

Whither News Shocks?

Robert B. Barsky, Susanto Basu, and Keyoung Lee*

Does news about future productivity cause business-cycle fluctuations? What other effects might it have? We explore the answer to this question using semi-structural VARs, where “news” is defined as the innovation in the expectation of TFP at a fixed horizon in the future. We find that systems incorporating a number of forward-looking variables, including stock prices, consumption, consumer confidence and inflation, robustly predict three outcomes. First, following a news shock, TFP rises for several years. Second, inflation falls immediately and substantially, and stays low, often for 10 quarters or more. Third, there is a sharp increase in a forward-looking measure of consumer confidence. Consumption typically rises following good news, but investment, consumer durables purchases and hours worked typically fall on impact. All the quantity variables subsequently rise, as does TFP. Depending on the specification of the reduced form VAR, the activity variables may lead TFP to some extent – possibly lending some support to the hypothesis of news-driven business cycles – or they may move in lockstep with productivity. For the most part, the quantity and inflation responses are quite consistent with the predictions of a standard New Keynesian model augmented with real wage inertia.

* Federal Reserve Bank of Chicago and NBER, barsky@frbchi.org; Boston College and NBER, susanto.basu@bc.edu; and Federal Reserve Bank of Chicago, klee@frbchi.org. We thank Larry Christiano for many helpful discussions, and Larry Christiano, Franck Portier, conference participants and the editors for useful comments. The paper does not reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

I. Introduction

In the last decade or so, the old “Pigouvian” hypothesis that good news about the future may cause business cycle expansions - which in turn are reversed by less favorable news - has again become the subject of an important discussion within macroeconomics, due in large part to the pioneering work of Beaudry and Portier (2004, 2006). The investigation of news shocks has been driven substantially 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 induce the right business-cycle comovements¹, nor can technology shocks yield anything close to the degree of forecastable variation in hours and output seen in real data (Rotemberg and Woodford, 1996). Monetary shocks may 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.² 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 even after positive realizations of technology change, as long as the realizations were smaller than expected. Third, news-driven models of cycles do not 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 even the

¹ See, for example, Gali (1999), Francis and Ramey (2009), and Basu, Fernald and Kimball (2006).

² Justiniano, Primiceri and Tambalotti (2010, 2011).

Great Recession of 2008. There is considerable appeal to the story that agents' demand for durable goods (including housing) preceding each recession was predicated on high growth expectations that—whether or not rational at the time—proved to be excessively optimistic, requiring a reduction in subsequent purchases and (at least in the case of the Great Recession) a concomitant deleveraging process.

On the other hand, news-driven models have a fundamental theoretical difficulty 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.³

This paper focuses on the empirical identification of news shocks and their effects. As we proceed with the empirical work, we attempt to understand and rationalize what we find in the data in terms of “DSGE reasoning” with and without formal DSGE models.

As far as the identification issue is concerned, we focus the lion's share of our effort in the area of structural (or semi-structural) VARs, for two reasons. First, the bulk of the substantive debate about the empirical effects of news shocks has used the methods of VAR analysis. Second, we believe that the substantially (though certainly not completely) nonparametric nature of the VAR approach (in contrast to fully specified DSGE models) maximizes the likelihood that the answers will be driven by the data rather than by model structure. Having identified the properties of news shocks and shown that they are robust to

³ For example, Beaudry and Portier (2007), Jaimovich and Rebelo (2009)

reasonable variations in specification, we turn to interpreting the results through the lens of a fully specified DSGE model using Monte Carlo simulation methods.

Before asking whether news shocks cause business cycles, one ought to ask whether news shocks do in fact exist to a nontrivial extent. Somewhat to our surprise, we find that they do. More precisely, we find that an innovation to the expectation of TFP at $t+20$ quarters that is orthogonal to current TFP explains between 20 and 40 percent of the variance of TFP at a horizon of five years, and typically in excess of 50 percent of the variance at a horizon of 10 years. Since we take the information set to be orthogonal to current TFP, the innovation we identify is a news shock by a stringent definition. Thus, it appears that news shocks exist, and are quantitatively important. A second important preliminary observation is that news shocks have the character of anticipated growth shocks. In response to our nonparametrically estimated news shocks, TFP displays a rather smooth increase over a number of quarters. Though the error bands are regrettably large, there is no positive evidence that the effect of news shocks on the level of TFP reverts even in the long run. This is in contrast to the parametric news processes in the original modern theoretical news shock paper (Beaudry and Portier, 2004), versions of which are often employed in DSGE models. These processes typically feature an anticipated *jump* in productivity at more or more dates in the future. On the other hand, an important early empirical paper on news shocks (Beaudry and Portier, 2004) estimates a smooth TFP response to news shocks which is far more drawn out than the more detailed work in this present paper indicates.

The use of a battery of forward-looking variables, including consumer confidence, stock prices, and consumption is key to our ability to identify an important role for predictable fluctuations in the level of future TFP. Thus, even though the univariate representation of TFP is best characterized as a random walk, implying the growth in TFP is unpredictable, the same is

not true in a multivariate context. The Granger causality from multiple forward-looking variables to TFP makes TFP growth predictable. An analogous point was made in the context of GDP by Cochrane (1994) and Rotemberg and Woodford (1996), both of whom emphasized the role of the consumption share of GDP in predicting future output. Shocks to linear combinations of innovations in the forward-looking variables constitute our estimated news shocks. More specifically, according to our conceptualization the empirical analogue of a news shock is the linear combination of reduced form VAR innovations that constitutes the revision in the optimal forecast of TFP at some fixed future horizon.

Turning now to the substantive results that follow from our identification of news shocks, we find that the most statistically significant and robust results concern not the quantity variables that are the usual focus of business cycle analysis, but inflation and the expectations component of consumer confidence. Strikingly, the historical decompositions of our VAR system show that our identified news shocks mimic closely innovations in consumer confidence. Equally robust and of perhaps even greater interest is the finding that inflation declines by 30-40 basis points after a news shock and in a typical specification stays below its pre-shock level for 10 quarters or more. This is unusual behavior if one thinks of news shocks as “demand shocks”—that is, shocks that raise output but are not technological in nature - but it is rationalizable in at least some variants of the New Keynesian model, which stresses the role of future real marginal costs in the determination of current inflation.

We of course go on to ask whether news shocks are plausibly significant drivers of economic fluctuations. As noted earlier, few if any of the shocks that economists have identified so far (at least using data-driven procedures consistent with large classes of interesting models) can simultaneously match the comovement that is a defining characteristic of business cycles and

also explain a large fraction of economic fluctuations at frequencies identified with economic fluctuations. Can our identified news shocks fill this gap?

While the case is not completely open and shut, the short answer appears to be no. As previously demonstrated by Barsky and Sims (2011), the impact effects of news shocks clearly does not induce the kind of comovement that is characteristic of business cycles. In most of our specifications, we find that consumption rises when there is good news, but investment, consumer durables purchases and hours worked all fall on impact. As discussed above and in Barsky and Sims (2011), this is exactly what one would predict in neoclassical models where there is a representative consumer with time-separable preference, the case analyzed by Barro and King (1984). In such settings, news shocks have an income effect but no substitution effect, so consumption and work hours must move in opposite directions, as must consumption and investment. Our results concerning the negative impact effect of news on hours, durables, and investment are weaker than those of Barsky and Sims (2011), suggesting that these authors may have overstated their case – though Barsky and Sims also cautioned briefly against an overemphasis on impact effects. In this paper we go much farther in that direction. One might easily imagine that Beaudry and Portier, as well as other advocates of news-driven theories of business cycles, would be quite happy with an empirical model in which the business cycle expansion is somewhat delayed, perhaps due to adjustment costs or other sources of inertia.

