demand that we find an alternative; in fact, in the macro profession, you pretty much *have* to back up any empirical result or simple model with a fully specified mainstream-ish DSGE model in order to be taken seriously.

So instead of trying to find which set of models really works, everyone just makes more models and more models and more models and more models...

(Note: If you know basic math and want to learn what DSGE models are all about, start with this chapter from David Romer's *Advanced Macroeconomics*.)