Expectations Driven Business Cycles: An Empirical Evaluation

Eric R. Sims*

University of Notre Dame

June 16, 2009

Abstract

The hypothesis that business cycles are driven by changes in expectations about future fundamentals has recently been the subject of renewed interest. This paper implements a new approach to identify so-called news shocks about future technology. In the context of a VAR featuring several forward-looking variables and macroeconomic aggregates, the news shock is identified as the shock that best explains future movements in a utilization-adjusted measure of aggregate technology among all shocks that are orthogonal to technology innovations. In post-war US data, a favorable news shock is associated with an increase in consumption and declines in output, hours, and investment on impact. After the impact effects, aggregate variables largely track predicted movements in technology. These are roughly the predictions of a simple neoclassical model augmented with news shocks. The negative conditional comovement among macroeconomic aggregates on impact in response to a news shock stands in contrast to the strong positive unconditional comovement among these series in the data. Moreover, an historical decomposition indicates that news shocks fail to account for output declines in four out of the six most recent US recessions. These results suggest that news shocks are not a dominant source of business cycles.

*Contact information: 723 Flanner Hall, Department of Economics and Econometrics, University of Notre Dame, Notre Dame, IN 46556. Email: esims1@nd.edu. I thank John Fernald of the San Francisco Fed for providing me with his TFP data. I am especially indebted to Bob Barsky for many years of fruitful collaboration, discussion, and guidance. I am also grateful to Rudi Bachmann, Daniel Cooper, Erik Hurst, Lutz Kilian, Miles Kimball, Bernd Lucke, Matthew Shapiro, and numerous seminar participants for helpful comments and advice. Any remaining errors are my own.

1 Introduction

Despite much progress in our understanding of aggregate fluctuations, the underlying source of business cycles remains a mystery. Most modern theories assume that fluctuations are driven by changes in current fundamentals, such as aggregate productivity. The last several years have witnessed a resuscitation of a much older theory in which business cycles can arise without any change in fundamentals at all. The expectations driven business cycle hypothesis – originally advanced by Pigou (1927) and reincarnated in its modern form chiefly in Beaudry and Portier (2004) – posits that business cycles might arise on the basis of expectations of future fundamentals.[1] Often referred to as the news driven business cycle, theories of this sort are appealing for a number of reasons.[2] If favorable news about future productivity can set off a boom today, then a realization of productivity which is worse than expected can induce a recession without any actual reduction in productivity itself ever occurring. As such, this theory of business cycles immediately addresses several of the concerns with conventional theories of the cycle – booms and busts can happen absent large changes in fundamentals and no technological regress is required to generate recessions. This theory is also appealing in that it seems to be a plausible explanation for several boom-bust episodes, with the stock price acceleration of the 1990s and ensuing recession in 2001 touted as a good recent example.

The chief difficulty faced by proponents of the expectations driven business cycle hypothesis is that it has proven extremely challenging to make news shocks about future fundamentals work in the context of relatively standard business cycle models, a point first recognized by Barro and King (1984) and later emphasized in Cochrane (1994). In a standard neoclassical setting, the wealth effect of good news about future productivity causes households to desire more consumption of both goods and leisure. With no change in labor demand, the inward shift in labor supply leads to a reduction in equilibrium employment and output. Falling output and rising consumption necessitate a fall in investment. Not only does good news about the future tend to cause a recession today, the implied negative comovement among macroeconomic aggregates is difficult to reconcile with the strong positive unconditional comovement of these series in the data.

In sharp contrast to the predictions of standard neoclassical models, recent empirical ev-

[1] This theory of business cycles is not to be confused with "sunspot" theories (e.g. Farmer (1998)) in which non-fundamental shocks can induce fluctuations. Expectations driven business cycle models generally presume rational expectations and a unique equilibrium in which agents receive stochastic signals of future fundamentals which are, in expectation, correct.

[2] This terminology differs from the literature on the effect on macroeconomic news on economic aggregates (Anderson, Bollerslev, Diebold, and Vega (2003)). I will follow the terminology introduced by Beaudry and Portier in referring to signals about changes in future productivity as news shocks.

idence suggests that news shocks about future productivity do induce positive comovement among the major macroeconomic aggregates. In particular, Beaudry and Portier (2006), Beaudry, Dupaigne, and Portier (2008), and Beaudry and Lucke (2009) propose two alternative VAR-based schemes for identifying news shocks. In one, these authors associate stock price innovations orthogonalized with respect to total factor productivity (TFP) with the news shock. In the other, they combine short and long run restrictions, identifying the news shock as the structural shock orthogonal to TFP innovations which has a long run effect on TFP. Under either orthogonalization scheme, their identified shocks are associated with a large, broad-based economic expansion occurring in anticipation of future TFP improvement.

While suggestive, these empirical findings are far from conclusive. A problematic feature is that the effect of the identified shock on TFP tends to be very delayed, leaving a large share of the variation in TFP unexplained at business cycle frequencies. Another is that these authors rely solely on a measure of stock prices as an "information" variable to help forecast future movements in TFP. Related to the work in Barsky and Sims (2008, 2009), consumer confidence and inflation also convey information about future productivity growth, much of which is not in fact revealed immediately in stock prices.

In this paper I thoroughly examine the empirical evidence in favor of the hypothesis that business cycles are driven by expectations about future productivity. I extend the identification strategy of Barsky and Sims (2009) to study the business cycle implications of news shocks about future productivity. I estimate a VAR featuring a utilization-adjusted measure of total factor productivity (hereafter "technology"), stock prices, inflation, consumer confidence, output, consumption, and hours. I identify the news shock as the shock contemporaneously orthogonal to technology which best explains future movements in technology not accounted for by its own innovation. In practice, this involves finding the linear combination of reduced form innovations in the VAR orthogonal to technology which maximizes the sum of contributions to technology's forecast error variance over a number of horizons. This approach is an application of principal components, and is a straightforward extension of the maximum forecast error variance identification proposed by Francis, Owyang, and Roush (2007) in a different context. It is also similar to work by Uhlig (2003).

This approach to identifying news shocks has a number of advantages over more traditional approaches to identification in VARs. Most importantly, there is a one to one correspondence between theory and identification. The feature which most models of expectations driven business cycles share in common is that only a limited number of shocks ever impact technology. My identification strategy imposes this implication of theory directly, while placing no other restrictions on the dynamic relationships among the other variables in the empirical VAR. This identification strategy explicitly seeks to minimize the unexplained variation in technology at short and long horizons, and therefore directly addresses the difficulty with previous work that the identified shocks fail to account for important variation in technology. I provide an overview of the details of my empirical strategy in Section 2. There I also present simulation evidence that my approach is in fact capable of recovering news shocks and their effects on aggregate variables from data generated by DSGE models.

In Section 3, I apply my empirical strategy to post-war US data. The news shock I identify begins to affect technology soon after impact, and explains a large share of technology movements at both short and long horizons. This contrasts in an important way with the news shocks identified by Beaudry and Portier (2006), Beaudry, Dupaigne, and Portier (2008), and Beaudry and Lucke (2009), which only begin to have a noticeable effect on technology after a period of several years. The main result of the paper is that a favorable news shock is associated with an increase in consumption and modest declines in output, investment, and hours of work on impact. After the impact effects, the macroeconomic variables largely track, rather than anticipate, movements in technology. While the identified news shock does appear to account for important long run movements in measured technology, it accounts for only modest shares of the forecast error variances of aggregate variables at short horizons. An historical decomposition suggests that news shocks fail to account for output declines in four out of six US recessions since 1961.

These results have important implications for macroeconomic modeling. Motivated in large part by Beaudry and Portier's (2006) findings, a number of authors have searched for theoretical frameworks capable of subverting the contrarian predictions of a standard neoclassical model augmented with news shocks. In particular, Beaudry and Portier (2004), Christiano, Ilut, Motto, and Rostagno (2007), Den Haan and Kaltenbrunner (2006), and Jaimovich and Rebelo (2008) all produce variants of the standard neoclassical model in which news about future productivity is capable of replicating the salient business cycle fact of comovement. While the underlying mechanisms in these models are quite different, they share the common feature that in each output, hours, investment, and consumption all jump up in anticipation of future technological improvement.

The positive conditional comovement in response to a news shock implied by these models stands in stark contrast to the impulse responses I estimate in the data. In fact, my estimated responses to a news shock are in rough accord with the qualitative predictions of a basic real business cycle model augmented with these shocks. In Section 4, I show that a standard calibration of the parameters of this model generates theoretical responses to a news shock which are a surprisingly good fit with my estimated empirical responses. I simulate data from the standard RBC model driven only by news shocks and show that it yields unconditional second moments completely at odds with post-war US data. That the model roughly matches the conditional responses to a news shock but fails to match the unconditional moments of the data when driven only by news shocks further suggests that some other disturbance(s) must be the main driving force behind fluctuations.

The remainder of the paper is organized as follows. In the next section I lay out the details of my empirical strategy. Section 3 presents the main empirical evidence and provides a comparison of my results with those in the existing literature. Section 4 offers a brief discussion of the expectations driven business cycle hypothesis in light of the empirical evidence. The final section concludes.

2 Empirical Strategy

My identification of news shocks is based on a ubiquitous assumption in the literature on expectations driven business cycles – that a limited number of shocks account for variation in technology. In particular, I assume that there are two distinct technology shocks – one that affects a measure of technology contemporaneously and one which affects it with a lag. Letting A_t denote an empirical measure of technology, this assumption can be expressed in terms of the moving average representation:

$$\Delta \ln A_t = [B_{11}(L) \qquad B_{12}(L)] \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix}$$

$\varepsilon_{1,t}$ is the conventional surprise technology shock and $\varepsilon_{2,t}$ is the news shock. The only restriction is that $B_{12}(0) = 0$ – that the news shock have no contemporaneous effect on technology. I place no other restrictions on the shapes of the responses of technology to these shocks, or on whether one or both of these shocks permanently affect technology.

It would not be possible to separately identify these two shocks in a univariate context – identification of news shocks must come from surprise movements in variables other than observed technology. As such, I implement my identification strategy in the context of an unrestricted vector autoregression (VAR) featuring a measure of technology and a number of other variables. Relative to the existing literature in this area, I include a large number of variables in the VAR – seven in the benchmark system, as opposed to two to five in the work of Beaudry and Portier (2006), Beaudry, Dupaigne, and Portier (2008), and Beaudry and Lucke (2009). In particular, I include a utilization-adjusted measure of total factor productivity ("technology"), stock prices, inflation, consumer confidence, consumption, output, and hours. The inclusion of additional variables represents an improvement along at least two dimensions. First, rather than including only one "information" variable (stock prices), I also include measures of inflation and consumer confidence, as these series are shown to be informative of future movements in productivity in Barsky and Sims (2009).[3] Like stock prices and explicit survey measures, inflation is a forward-looking variable in many models with nominal rigidities. Second, I am able to jointly estimate the responses of aggregate variables to a news shock, as opposed to doing so one or two at a time.

I identify the surprise technology shock ($\varepsilon_{1,t}$) as the reduced form innovation in technology. I then identify the news shock ($\varepsilon_{2,t}$) as the structural shock from the system that comes as close as possible to satisfying the identifying assumption that two shocks alone account for variation in technology.[4] In particular, the news shock is identified as the structural shock that best explains future variation in technology not explained by its own innovation. This identification strategy is an application of principal components, identifying the news shock as the linear combination of reduced form innovations orthogonal to technology that maximizes the sum of contributions to technology's forecast error variance over a finite horizon. It is similar to the maximum forecast error variance strategy proposed by Francis, Owyang, and Roush (2007), which in turn builds on Faust (1998), and is closely related to Uhlig's (2003) strategy of finding a small number of shocks to which to attribute movements in GDP.

The details of my empirical strategy are found in Appendix 6.1. The existing VAR-based identification strategies in the literature on expectations driven business cycles are special cases of mine which hold under more restrictive assumptions. Beaudry and Portier (2006) and Beaudry, Dupaigne, and Portier (2008) associate the news shock with the stock price innovation orthogonalized with respect to technology. Were the conditions for this restriction to be valid satisfied, my approach would (asymptotically) identify the same shock and impulse responses. As shown in Barsky and Sims (2009), there appear to be important movements in stock prices unrelated to technology shocks altogether.[5] Beaudry and Portier (2006) and Beaudry and Lucke (2009) also propose identifying news shocks with a combined recursive and long run restriction. This identification strategy is valid under the same assumption I make, but is problematic for two reasons. First, long run restrictions often perform poorly in finite samples (Faust and Leeper (1997)); Francis, Owyang, and Roush (2007) show that medium run identification similar to what I pursue performs significantly better in simulations. Second, a long run restriction fails to exploit the stronger implication that news and contemporaneous technology shocks completely explain variation

[3] The inclusion of additional forward-looking variables will also help to ameliorate any potential invertibility issue – see Watson (1994) or the discussion below.

[4] Imposing this restriction exactly would require highly restrictive assumptions on the VAR coefficients.

[5] There are numerous reasons why stock prices might move for reasons unrelated to technology – taxes, leverage, bubbles (either rational or irrational), and the interaction of inflation with taxes.

in technology at all horizons, not just in the long run. In practice, the news shocks identified by Beaudry and Portier (2006) and Beaudry and Lucke (2009) fail to account for important variation in technology at medium horizons (as much as 40 percent), making their interpretations problematic. In contrast, my identification strategy explicitly seeks to minimize the unexplained variance in technology over these horizons, and produces a news shock which leads to increases in technology soon after impact.

Recent work has questioned the ability of structural VARs to adequately recover shocks from economic models. Following the recommendation in Chari, Kehoe, and McGrattan (2008), I simulate data from a dynamic stochastic general equilibrium (DSGE) model to examine the performance of my empirical approach. I consider a neoclassical model with real frictions (habit formation and investment adjustment costs) augmented with both news shocks and surprise technology shocks.[6] The full description of the model and the details of the simulation exercise can be found in Appendix 6.2. I estimate VARs featuring technology, consumption, output, and hours on the simulated data. These are the same variables that are included in my empirical VARs in Section 3 (without the "information" variables).

