Comment on Christiano, Motto and Rostagno

Christopher A. Sims
Princeton University
sims@princeton.edu

June 12, 2009

1 What's new in the model

New variables

- Financial index: External finance premium, (AAA minus BAA rates in US, bank loan rate minus corporate/gov't bond rate in EA)
- Banking: M1, M3, reserves

New mechanisms

- BGG financial frictions
- Expanded monetary sector
- Risk shocks
- Private info about risk shocks better than modeler's

2 What's missing or problematic in the model

2.1 Generic problems with policy DSGE's

Fiscal policy

- What's the most important and most uncertain aspect of the current policy environment?
- The question of what happens when the irresistible force (the fiscal deficit) meets the immovable object (the Fed's inflation-control policy).
- This model, like most (all?) policy DSGE's in use in central banks, has no serious treatment of fiscal policy. It recognizes neither that inflation, and expectations of future inflation, influence private sector wealth in the form of government bonds, nor the reverse, that changes in the outstanding volume of bonds affect actual and expected future inflation.
- We can see these channels operating in the current environment. Models that omit them or assume them away are seriously deficient.

Limitations of the BGG financial friction

- It has risk residing directly in real returns at the level of the producing firm.
- Our recent problems (and those in the 1930's) seem to have originated in the financial system, not so much in the riskiness of production itself.
- The model does not contain layers of cross-leveraged purely financial firms that could be the locus of a financial shock.
- This does not mean that the BGG framework won't work pretty well as an approximation in a macro model.

2.2 Specific issues with this model

Household objective function

- Despite the amazingly persistent popularity of inserting real balances in the utility function, this is hazardous, and this paper is a nice illustration of the hazards.
- Putting real balances in the budget constraint, as generating transactions costs, enforces realistic treatment of the transactions value of real balances. Furthermore, in this case we can do without the (also amazingly persistent) "cash in advance" fiction.

Household objective, cont.

This model has an additive term in the objective function of the form

$$-\nu \frac{\left(\frac{PC}{M}\right)^{1-\sigma_q}}{1-\sigma_q} \,.$$

The estimated value of σ_q is negative and large. This implies that the curvature in this term dominates that in the (logarithmic) direct utility of C for large C, so that there is satiation in consumption, because transactions costs start to absorb more than the entire increase in consumption spending.

Also [this point not in the presentation] this component of the utility function is not concave. From a couple of numerical experiments it seems to be quasi-concave, so second-order conditions are probably ok. But the paper should note the non-monotonicity and verify that solutions are well away from the region where marginal utility of consumption turns negative.

Household objective, cont., cont.

- Why a utility cost to changing real balances?
- Even when the change occurs passively, via lower than trend growth in prices?
- This is so counterintuitive that the reason for it needs explanation. Does the model misbehave without it? Why not make the cost a function of $(M_t - M_{t-1})/P_{t-1}$, so the disutility of price changes is eliminated?

News about risk

- It is specified that the shock u_t to the risk variable σ_t can be represented as

$$u_t = \sum_{s=0}^{8} \xi^s_{\sigma,t-s} \,,$$

  with σ^s, the standard deviation of $\xi^s_{\sigma,t}$, constant across $s > 0$.

- The $\xi^s_{\sigma,t-s}$ are specfied to be i.i.d. across s with a diagonal covariance matrix, which follows from the martingale property of optimal forecasts.

- But they are also specifed to be i.i.d. across t, which is a strong, even unreasonable, additional restriction. It rules out, e.g., a situation where each period we see not only current u_t, but a noisy indicator $u_{t+1} + u_{t+2} + \varepsilon_t$. I.e., it rules out people realizing that u is likely to rise more than usual in the next few quarters, while being unsure about the timing of the rise.

3 Assessing results and fit

News about risk, cont.

- That the assumptions about the risk shock components are problematic is borne out in Table 6.
- That table shows correlations of smoothed estimates of shocks. There are many large off-diagonal elements in the submatrix corresponding to the $\xi^s_{\sigma,t}$ shocks.
- This could reflect large errors in the smoothed estimates of the shocks, but the correlations look too big to be explained this way.

Risk spread

Since one of the main empirical advances here is bringing in the external finance premium variable, it is particularly disappointing that the variable on which the CMR model does worst in relative RMSE is the risk spread (= external finace premium

RMSE vs. MCMC

- RMSE's from recursively estimated models are expected to shrink between small early samples and larger later ones.
- More generally, there may be periods when variables that have moved little to date start moving sharply. Their coefficients are likely to be poorly estimated, and expected squared forecast error is therefore high in such periods.
- MCMC measures of fit are basically recursively estimated RMSE's, except properly weighted to reflect parameter uncertainty. They are easy to do these days. They are strictly better than the plain RMSE's.

- The paper claims that it is easy to produce frequentist confidence intervals for the RMSE's. This is not true, since by construction the expected squared errors change over time in a complicated way. Furthermore, multi-step forecast RMSE's can given an (asymptotic) frequentist distribution theory only via estimation of large numbers of nuisance parameters, since the multi-step forecasts are serially correlated. (Was this done here?)

- [Added post-conference] I gather from conversation that the large number of nuisance parameters was in fact estimated, so that the confidence intervals have asymptotic justification. But that they make little sense is apparent from the plots in Figure 7a, where half the plots show the confidence interval for RMSE tighter at 12 quarters than at 1 quarter.