Unsatisfactory impact effects of news shocks, however, are only the first stage in the uphill battle that expectations-driven business cycle theory faces in light of the empirical evidence. As such a theory predicts, in subsequent periods consumption, investment and durables purchases all rise strongly, and in some specifications hours of work do as well. However, unlike Beaudry and Portier (2006), we find that TFP begins rising markedly one or

two periods after the news shock. Thus, it is difficult to argue convincingly that the expansion in business-cycle variables following a news shock is due to the effects of news per se rather than to the change in fundamentals that is heralded by a news innovation. The hallmark of a news-driven business cycle theory must be that quantity variables – particularly hours and investment in consumer and producer durables – rise *in advance* of a predictable rise in TFP. It is certainly possible to find specifications of the reduced form VAR in which there is a significant element of this phenomenon, and large confidence bands further militate against a definitive rejection of its existence. However, the canonical case appears to be one in which productivity and quantity variables move largely in tandem.

How well can a New Keynesian DGSE model explain the delayed response of real variables alongside the immediate negative response of inflation to a news shock? We take up this topic in the last section, constructing a New Keynesian model that is quite standard except for the addition of real wage rigidity. As stressed by Christiano et al. (2010) and Barsky and Sims (2011), expected future productivity improvements, *ceteris paribus*, lower expected real marginal costs. If the real wage does not rise too sharply (either on impact or over time), inflation will jump down and stay down, potentially for a number of quarters – just as in the data. Whether this happens or not depends on the persistence of the news process, the monetary policy rule (particularly the extent to which the monetary authority observes a good proxy for the natural real interest rate), and the degree of price and wage stickiness.

II. Semi-Structural VAR Identification of News Shocks

A. Basic Method

Let's turn to VAR identification. This is of course intimately tied up with the question of the extent to which quantitatively important news shocks exist in the first place. If the data support a finding that indeed news shocks do exist, we will be interested in the shape of the impulse responses of standard macroeconomic variables (including the usual real variables as well as inflation and asset prices) to the news shock, as well as to the surprise technology shock. We will also study variance decompositions, and historical simulations. At this point, the reader should regard what we are doing as establishing facts regarding news shocks in a relatively nonparametric way, not constrained by a structural model. In a later section of the paper, we will compare the impulse responses from the empirical VAR to those from Monte Carlo simulations of a simple DSGE model that illustrates some important mechanisms underlying the observed effects of news shocks, and turns out to fit a number of the facts rather well.

A structural news shock will be defined as an advance signal that agents receive about future productivity. One example of a stochastic process for log TFP, a , that contains a news shock would be:

$$a_t = \rho a_{t-1} + v_{t-1} + \varepsilon_t \quad 0 < \rho < 1$$

This is the baseline news specification of Christiano, et al (2010). A second example (Barsky and Sims, 2009, 2011; Jinnai, 2013) is:

$$\begin{aligned} a_t &= a_{t-1} + g_{t-1} + \varepsilon_{1t} \\ g_t &= (1 - \rho) \bar{g} + \rho g_{t-1} + \varepsilon_{2t} \end{aligned}$$

In this specification productivity has a time-varying trend, and agents observe the innovation in that trend one period in advance. Hence the news shock is $\varepsilon_{2,t-1}$. Since the impulse response to

the news shocks we identify via vector autoregression typically displays a rather smooth increase in productivity over several quarters, the latter process is closer to what is nonparametrically estimated in the data.

How do we identify news shocks via vector autoregressions? For the purposes of this paper, an identified news shock will be a linear combination of reduced-form innovations that 1) predicts future productivity; and 2) is not “excessively” correlated with current productivity (more particularly, with a productivity measure stripped of the endogenous, utilization-driven component). Criterion (1) needs no discussion (except for issues about the forecast horizon) – it is certainly a *sine qua non* for a news shock. Nor is it difficult to see why criterion (2) is critical. Consider the extreme case where productivity is a univariate random walk that is also not Granger-caused by any known variables. In that case, the best forecast of future productivity is current productivity; the innovations in the two are equivalent at all horizons, and a shock to current technology and a shock to the forecast of future technology are one and the same. Such an innovation is no doubt a kind of news (*all* innovations are), but hardly the kind of advance information about future technology with which the expectations-driven business cycle literature is concerned.

It is well-known that TFP growth is approximately white noise. Since the univariate TFP process cannot predict future TFP, the possibility of identifying news shocks from the VAR arises if and only if there are observable variables that Granger-cause TFP. These variables may or may not be the actual signals seen by the agents in the model, but they serve as *indicators* of the signals agents observe. They may be asset prices, survey measures of expectations, or macroeconomic variables such as consumption, investment, and hours that reflect – and hopefully reveal – the information possessed by the agents.

A positive innovation in the stock price orthogonal to the current productivity innovation but Granger-causing future productivity was *the* operationalization of a news shock in the pioneering empirical paper of Beaudry and Portier (2006). While there are a number of reasons not to rely too heavily on the stock market as an indicator of technological news (some of them discussed in Barsky and Sims, 2011), we retain the stock price as one indicator of a news shock. Consumer confidence also turns out to be a surprisingly valuable indicator (Barsky and Sims, 2009, 2012).

As a preliminary exercise (not shown here) we computed impulse responses from four separate bivariate VARs, each involving TFP and one forward-looking variable. In each case, TFP is ordered first. We first reproduced the key result in Beaudry and Portier (2006) – an innovation in stock prices orthogonal to current productivity presages a rather long period of increased productivity growth without any apparent subsequent reversion. Almost precisely the same pattern – with a very similar estimate of the long-run increase in the level of productivity – holds when the forward-looking variable is E5Y, the measure of five year business expectations from the Michigan Survey of Consumers. Next, inspired by the simple Permanent Income Hypothesis and Cochrane (1994), we considered the implications for productivity growth of the orthogonalized innovation in consumption of nondurables and services. Once again, we observe more or less the same pattern, with a somewhat more rapid step-up in productivity and a slightly smaller long-run effect (though the difference is clearly not statistically significant). Finally, informed by the canonical New Keynesian model in which inflation is a jump variable that presages future real marginal costs, we choose inflation for our forward-looking variable. The result is a mirror image of what we saw in the previous three. Thus stock prices are joined by

consumer confidence, consumption, and inflation as forward-looking indicators of future productivity.

Barsky and Sims (2012), used an agnostic VAR identification with a medium-sized VAR that includes consumption, GDP, hours, inflation, stock prices, consumer confidence and interest rates. Their identification strategy considers all shocks that are orthogonal to the innovation in current productivity, and among these - following Uhlig (2004) - chooses the shock that maximally explains a weighted average of future levels of productivity. They find that their identified news shocks raise consumption, but reduce hours, investment, and GDP in the short run. This finding holds across different VAR specifications in their paper, and is consistent with the results of standard neoclassical models. Additionally, news shocks seem not to account for historical recessions.

The maximization-based identification in Barsky and Sims is not entirely transparent, and there is an arbitrariness (inherited from Uhlig (2004)) about the weights attached to the various horizons over which technology shocks are to be explained. Furthermore, it is not absolutely clear that we want to impose *a priori* that the news shock be strictly uncorrelated with the contemporaneous technology innovation.

In this paper, we focus our attention on the object:

$$[E_t - E_{t-1}]TFP_{t+k}$$

where E is the expectations operator and the forecast variable is quarterly utilization-adjusted TFP (from Fernald (2012), who uses a subset of the corrections to standard TFP proposed by Basu, Fernald, Kimball (2006) to create a quarterly TFP series purged of its endogenous utilization component). This object is the innovation in the optimal VAR forecast of TFP at

some *fixed* point in the future. We vary this forecast horizon k sequentially, though we focus on what appears to be the reasonable benchmark of five years (20 quarters).

Is it clear that we ought to orthogonalize $[E_t - E_{t-1}]TFP_{t+k}$ with respect to the current productivity innovation, as is standard in the new shock literature? On the one hand, one might think that a “pure news shock,” by definition, should not affect current TFP. The current line of DSGE modeling for news shocks reflects this idea that news shocks are shocks that affect productivity in the future while being unrelated to current (true) productivity. On the other hand, a “news shock” in the real world may not have this kind of feature. It is possible that news about future productivity arrives along with innovations in productivity today. Positive technological innovations today may be introduced today with the understanding that significant improvements to technology will occur in the years to come. Or there may be gradual diffusion of a new general purpose technology across sectors. In these cases, a single structural shock raises TFP on impact and then raises it further over time. In such cases, orthogonalizing the news shock with respect to innovation in current TFP is simply the wrong identification. Therefore, it would be nice to examine both kinds of “news shocks,” one where the innovation in the k -horizon-ahead forecast is orthogonalized with respect to the reduced form innovation in current TFP and one where it is not.