Figure 1 depicts both theoretical and estimated impulse responses averaged over the simulations to a news shock that technology will be permanently higher. The theoretical responses are solid black and the average estimated responses over the simulations are depicted by the dotted line, with the dashed lines depicting the 10th and 90th percentiles of the distribution of estimated impulse responses. Although investment does not appear directly in the system, I impute its response as the output response less the share-weighted consumption response. A number of features from the simulations stand out. The estimated empirical impulse responses are roughly unbiased on impact and for most horizons thereafter. A favorable news shock leads to rising consumption but falling output, hours, and investment on impact in the model. After impact, the aggregate variables track movements in technology. My empirical identification captures these features quite well. The estimated responses to a news shock are only slightly downward biased at long horizons, and the estimated dynamics are very close to the true dynamics at all horizons. Figure 2 shows results for the identification of the surprise technology shock. Similarly, the estimated impulse responses are roughly unbiased on impact and for a number of quarters. The long horizon biases in the responses are larger here than for the identification of the news shock, but the responses nevertheless do a good job at capturing the model's dynamics.[7]

The average correlation between the identified news shock and the true news shock across

[6] Similarly good simulation results obtain in a model with nominal frictions.

[7] The simulation results are similarly good when a reasonable amount of noise is added to the observation of the series entering the estimated VAR.

simulations is 0.73, with the median correlation 0.81 and the 10th and 90th percentile correlations 0.55 and 0.88, respectively. The average correlation between the identified and true surprise technology shock is even higher at 0.92. The results improve even further as I let the size of the simulated samples become arbitrarily large. While very small biases persist in large samples, the estimated impulse responses to both kinds of technology shocks are extremely close to the true responses at all horizons, the distributions of responses collapse on a point, and the correlation between the identified and true shocks exceeds 0.95. My empirical procedure does not identify a large statistically significant news shock when the simulated data contain no such shock. In particular, the average estimated responses of aggregate variables to a news shock are all zero at short horizons when the data are generated without news shocks.[8] The simulations are of similar high quality under alternative specifications of the model and over a variety of different parameterizations.[9] Taken as a whole, the simulations suggest that my approach is capable of doing a good job in identifying both news shocks and surprise technology shocks and their effects on macroeconomic variables.

I close this section of the paper by addressing the implications of news shocks for VAR invertibility. Fernandez-Villaverde, Rubio-Ramirez, Sargent, and Watson (2007) discuss the conditions under which DSGE models produce moving average representations in the observables which can be inverted into a VAR representation in which the VAR innovations correspond to economic shocks. Invertibility problems potentially arise when there are unobserved state variables which do not enter the estimated VAR (Watson (1994)). If the observables do not fully reveal the values of the unobserved states, then the VAR innovations will be linear combinations of structural shocks and measurement errors, potentially invalidating conclusions drawn from structural impulse response analysis. Leeper, Walker, and Yang (2008) stress that anticipated shocks to future state variables are potentially pernicious for VAR invertibility. The essential difficulty is that when shocks are anticipated by agents several periods in advance, the shocks themselves become unobserved state variables.

It is straightforward to verify that the condition for invertibility in Fernandez-Villaverde, et al (2007) is not strictly satisfied in the model as presented in Appendix 6.2. The intuition for the failure of invertibility is that the news shock is both a shock and a state variable

[8] The procedure does identify a (small) spurious response of the non-stationary aggregate variables (technology, output, consumption, and investment) at low frequencies when no news shock is present. This spuriousness disappears as the size of the sample increases. Nevertheless, the point estimates remain unbiased at high frequencies when there are no news shocks in the simulations.

[9] In particular, the quality of the simulations is roughly invariant to alternative parameterizations of the variance of non-technology shocks. This differs from the conclusions in Chari, Kehoe, and McGrattan (2008), who find that the small sample performance of long run restrictions to identify technology shocks depends heavily on the contribution of non-technology shocks in the model. One reason I do not reach the same conclusion is because the measure of technology is not influenced by non-technology shocks, whereas average labor productivity (which is what enters their VARs) is.

in the model – agents must keep track of its value for several periods until it loads onto the level of technology. Nevertheless, as noted by Sims and Zha (2006) and Sims (2009), non-invertibility is not an either/or proposition, and structural VAR techniques applied to data generated from a model with a non-invertibility may nonetheless perform quite well. The simulation results here indicate my VAR-based procedure does well in practice, even though the non-invertibility leads to small asymptotic biases. As stressed by Watson (1994) and Sims (2009), the inclusion of forward-looking variables in the system helps to forecast the missing state variables, and mitigates the role of the non-invertibility in practice.

3 Empirical Evidence

In this section I present the main results of the paper. I find that a favorable news shock is associated with a slight rise in consumption and modest declines in output, investment, and hours of work on impact. In the next section I will argue that this robust feature of the data poses problems for the expectations driven theory of business cycles. Before proceeding I begin with a brief discussion of the data.

3.1 Data

The most critical data series needed to proceed is the technology series itself. The Solow residual is not a particularly appealing measure of technology, primarily due to the fact that standard growth accounting techniques make no attempt to control for unobserved input variation (labor hoarding and capital utilization). Since identification of the news shock requires orthogonalization with respect to technology, it is critically important that the empirical measure of technology adequately control for unobserved input variation. To address these issues, I employ a quarterly version of the Basu, Fernald, and Kimball (2006) technology series, which arguably represents the state of the art in growth accounting. Their essential insight is to exploit the first order condition which says that firms should vary intensity of inputs along all margins simultaneously. As such, they propose measuring unobserved input variation as a function of observed variation in hours per worker. They also make use of industry level data to allow for non-constant returns to scale in the production function. As the industry level data is only available at an annual frequency, it is not possible to construct a quarterly technology series with both the unobserved input and returns to scale corrections. What I use in this paper is a quarterly measure using only the utilization correction.[10]

[10] This series was constructed and given to me directly by John Fernald.

Formally, the quarterly version of this technology series presumes a constant returns to scale production function of the form: $Y = AF(ZK, EQH)$, where Z is capital utilization, E is labor effort, H is total labor hours, and Q is a labor quality adjustment. The traditional uncorrected TFP is then $\Delta A = \Delta Y - \theta\Delta K - (1-\theta)\Delta QH$, where θ is capital's share. The utilization correction subtracts from this $\Delta U = \theta\Delta Z + (1-\theta)\Delta E$, where observed labor variation is used as a proxy for unobserved variation in both labor and capital. The standard Solow residual is both more volatile and procyclical than the resulting corrected technology measure. In particular, the standard deviation of the HP detrended Solow residual is roughly 33 percent larger than for the adjusted series. The correlation between HP detrended output and uncorrected TFP is roughly 0.8, while the output correlation with corrected technology is about half that at 0.4.

The output measure I use is the log of real output in the non-farm business sector at a quarterly frequency. The consumption series is the log of real non-durables and services. The hours series is total hours worked in the non-farm business sector. I convert these series to per capita terms by dividing by the civilian non-institutionalized population aged sixteen and over. The results are not sensitive to this transformation. The consumption data are from the BEA; the output, hours, and population data are from the BLS. The population series in raw form is at a monthly frequency. I convert it to a quarterly frequency using the last monthly observation of each quarter.

The measure of stock prices which I use is the log of the real S&P 500 Index, taken from Robert Shiller's website. The measure of inflation is the percentage change in the CPI for all urban consumers. Use of alternative price indexes produces similar results. Both the stock price and inflation series are at a monthly frequency. As with the population data, I convert to a quarterly frequency by taking the last monthly observation from each quarter. The consumer confidence data are from the Michigan Survey of Consumers, and summarize responses to a forward-looking question concerning aggregate expectations over a five year horizon. For more on the confidence data, see Barsky and Sims (2008).[11] The confidence data are available beginning in 1960; all other series begin in 1948.

3.2 Benchmark Results

I include seven variables in the benchmark system: the technology series, stock prices, inflation, consumer confidence, non-durables and services consumption, total hours worked, and

[11] The specific survey question is: "Looking ahead, which would you say is more likely – that in the country as a whole we'll have continuous good times during the next five years, or that we'll have periods of widespread unemployment or depression, or what?" The series is constructed as the percentage of respondents giving a favorable answer less the percentage giving an unfavorable answer plus 100.

real output. Given the limited availability of the confidence data, the data in the VAR run from the first quarter of 1960 to the third quarter of 2007. As a benchmark, I estimate the system as a vector error correction model (VECM); I obtain very similar results when estimating the VAR in levels. Standard unit root tests fail to reject the hypotheses that technology, stock prices, consumption, and output are $I(1)$; tests are inconclusive for consumer confidence, hours per capita, and inflation, though they tend to indicate that these series are stationary. I allow for and estimate three cointegrating relationships among the four assumed trending series; confidence, hours, and inflation enter the system in levels. My results are robust to alternative assumptions about cointegration.

The median suggestion of a variety of popular information criteria, I estimate the benchmark system with three lags. I consider robustness to lag length in the next subsection. In terms of the identification strategy detailed in the Appendix, I set the truncation horizon at $H = 50$. In words, then, I identify the news shock as that structural shock orthogonal to technology which maximally explains movements in technology over a twelve year horizon. A truncation horizon of twelve years is both long enough to capture medium run forces and short enough to provide fairly reliable results. It also focuses in on forecastable movements in technology at the frequencies typically studied in the theoretical literature on expectations driven business cycles. As with lag length, I discuss robustness along this dimension in detail below.

Figure 3 shows the estimated impulse responses of technology, consumption, investment, output, and hours to a favorable news shock from the benchmark VAR, with the dashed lines representing 5th and 95th percentile confidence bands.[12] These bands are constructed from the bias-corrected bootstrap procedure proposed by Kilian (1998). Following a favorable news shock, technology grows rapidly for about a year and a half before leveling off approximately one third of percent higher than its pre-shock value. Consumption jumps up only very slightly on impact. The most striking features of the estimated responses are the point estimates of the impact effects of a favorable news shock on output, investment, and hours, all of which are negative. In particular, the favorable news shock leads to an immediate reduction in hours worked of more than one third of percent and in output of slightly more than 0.2 percent. Both of these effects are statistically significant at better than the 5 percent level. The estimated impact effect of a news shock on investment is negative and statistically significant as well. The negative conditional comovement among aggregate variables on impact is broadly consistent with the implications of standard neoclassical business cycle models; it is incompatible with news shocks being the primary source

[12] As in the simulations of the previous section, the investment response is imputed as output less the share-weighted consumption response.

of fluctuations.

The estimated dynamic paths of the macroeconomic variables following a news shock largely track that of technology. In particular, following the small impact effect, consumption grows smoothly and slowly for a number of quarters, with a peak response of roughly one half percent and a long run response of slightly less than that. The large predictable increase in consumption following impact would be associated with an increase in real interest rates in most equilibrium models, and it is also a feature of the data. In particular, the identified news shock series is positively and significantly correlated with the ex-post real rate of return, measured as the three month T-Bill rate less one quarter ahead inflation.[13] Similarly to consumption, after the initial negative impact effects, output, investment, and hours all grow for a number of quarters. The estimated swings in these series are large, with both output and investment slightly overshooting their long run values. After the initial negative response, hours recover strongly, with the response positive several quarters after the shock before reverting to zero. The peak responses of these aggregate variables all occur a couple of quarters after the maximal response of technology. In short, there is no evidence that these main macroeconomic variables strongly anticipate technology improvements with large, broad-based comovement.

As the primary focus of this paper is the business cycle relevance of news shocks, the impulse responses of the "information" variables are omitted. A positive news shock leads to an increase in stock prices. While relatively small in magnitude, it is nevertheless not possible to reject the hypothesis that stock prices obey a random walk following a news shock. That news shocks explain a relatively small component of variation in stock prices (less than 20 percent at high frequencies) is helpful in understanding the differences between my results and others in the literature. The news shock is associated with large declines in inflation and increases in consumer confidence. For further discussion of the responses of these "information variables" see Barsky and Sims (2009).

Table 1 depicts the share of the forecast error variance of several of the variables in the VAR attributable to the news shock at a number of horizons. The numbers in brackets are the 5th and 95th percentiles from the same bootstrap procedure used to construct the confidence bands for the impulse responses. News shocks account for about 40 percent of the forecast error variance share of technology at a horizon of five years and almost 60 percent at ten years. The final row of the table shows the total contribution to technology's forecast error variance of the news shock and the surprise technology shock. The news shock and the contemporaneous technology innovation combine to account for 90 percent or more

[13] Similar results obtain when the interest rate is included directly in the system. See also the discussion in Barsky and Sims (2009).

of the forecast error variance of technology at all horizons up to ten years. That so little of technology remains unexplained at most horizons validates the assumption underlying identification that most of the movements in technology can be attributed to only two shocks, and suggests that my approach has done a good job at identifying the news shock.

The news shock accounts for a relatively small share of the forecast error variances of consumption at short horizons, and a somewhat larger share of the forecast error variance of output. The shock contributes more significantly to the variance decomposition of hours at high frequencies and much less so at lower frequencies. At longer horizons the news shock contributes more significantly to the variance decomposition of the aggregate variables excluding hours, explaining between ten and forty of the variance of output at business cycle frequencies. While news shocks thus appear to be a non-negligible feature of the data, I argue below in Section 4 that the negative conditional comovement at high frequencies limits the extent to which such shocks are a major source of fluctuations.

Figure 4 plots the cross correlogram between the identified news shock series and HP detrended output. The figure shows both the contemporaneous correlation and the correlation between the news shock and detrended output led over the span of three years. The news shock is negatively correlated with detrended output contemporaneously and for a few quarters and positively correlated with detrended output led several more quarters. This plot corroborates the estimated impulse responses. The correlation between the news shock and the cyclical component of output is modest, with good news about future technology associated with output falling below trend for a number of quarters before picking up.

