Whither News Shocks?

A Proposal for the 2014 NBER *Macroeconomics Annual*

Robert B. Barsky and Susanto Basu[1]

Overview

After the original suggestion of Pigou (1927), the idea that macroeconomic fluctuations might be driven by news about future events languished for the better part of a century. In the last decade, however, the hypothesis has become the subject of an important discussion within macroeconomics, due in large part to the pioneering work of Beaudry and Portier (2004, 2006). Apart from interest in the question for its own merits, the investigation of news shocks has also been driven by the growing realization that it is difficult to identify structural shocks that might be responsible for the bulk of business cycle fluctuations. Identified shocks to technology do not appear to induce the right business-cycle comovements.[2] Monetary shocks induce the right comovements, but their variance is too small. The importance of investment-specific technology shocks appears to decline when the shock is asked to explain the relative price of investment.[3] Other shocks are often either implausible (for example, high-frequency labor supply shocks) or difficult to identify without conditioning on a full DSGE model (for example, markup shocks).

News shocks are appealing for several reasons. First, they are *a priori* plausible sources of fluctuations. Second, they fix some of the unappealing features of technology-driven business-cycle models. For example, recessions can occur without technological regress in models with news shocks; downturns can take place after even positive realizations of technology change, as long as the realizations were smaller than expected. Third, news-driven models of cycles don't require large and frequent changes in current fundamentals, since news about the same realization of future technology can change many times. Thus, it is unsurprising that news-driven stories feature in casual explanations of the recession of 2001, and to some extent even the Great Recession of 2008.

On the other hand, news-driven models have problems of their own. Notably, in standard neoclassical models, news shocks do not produce the correct business-cycle comovements. The reasoning follows directly from the insight of Barro and King (1984): without a current increase in the marginal product of labor, consumption and labor supply must move in opposite directions. So good news shocks may raise consumption, but they will then lower investment, labor input and output. Much theoretical work has been devoted to augmenting neoclassical models so that good news shocks will be expansionary in these models.[4]

In our proposed paper, we wish to pursue the empirical question of what news shocks actually do using a variety of VAR identification schemes, and then ask what models are capable of

---

[1] Federal Reserve Bank of Chicago and NBER and Boston College and NBER, respectively. Depending on our findings, the title may be changed to "Wither News Shocks."

[2] See, for example, Gali (1999), Francis and Ramey (2005), and Basu, Fernald and Kimball (2006).

[3] Justiniano, Primiceri and Tambalotti (2010, 2011).

[4] For example, Beaudry and Portier (2007) and Jaimovich and Rebelo (2009).

explaining the empirical effects of news shocks. The first set of questions we ask are, Can we identify news shocks from VARs, while taking account of the important issue of invertibility? If so, how? How can we be sure that the shocks we identify are really news shocks? These are the questions we propose to address in the empirical portion of our paper.

Based on our preliminary results, we think we can indeed identify news shocks. We find that these shocks have some properties consistent with simple neoclassical models: they raise consumption, but lower investment and hours worked, just as Barro and King (1984) might have predicted. Furthermore, they raise real interest rates, as they should if consumption is expected to rise in the future. However, we also find results that are not so simple to explain in simple neoclassical models. First, we find that stock prices rise. While this may seem intuitive—after all, the shock represents good news about future marginal products of capital—it is not straightforward to have stock prices and investment move in opposite directions in a neoclassical setting. Second, we find that perhaps the strongest effect of a news shock is that it causes an immediate and persistent disinflation. While the disinflation gives us further evidence that we are indeed identifying a news shock, which on average raises future full-employment output, the fact that it has an immediate effect on inflation is difficult to square with a neoclassical model. New Keynesian models might have this property, since the shock might imply that expected future marginal costs will be lower, but will New Keynesian models rationalize the inflation findings, while simultaneously preserving the appealing short-run implications of the simple neoclassical model for other variables? Furthermore, will a model that can rationalize the full set of findings for news shocks also be consistent with the evidence on the effects of other identified shocks, particularly technology shocks? These are the questions we propose to address in the theoretical sections of our paper.

Finally, we wish to step back and in the light of our findings assess what the news shock literature might contribute to macroeconomics, and suggest the directions that new research on the topic should take.