The forecast innovation approach of this paper allows us to study both of these cases. If we find that the unorthogonalized news shock looks quite similar to the orthogonalized one, so that (subject to small sample bias) the *data alone* tell us the effects of TFP news unrelated to current TFP, then we would be in an ideal world in terms of defining and identifying news shocks – no identifying restrictions would be needed. Indeed, this happy coincidence obtains in the levels specification for the 40-quarter horizon. Unfortunately, the same is not the case for the

shorter horizons of 12 and 20 quarters.⁴ Indeed, three years ahead much of the forecast innovation is due to the contemporaneous innovation. We regard five years as a reasonable benchmark horizon for predictions of future productivity, and this is the horizon on which we focus. Here we do orthogonalize with respect to current TFP so that the resultant median impulse response looks like what the literature has come to regard as a news shock. Statistically, however, even if we do not impose the orthogonality, it is not rejected by the data, and in this sense can be seen as an *overidentifying* restriction.

We see this forecast innovation approach as preferable to the scheme of Barsky and Sims because it is more closely connected with the fundamental definition of a news shock.⁵ Note that the forecast is computed entirely from the reduced form VAR. We see this as a major virtue, because it allows us to go as far as we can with the data alone before imposing any *a priori* restrictions. In the case where the VAR is entirely in differences and k goes to infinity it amounts to the Beveridge-Nelson procedure (which is not to say that either the differencing or the infinity are advisable choices.)

To see how we compute the above object, denote a VAR in companion form as

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

Note that in this case, $[E_t - E_{t-1}]Y_{t+k} = A^k \varepsilon_t$. We choose the vector of the matrix A^k associated with TFP to arrive at the linear combination of reduced form residuals. At a mechanical level, one can always regard the shock as one of a number of shocks in a structural or semi-structural

⁴ A finding that at the 40 quarter horizon the unorthogonalized forecast variance is more or less identical to the orthogonalized one is equivalent to saying that the effect on TFP of the contemporaneous innovation in productivity has largely disappeared after ten years. One must be concerned that the 40 quarter result is due in part to downward bias in the largest autoregressive root in the productivity equation.

⁵ Our method is of course conceptually related to Barsky-Sims (after all, optimal forecasts are the solution to a maximization problem) and overall it gives rather similar – though certainly not identical – results.

VAR. We then compute the impulse responses to this shock after orthogonalizing with respect to current TFP. We focus on k equal to 20 quarters (5 years). Such a “medium run” identification approach has been shown to have considerably better small sample properties in some controlled settings (Francis et al., 2005).

B. A Brief and Informal Digression on Invertibility

Much has been written about the “invertibility problem” (see, for example, Fernandez-Villaverde et al., 2007). A DSGE model (which typically has a state space representation in the form of a VARMA) is *invertible* if that VARMA can be reduced to a VAR – i.e. if the structural shocks in the DSGE model are linear combinations of reduced-form VAR residuals. At this level, invertibility is about the suitability of the *reduced-form* VAR for studying the underlying DSGE model, and prior to the problem of imposing identifying restrictions on the VAR. If it is known that the news shock cannot be represented as a linear combination of VAR residuals, there is no point in trying to “structuralize” the VAR. In our context the underlying economic reason for potential noninvertibility is that agents in the economy might well receive signals about the future that are not innovations with respect to variables observed by the econometrician.

Earlier writers often thought that the invertibility problem is a disaster for structural inference from VARs. We take the (still controversial but increasingly respectable) position that invertibility is not a devastating problem. This is one place in which we are consonant with Beaudry and Portier (2013). The reason has to do with recent refinements in understanding invertibility issues at both the theoretical and practical level. The new understanding relates invertibility to insufficient richness either in the variables included in the underlying DSGE

model or in the data available as indicators of news shocks rather than to the existence of news itself.

The whole idea of a model of news-driven business cycles revolves around concrete economic decisions that agents make in response to news. Signals that are observed by economic actors but not reflected in their behavior are of limited interest. If news shocks are quantitatively important, and a rich enough set of observable indicators is included in the model and available to the econometrician, then news should be reflected in such activity variables as consumption, investment, and hours as well as in stock and bond prices. News may also be partially revealed in measures of survey confidence.

Roughly speaking, the richer the underlying DSGE model is in observable variables, the more likely it is that we will be able to uncover the news shock and their impulse responses, either exactly (in large samples) or approximately. At an analytical level, this point is reflected in the example given at the end of Fernandez-Villaverde et al. (2007). These authors exhibit a permanent-income model that is noninvertible because the econometrician does not have enough information to infer news that agents receive about permanent income. In a somewhat anticlimactic turn of events at the end of the paper, it is noted that the information set of the econometrician did not include consumption. Once consumption is observed, the system becomes invertible and once again a VAR can be used to recover the underlying shocks of the model. Not every case will be this straightforward, but the general lesson is clear.

Moreover, at a practical level, it may not even be critical that the DSGE model be invertible in the analytical sense. In an important paper, Sims (2012) argues - and demonstrates via Monte Carlo simulations - that VARs with news shocks and rich information sets may give good estimates of theoretical impulse response from DSGE models that are technically not invertible.

In this sense, invertibility is not an “on/off” issue. Simulations of a New Keynesian model without capital in Barsky and Sims (2009) and a real business cycle model with capital in Barsky and Sims (2011) are favorable to the VAR identification method (conditional, of course, on the “true” model being of the same variety as the DSGE model studied).

III. Data

We use standard national income accounts data on gross investment, purchases of consumer durables, and consumption of non-durables and services (aggregated into a single Divisia index). We express each variable in per-capita terms by dividing by the civilian, non-institutional population. Hours worked are the BLS measure of aggregate nonfarm payrolls hours, again put on a per-capita basis. The stock price variable is Shiller’s real S&P 500 index; the interest rate is the 3-month Treasury bill rate; inflation is measured by the CPI-U. The consumer confidence measure is from the Michigan Survey of Consumers. Data on quarterly, utilization-adjusted TFP are from Fernald (2012), who uses a subset of the procedures proposed by Basu, Fernald and Kimball (2006) to create a quarterly TFP series purged of the endogenous utilization component. When necessary, we convert the growth rates in Fernald’s TFP series to an index in log levels. We take logs of the quantity variables and the stock price. See Appendix I for details.

IV. Empirical News Shocks

We have run a number of specifications in our VAR exercises with news shocks. The differences in specification concern the handling of nonstationarity (and the associated issue of cointegration), and the forecast horizon. The *levels* specification needs no explanation, except to point out that it is often thought to be desirable because of its robustness to specification error concerning the number of cointegrating vectors, and that it can be given a Bayesian justification

even when some of the variables are neither stationary nor cointegrated with others. The *hybrid* specification puts nonstationary variables in differences (thus assuming an absence of cointegration - a feature which is in fact not particularly supported by statistical tests), and presumptively stationary variables (consumer confidence, inflation, and interest rates) in levels. Of course, if cointegration holds then there are potential significant benefits to imposing it. However, evidence has been accumulating in recent decades that the US economy might be better understood through the lens of a two-sector growth model rather than a one-sector model. In a two-sector model, aggregate TFP would no longer be cointegrated with either nondurables consumption or with investment and consumer durables purchases, so that the hybrid specification might be preferable. That said, we also present results of an error correction model in which we impose cointegrating factors of unity relating TFP to nondurables and services, durables, and investment, but let the data determine the cointegrating relationship between TFP and the stock market.

Figures 1 and 2 display the impulse responses to news shocks and unanticipated technology shocks, respectively, for the levels specification. Figure 3 shows the impulse responses to news shocks for the hybrid specification, with Figure 4 showing the responses to news shocks for the ECT specification. As previously noted, our baseline model is run with a 5-year forecast horizon in identifying the news shock. We have also run the model with alternative forecast horizons, identifying news shocks as orthogonalized innovations in the expectation of TFP three and ten years in the future. To save space, we do not report the results here, but we discuss which of our findings are sensitive to the choice of horizon. The full results for the different horizons are available from the authors on request.