Figure 5 plots the time series of identified news shocks from the benchmark VAR. The shaded areas represent recession dates as defined by the NBER. So as to make the figure more readable, I show the one year moving average of the identified shock series as opposed to the actual series.[14] There are a preponderance of negative shock realizations throughout the 1970s (corresponding with the productivity slowdown) and a series of positive realizations in the first half of the 1990s, corresponding with the productivity speed-up. There are large negative shock realizations which are not associated with recessions at all (particularly in the late 1970s and mid 1980s). It is also interesting to note that the smoothed shock series is persistently positive immediately after both the 1990-1991 and 2001 recessions. The wealth effect associated with positive news shocks may help to explain the "jobless recoveries" associated with both of these episodes.

An historical decomposition of the variables in the VAR attributable to the news shock

[14] To be clear, the smoothed version of the series is $\varepsilon_t^s = (\varepsilon_{t-2} + \varepsilon_{t-1} + \varepsilon_t + \varepsilon_{t+1} + \varepsilon_{t+2})/5$. The series begins in 1961:3 and ends in 2007:1. I lose four observations at the beginning of the sample due to the lag length and two additional observations at the beginning and end of the sample due to the moving average.

is likely to be more informative about the business cycle relevance of news shocks than is a plot of the shocks alone. Figure 6 plots simulated and actual values of output per capita in the six NBER defined US recessions since 1961, with the simulated values constructed using the VAR estimates and identified news shocks. In particular, the decomposition shows the simulated time path of real output per capita as if the news shock were the only stochastic disturbance impacting the system beginning in the first quarter of 1961. The decomposition fails to predict output declines in four out of the six US recessions in the sample period (1969-1970, 1981-1982, 1990-1991, and 2001). For example, the cumulative effects of news shocks suggest that output per capita should have risen by some two percent during the 2001 recession, whereas in actuality it fell by about one percent. While the decomposition does show output falling slightly during the 1973-1975 and 1980 recessions, the magnitudes of the predicted declines are much smaller than observed in the data. Taken as a whole, the historical decomposition suggests that news shocks have not been an important source of US recessions.

Figure 7 shows the impulse responses of aggregate variables to the surprise technology shock (i.e. ε_1 above). Technology's response to its own innovation, while large and significant on impact, is quite transitory. Output and investment show a significant transitory response to the surprise technology shock, with hours rising slightly and following a hump-shape, and a small increase in consumption.[15] These responses are roughly consistent with the transitory but persistent productivity disturbances emphasized in the real business cycle literature (Kydland and Prescott (1982)), though the hours response to the surprise technology shock is small and suggests only weak amplification. Given technology's apparently permanent response to the news shock, my results indicate that the bulk of the low frequency component of productivity is attributable to news shocks. Beaudry and Portier (2006) and Barsky and Sims (2009) reach similar conclusions. These results also accord well with the analysis in Rotemberg (2003).

These results have implications reaching beyond the study of expectations driven business cycles. Many VAR identifications based on long run restrictions find that the shock responsible for the unit root in labor productivity leads to an impact reduction in hours (Shapiro and Watson (1989), Gali (1999)). Some have argued that this finding lends support to sticky price models (Gali (1999), Basu, Fernald, and Kimball (2006)).[16] My finding

[15] The contribution of the surprise technology shock to the forecast error variance of output and technology can be found by subtracting the appropriate rows in Table 1. The surprise technology shock makes only small contributions to the forecast error variance of both consumption and hours.

[16] Basu, Fernald, and Kimball (2006) find that innovations to their annual technology series are negatively correlated with the change in hours. Since these authors assume that technology is exogenous (i.e. there is no Granger causality from other variables to technology), the innovation in their technology series is a combination of surprise and news shocks from previous periods, offering a potential reconciliation with the

that the low frequency component of technology is mainly driven by news shocks offers a potential reconciliation of these results without relying upon nominal frictions. As argued below in Section 4, a negative conditional correlation between hours and the "technology shock" obtained from a long run restriction is exactly the qualitative prediction of a flexible price model when the low frequency component of productivity is mainly attributable to news shocks.

My results nevertheless do suggest that non-technology shocks are an important source of fluctuations. The final row of Table 1 shows the total variance of output explained by the surprise technology shock and news shock combined. Surprise technology shocks explain most of output fluctuations at very high frequencies, while news shocks explain large movements in output in the long run. At medium horizons, however, some other disturbance(s) explains between half and three-quarters of output fluctuations. Further understanding of these other shock(s) remains an important task for future research.

3.3 Sensitivity Analysis

My main result that news shocks induce negative comovement is robust to alternative lag structures in the reduced form system as well as to various different assumptions and/or specifications concerning the long run relationships among the series.[17] At all tested lag lengths, output, investment, and hours decline on impact in response to a favorable news shock, while consumption rises. With more lags in the reduced form system the impulse responses are less smooth and there is more evidence of reversion in all series at longer horizons, but the basic qualitative nature of the responses is unchanged.

Earlier versions of this paper reported results with all systems estimated as VARs in levels. The impulse responses are nearly identical under a levels specification compared with the VECM specifications reported here; the primary differences lie in the estimated responses at longer horizons, with more evidence of reversion evident in the levels specification. The qualitative nature of the responses is unaffected by different assumptions concerning the number of cointegrating relationships – with fewer cointegrating vectors, the long run responses are quantitatively larger, but the high frequency impulse responses of aggregate variables to a news shock are virtually identical to the benchmark results. In the

results presented here.

[17] My results are also qualitatively robust with alternative measures of technology. Application of my benchmark identification to a system with the Solow residual in place of the utilization-adjusted technology measure again finds negative impact comovement, with output, hours, and investment all declining in anticipation of good news. The main difference is that the response of the Solow residual itself to the news shock appears far more transitory than is the response of technology in Figure 3. Similar results to those in Figure 3 obtain with alternative utilization corrections to the Solow residual – see Footnote 20.

benchmark system I assume that aggregate hours per capita is stationary, and thus it enters system in levels and does not appear in any of the cointegrating relationships. There is a large debate in the VAR literature over the stationarity of hours (for a review see Christiano, Eichenbaum, and Vigfusson (2004)). I obtain very similar impulse responses to a news shock whether hours enter the system in levels, in first differences, or as deviations from a trend, as well as whether or not hours enter the estimated cointegrating relationships.

In Figure 8 I show estimated impulse responses to a news shock from a smaller system than my benchmark. In particular, I omit consumer confidence and inflation from the system, leaving a five variable system featuring technology, stock prices, consumption, output, and hours, which is very similar to the larger systems estimated in Beaudry and Portier (2006) and Beaudry and Lucke (2009). The estimation and identification are otherwise similar to above. While the quantitative response of technology to the news shock is somewhat smaller than before, the qualitative results are otherwise the same. Output and hours both decline on impact and for a number of quarters following a favorable news shock, while consumption rises. As earlier, output and hours appear to track movements in technology as opposed to anticipating them. These results for the smaller system obtain regardless of the sample period or the measure of technology.

I also consider how sensitive my results are to the specification of the maximization problem underlying identification. Figure 9 shows impulse responses from the benchmark system when the truncation horizon is 100 years, which effectively makes my identification identical to the combined recursive-long run restriction of Beaudry and Lucke (2009). As before, I find that favorable news about future technology is associated with impact declines in output, hours, and investment and an increase in consumption.[18] The impact increase in consumption is larger here than in Figure 3, the impact declines in output, hours, and investment are smaller, and the shock accounts for a larger share of output fluctuations at business cycle frequencies. The response of technology itself also changes in an important way. In particular, similarly to the results in Beaudry and Portier (2006) and Beaudry and Lucke (2009), the response of technology to news is far more delayed, with most of the productivity improvement occurring much further off in the future than what I estimate in the benchmark. This fact becomes important in understanding the differences between our conclusions, and I discuss it in more depth in the next subsection.

While the utilization-adjusted technology series I use in this paper arguably represents the state of the art in growth accounting, one may nevertheless object to the notion that

[18] My results are also qualitatively similar when identification is achieved from maximizing the variance share of technology at lower truncation horizons. The main difference is that the responses of technology, output, etc. appear slightly hump-shaped and more transitory. The impact effects of the identified news shocks are nevertheless very similar to my benchmark.

the resulting series accurately reflects true technology. To the extent to which it is a poor measure of true technology, the analysis conducted thus far may be invalid. Indeed, most of the work in the structural VAR literature assumes that technology is unobservable, and attempts to identify technology shocks off of the time series properties of observed labor productivity.

As a final robustness check, as well as to provide a closer link to some of the existing VAR literature, I estimate a system with output per hour of work in place of the technology series. As this necessitates dropping the measure of output, the resulting system features output per hour, stock prices, consumption, hours, inflation, and consumer confidence. In most DSGE models news shocks will lead to movements in output per hour on impact and non-technology shocks will affect labor productivity in the short run. As such, the identifying restrictions that the news shock is orthogonal to the measure of productivity and that surprise and news shocks completely explain variation in productivity are no longer theoretically valid when output per hour is used in place of a measure of technology. A different orthogonalization strategy to identify news shocks is therefore required.

I employ a shape restriction to identify a set of candidate impulse responses to a news shock in the system with labor productivity. In particular, I impose that a favorable news shock results in a "small" impact effect on labor productivity followed by sustained growth. While most DSGE models would have the implication that a favorable news shock raises labor productivity on impact (because real wages rise), I do not explicitly impose this. The only restriction on the data is that the response of labor productivity be growing – in particular that the response is larger after ten quarters than it is on impact, and in turn that the response is larger after twenty-four quarters than it is at ten.

The shape restriction methodology has much to recommend it in this context. As noted, it does not rely upon an explicit and potentially controversial measure of technology. Furthermore, it is valid under less restrictive assumptions than is my benchmark identification. In particular, it does not assume that the news shock is contemporaneously orthogonal to true technology, nor does it assume that the stochastic process for true technology is driven only by two shocks. It merely imposes that news shocks lead to increasing movements in measured labor productivity over time. The details of the shape restriction identification are in Appendix 6.3. See also Faust (1998).

The impulse responses of aggregate variables for this identification are in Figure 10. The solid line shows the median response, while the dashed lines depict the 16th and 84th percentiles of the distribution of candidate responses. These responses are both qualitatively and quantitatively in line with the responses in Figure 3. Labor productivity moves little on impact, followed by rapid growth for the first several quarters, and slower, prolonged growth thereafter, with a long run response slightly below 0.5 percent. Consumption rises slightly on impact, followed by sustained growth. As in Figure 3, output, hours, and investment all fall on impact and then track movements in productivity.[19] The impact declines in both output and investment are virtually identical to the benchmark estimates; the impact decline in hours is somewhat smaller here than in the system with technology. The median long run responses of output and consumption to the news shock are slightly smaller here than before, but the overall qualitative nature of the responses is very similar to my baseline estimates.

3.4 Comparison with Existing Literature

As noted in the Introduction, the existing empirical literature on news shocks is somewhat limited. The most well-known of these papers are by Beaudry and Portier (2006), Beaudry, Dupaigne, and Portier (2008), and Beaudry and Lucke (2009). These authors estimate two to five variable systems featuring measures of technology, stock prices, and other macroeconomic aggregates. They propose two alternative orthogonalization schemes aimed at isolating news shocks – the first is to associate the news shock with the stock price innovation orthogonalized with respect to technology, and the second combines short and long run restrictions to identify the news shock. These authors argue that both orthogonalization schemes yield very similar results. They find that news shocks lead to positive conditional comovement among macroeconomic aggregates on impact, that aggregate variables strongly anticipate movements in technology, and that news shocks account for the bulk of the variance of aggregate variables at business cycle frequencies.

The conditions under which either of these orthogonalization schemes are valid are encompassed by my empirical identification strategy. In particular, were the conditions required for the pure recursive identification satisfied, my identification would (asymptotically) identify the same shock and impulse responses. Likewise, their long run identifying assumption in the second orthogonalization scheme rests on the same implicit assumption underlying my identification – that a limited number of shocks account for variation in measured technology. I have generated data from both economic and econometric models satisfying their assumptions and my identification strategy routinely does an excellent job in identifying the news shock and its associated impulse responses. Further, my identification consistently outperforms identification based on long run restrictions in finite samples; the model-based simulation results described in Section 2 become significantly less reliable for very high trun-

[19] Output does not enter the system directly, but is imputed as the output per hour response plus the hours response. As before, the investment response is then calculated as the output response less the share-weighted consumption response.

cation horizons.[20]

There is a large quantitative and qualitative difference between my results and theirs in the estimated effects of news shocks on technology itself. When using a corrected measure of technology similar to the one used in this paper, the shock identified by these authors typically does not have any noticeable effect on technology for several years.[21] Indeed, Beaudry and Portier (2006) note that "growth beyond its [TFP's] initial level takes somewhere between 12 and 16 quarters" (p. 1303) following a news shock, while in Beaudry, Dupaigne, and Portier (2008), they state "it [news shock] has almost no impact on TFP during the first five years" (p. 3). In contrast, the news shock I identify begins to affect technology soon after impact, and explains technology movements well at both short and long horizons.

That these authors' identified shock has such a delayed effect on technology makes its interpretation as a news shock problematic. I estimated my benchmark seven variable system and identified a shock using both a short and a long run restriction as in Beaudry and Portier (2006) and Beaudry and Lucke (2009) (similar results obtain when applying this identification to smaller systems than my benchmark). As noted above in Section 3.3, I find that this identification yields a shock more closely aligned with their results. In particular, the evidence in support of negative impact comovement is less drastic, and the shock accounts for a much larger share of the variance of aggregate macro variables at business cycle frequencies.