- The frequentist distribution theory used here invokes asymptotics to ignore the shrinkage in forecast errors between early and later forecasts. This follows from the fact that in a large enough sample all parameters will be estimated with high accuracy, so parameter estimates contribute little to forecast error. But in realistic sample sizes and models, parameter estimation error is the dominant source of forecast error, not a negligible source. The CMR model estimates 46 parameters by formal econometric methods and 26 more by educated guessing. Its total sample size is 93. It is unlikely that frequentist asymptotics are a useful approximation to small sample distribution theory here.

- Even within the frequentist framework, attaching an error band to one model's RMSE's (here the BVAR) is not very useful when, as in forecast comparisons, it is expected that forecast errors are highly correlated across models. One could, e.g., plot the differences between the other model RMSE's and the BVAR RMSE and put error bands on those.

Autocorrelation functions [added post-conference]

- The exhaustive display of autocorrelation functions is wasteful of space. ACF's tend to be dominated by the largest root of the system, so they give less useful information to the eye than the same information presented as an exhaustive set of impulse responses to orthogonalized innovations. It's like staring at a raw covariance matrix vs. looking at its eigenvector/eigenvalue decomposition.

- How were the error bands on the autocorrelation functions generated? I could easily understand if the paper involved MCMC posterior simulation, but it seems to involve only maximization of the posterior pdf and use of local expansions of the log likelihood around that point. So are the error bands based on draws from the Gaussian approximation to the posterior? However it was done, this, like the methods for error bands on the RMSE's, need to be explicit.

Laplace expansion for model comparison

- Laplace expansions are a useful shortcut that sometimes, even often, give results similar to what is obtained by Monte Carlo posterior simulation.
- But in large models Laplace expansions raise numerical difficulties of their own, and
- Sometimes Laplace expansions are far off the mark, and this is especially likely in weakly identified models with many parameters.
- MCMC simulation is standard now. The paper should include it.

Priors

- Priors should reflect what we know about parameters before looking at data.
- It is particularly important, when our work is part of a stream of research modeling the same data in different ways, that priors not be based on posteriors from previous studies.
- (If a sequence of researchers estimate the same model on the same data, each researcher j basing his prior on the posterior of researcher $j-1$, the posterior converges toward a spike at the MLE.)
- This paper pays uncomfortably much attention to "previous research" in justifying its prior.
- Micro data based estimates of parameters seldom are estimating exactly the parameter that enters the aggregate model, and in any case it is well known that micro and macro based parameter estimates often differ, for a variety of reasons, with neither being obviously more correct.

Posteriors in tails

- When the posterior for a parameter is concentrated in the distant tail of the prior, there are two possible interpretations.
  - The parameter estimates are unreasonable from the point of view of a priori knowledge. This suggests the model is deficient, and indeed posterior odds will be pushed down by this kind of result.
  - The prior was not carefully thought out. These estimates actually aren't very unreasonable, now that we think about it. Then the prior should be adjusted, since the earlier, unreasonable prior will have had strong effects on the estimates.
- This paper's estimates have posteriors in the distant tails for quite a few parameters.

Fisher effect

- Figures 12-14 show impulse responses to a risk shock, a monetary policy shock, and a technology shock.
- It is difficult to interpret them because they are accompanied by no error bands. We we don't know whether the differences we see are well determined by the data or not.
- But on the face of it, they show that the Fisher debt deflation effect is an important source of non-neutrality. Responses to monetary policy shocks are substantially larger in the EA with this effect (but not in the US), while responses to technology shocks with the effect are larger in both regions.

Preprocessing of the data

- The model is used in a form that can be reduced to stationary by removing some common trends.
- The data that might trend is all first-differenced.
- The sample means of these differenced data are removed in advance.

- Recursive "RMSE" based on data demeaned using whole sample?
- This sleight of hand is what let Smets and Wouters claim a better fit than BVAR's in their initial Euro area paper.
- If BVAR's were fit to differenced, demeaned, data, Minnesota prior is inappropriate.
- If BVAR's were fit to levels, recursive RMSE's problematic — BVAR will spend "fit energy" on matching divergent trends among the series.

Splitting of parameter vector

- As in many such models, this paper calibrates many of the model parameters to match steady states to sample means or on the basis of micro studies, then estimates remaining parameters from prefiltered data that removes low frequencies.
- Models that can match the actual divergent trends of many real variables are complicated to specify and difficult to linearize.
- Nonetheless, much of the promise of the RBC, and hence of its successor DSGE, modeling approach is that in principle the models are *growth* models. It is a shame that we lose this promise in applied models.

Short sample

- The EA has only existed for a relatively short period, and the paper uses a similarly short time series for both the EA and the US, to make results comparable.
- This is throwing away information.
- Very important information. We may be about to repeat the history of the late 70's and early 80's. Throwing out the data from this period is not the way to look forward from our current situation.

4 Conclusion

Summing up

- The paper is an important step toward a policy model that can realistically model financial frictions.
- The importance of the Fisher debt deflation effect is interesting and should be explored more widely by modelers generally. We need competitors for Calvo pricing as a source of non-neutrality.
- The empirical implementation leaves a lot of room for improvement.