A. Discussion of Empirical Results

Across all our specifications for the effects of news shocks, our most consistent finding of the effects of news on a standard business-cycle variable is that inflation falls immediately, substantially and persistently in response to a “good” news shock. The results in the level specification of Figure 1 are representative. A current news shock that by assumption has no effect on TFP today but raises TFP by about 0.2 percent in three quarters and an additional 0.2 percent 20 quarters after the shock, lowers inflation on impact by more than 0.25 percentage points. Inflation continues to fall over the next year, reaching a maximum drop of 0.5 percentage points four quarters after the shock. Inflation then rises slowly from this trough, but is still 0.1 percentage point below its pre-shock value 10 quarters after the news shock.

The literature on news shocks has not focused much attention on this very robust finding, with Barsky and Sims (2009), Christiano et al. (2010) and Jinnai (2013) being notable exceptions. We feel this result merits substantially more discussion, in part because it is difficult to understand without a model where there is substantial interplay between real and nominal variables. We thus devote a separate section to understanding and explaining this finding in the context of popular DSGE models of business cycle.

The next most robust findings are the effects of news shocks on consumer confidence and the stock market (though results for the latter are not always statistically significant). In every specification in our baseline VAR model, stock prices jump up when there is a positive news shock, and typically continue to rise for about two years. At its peak, the stock market reaches a level between 2 and 4 percent above its pre-shock value.

Unlike the stock market, which takes between five and 10 quarters to reach its peak, consumer confidence jumps on impact and either is at its peak immediately, or reaches that peak

within a quarter or two. Confidence then remains significantly above its pre-shock level for years, typically 15 quarters or more.

In Figure 5 we plot the time series of our variables and the fitted values that can be attributed to news and unanticipated technology shocks. We use results from the levels specification to generate the figure. To isolate business-cycle comovements, after computing the decomposition, we filter both the data and the fitted values using the bandpass filter, set to isolate cycles with frequencies between 6 and 32 quarters.

The historical decomposition for inflation shows that the news shock explains a large fraction of cyclical inflation movements. The comovement is particularly impressive in the 1970s and early 1980s, but unlike the case for other variables, the news shocks continue to play an important role in explaining inflation fluctuations in the 1990s and 2000s.

If the decomposition of inflation is striking, that of consumer confidence is even more so. It is clear that to a large extent a news shock *is* a consumer confidence shock. But whereas confidence is often discussed as a measure of animal spirits, an exogenous driver of aggregate demand, it is clear from our results that confidence is driven heavily by fluctuations in expectations of TFP. This is a welcome confirmation of Barsky and Sims (2012), who come to this conclusion using a quite different approach involving a New Keynesian structural model. The decomposition for the stock market, by contrast, shows that while both news and unanticipated TFP are positively correlated with changes in the real value of equities, the relationship is not a close one, and is especially weak in recent decades.

We now turn to exploring the effects of news shocks on quantity variables. First is the effect of the shock on the future path of TFP. Strikingly, news shocks predict future TFP growth at high levels of significance. Furthermore, once realized, the path of TFP following a news

innovation resembles a growth rate shock, rising predictably over time. As is well known, growth rate shocks create different incentives for intertemporal substitution than level TFP shocks. Thus, in addition to the fact that (orthogonalized) news shocks create a wealth effect with no substitution effect on impact, even when the news shock leads on average to an increase in future TFP, the wealth effects in the early periods after TFP begins to rise should be strong relative to the substitution effects. The historical decomposition shows that news shocks are definitely correlated with TFP changes, although the bulk of changes in TFP at the included frequencies are due to the unanticipated shocks.

We now turn to the more controversial results regarding the impulse responses of quantities to news. This of course leads to the most frequently posed question about news shocks – are they an important cause of business cycles? In almost all of our specifications, we find that investment, hours and consumer durables purchases do not rise on impact of a news shock. Thus, news shocks do not on impact move these endogenous variables in the direction that would be expected by an advocate of news-driven business-cycles. This is particularly true for hours worked. The failure of consumption to rise more strongly may appear to be at odds with the notion of a news shock as a positive wealth shock or with the idea that consumption should also, as a forward-looking variable, react strongly to news. Interestingly, the historical decomposition shows a pattern that will become familiar as we discuss the other quantity variables. It is clear that news does drive consumption to some extent, but it is also apparent that this pattern was much stronger in the early part of the sample, especially around the time of the dramatic oil price increases of the 1970s. However, consumption is expected to rise over time following a positive news shock, which in turn requires that real interest rates rise at the time of the news. Indeed,

this is what happens: the nominal interest rate falls, but expected inflation falls further, so real interest rates rise when there is good news.

As noted, in two out of our three specifications hours worked decline when a news shock hits (although the decline is never statistically significant). In no case do hours worked rise significantly on impact of a news shock. This pattern of results is consistent with the idea that news shocks constitute a positive wealth shock for consumers. If consumption and leisure are both normal goods, then an increase in wealth with no offsetting substitution effect should cause hours worked to decline and consumption to rise, which is the pattern we generally see in the data. It then follows that investment and/or purchases of durable goods must fall. We see this qualitative pattern in Figures 1, 3 and 4.

Turning to the historical decompositions, we see that for hours worked news shocks were important early in the sample, but cease to be so around the early 1980s. Similar results hold for investment and, to a lesser extent, for consumer durables.

Finally, we examine the behavior of unanticipated TFP shocks. For the sake of brevity, we report only the levels specification, in Figure 2. We see that the time path of TFP following the shock is mean-reverting, although in the hybrid specification (not shown) TFP follows basically a random walk. In general we see that consumption, investment, worked and durables purchases all rise following a positive TFP shock. Hours first rise and then fall, and are just significant at the 90 percent level in the impact period. Interestingly, it appears that inflation rises after an increase in TFP (significantly so after 8 quarters), which appears difficult to reconcile with sticky-price models, where a mean-reverting TFP shock typically lowers marginal cost for several periods. It is also inconsistent with the findings of Basu, Fernald and Kimball (2006), who found that inflation fell significantly for several years after a surprise improvement

in technology. Stock prices either decline or fail to rise significantly, while consumer confidence typically rises by a small amount on impact and then falls. (We found the same results for the hybrid and ECT specifications.) In general, the results following unanticipated TFP shocks are difficult to reconcile with the theoretical predictions for positive technology shocks, especially in sticky-price models with limited monetary accommodation (see, for example, Gali, 2008).

These results lead us to question whether our “purified” TFP series has really succeeded in purging the Solow residual of its non-technological component. As Fernald (2012) discusses carefully, the quarterly series that he constructed and we use cannot make all the corrections that Basu, Fernald and Kimball (2006) were able to make using annual data, because some of the source data are not available at a quarterly frequency. But if more non-technological components remain in Fernald’s TFP series, then an unanticipated rise in his measure may be a composite of a change in technology and an endogenous change in utilization coming from a non-technological “demand” shock. This hypothesis would explain why we find, contrary to Basu, Fernald and Kimball (2006), that an innovation in TFP does not lower inflation and raises hours worked.

Since our identification of news shocks relies on predicting changes in Fernald’s TFP measure, one might ask whether this problem also calls our identification of news shocks into question. We believe it does not, because our identification rests on forecasting TFP k quarters in the future, where we take k to be a large value (such as 20, or even 40, quarters). Over such long periods of time, predictable changes in utilization should be essentially zero. This can be verified by checking the predictability of hours per worker at such horizons, which according to the model of Basu, Fernald and Kimball (2006) is an observable proxy for unobserved factor utilization.

Looking further at the impulse responses to the news shocks, we see that hours, though contracting a bit on impact (less than in Barsky and Sims 2011), *expand* quite a bit at medium frequencies, especially as do the various components of output. We confirm this visual impression with the variance decompositions for the hybrid specification in Table 1, which show that at a 20-quarter horizon news shocks explain about 30 percent of the variance in consumption and purchases of consumer durables and 16 percent of investment, although only 5 percent of hours. Table 2 shows variance decompositions for the cointegrated specification, which appear more favorable for the news-driven business cycle hypothesis. At a 20-quarter horizon, news shocks account for 68 percent of the variance of consumption, about 30 percent of the variation in investment and consumer durables, and 40 percent of the variance of hours worked. This looks quite a bit more like a news-driven business cycles, albeit with a delay. The results for the level specification, in Table 3, are if anything even more favorable for the “news shocks” view of economic fluctuations.