Table 2 shows the fraction of the forecast error variance of technology attributable to this shock at various horizons as well as the total variance in technology accounted for by this shock along with the surprise technology shock. A comparison with the corresponding rows in Table 1 is instructive. Whereas the news shock identified using my empirical strategy explains between 20 and 60 percent of the variance of technology at business cycle frequencies, the shock identified using the long run restriction explains only 5 to 25 percent of the technology variance at horizons from one to ten years. Importantly, the long run identification leaves up to 40 percent of the variance of technology unexplained at business cycle frequencies. In other words, some other structural shock orthogonal to technology's innovation potentially accounts for twice as much variation in technology at these frequencies than does what these authors deem the news shock. In comparison, my identification leaves less than 10 percent of technology's variance unaccounted for at business cycle frequencies.

That the shock identified by these authors leaves so much of the variance of technology

[20] Francis, Owyang, and Roush (2007) report similar findings in that their medium run identification of technology shocks performs significantly better in finite samples than does a long run identifying restriction.

[21] Discrepencies in our results do not result from different data, and in particular from different measures of TFP. I have conducted my empirical analysis using Beaudry and Portier's (2006) TFP data (available from the *American Economic Review* website) and obtain very similar results.

unexplained at business cycle frequencies is obviously problematic. In particular, it leaves unanswered the question of what the business cycle implications are of the shock orthogonal to technology's innovation which explains the remaining variance in technology. In systems small and large, under a variety of different assumptions, I robustly find that a shock orthogonal to technology's innovation which accounts for the bulk of technology movements at horizons up to fifteen to twenty years leads to negative comovement among macroeconomic aggregates at high frequencies.

An alternative approach to the VAR-based methodology of estimating the implications of news shocks for aggregate variables would be the estimation of a fully specified DSGE model. This is the approach taken by Schmitt-Grohe and Uribe (2008), who argue that news shocks about future technology are quantitatively important for understanding fluctuations (though they also find that favorable news shocks lead to an immediate reduction in hours), and Kahn and Tsoukalas (2009), who reach conclusions similar to mine. Kahn and Tsoukalas show that Schmitt-Grohe and Uribe's results are highly sensitive to model structure, and in particular to the range of potential shocks taken into consideration. The advantage of the VAR methodology pursued here is that it is highly flexible, and reliably identifies news shocks from a variety of different model structures, including those in which news shocks are a quantitatively important feature of the data generating process. In practice, estimation of a fully specified model in this context is problematic, as there is no consensus on what the appropriate theoretical structure is in which news shocks have a chance to be an important component of the data generating process.

4 Evaluating Expectations Driven Business Cycles

During the last several years, a number of authors have advanced various features capable of overturning the prediction of negative comovement in response to a news shock in most neoclassical models. Beaudry and Portier (2004), Den Haan and Kaltenbrunner (2006), Christiano, Ilut, Motto, and Rostagno (2007), and Jaimovich and Rebelo (2008) all propose models in which consumption, investment, hours, and output all rise significantly on impact in anticipation of future technological improvement. My results suggest that this research may have been misguided. In fact, my empirical finding of an increase in consumption but declines in hours, investment, and output on impact in response to good news is exactly the qualitative implication of a standard neoclassical model augmented with news shocks.

To illustrate the good fit of the basic neoclassical model, Figure 11 shows the theoretical responses from the model and my estimated empirical responses to a good news shock placed together. The model responses are generated from the same theoretical structure as described in Appendix 6.2, with a slightly different calibration and a different specification of the process for aggregate technology. Rather than assuming that news shocks portend discreet jumps in technology j periods into the future, I model news as diffusing slowing into the level of the permanent component of technology as follows:

$$\begin{aligned}
\ln A_t^p &= g_{t-1}^A + \ln A_{t-1}^p \\
g_t^A &= (1-\lambda)g^A + \lambda g_{t-1}^A + \varepsilon_{2,t}
\end{aligned}$$

Given the timing assumption, a shock to $\varepsilon_{2,t}$ has no immediate effect on the level of technology but portends a period of sustained, smooth growth. I calibrate the parameters of this specification so as to match the estimated empirical response of technology to a news shock. So as to make the point as stark as possible, I calibrate $b = \gamma = 0$ (no habit formation and no adjustment costs), so that the model reduces to the standard real business cycle model. The remaining parameters of the model are calibrated as detailed in Appendix 6.2.

While far from perfect, it is clear that the simple RBC model provides a reasonably close characterization of the estimated impulse responses. The theoretical impact effects of a favorable news shock on macroeconomic aggregates are the same signs as the estimated ones. After the impact effects, the aggregate variables rise smoothly in tandem with the predicted increase in technology in both the model and in the data. The impact jump in consumption is smaller in the data than in the model, while the impact decline in hours is larger in the data than in the model. The lack of a strong internal propagation mechanism in the model results in the theoretical responses failing to fully match the large estimated swings in the data after a number of quarters. The responses to a surprise transitory technology shock in the model are also qualitatively similar to what is estimated and shown in Figure 7. The main inconsistency between the responses to the two kinds of technology shocks and the simple RBC model is in the response of hours. Hours decline sharply in response to good news in the data, which would be consistent with a high labor supply elasticity in the model. In comparison, the hours response (though positive at high frequencies) is quite modest to the surprise technology shock, consistent with a low labor supply elasticity. In spite of this inconsistency, the overall nature of the estimated impulse responses to the shocks are qualitatively consistent with the predictions of the stylized model.

That the simple RBC model provides a fairly good qualitative fit with the estimated empirical responses to a news shock is somewhat surprising. A number of realistic and common features would likely improve the fit even further. In particular, features mitigating the positive wealth effect of good news on consumption would likely serve to improve the fit of the consumption response. In particular, habit formation or "rule of thumb consumers" (Campbell and Mankiw (1990)) would limit the impact jump of consumption in the model; the empirical response of consumption to a news shock appears to strongly track that of output, suggesting that a rule of thumb specification may provide the better fit. Liquidity constraints may help to reconcile both the consumption and hours responses to the two kinds of technology shocks – if constrained, the only way to increase utility in response to good news is to decrease hours of work, whereas surprise technology shocks would ease constraints and allow increased consumption without having to work more. Common modifications in factor markets (such as variable utilization or a higher labor supply elasticity) might also help to improve the fit of the hours response.

The most salient feature of the business cycle is broad-based comovement among macroeconomic aggregates. This comovement typically refers to the unconditional correlations of filtered consumption, investment, and hours with output. These correlations for HP filtered post-war US data are in the far right column of Table 3. Though positive conditional comovement on impact is neither necessary nor sufficient for unconditional comovement, given how high these correlations are, it is unlikely that a shock inducing negative conditional comovement on impact could be the main driving force behind the data.

The middle column of Table 3 makes this point clear. There I show the average correlations between HP filtered macro aggregates and output from 2000 simulations of the simple RBC model in which the news shock is the only stochastic disturbance. While the model correlations between investment, technology, and hours with output are close to those in the data, the model correlation between consumption and output, though still positive, is roughly one quarter its value in the data. Alternative calibrations of the model's parameters do little to improve the fit along this dimension. That the model provides a close fit with the estimated impulse responses but is unable to match the unconditional correlations suggests that some other shock must be the primary driving force behind the data. This finding is consistent with the conclusions from the variance decompositions of Section 3, which suggested that news shocks play only a modest role in accounting for variation in output at business cycle frequencies.

5 Conclusion

The expectations driven business cycle hypothesis has been advanced as an alternative to business cycle models based on aggregate productivity shocks. In particular, it offers the tantalizing possibility that business cycles could emerge absent any (ex-post) change in fundamentals. If good news about the future can set off a boom today, then a realization worse than anticipated can set off a bust. For this story to work, however, good news about the future must induce broad-based comovement, which is not the prediction of standard macro models. Existing empirical evidence suggesting that news shocks do lead to broad-based comovement has spawned a new literature searching for theoretical frameworks capable of delivering business cycle-like behavior when driven by news shocks about future technology.

This paper has taken a closer look at the empirical evidence in favor of this theory of fluctuations. I implemented a new empirical approach of identifying news shocks that is directly based on the implications of theoretical models of expectations driven business cycles, and I showed that my approach performs well on model generated data. While I corroborate earlier evidence that agents do receive advance signals about future productivity, I find that good news is associated with an increase in consumption and impact declines in output, hours, and investment. After impact, aggregate variables largely track, as opposed to anticipate, predicted movements in measured technology. The impulse responses I estimate are broadly consistent with the implications of standard macro models. The estimated negative conditional comovement on impact is difficult to reconcile with the salient business cycle fact of strong broad-based comovement among macroeconomic aggregates. As such, my results suggest that news shocks about future productivity are not a dominant source of business cycles.

6 Appendix

6.1 Identifying News Shocks

Let $\mathbf{y}_t$ be the $N \times 1$ vector of observables. One can form the reduced form moving average representation in the levels of the observables either by estimating an unrestricted VAR in levels or by estimating a stationary vector error correction model (VECM):

$$\mathbf{y}_t = \mathbf{B}(\mathbf{L})\mathbf{u}_t \tag{A.1}$$

Assume there exists a linear mapping between reduced form innovations ($\mathbf{u}$) and structural shocks ($\boldsymbol{\varepsilon}$):

$$\mathbf{u}_t = \mathbf{A}_0\boldsymbol{\varepsilon}_t \tag{A.2}$$

This implies the following structural moving average representation:

$$\mathbf{y}_t = \mathbf{C}(\mathbf{L})\boldsymbol{\varepsilon}_t \tag{A.3}$$

Where $\mathbf{C}(\mathbf{L}) = \mathbf{B}(\mathbf{L})\mathbf{A}_0$ and $\boldsymbol{\varepsilon}_t = \mathbf{A}_0^{-1}\mathbf{u}_t$. The impact matrix must satisfy $\mathbf{A}_0\mathbf{A}_0' = \boldsymbol{\Sigma}$ after normalizing the variances of structural shocks to be unity, but it is not unique. Letting $\mathbf{D}$ denote an orthonormal matrix of conformable size and $\widetilde{\mathbf{A}}_0$ be an arbitrary orthogonalization of the reduced form, then the matrix $\widetilde{\mathbf{A}}_0\mathbf{D}$ spans the space of possible orthogonalizations (see Faust (1998)).

The h step ahead forecast error in terms of the structural shocks over the space of possible orthogonalizations is:

$$\mathbf{y}_{t+h} - E_{t-1}\mathbf{y}_{t+h} = \sum_{\tau=0}^{h} \mathbf{B}_\tau \widetilde{\mathbf{A}}_0 \mathbf{D} \boldsymbol{\varepsilon}_{t+h-\tau} \tag{A.4}$$

$\mathbf{B}_\tau \widetilde{\mathbf{A}}_0 \mathbf{D}$ is the matrix of structural moving average coefficients at horizon τ. The share of the forecast error variance of variable i attributable to structural shock j at horizon h is then:

$$\Omega_{i,j}(h) = \frac{\mathbf{e}_i' \left(\sum_{\tau=0}^{h} \mathbf{B}_\tau \widetilde{\mathbf{A}}_0 \mathbf{D} \mathbf{e}_j \mathbf{e}_j' \mathbf{D}' \widetilde{\mathbf{A}}_0' \mathbf{B}_\tau' \right) \mathbf{e}_i}{\mathbf{e}_i' \left(\sum_{\tau=0}^{h} \mathbf{B}_\tau \boldsymbol{\Sigma} \mathbf{B}_\tau' \right) \mathbf{e}_i} = \frac{\sum_{\tau=0}^{h} \mathbf{B}_{i,\tau} \widetilde{\mathbf{A}}_0 \boldsymbol{\varsigma}\boldsymbol{\varsigma}' \widetilde{\mathbf{A}}_0' \mathbf{B}_{i,\tau}'}{\sum_{\tau=0}^{h} \mathbf{B}_{i,\tau} \boldsymbol{\Sigma} \mathbf{B}_{i,\tau}'} \tag{A.5}$$

The $\mathbf{e}$s are selection vectors with one in the ith or jth places and zeros elsewhere. The selection vectors inside the parentheses in the numerator pick out the jth column of $\mathbf{D}$, which I denote by $\boldsymbol{\varsigma}$. $\widetilde{\mathbf{A}}_0\boldsymbol{\varsigma}$ is then the $N \times 1$ column vector corresponding to the jth column of a possible orthogonalizing matrix. The selection vectors outside the parentheses in both numerator and denominator pick out the ith row of the matrix of moving average coefficients, which I denote by $\mathbf{B}_{i,\tau}$.

My identifying assumption implies that ε_1 and ε_2 should account for all of the forecast error variance of technology at all horizons. Formally:

$$\Omega_{1,1}(h) + \Omega_{1,2}(h) = 1 \quad \forall\ h$$

With the unanticipated shock identified as the innovation in technology, $\Omega_{1,1}(h)$ will be invariant at all h to alternative identifications of the other $N-1$ structural shocks. As such, choosing elements of $\mathbf{A}_0$ to come as close as possible to making the above expression hold is equivalent to choosing the elements of $\mathbf{A}_0$ to maximize contributions to $\Omega_{1,2}(h)$ over h. Since the contribution to the forecast error variance depends only on a single column of $\mathbf{A}_0$, this suggests choosing the second column of the impact matrix to solve the following optimization problem:

$$\boldsymbol{\varsigma}^* = \arg\max \quad \sum_{h=0}^{H} \Omega_{1,2}(h) = \frac{\sum_{\tau=0}^{h} \mathbf{B}_{i,\tau}\widetilde{\mathbf{A}}_0\boldsymbol{\varsigma}\boldsymbol{\varsigma}'\widetilde{\mathbf{A}}_0'\mathbf{B}_{i,\tau}'}{\sum_{\tau=0}^{h} \mathbf{B}_{i,\tau}\boldsymbol{\Sigma}\mathbf{B}_{i,\tau}'}$$

s.t.

$$\begin{aligned} \widetilde{\mathbf{A}}_0(1,j) &= 0 \ \ \forall\ j > 1 \\ \boldsymbol{\varsigma}(1,1) &= 0 \\ \boldsymbol{\varsigma}'\boldsymbol{\varsigma} &= 1 \end{aligned}$$

Expressing the problem in terms of choosing $\boldsymbol{\varsigma}$ conditional on $\widetilde{\mathbf{A}}_0$ ensures that the resulting identification belongs to the space of possible orthogonalizations. The first two constraints impose that the news shock has no contemporaneous effect on technology. The third restriction (that $\boldsymbol{\varsigma}$ have unit length) ensures that $\boldsymbol{\varsigma}$ is a column vector belonging to an orthonormal matrix. H is an arbitrary, finite truncation horizon. This maximization problem can be rewritten as a quadratic form in which the non-zero portion of $\boldsymbol{\varsigma}$ is the eigenvector associated with the maximum eigenvalue of a weighted sum of the lower $(N-1) \times (N-1)$ submatrices of $\left(\mathbf{B}_{1,\tau}\widetilde{\mathbf{A}}_0\right)'\left(\mathbf{B}_{1,\tau}\widetilde{\mathbf{A}}_0\right)$ over τ. In other words, this procedure essentially identifies the news shock as the first principal component of technology orthogonalized with respect to its own innovation. Given the estimate of $\boldsymbol{\varsigma}$, the structural impulse response function to the news shock is given by $\mathbf{B}(\mathbf{L})\widetilde{\mathbf{A}}_0\ \boldsymbol{\varsigma}^*$, while the news shock itself is $\boldsymbol{\varepsilon}_{2,t} = \boldsymbol{\varsigma}^{*\prime}\ \widetilde{\mathbf{A}}_0^{-1}\mathbf{u}_t$.