Identifying News Shocks with VARs

Our method will be to focus on the object

$$\left[E_t - E_{t-1}\right] TFP_{t+k},$$

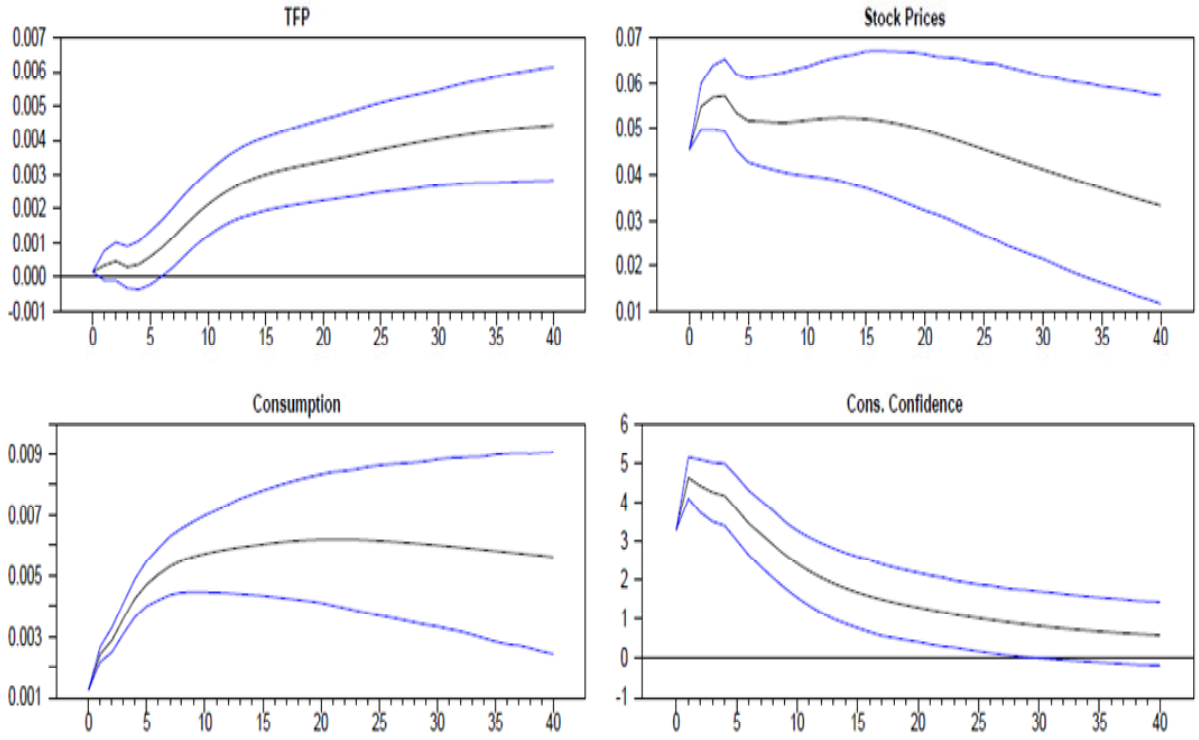
where E is the expectations operator and the forecast variable is utilization-adjusted TFP (from Fernald, 2012). This method is different from those used by either Beaudry-Portier (2006) or Barsky-Sims (2011). It is preferable to the scheme of Barsky and Sims because it is more closely connected with the fundamental definition of a news shock, but it is related to that method conceptually and turns out to give similar empirical results.