To what extent are the variance decompositions in Table 2 evidence for the news theory of business cycles? To answer this question, it is important to take into consideration the behavior of TFP following a news shock. Recall that *today's news shock is (on average) tomorrow's rise in actual productivity*. Higher productivity, whether anticipated or not, *will* raise output permanently and (probably) hours temporarily—although perhaps after a delay, as in Basu, Fernald, and Kimball (2006) and Gali (1999). That hours rise at medium frequencies may not be the effect of news *per se*, but rather the effect of the subsequent actual improvement in technology. Note from Table 1 that at a 20-quarter horizon, the news shock “explains” 40 percent of the variance of TFP (rising ultimately to 55 percent at 40 quarters) while only accounting for a relatively smaller fraction of hours (5 percent) and investment (16 percent). So

on one hand, the rise of hours and investment at medium frequencies is an important result. But on the other hand, if those variables are merely rising in conjunction with TFP, then it is difficult to say that news *per se* is causing the business cycle movements.

Thus the critical issue for the evaluation of the news-driven business cycle hypothesis is the extent to which hours, investment, and durables rise ahead of TFP, or only contemporaneously with it. To the extent that activity variables rise sharply before cumulative TFP growth has been substantial, we have evidence of news-driven cycles. In the hybrid specification (Table 1) this is clearly not the case. In the levels and error correction specifications, the case for news-driven cycles is somewhat stronger. Some substantial movement in investment and hours leads productivity. Yet these activity variables continue to rise as productivity growth picks up. This is perhaps suggestive of some combination of a pure news effect and a direct productivity effect on activity.

In the case of the hybrid specification, noting that activity variables pick up only as does productivity, we propose a heuristic exercise to decompose the effects of a news shock into the “pure news” effect and the effect of actual (predictable) technology change. Since this is a counterfactual experiment, it is difficult to do without a full structural model. However, we can get an idea of the importance of “pure news” versus the realization of technology by combining results from our estimated impulse responses to news shocks and to unanticipated technology shocks. The idea is to use the impulse response to unanticipated technology to uncover the “direct effect” of technology on the various outcome variables. Then the pure news effect is computed as the impulse response to news subtracting the imputed direct effect of the technology change.

This procedure would be rigorous if the path of TFP following a news shock (after TFP began to increase) followed exactly the same stochastic process as an unanticipated TFP shock. To see why, consider an RBC model, where the state variables are just capital and TFP, and the current level of TFP summarizes all the expectations for the future path of TFP. Suppose that TFP always follows the same stochastic process whether it is anticipated or unanticipated. However, some TFP changes are anticipated j periods in advance. Then the response of economic variables after anticipated and unanticipated TFP changes would differ only because the starting values of the capital stock would be different: an anticipated news shock would create incentives to accumulate or de-cumulate capital over those j periods, while an unanticipated shock by definition would not. Assuming that the effects of capital and technology are additive (which would be true in a linearized model), the impulse responses to unanticipated technology would allow us to uncover the effects due to the change in the capital stock from the reaction of agents to the anticipated shock, which would be the effects of news *per se*. The same result would hold in more complex linearized models with more state variables (for example, nominal prices or wages).

Unfortunately, the impulse responses to anticipated and unanticipated TFP shocks are not similar in shape. The path of TFP following a news shock looks like the response to a stationary growth rate shock, with the univariate time series process for technology resembling an AR(1) in first differences. The response of TFP to its own unanticipated shock looks like a stationary AR(1) in levels or, in the hybrid specification, resembles a random walk. Assuming rational expectations on the part of economic agents, the differences in shape imply different incentives for intertemporal substitution in consumption and labor supply. Thus, our calculations can only be suggestive.

However, to mitigate the problems just outlined, we take as our baseline the effects of an unanticipated TFP shock 10 quarters after the shocks occurs. This is roughly when TFP is halfway back to its pre-shock value, and where TFP and other variables are roughly at their long-run values in the hybrid specification shown in Figure 3.

Thus, we implement the following calculation:

$$\text{"Pure News Effect"}_k = \text{IRF|News}_k - \frac{\text{TFP|News}_k}{\text{TFP}_{10}} \text{IRF|TFP}_{10}.$$

We implement this imputation of the pure news effect using the hybrid specification. The results are shown in Figure 6.

We find that in the hybrid case this correction is significant. The news responses of consumption, investment and durables are cut roughly in half by correcting for the TFP change following a news shock. The hours response is reduced even further, and hours are now negative for the first 6 quarters following a news shock. A news shock that ultimately raises TFP by about 0.7 percent has a maximum positive effect on hours of about 0.1 percent. On the other hand, the correction actually increases the decline in inflation and interest rates, leaves consumer confidence basically unaffected, and has only a modest effect on the stock price response. These results cast some doubt on the claim that the data support the news-driven business cycle hypothesis. On the other hand, the corresponding corrections for the level and error correction models (not shown) are - as expected - smaller. Once again, these latter two specifications are somewhat more supportive of the notion of purely news-driven cycles.

C. Robustness

We extend our basic results in a variety of ways to check the robustness of our core findings. First, we see whether our results change if we allow for a two-sector model of technical change rather than the one-sector model we have implicitly employed so far. Second, we assess

robustness to using different measures of news, especially to the idea that news about the second moment of technology (“risk shocks”) is more important than news about the first moment shocks we have considered so far.

i. Two-sector results

In our earlier discussion of the potential pitfalls of cointegration – and the concomitant case for the hybrid specification – we noted that evidence supports idea that US data are better modeled as a two-sector growth model. In particular, relative price data strongly indicate that technology in the production of investment goods and consumer durables has been advancing at a rate faster than the technology for producing other forms of output. Indeed, Ben Zeev and Khan (2013) found that investment-specific technology shocks, measured as unanticipated drops in the relative price of investment, trump neutral technology shocks in the explanation of business cycle movements. Thus, we checked the robustness of our previous results by estimating a specification that allows for two different shocks to technology, and therefore news about these shocks separately. Thus, in addition to our basic specifications, we estimated a model where we included consumption-specific and investment-specific TFP as separate variables, TFPC and TFPI. The TFP series used in our previous specifications is a weighted average of the two, but standard economic models imply that the effects of the two shocks should be quite different, so it makes sense to separate the two measures.⁷

The measure of investment technology is Fernald’s (2012) utilization-corrected TFP in equipment and consumer durables. Our TFPC is also from Fernald, and is intended to measure technology in the other industries. Thus, we can estimate impulse responses to investment or

⁷ For discussion of the business-cycle implications of investment-specific technical change, see Greenwood, Hercowitz and Krusell (2000) and Basu, Fernald, Fisher and Kimball (2011).

consumption TFP news shocks, which are the forecast revisions in TFPI or TFPC at a horizon of 20 quarters, orthogonalized with respect to the unanticipated TFPC and TFPI shocks.

We found that news about consumption TFP has an interesting constellation of effects. It predicts a rise in consumption TFP of 0.2 percent starting about a year after the shock, but an even larger and more significant increase in investment TFP, almost 1 percent after 20 quarters. Consumption is unchanged on impact, but higher starting about 3 quarters after the shock. Durables consumption jumps down, but becomes positive four quarters after the shock. Hours worked are basically unchanged on impact, and also becomes positive four quarters after the shock. Investment, by contrast, jumps up and generally keeps increasing. As is our earlier results, investment falls on impact, and reaches its trough several quarters after the shock. The nominal interest rate also declines. Stock prices rise until four quarters after the shock. Consumer confidence, by contrast, rises on impact, and stays high for an extended period. The confidence bands for these results are unfortunately very wide.