6.2 Simulations

The simulations discussed in Section 2 are from a neoclassical model with real frictions and augmented with news shocks. The model can be expressed as a planner's problem:

$$\max \quad \sum_{t=0}^{\infty} \beta^t E_0 \left(\ln\left(C_t - bC_{t-1}\right) - \psi_t \frac{N_t^{1+1/\eta}}{1+1/\eta} \right)$$

s.t.

$$\begin{aligned}
K_{t+1} &= (1-\delta)K_t + \left(1 - \phi\left(\frac{I_t}{I_{t-1}}\right)\right) I_t \\
Y_t &= A_t K_t^{\theta} N_t^{1-\theta} \\
Y_t &= C_t + I_t + G_t \\
G_t &= g_t Y_t \\
\ln A_t &= g_A + \ln A_{t-1} + \varepsilon_{1,t} + \varepsilon_{2,t-j} \\
\ln g_t &= (1-\rho)\ln \overline{g} + \rho \ln g_{t-1} + \varepsilon_{3,t} \\
\ln \psi_t &= \nu \ln \psi_{t-1} + \varepsilon_{4,t}
\end{aligned}$$

C is consumption, N is employment, Y is output, K is capital, A is the level of technology, and G is government spending. β is the subjective discount factor, b is the degree of habit persistence, η is the Frisch labor supply elasticity, θ is capital's share of income, and δ is the depreciation rate on capital. ψ_t is a time-varying preference parameter with mean one and autoregressive coefficient ν. I assume that the government consumes a stochastic fraction of output, g_t. The log government share of output follows a stationary autoregressive process, with autoregressive parameter ρ. The government spending and preference shocks are included so as to introduce sufficient variation to be able to estimate a VAR with more than a few variables. $\phi(\cdot)$ is a convex function describing costs associated with adjusting investment. I assume that it has the following properties: $\phi(1) = 0$, $\phi'(1) = 0$, and $\phi''(\cdot) = \gamma \geq 0$. Log technology follows a random walk with drift with both an unanticipated shock and a news shock, with j describing the number of periods of anticipation in the news process. This specification of the process for technology is consistent with the more general specification in the text of Section 2.

The model as presented has the desirable property that there exist parameterizations in which news shocks induce positive comovement among aggregate variables on impact. In particular, for sufficiently high degrees of habit persistence and adjustment costs to investment, output, consumption, hours, and investment can all rise upon news of future technological improvement. In the case with $b = \gamma = 0$, the model converges to the simple real business cycle model in which good news about the future leads to falling output, hours, and investment on impact.

The model is solved by log-linearizing the first order conditions about the balanced growth path. As a baseline, I calibrate the parameters as follows: $\beta = 0.99$, $b = 0.8$, $\psi = 1$, $\eta = 1$, $\delta = 0.025$, $\theta = 0.33$, $\gamma = 0.3$, $\psi = 1$, $\overline{g} = 0.2$, $g_A = 0.25$, $\nu = 0.8$, and $\rho = 0.95$. This calibration implies that, along the balanced growth path, government consumption is 20 percent of output, private consumption is 56.5 percent of output, and investment is 23.5 percent of output. These numbers are in line with US data when durable consumption is included as a component of investment. Technology grows at the annualized rate of one percent per year, with output, consumption, and investment per capita growing at 1.5 percent per year. I assume three periods of anticipation for news shocks (i.e. $j = 3$).

I simulate 2000 sets of data with 200 observations each, drawing all four exogenous shocks from normal distributions. I set the standard deviation of the unanticipated technology shock to 0.66 percent and the standard deviation of the news shock at 0.33 percent. I calibrate the standard deviations of the remaining two shocks at 0.15. Similar results obtain for alternative calibrations of the non-technology shocks. For each simulation, I estimate a VECM with technology, consumption, output, and hours with three lags. I allow for and estimate two cointegrating relationships among the three trending series (technology, consumption, and output); hours is stationary in the model and enters the VECM in levels. Very similar results obtain when I estimate an unrestricted VAR in levels. I identify the contemporaneous technology shock as the innovation in technology and the news shock by maximizing the variance share of technology over a horizon of twenty quarters.

6.3 Shape Restrictions

For the shape restriction results in Section 3.3, the vector of observables, $\mathbf{y}_t$, contains output per hour, stock prices, consumption, hours, inflation, and consumer confidence. After estimating either a VAR in levels or a VECM, the reduced form moving average representation can be formed as in equation A.1. The structural impulse response function over the entire space of possible orthogonalizations is:

$$\mathbf{y}_t = \mathbf{B}(\mathbf{L})\widetilde{\mathbf{A}}_0\mathbf{D}\boldsymbol{\varepsilon}_t \tag{A.6}$$

As above, $\widetilde{\mathbf{A}}_0$ is an arbitrary orthogonalization of the reduced form and $\mathbf{D}$ is an orthonormal matrix. The impulse response to a particular structural shock depends only on a column of $\mathbf{D}$, again denoted by $\boldsymbol{\varsigma}$, which must be unit length.

I compute an arbitrary Choleski decomposition of the reduced form, $\widetilde{\mathbf{A}}_0$. I then draw 100,000 unit length vectors of conformable size (i.e. 100,000 candidate $\boldsymbol{\varsigma}$) from a normal distribution. These random vectors are then rescaled to be of unit length. For each of these vectors, I compute the implied impulse response of output per hour to the shock defined by the impulse vector $\widetilde{\mathbf{A}}_0\boldsymbol{\varsigma}$. I keep all candidate $\boldsymbol{\varsigma}$ which produce an impulse response of output per hour that satisfies the shape restriction described in the text. In particular, I impose that the candidate $\boldsymbol{\varsigma}$ yield an impact effect on output per hour which is less than 0.25 percent in absolute value, for which the impulse response at a horizon of 10 quarters is greater than the impact response, and for which the impulse response at 24 quarters is greater than the impulse response at 10 quarters. This procedure identifies a set of candidate impulse response functions to a shock which leads to a small impact effect on labor productivity but a growing response over time. Of the 100,000 candidate $\boldsymbol{\varsigma}$, roughly 30 percent produce impulse responses satisfying the shape restriction. The results in Figure 10 are presented where the solid line is the median response over the candidate $\boldsymbol{\varsigma}$, while the the dashed lines are the 16th and 84th percentiles of the distribution of candidate responses.

References

[1] Anderson, Torben, Tim Bollerslev, Francis Diebold, and Clara Vega. "Micro Effects of Macro Announcements: Real Time Price Discovery in Foreign Exchange." *American Economic Review* 93: 38-62, 2003.

[2] Barro, Robert and Robert King. "Time Seperable Preferences and Intertemporal Substitution Models of Business Cycles." *Quarterly Journal of Economics* 99: 817-839, 1984.

[3] Barsky, Robert and Eric Sims. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." Working paper, University of Michigan, 2008.

[4] Barsky, Robert and Eric Sims. "News Shocks." Working paper, University of Michigan, 2009.

[5] Basu, Susanto, John Fernald, and Miles Kimball. "Are Technology Improvements Contractionary?" *American Economic Review* 96: 1418-1448, 2006.

[6] Beaudry, Paul and Bernd Lucke. "Letting Different Views About Business Cycles Compete." Working paper, University of British Columbia, 2009.

[7] Beaudry, Paul and Franck Portier. "An Exploration into Pigou's Theory of Cycles." *Journal of Monetary Economics* 51: 1183-1216, 2004.

[8] Beaudry, Paul and Frank Portier. "News, Stock Prices, and Economic Fluctuations." *American Economic Review* 96: 1293-1307, 2006.

[9] Beaudry, Paul, Martial Dupaigne, and Franck Portier. "The International Propagation of News Shocks." Working paper, University of British Columbia, 2008.

[10] Campbell, John and Gregory Mankiw. "Permanent Income, Current Income, and Consumption." *Journal of Business and Economic Statistics* 8: 265-279, 1990.

[11] Chari, VV, Patrick Kehoe, and Ellen McGrattan. "Are Structural VARs with Long Run Restrictions Useful in Developing Business Cycle Theory?" *Journal of Monetary Economics* 55: 1337-1352, 2008.

[12] Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson. "The Response of Hours to a Technology Shock: Evidence Based on a Direct Measure of Technology." *Journal of the European Economic Association* 2: 381-395, 2004.

[13] Christiano, Lawrence, Cosmin Ilut, Robert Motto, and Massimo Rostagno. "Monetary Policy and Stock Market Boom Bust Cycles." Working paper, Northwestern University, 2007.

[14] Cochrane, John. "Shocks." *Carnegie-Rochester Conference Series on Public Policy* 41: 295-364, 1994.

[15] Den Haan, Wouter and Georg Kaltenbrunner. "Anticipated Growth and Business Cycles in Matching Models." Working paper, Euro Area Business Cycle Network, 2006.

[16] Farmer, Roger. *The Macroeconomics of Self-Fulfilling Prophecies.* Boston: MIT Press, 1998.

[17] Faust, Jon. "The Robustness of Identified VAR Conclusions about Money." *Carnegie-Rochester Conference Series on Public Policy* 49: 207-244, 1998.

[18] Faust, Jon and Eric Leeper. "When Do Long Run Identifying Restrictions Give Reliable Results?" *Journal of Business and Economic Statistics* 15: 345-353, 1997.

[19] Fernandez-Villaverde, Jesus, Juan Rubio-Ramirez, Thomas Sargent, and Mark Watson. "ABCs (and Ds) of Understanding VARs." *American Economic Review* 97: 1021-1026, 2007.

[20] Francis, Neville, Michael Owyang, and Jennifer Roush. "A Flexible Finite Horizon Identification of Technology Shocks." Working paper, Federal Reserve Bank of St. Louis, 2007.

[21] Gali, Jordi. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89: 249-271, 1999.

[22] Jaimovich, Nir and Sergio Rebelo. "Can News about the Future Drive the Business Cycle?" Working paper, Northwestern University, 2008.

[23] Kahn, Hashmat and John Tsoukalas. "The Quantitative Importance of News Shocks in Estimated DSGE Models." Working paper, Carleton University, 2009.

[24] Kilian, Lutz. "Small Sample Confidence Intervals for Impulse Response Functions." *Review of Economics and Statistics* 80: 218-230, 1998.

[25] King, Robert and Sergio Rebelo. "Resuscitating Real Business Cycles." In John Taylor and Michael Woodford (editors) *Handbook of Macroeconomics,* 2000.

[26] Kydland, Finn and Edward Prescott. "Time to Build and Aggregate Fluctuations." *Econometrica* 50: 1345-1370, 1982.

[27] Leeper, Eric, Todd Walker, and Shu-Chun Yang. "Fiscal Foresight: Analytics and Econometrics." Working paper, Indiana University, 2008.

[28] Pigou, Arthur. *Industrial Fluctuations*. London: MacMillan, 1927.

[29] Schmitt-Grohe, Stephanie and Martin Uribe. "What's News in the Business Cycle?" Working paper, Duke University, 2008.

[30] Rotemberg, Julio. "Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles." *American Economic Review* 93: 1543-1559, 2003.

[31] Shapiro, Matthew and Mark Watson. "Sources of Business Cycle Fluctuations." *NBER Macroeconomics Annual* 3:111-148, 1988.

[32] Sims, Christopher and Tao Zha. "Does Monetary Policy Generate Recessions?" *Macroeconomic Dynamics* 10: 231-272, 2006.

[33] Sims, Eric. "Non-Invertibilities and Structural VARs" Working paper, University of Michigan, 2009.

[34] Uhlig, Harald. "What Drives GNP?" Working paper, Euro Area Business Cycle Network, 2003.

[35] Uhlig, Harald. "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association* 2: 361-371, 2004.

[36] Watson, Mark. "Vector Autoregressions and Cointegration." in Robert Engle and Daniel McFadden (editors) *Handbook of Econometrics* 4, 2843-2915, 1994.