Denote a VAR in companion form as

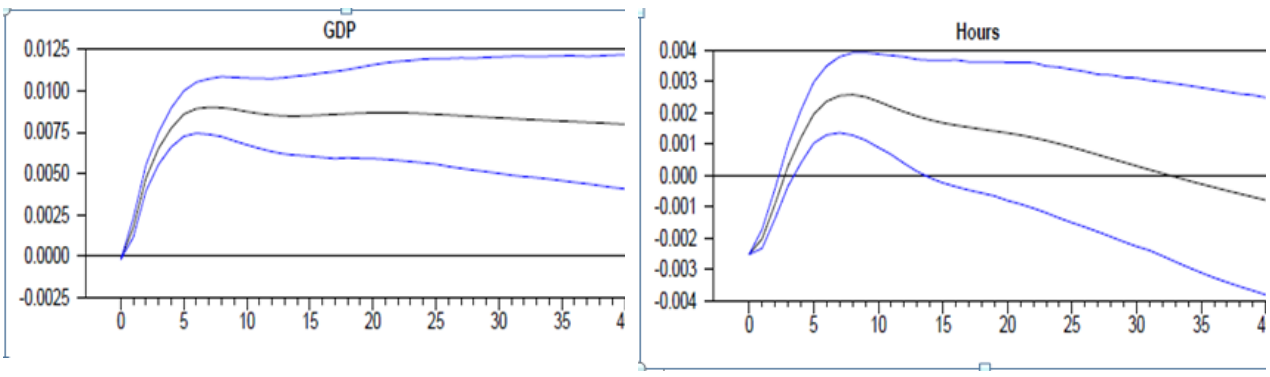
$$Y_t = AY_{t-1} + \varepsilon_t .$$

Note that in this case $\left[E_t - E_{t-1}\right]Y_{t+k} = A^k \varepsilon_t$. Choose the element of this vector associated with TFP. At a mechanical level, one can always regard the shock as one of a number of shocks in a structural or semi-structural VAR. We then compute the impulse responses to this shock (possibly after orthogonalizing with respect to current TFP). We begin by setting *k* to 10 years.

The new method produces impulse responses that look like news shocks: TFP is basically unchanged for several quarters after the news is received, but consumption rises steeply, and stock prices and consumer confidence both jump:

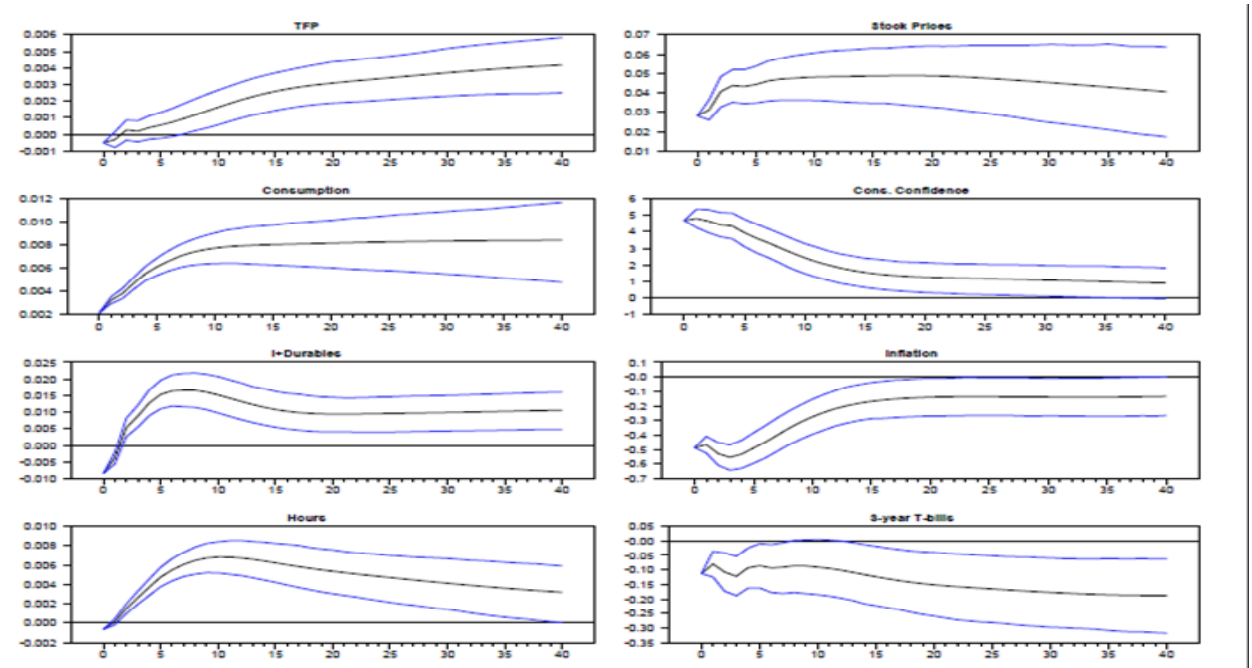


However, the business-cycle implications of this new identification scheme are not strongly in favor of a business-cycle model driven heavily by news and technology shocks. The shocks we identify as "news" still have a short-run contractionary impact effect on hours and investment (not shown), with essentially no impact effect on GDP:



However, there is ultimately a substantial expansionary effect on hours at a horizon of 1-3 years. We discuss this issue further below.