From 5 to 15 quarters after the shock, investment, consumption, hours and durables purchases are higher, so over this horizon the fluctuations in these variables exhibit the comovement one expects over the business cycles. But of course this is also the period when consumption and investment TFP are noticeably higher following the news shock. Thus, the 2-sector results basically give results similar to those of the levels and ECT specifications, which gives us more confidence in these results. Thus, once one takes into the account the distinction between true news effects and the effect of realizations of anticipated investment-specific technology improvements, the results of Ben Zeev and Khan (2013) that investment-specific news shocks have much larger business than neutral TFP shocks ceases to be compelling.

ii. Other Measures of News

We also experimented with a third approach to the empirical identification of news shocks, an approach based on publication counts of science and technology titles, based on the pioneering work of Alexopoulos (2011). While developed to discuss contemporaneous technology shocks, her time series data on the number of science and technology titles published seemed to us to be potentially suited for interpretation in terms of the news shock concept, albeit at a relatively short horizon. Interestingly, Alexopoulos (2011, table 4) finds that her measure of computer publications Granger-causes the utilization-adjusted technology measure of Basu, Fernald and Kimball (2006) (p-value of 0.02). The correlation between the two series at one- and two-year horizons is very small (on the order of two percent). Alexopoulos also finds that an innovation in her measure is positively correlated with changes in GDP, investment and hours worked. After two years, the correlation between the measure of computing-related publications and utilization-adjusted TFP jumps to about 0.20 (her table 3B). This suggested to us that the effects on GDP and other business-cycle variables in the first two years following an innovation in Alexopoulos's publication series might be fruitfully interpreted as being due to the effects of news *per se*. If so, her measure, interpreted as a news shock, has potential for explaining business cycles as traditionally defined.

It thus seemed natural to run our reduced form VARs with the addition of the Alexopoulos' time series of computer titles as an additional variable, and interpret the innovation in that measure (or perhaps the innovation in the first principal component of that measure and stock prices) as a proxy for news shocks. Unfortunately, these data are available only at an annual frequency. The error bands in all of the specifications we tried were far too wide to make any inference that would add to what Alexopoulos already found, and we do not include the impulse responses in the paper, though they – along with the data – are available on request.

In recent work, Christiano et al. (2010) and Christiano, Motto, and Rostagno (2013) (henceforth CMR) find that in an estimated DSGE model with financial frictions along with the now-standard features of sophisticated New Keynesian empirical models, “risk news shocks” drive out technology news shocks. Risk shocks as viewed by CMR manifest themselves as increases in the cross-sectional standard deviation of TFP at the firm level. Since the underlying theoretical model is closely related to that of Bernanke, Gertler, and Gilchrist (1996), the ultimate significance of the CMR measure is presumably that those firms that do relatively poorly have a high probability of defaulting on their obligations. CMR find that *news* about *future* risk is a dominant source of the business cycle, while TFP news is of relatively trivial importance.

It is important to note that CMR have a quite different aim than do we in the present paper. They are seeking a shock (or at most, a few shocks) that maximize(s) the ability to explain business cycle variation. We, on the other hand, are seeking to identify and understand the effects – large or small – of a particular pre-specified shock (namely, the TFP news shock) on real and nominal macroeconomic variables. Nonetheless, in light of the CMR results in the DSGE context, it seems appropriate to ask whether the inclusion of a proxy for risk news shocks in a VAR changes the magnitude of the identified TFP news shocks, or the behavior of their impulse responses.

To address this question we chose as our proxy for risk news shocks the interest rate spread between Moody’s Aaa bonds of twenty or more years of maturity and the corresponding Baa bonds. Since the duration of these bonds (at least at recent interest rates) is quite long, this spread responds to information about future bankruptcy risk over a long horizon, and thus represents news about future risk.

We included this spread variable in addition to our other variables in both the levels and hybrid specifications (not shown, see web appendix), ordering it first so as to maximize its potential to dwarf the importance of TFP news shocks. In all of our specifications, we found that the risk news shock has essentially no effect on the impulse responses or variance decompositions of the TFP news shock that we presented earlier. This result holds whether or not the TFP news shock is orthogonalized with respect to the risk spread and/or the unanticipated TFP shock. As expected, we found in both the levels and hybrid specifications that durables, investment, hours, inflation, treasury bill yields fall substantially on impact in response to a spread shock.⁸

V. Can a DSGE Model Explain the Nominal and Real Effects of News Shocks?

By far the most robust result that we obtain across different specifications is that good news shocks are highly disinflationary. The strong disinflationary response – on the order of half a percentage point or more for a one standard deviation news shock – appears immediately on impact and in most specifications lasts for a number of quarters. Whether or not they are a cause of business cycles, news shocks are very important for understanding inflation.

The disinflationary effects of news shocks appeared in previous literature in two guises. Barsky and Sims (2009, 2011) included inflation in their VAR systems – just as in the present paper – and observed that the identified news shock was associated with a sharp drop in inflation on impact. Christiano et al. (2010) took a more indirect approach, in which the empirical observation and its theoretical rationale were more closely intertwined. They looked at a number of stock market booms, and noted that they were consistently associated with low inflation. They then constructed a model in which news shocks – in combination with a monetary policy in

⁸ Somewhat surprisingly, confidence and stock prices respond positively in the levels specification, though not in the hybrid specification.

which the Fed did not properly account for changes in the natural rate of interest (e.g. by including it as an intercept in the Taylor Rule as in Woodford (2003)) – accounted for both the rise in the stock market and the drop in inflation. The underlying problem is that news shocks cause a rise in the natural rate of interest, while the Taylor rule without the time-varying intercept leads to a drop in the interest rate in response to falling inflation.

Both Barsky and Sims (2009, 2011) and Christiano et al (2010) regard the disinflationary effect of news shocks as a manifestation of the forward-looking nature of inflation in the New Keynesian model. Consider the now-classic New Keynesian relationship between inflation and future expected real marginal costs (Gali and Gertler, 1999; Sbordone, 2002):

$$\pi_t = \frac{(1-\theta)(1-\beta\theta)}{\theta} \sum_{s=0}^{\infty} \beta^s E_t mc_{t+s}$$

where θ is the Calvo price adjustment parameter and β is the discount factor. In the benchmark model, the log deviation of real marginal cost is equal to the log difference between real wages and labor productivity: $mc_t = w_t - p_t - y_t + n_t$. Expected future productivity improvements, *ceteris paribus*, lower expected real marginal costs. If the real wage does not rise too sharply (either on impact or over time), inflation will jump down and stay down, potentially for a number of quarters – just as in the data. Whether this happens or not depends on the persistence of the news process, the monetary policy rule (particularly the extent to which the monetary authority observes a good proxy for the natural real interest rate), and the degree of price and wage stickiness. Figure 7, which simply augments our level specification with real compensation per hour and labor productivity, shows that in fact real wages do lag labor productivity considerably following the news shock. This lends empirical support to the focus on the forward-looking New

Keynesian relationship and illustrates the importance of finding some mechanism to limit the extent to which real wages rise along with labor productivity.

Our objective is to present a DSGE model in which the impulse responses of key variables to a news shock match the empirical responses we estimate. Our criterion for “matching” is that the theoretical impulse responses lie within the confidence intervals for the empirical responses. Thus, our exercise is akin to that of Christiano, Eichenbaum and Evans (2005), although we fit the impulse responses “by eye,” rather than choosing the model parameters to maximize a formal goodness-of-fit criterion. However, recent research has shown that VAR impulse responses can be subject to significant small-sample bias. That being the case, Kehoe (2006) argues that one should compare VAR impulse responses estimated from actual data to impulse responses estimated from data of the same sample size generated by simulating the DSGE model. This is the comparison that we perform.

Our assumed law of motion for technology is once again:

$$a_t = a_{t-1} + g_{t-1} + \varepsilon_{a,t}$$

$$g_t = \rho_g g_{t-1} + \varepsilon_{g,t}$$

where a_t is the log of technology, and g is the steady-state technology growth rate. In this framework, $\varepsilon_{a,t}$ is a surprise technology shock and $\varepsilon_{g,t-1}$ is the news shock, which has no effect on actual technology in the period that the shock hits.⁹ All lower-case variables denote log deviations from the steady state. The process is chosen so that the news shock leads to a sustained growth in technology starting from the following period, as in our VAR results.