Table 1: Share of Forecast Error Variance Attributable to News Shock

	$h = 0$	$h = 4$	$h = 8$	$h = 16$	$h = 24$	$h = 40$
Tech.	0.0	5.5	12.4	32.1	44.0	56.3
	[0.0,0.0]	[0.3,18.2]	[2.0,32.1]	[15.8,55.0]	[27.8,65.4]	[38.7,74.1]
Consumption	0.6	16.5	31.2	46.5	53.1	56.2
	[0.0,13.5]	[1.4,40.8]	[3.1,59.8]	[8.3,75.4]	[11.2,78.1]	[12.1,80.1]
Output	9.6	7.5	21.3	38.7	47.0	50.7
	[0.5,22.4]	[3.0,24.0]	[4.2,46.3]	[8.2,63.5]	[11.4,69.8]	[11.3,73.4]
Hours	65.0	20.6	10.1	8.8	9.3	8.5
	[22.7,83.0]	[5.7,48.3]	[4.8,33.0]	[4.1,31.0]	[3.7,32.5]	[34.1,31.7]
Total Tech.	100	91.6	90.5	92.8	93.3	91.0
Total Output	71.5	25.2	32.0	44.9	53.2	58.1

The numbers in brackets are the 5th and 95th percentiles of the bootstrapped distribution of variance decompositions. The second to last row shows the fraction of the total technology variance explained by the news shock and the surprise technology shock combined, while the final row shows the total fraction of the forecast error variance of output explained by the two technology shocks.

Table 2: Share of Forecast Error Variance Attributable to News Shock
Long Run Identification

	$h = 0$	$h = 4$	$h = 8$	$h = 16$	$h = 24$	$h = 40$
Tech.	0.0	5.3	9.9	18.6	21.3	27.9
Total Tech.	100	91.3	87.9	79.2	70.5	62.6

This table shows the variance decomposition of technology using a long run restriction similar to Beaudry and Lucke (2009). The final rows show the total technology variance explained by the identified news and surprise technology shocks combined under this identification.

Table 3: HP Filtered Correlations with Output

	RBC Model	US Data
Consumption	0.20	0.88
Hours	0.88	0.88
Investment	0.93	0.80
Tech.	0.90	0.78

This table shows the HP filtered correlations among the variables in the lefthand column with output. The numbers in the column "RBC Model" show the correlations from the standard RBC model with news shocks as the only stochastic shock. The numbers under "US Data" are correlations from postwar US data, and are taken from Table 1 in King and Rebelo (2000).

Figure 1
Model and Monte Carlo Estimated Impulse Responses to News Shocks

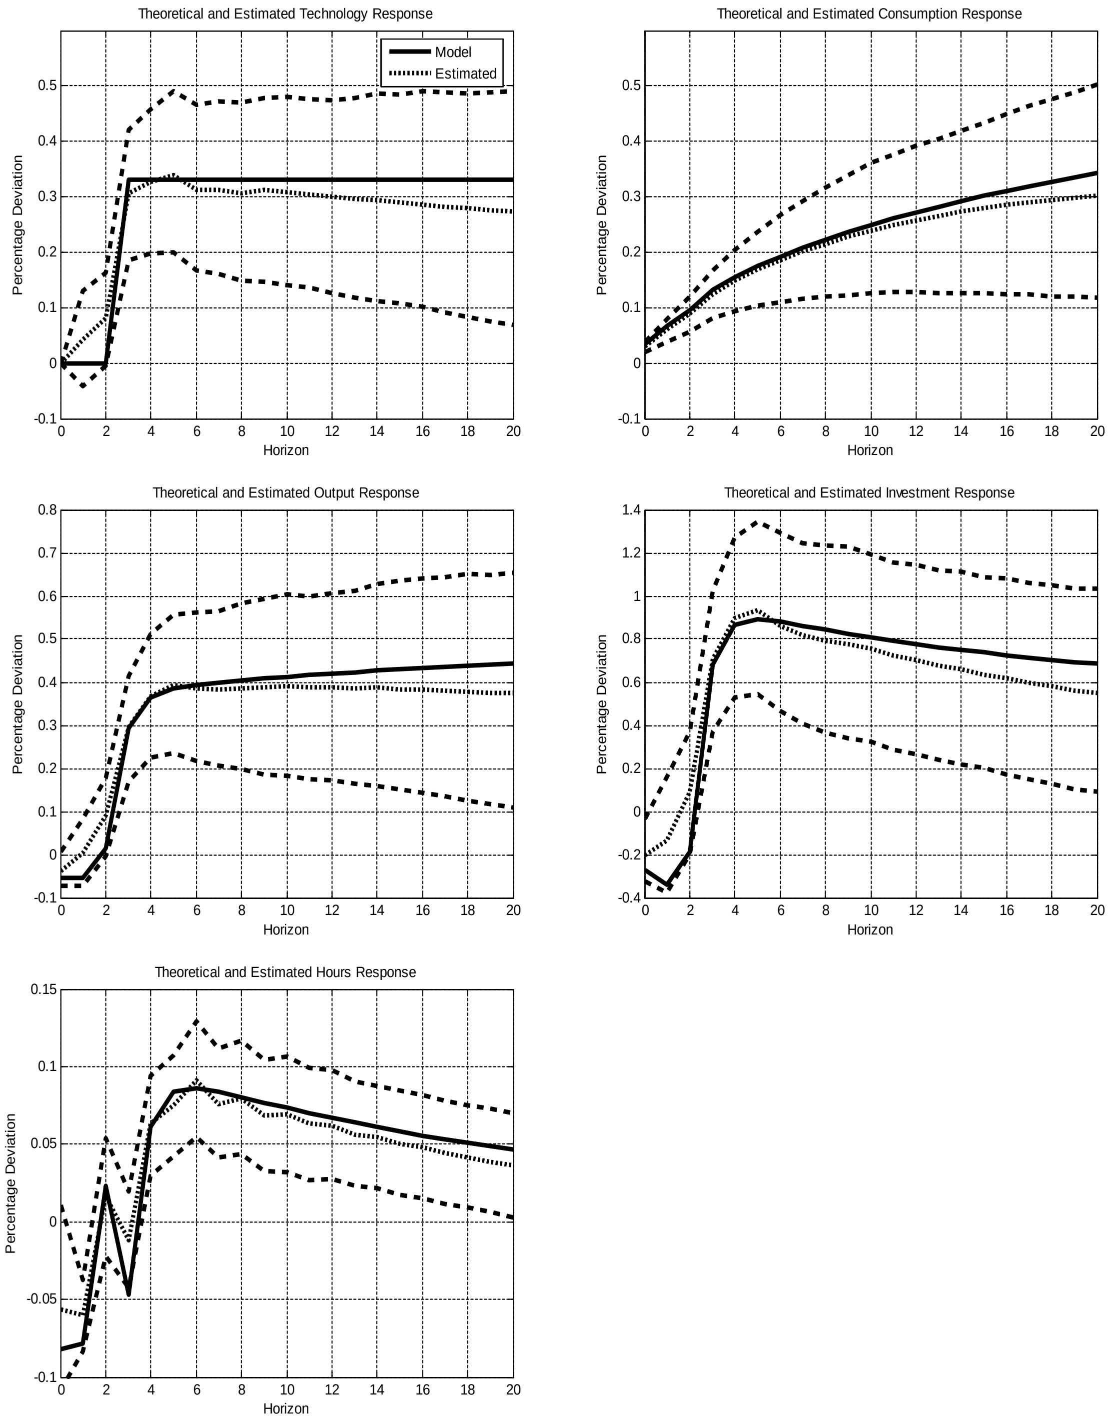

The solid lines show the theoretical impulse response to a news shock from the model of the Appendix, The dotted lines depict the average estimated impulse responses over 2000 Monte Carlo simultions, with the dashed lines representing the 10th and 90th percentiles of the distribution of estimated impulse responses.

Figure 2
Model and Monte Carlo Estimated Impulse Responses to Surprise Technology Shock

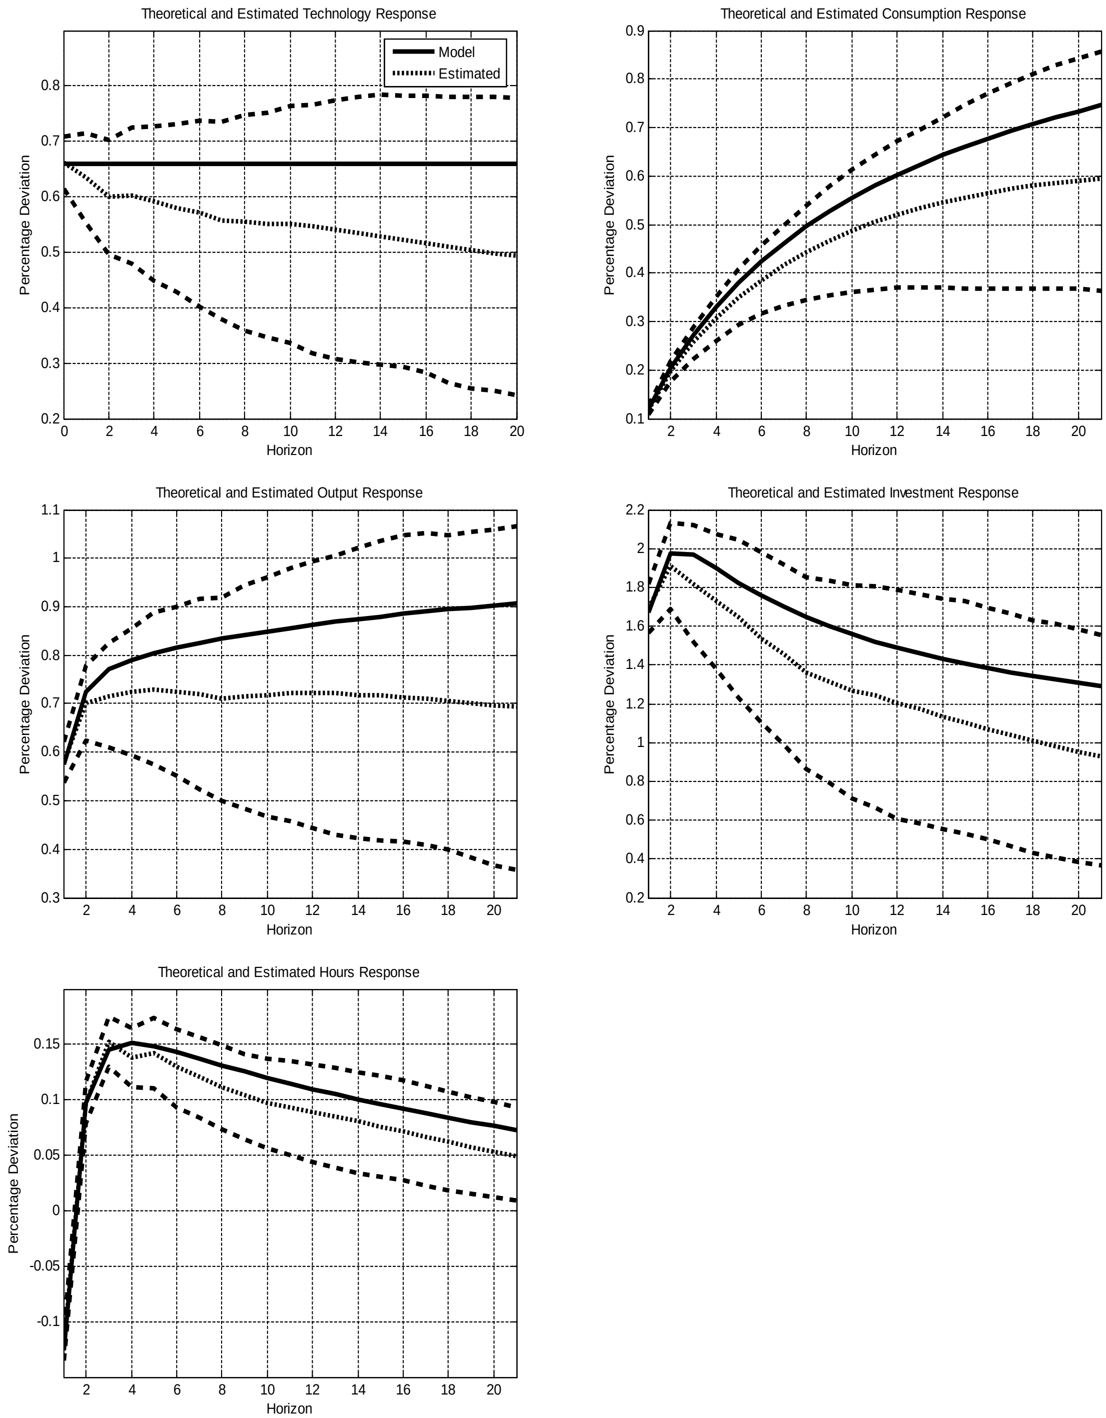

The solid lines show the theoretical impulse response to a surprise technology shock from the model of the Appendix, The dotted lines depict the average estimated impulse responses over 2000 Monte Carlo simultions, with the dashed lines representing the 10th and 90th percentiles of the distribution of estimated impulse responses.