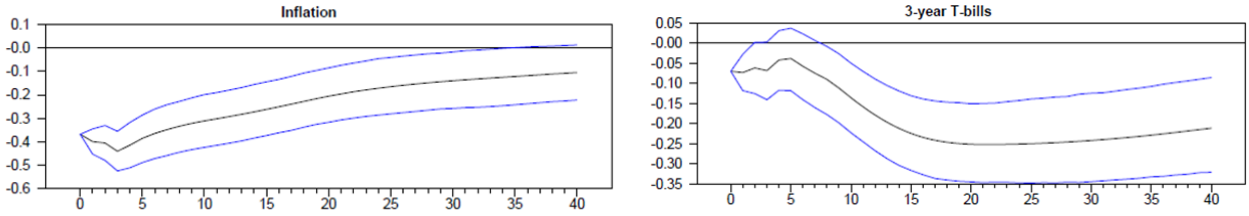
It is by no means obvious that one should orthogonalize with respect to current TFP. The arguments in favor of so doing are that (1) pure news shocks should by definition not affect current TFP, and (2) orthogonalization works well in simulations of DSGE models in which the news shock is constrained to be orthogonal to the current TFP shock. The arguments against orthogonalization are strong. First, the assumption might simply be wrong. For example, suppose the truth is that a single shock raises TFP on impact and then raises it further over time. Second, it is helpful to be able to look at various kinds of shocks simultaneously without having to order one against the others. Fortuitously, the shock that explains TFP 10 years ahead is approximately uncorrelated with the current innovation in TFP, just as Beaudry and Portier (2006) hypothesized, so the issue turns out not to be too important empirically. We show the impulse responses without orthogonalization below for comparison:



Responses to News Shock (unorthogonalized)

We also plan to investigate the important issue raised by Blanchard, L'Huillier and Lorenzoni (2012), on subject of whether news can be distinguished from "noise" using only VAR methods. We agree with them and with Barsky-Sims (2012) that one needs structural methods to disentangle the two completely. However, the VAR alone is enough to reject null that there is *only* noise, since in this case our shock would not predict the future evolution of TFP, at least in a setting where (utilization-adjusted) TFP is exogenous. We plan to explore this issue further, perhaps with the aid of Monte Carlo simulations.

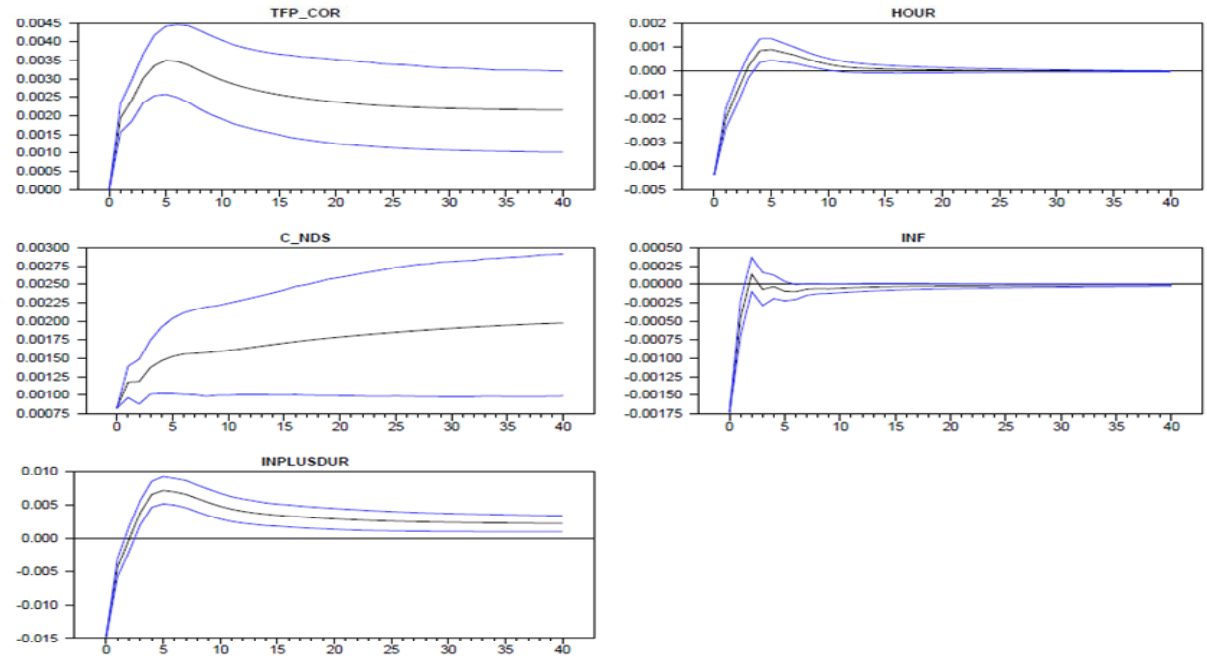
Finally, we find the responses of inflation and interest rates to our news shocks both interesting and challenging. Note that the combination of the two shows that real interest rates rise in response to good news:



News shocks in DSGE models

We plan to explore the effects of news shocks in simple DSGE models, to see which models are capable of reproducing our empirical findings. We shall also use Monte Carlo simulations to investigate the robustness of our identifying assumptions.

Early indications are that the method works well in Monte Carlo studies. Jinnai (2011) extends the New Keynesian model of Barsky and Sims (2011) to include capital and sticky nominal wages. We plan to do more extensive checks with a variety of models, but his model captures the basic findings well. In addition to predicting a positive effect on consumption, his model predicts declines in investment, hours worked and inflation, just as we find:



Responses to News Shock

Note from both the calibrated results and the empirical impulse responses discussed earlier that there is significant positive comovement between output and hours in response to a news shock at medium frequencies. Is this a victory for the news theory of business cycles? We think not. Recall that today's news shock is (on average) tomorrow's rise in actual productivity. We would expect higher productivity—whether anticipated or not—to raise output (permanently) and hours (temporarily). It *can* be said that a substantial part of hours and activity variation is due to anticipated productivity changes: this result shows up clearly in variance and historical decompositions. But this is not an effect of news *per se*, since the pure effects of news pertain to the periods when the TFP increase is anticipated but has not yet occurred.

Non-technological shocks and endogenous technology

While almost all models of news shocks assume that the news regards future exogenous technology changes, we also plan to investigate two alternative hypotheses. First, we ask how the previous results and methods change if the news is about something other than technology. Second, we plan to investigate what happens if future technology is endogenous.

Suppose that the news is not about future technology but about a non-technological shock. We are still discussing how best to represent such a shock, but for now take an exogenous shock to the markup of price over marginal cost as a prototypical "demand" shock. News about a future markup decline may have some of the effects we find for news shocks: in particular, investment is likely to decline when the news is received (due to intertemporal substitution). The same effect is also present for non-durable consumption, but is weaker and the wealth effect goes the other way, so consumption may rise. Finally, intertemporal substitution in labor supply may also explain the behavior of hours worked. However, one variable that is not likely to respond as we find in the impulse responses to news shocks is TFP, at least if we use utilization-adjusted TFP.

Now suppose we extend the model to allow technology to be endogenous. For example, suppose that TFP depends on R&D investment, which, due to credit market imperfections, depends partially on firms' cash flows. Then the future markup shock will also raise TFP endogenously. (This is essentially the model of Comin and Gertler (2006), augmented with a news shock.) The endogenous TFP improvement will make it more likely that the responses of variables in the model will match the empirical impulse responses to news shocks discussed earlier.

Thus, we conjecture that the impulse responses we discussed earlier may arise from news about exogenous future TFP or from news about future "demand" shocks, particularly when the latter are combined with models of endogenous technological change. We plan to investigate how the importance of these two sources of news can be disentangled from one another.