Technology finally asymptotes to a permanently higher level following a news shock. Since the

⁹ Some of the VAR impulse responses suggest that the main rise in TFP begins two quarters after the shock. It would obviously be easy to modify the stochastic process for technology to introduce a two-quarter delay, rather than our current choice of one quarter.

growth-rate shocks are stationary, the model can be solved by log-linearization around the non-stochastic balanced growth path.

The real block of the model is a fairly standard one-sector model with endogenous capital accumulation. Here we list the main model features; the log-linearized model equations and the parameter values are listed in Appendix II. The model has habit formation in consumption, capital adjustment costs as in Hayashi (1982) (as opposed to the “higher-order” adjustment costs of Christiano et. al (2005)), Cobb-Douglas production, and government purchases financed by lump-sum taxes. In order to estimate a VAR for a number of variables using simulated model data while avoiding a stochastic singularity, the model has a number of structural shocks. In addition to the unanticipated level technology shock and the news shock discussed above, there is a government purchase shock, a labor supply preference shock, a shock to the Euler equation (which may be interpreted as a financial friction shock), and a monetary policy shock.

In order to obtain a definite prediction for inflation, we need to add nominal features to the model. We begin by assuming price stickiness, governed by the standard New Keynesian Phillips curve discussed above and a Taylor-type rule for monetary policy:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \varphi_\pi \pi_t + (1 - \rho_i) \varphi_y (y_t - y_{t-1}) + \varepsilon_{i,t}$$

This rule has the monetary authority respond only to variables that can be observed readily, without the need to construct an entire counterfactual model. The rule also assumes interest rate smoothing, which is generally acknowledged to be a realistic feature of monetary policy-making.

Our prior belief, based on Barsky and Sims (2009, 2011) and Jinnai (2013) was that this model would not be able to explain the persistent period of low inflation that follows a news shock in the data. Our reasoning was as follows: According to the NKPC combined with Cobb-Douglas production, labor productivity needs to grow faster than the real wage for real marginal

costs to decline, which is necessary for inflation to fall. But a news shock is unlikely to achieve this outcome, for two reasons. First, since the shock is to the growth rate of technology rather than the level of technology, labor productivity will not rise much in the short or medium term. Second, exactly because the shock is to the growth rate and will thus have a large effect on consumption in the long run, the wealth effect on labor supply should imply that real wages rise sharply on impact. Thus, inflation should rise rather than falling. Barsky and Sims (2009, 2011) and Jinnai (2013) argue that one needs some form of wage rigidity to counteract the wealth effect and stop wages from rising on impact of a positive news shock.

This reasoning, while correct, is incomplete. It leaves out the demand side of the labor market, and it turns out that this channel is sufficient to explain a long period of low inflation. Since purchasers of output also recognize that there has been a positive growth rate shock, they know that prices will ultimately decline significantly to reflect higher labor productivity. However, the combination of sticky prices and a shock to the growth rate rather than the level of technology implies that most of the price decline will come in the future, not on impact. Thus, there is an incentive to delay purchasing output until the price falls. This intertemporal substitution in the timing of purchases is particularly strong for investment, which has a very high elasticity of intertemporal substitution. Thus, on impact of a news shock, labor demand can fall more than labor supply, leading to a decline in real wages and thus in real marginal cost. (In other words, the short-run decline in demand interacts with nominal price rigidity to raise the equilibrium markup, which is equivalent to lowering real marginal cost.) In our model, the net effect is to keep inflation below trend for about 12 quarters after a shock, as we find in the data.

It turns out that a simple model with just nominal price rigidity misses mostly the medium-term surge in real, rather than nominal, variables that can be seen in the “levels”

impulse responses of Figure 1. For the reasons discussed above, in our calibrated model with only sticky prices, investment and hours either shrink or grow slowly for the first 15 quarters or so following a news shock. This contrasts with the empirical impulse responses, which show predictable, transitory booms in these variables from roughly 5 to 15 quarters following a positive news shock. It is to match this boom in real variables in the middle frequencies that we introduce our major deviation from the standard New Keynesian model, and augment our model with real wage inertia, as discussed by Blanchard and Gali (2007) and used in a similar context by Barsky and Sims (2009). We assume that wages are determined by the time-series process

$$w_t = (1 - \rho_w)w_t^* + \rho_w w_{t-1},$$

where w^* is the wage determined by the Frisch labor supply relationship $w_t^* = \frac{1}{\eta} n_t - \lambda_t$, where λ_t is the marginal utility of wealth. Real wage inertia creates a second reason why inflation stays low, as discussed by Barsky and Sims (2009). But it also creates a second reason for intertemporal substitution, this one leading purchasers to buy more in the short run. As discussed above, the wealth effect on labor supply leads real wages to rise in the long run following a positive news shock. Thus, purchasers of output need to balance two forces pulling in opposite directions. On the one hand, they are tempted to wait for the level of technology to improve further following a positive growth rate shock. This force suggests they postpone their purchases, and keeps output low. On the other hand, they want to purchase output before wages, and thus prices, rise substantially. Thus real rigidity acts to pull purchases forward in time. The net effect of the two is to keep output and labor input low on impact, but raise output and hours 5 to 15 quarters after a positive news shock, better matching the data.

The dashed lines in Figure 8 show the median of the impulse responses estimated from a VAR run on data from the simulated DSGE model. The solid lines reproduce the empirical

impulse responses from the levels specification of Figure 1, and the shaded areas show 90 percent confidence intervals around the empirical impulse response. For consumption and hours, the theoretical impulse responses agree closely with the empirical ones. The model responses for the nominal interest rate and inflation, on the other hand, while remaining within the confidence intervals, are generally higher than their empirical counterparts, while investment undershoots its empirical counterpart, especially 1-3 years after a news shock.

The problems with the nominal interest rate and investment may be linked: if the equilibrium of the model had interest rates being temporarily low, as they are in the data, then investment would probably be temporarily high, also matching the data. On the other hand, the extra demand for output might lead to counterfactual predictions for hours worked and lead to higher predicted inflation, thus worsening the model's fit for inflation even further. Overall, however, by the conventional criterion of producing model impulse responses that remain within the confidence intervals for the estimated responses, our modeling exercise is a success.

The model can be used to shed some interesting light on the claim of Christiano et al (2012) that the Fed moves the real interest rate in the wrong direction in response to news shocks, creating an inefficient boom. Our empirical VAR results show that the actual real rate *rises* somewhat following a good news shock. By turning off price rigidity (but not the real wage rigidity, which here is a real feature of the economy), we find that in our real wage rigidity model the response of the Fed to a news shock is in fact somewhat too contractionary, leaving an output gap that both contributes to the disinflation process and explains why, unlike Christiano et al (2012) whose model produces an immediate output boom in response to a news shock, we find slightly contractionary impact effects on hours, investment, and output.

VI. Conclusion

We found that systems incorporating a number of forward-looking variables, including stock prices, consumption, consumer confidence and inflation, robustly predict three outcomes. First, following a news shock, TFP rises for several years. Second, inflation falls immediately and substantially, and stays low, often for 10 quarters or more. Third, there is a sharp increase in a forward-looking measure of consumer confidence. The business cycle implications of news shocks are far more ambiguous. Although it is hard to definitively reject the hypothesis that news about future productivity might have a causal effect on hours and investment even before the news is realized, the modal tendency is for quantity variables to rise only after substantial realizations of productivity improvements. Thus a robust affirmative argument for the news-driven business cycle hypothesis will require a new line of argument that goes considerably beyond the both VAR methodology and current DSGE research.

Appendix I: Data

Productivity: We use the TFP data constantly updated and maintained by John Fernald on his website. For TFP, we use utilization-corrected TFP at a quarterly frequency, which is provided in growth rates, and convert it into log-levels. We also use utilization-adjusted TFP in producing equipment and consumer durables for our measure of investment TFP, and utilization-adjusted TFP in producing non-equipment output for our measure of consumption TFP. The TFP measures were retrieved on 3/12/2013.

Civilian Non-institutional Population: We convert many of our series into per capita terms by dividing them by the Civilian Non-institutional Population (not seasonally adjusted) for each quarter. This data is from Haver Analytics (PN16@EMPL).