Figure 3

Empirical Impulse Responses to News Shocks

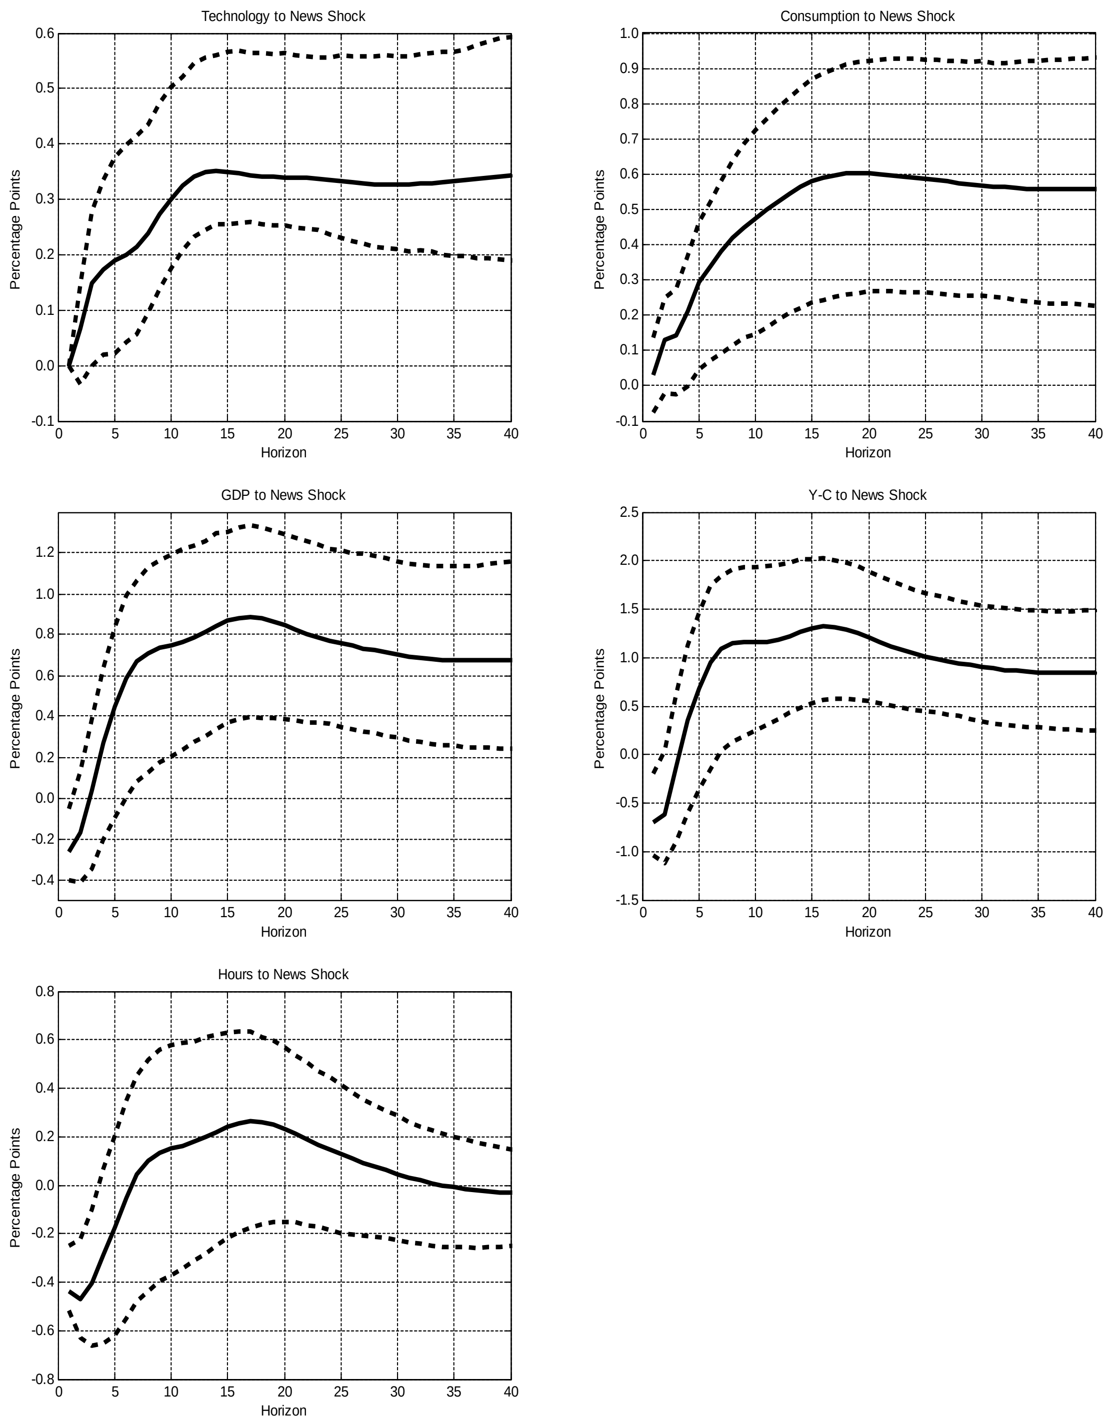

The dashed lines depict the 5th and 95th percentiles of the empirical distribution of impulse responses from a bias-corrected bootstrap procedure.

Figure 4

Cross Correlogram Between News Shock Series and HP Detrended Output

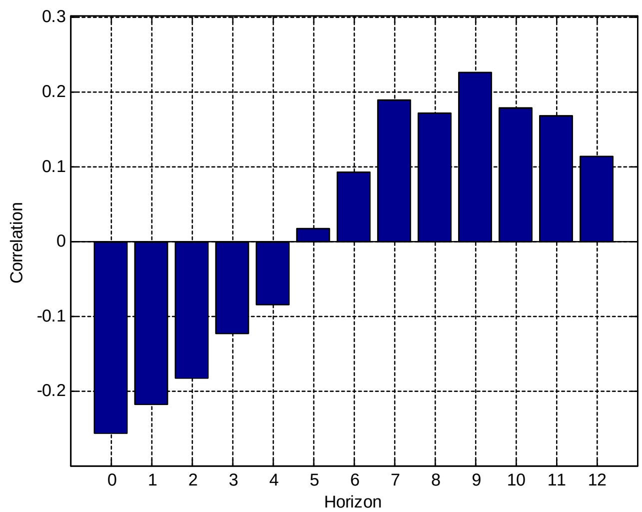

The above figure plots the correlation between the identified news shock series from the benchmark VAR with HP detrended output led over the specified horizon. As such, the number for horizon 0 is the contemporaneous correlation, the number for horizon 1 is the correlation between the shock series and detrended output led one period, and so on.

Figure 5

Identified News Shock Time Series and US Recessions

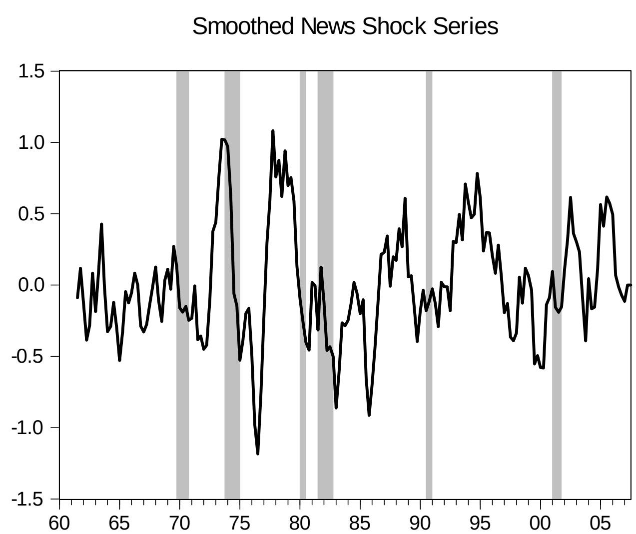

This figure plots the time series of identified news shocks from the benchmark VAR. So as to render the figure more readble, the plotted data is smoothed using a one year moving average.

Figure 6

Historical Decomposition of Output per Capita in US Recessions

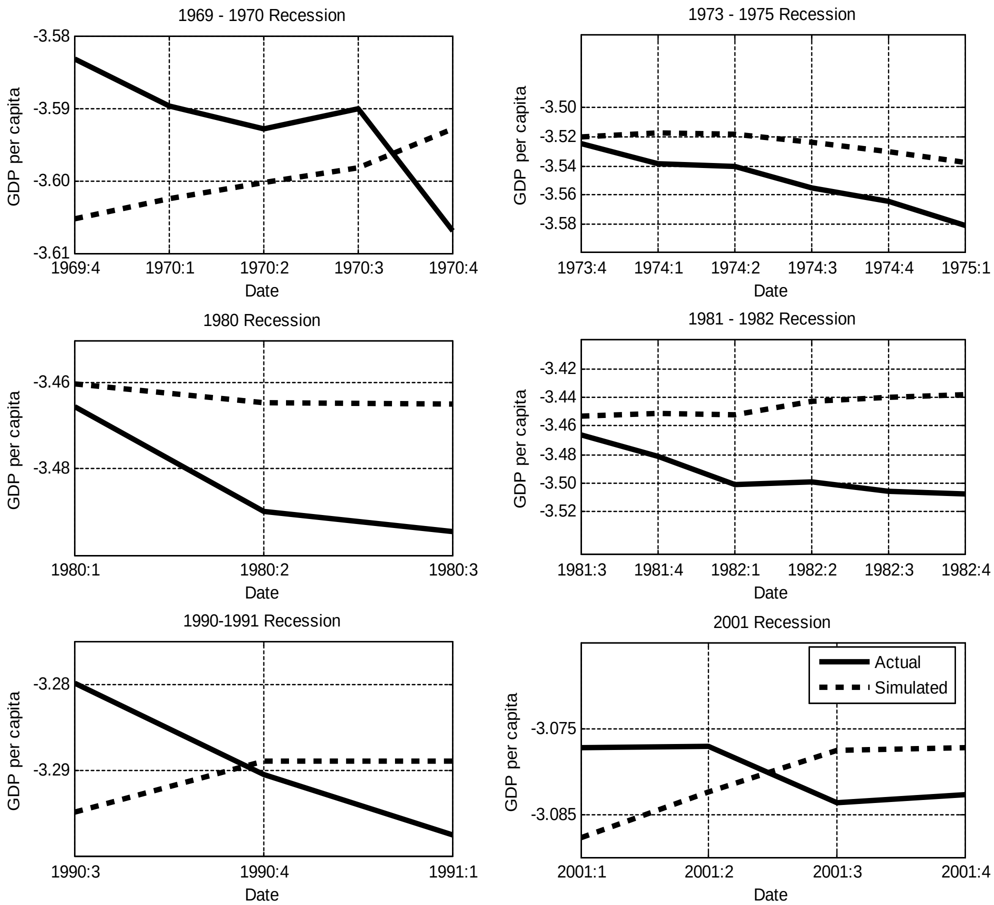

The above shows the actual (solid line) and simulated (dashed line) paths of GDP per capita during the six NBER dated US recessions since 1961.