Investment: We take gross private domestic investment (I@USNA) and make it into real terms using Gross Private Domestic Investment: Chain Price Index (JI@USECON). The series converted into per capita terms and log-levels. The data is from Bureau of Economic Analysis/Haver Analytics.

Durables Consumption: We take Real Personal Consumption Expenditures: Durables Goods from Bureau of Economic Analysis/Haver Analytics (CDH@USECON). The series is converted into per capita terms and log-levels.

Hours: We use aggregate nonfarm payrolls hours series from Bureau of Labor Statistics and retrieved from Haver Analytics (LHTNAGRA@USECON). The series is converted into per capita terms and log-levels.

Consumption: For consumption, we use non-durables and services consumption. Since consumption series for non-durables and services from Bureau of Economic Analysis are provided separately for consumption of non-durables and consumption of services, we use the two series to construct our own non-durables and services consumption measure. We get the data from BEA/Haver Analytics for real nondurable consumption (CNH@USNA) and real services consumption (CSH@USNA) and weight the two series using nominal shares. In particular,

$$d\tilde{c}_t = s_{n,t}dn_t + (1 - s_{n,t})ds_t,$$

$$\text{where } s_{n,t} = \frac{\text{nominal nondurables at time } t}{\text{nominal nondurables at time } t + \text{nominal services at time } t}$$

where $d\tilde{c}$ is the log difference in our consumption measure, dn is the log difference in real non-durables consumption, ds is the log change in real services consumption and s_n is the nominal share of non-durables consumption. We weight the share of real non-durable and services consumption by their respective nominal shares. We take the final consumption measure and convert it into per capita terms and log-levels.

Stock Prices: The stock price measure is Robert Shiller's real S&P500 Index. The data is available on Shiller's website. It is converted into per capita terms and log-levels. The data was retrieved 12/10/2012.

Consumer Confidence: The Consumer Confidence data is from the Michigan Survey of Consumers. We take the series *Relative* from Business Conditions Expected the Next 5 Years. The data is available online and was retrieved 1/2/2013.

Inflation: Inflation data comes from Bureau of Labor Statistics/Haver Analytics (PCUY@USECON). We take the monthly series for CPI-U: All Items in year-over-year percent change and take the last month's data from each quarter to convert it into quarterly series.

Treasury Bills: We take the secondary market 3-month treasury bills data from Federal Reserve Board/Haver Analytics (FTBS3@USECON) at a quarterly frequency.

Spread: We take Moody's Seasoned Aaa Corporate Bond Yield (FAAA@USECON) and Moody's Seasoned Baa Corporate Bond Yield (FBAA@USECON) from Haver Analytics, both with maturities of at least 20 years. Then we subtract the Baa from Aaa to create our spread variable.

W/P (Real Compensation): We take the quarterly, seasonally adjusted compensation per hour index (LXNFC@USECON) from the Bureau of Labor Statistics, "Productivity and Costs," retrieved from Haver Analytics. We convert the series into real terms by dividing it by the implicit price deflator (LXNFI@USECON), which comes from dividing the current dollar GDP by the output index provided in the "Productivity and Costs" table by the Bureau of Labor Statistics. Then, we convert the series into log-levels.

Y/L (Labor Productivity): We take the quarterly, seasonally adjusted real output per hour of all persons index (LXNFA@USECON) from the Bureau of Labor Statistics, retrieved from Haver Analytics. We convert the series into log-levels.

Appendix II: A DSGE Model with News Shocks

This appendix collects the equations of the log-linearized DSGE model used in Section V.

Lower-case letters represent log deviations from the steady state. Government spending, G , is also a log deviation from the steady state, but the lowercase g is already used to represent the stochastic growth rate of the economy. For similar reasons, we use x to represent investment, since i is used to denote the nominal interest rate.

$$a_t = a_{t-1} + g_{t-1} + \varepsilon_{a,t} \quad (1)$$

$$g_t = \rho_g g_{t-1} + \varepsilon_{g,t} \quad (2)$$

$$\lambda_t = \left(\frac{1}{(1-\tau)(1-\beta\tau)} \right) (\tau c_{t-1} - (1+\beta\tau^2)c_t + \beta\tau c_{t+1}) \quad (3)$$

$$0 = \lambda_{t+1} - \lambda_t + E_t[i_t - \pi_{t+1} + \varepsilon_{b,t}] \quad (4)$$

$$w_t = (1 - \rho_w) \left(\frac{1}{\eta} n_t - \lambda_t - \psi_t \right) + \rho_w w_{t-1} \quad (5)$$

$$\psi_t = \rho_\psi \psi_{t-1} + \varepsilon_{\psi,t} \quad (6)$$

$$y_t = \frac{C^*}{Y^*} c_t + \frac{X^*}{Y^*} x_t + \frac{G^*}{Y^*} G_t \quad (7)$$

$$k_t = (1 - \delta) k_{t-1} + \delta x_t \quad (8)$$

$$G_t - y_t = \rho_G (G_{t-1} - y_{t-1}) + \varepsilon_{G,t} \quad (9)$$

$$y_t = a_t + (1 - \alpha) k_{t-1} + \alpha n_t \quad (10)$$

$$w_t = mc_t + a_t + \alpha (k_{t-1} - n_t) \quad (11)$$

$$q_t = \gamma x_t - \gamma k_{t-1} \quad (12)$$

$$E_t[i_t - \pi_{t+1}] = E_t[(1 - v^*) mc_{t+1} + (1 - v^*) y_{t+1} - (1 - v^*) k_t + v^* q_{t+1} - q_t], \quad (13)$$

$$\text{where } v^* = \frac{1 - \delta}{\alpha \left(\frac{1}{\mu^*} \right) \left(\frac{Y^*}{K^*} \right) + (1 - \delta)}$$

$$\pi_t = \frac{(1 - \theta)(1 - \beta\theta)}{\theta} mc_t + \beta E_t[\pi_{t+1}] \quad (14)$$

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) \varphi_\pi \pi_t + (1 - \rho_i) \varphi_y (y_t - y_{t-1}) + \varepsilon_{i,t} \quad (15)$$

Equations (1) and (2) comprise the laws of motion for technology. $\varepsilon_{a,t}$ is a level technology shock, and $\varepsilon_{g,t}$ is the news shock, which does not affect technology at time t but changes the growth rate of technology for subsequent periods. Equation (3) defines λ_t , the marginal utility of wealth, in the presence of habit formation in consumption, where τ is the habit parameter, and β is the discount rate. Equation (4) is the consumption Euler equation, with ε_b representing a white-noise wedge between the real interest rate and the return on assets held by households. Equation (5) is the partial-adjustment equation for real wages, where ψ is a labor supply shock; equation (6) is the law of motion for ψ . Equations (7) and (8) are the intra- and inter-temporal accounting identities, and (9) defines a partial-adjustment equation for G that keeps the ratio of government spending to GDP constant in the long run with ε_G representing a government spending shock. Equation (10) is the production function, and (11) is the labor demand equation. (mc denotes marginal cost, the inverse of the markup.) Equation (12) defines marginal q as a function of the convex cost of adjusting the capital stock, where γ is the elasticity of the adjustment cost function. (13) is the capital Euler equation, where δ is the depreciation rate and μ^* is

the steady-state mark-up. Equation (14) is the conventional New Keynesian Phillips curve, with θ the Calvo parameter, and (15) is the monetary policy rule.

We set the following parameter values:

$$\begin{aligned} \rho_i &= 0.9, \phi_\pi = 1.3, \phi_y = 0.94, \theta = 0.76, \eta = 1.32, \gamma = 0.17, \tau = 0.7, \mu^* = 1.08, \alpha = 0.36, \delta = 0.03, \beta = 0.985, \\ \rho_G &= 0.95, \rho_g = 0.74, \rho_w = 0.95, \rho_\psi = 0.6, \\ \sigma_a &= 0.57, \sigma_g = 0.15, \sigma_i = 0.21, \sigma_G = 0.0025, \sigma_b = 0.0025, \sigma_\psi = 0.0025 \end{aligned}$$