Figure 7
Estimated Impulse Responses to Surprise Technology Shocks

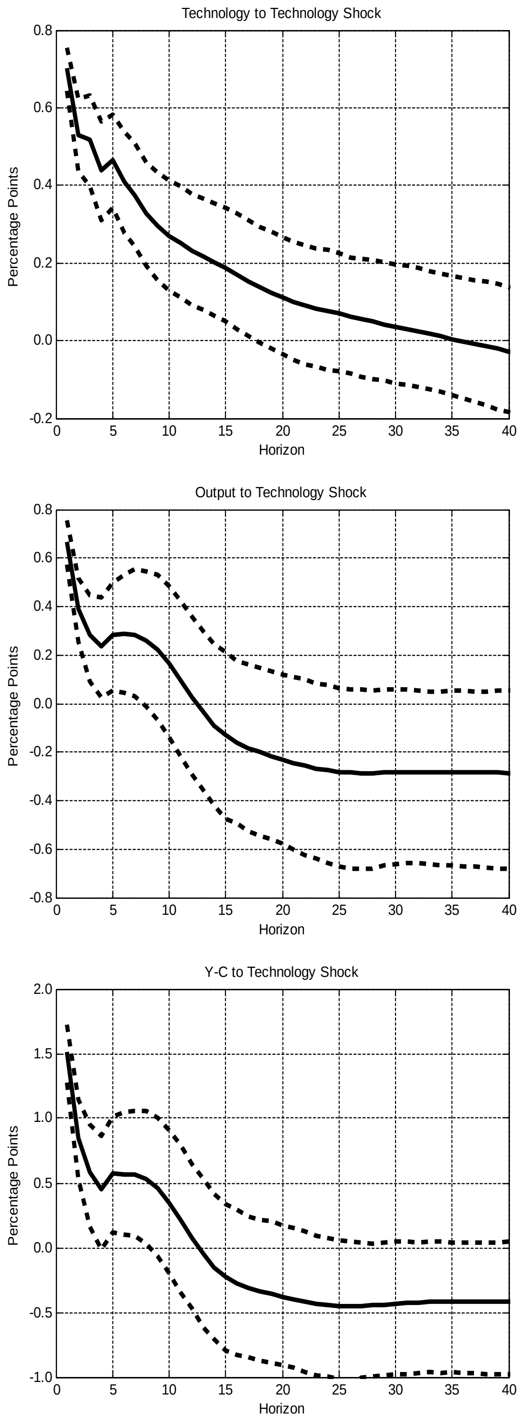

Figure 8
Impulse Responses to News Shocks: Smaller System

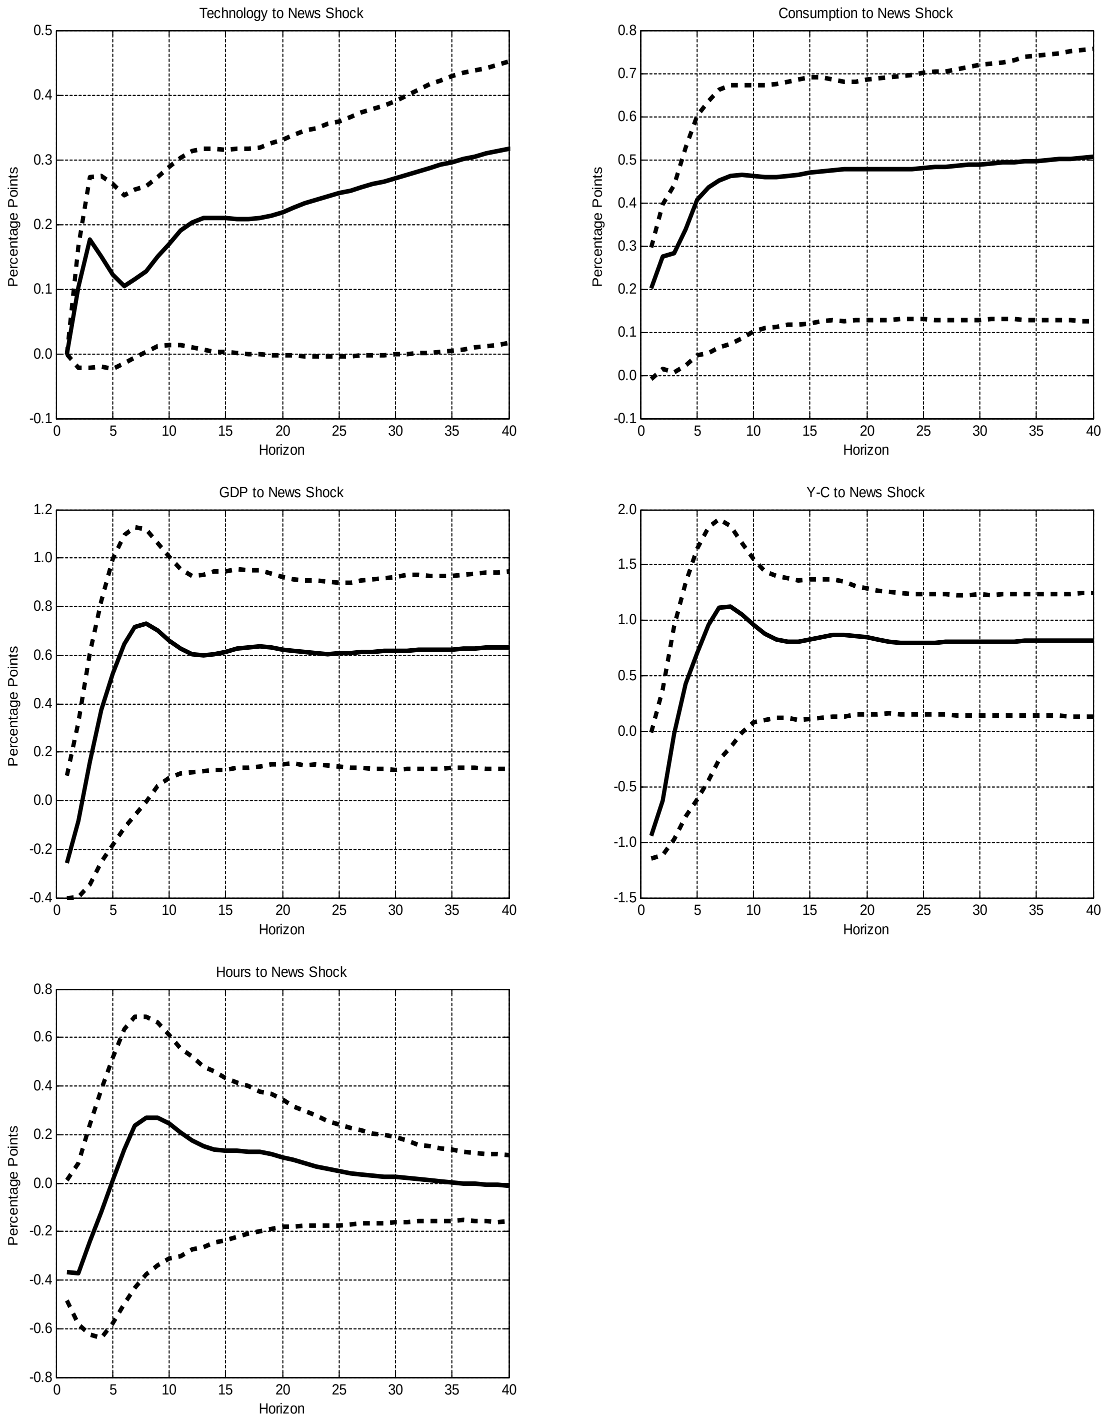

The above are responses from a VAR with technology, stock prices, consumption, output, and hours.

Figure 9
Impulse Responses to News Shocks
Long Run Identification

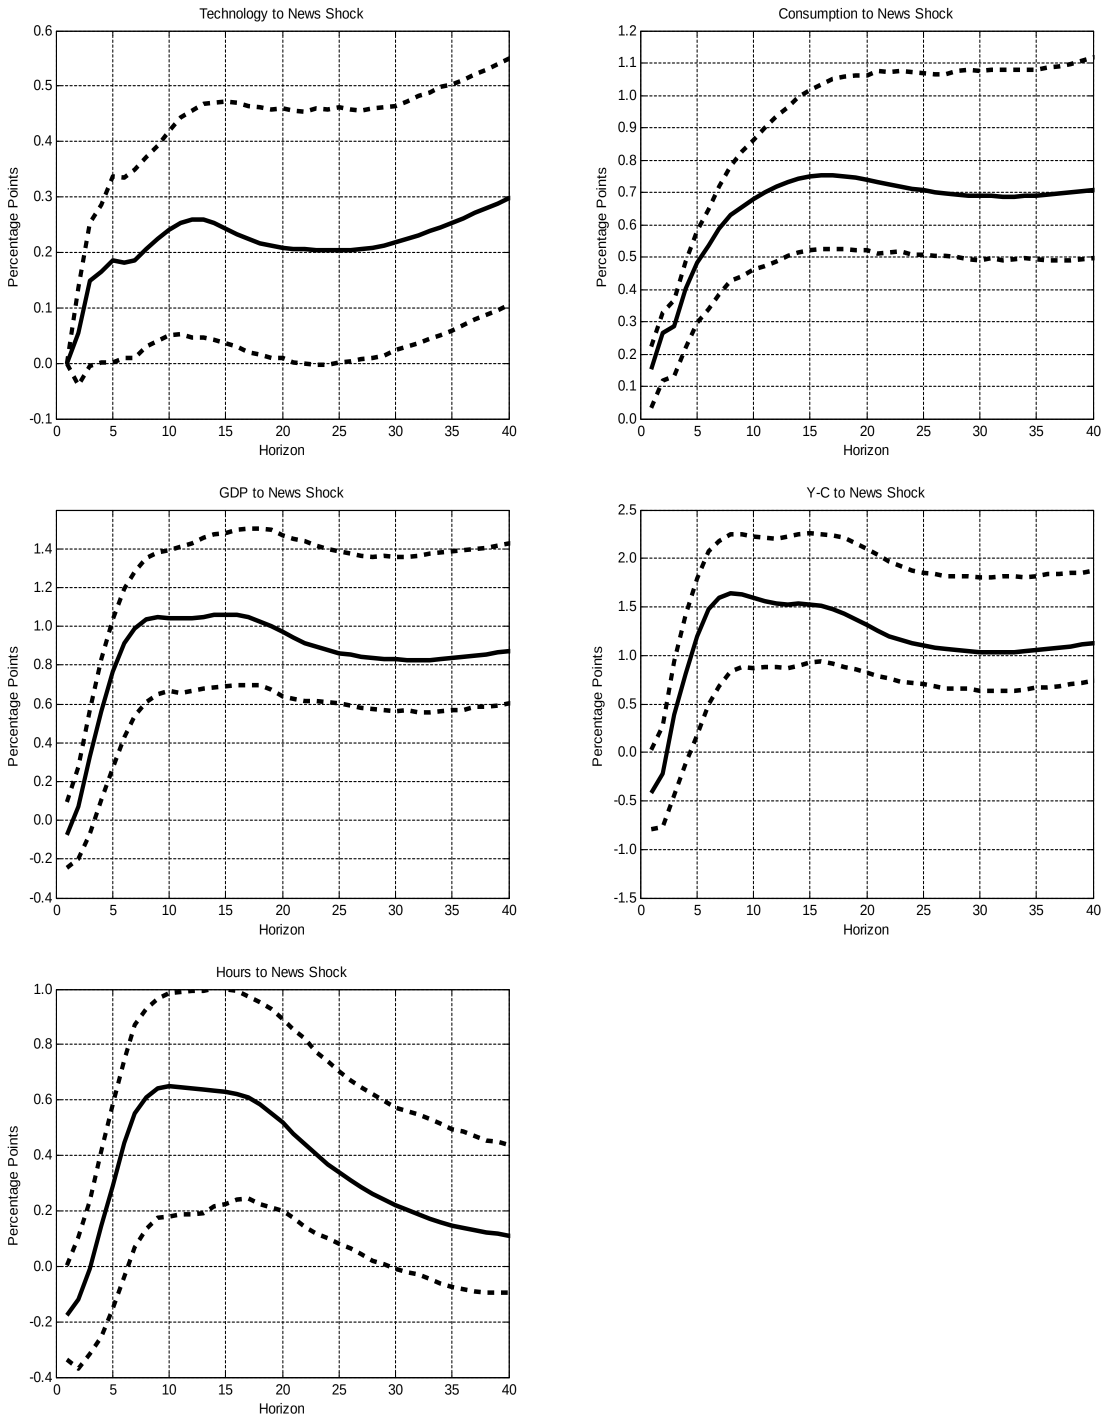

Figure 10

Impulse Responses to News Shocks:
System with Labor Productivity, Shape Restriction Identification

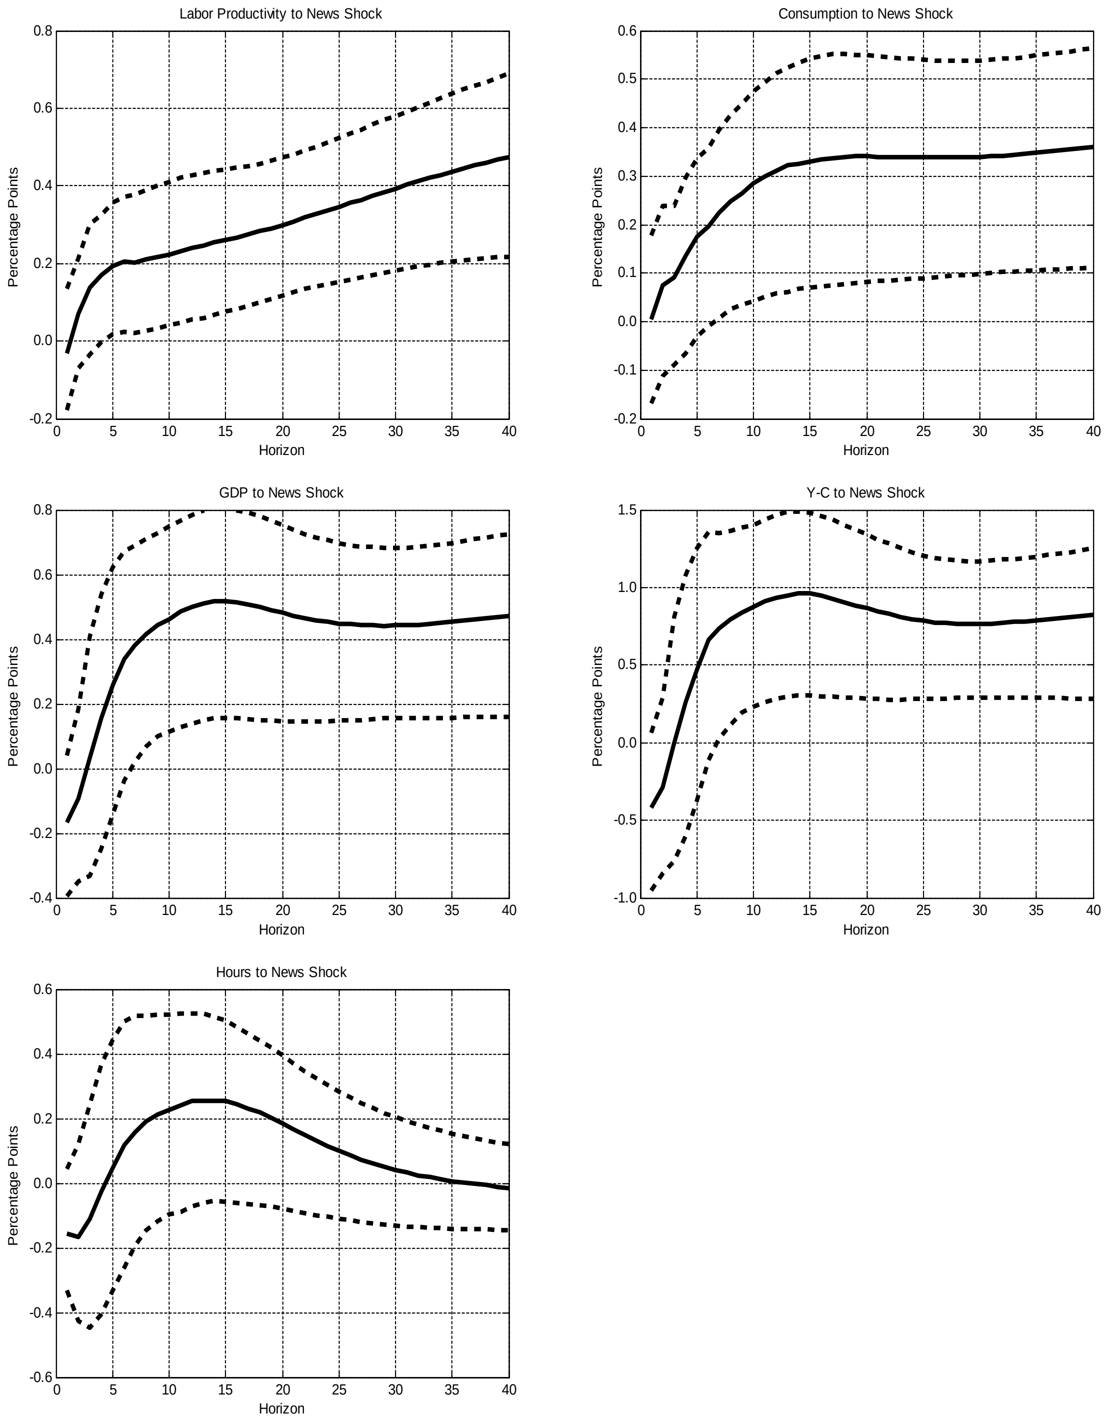

The dark lines are the median impulse responses of the variables to the news shock identified via a shape restriction as described in Section 3.3 and Appendix 6.3. The dashed lines are the 16th and 84th percentiles of the distribution of responses.

Figure 11
Estimated and Theoretical Responses to News Shocks

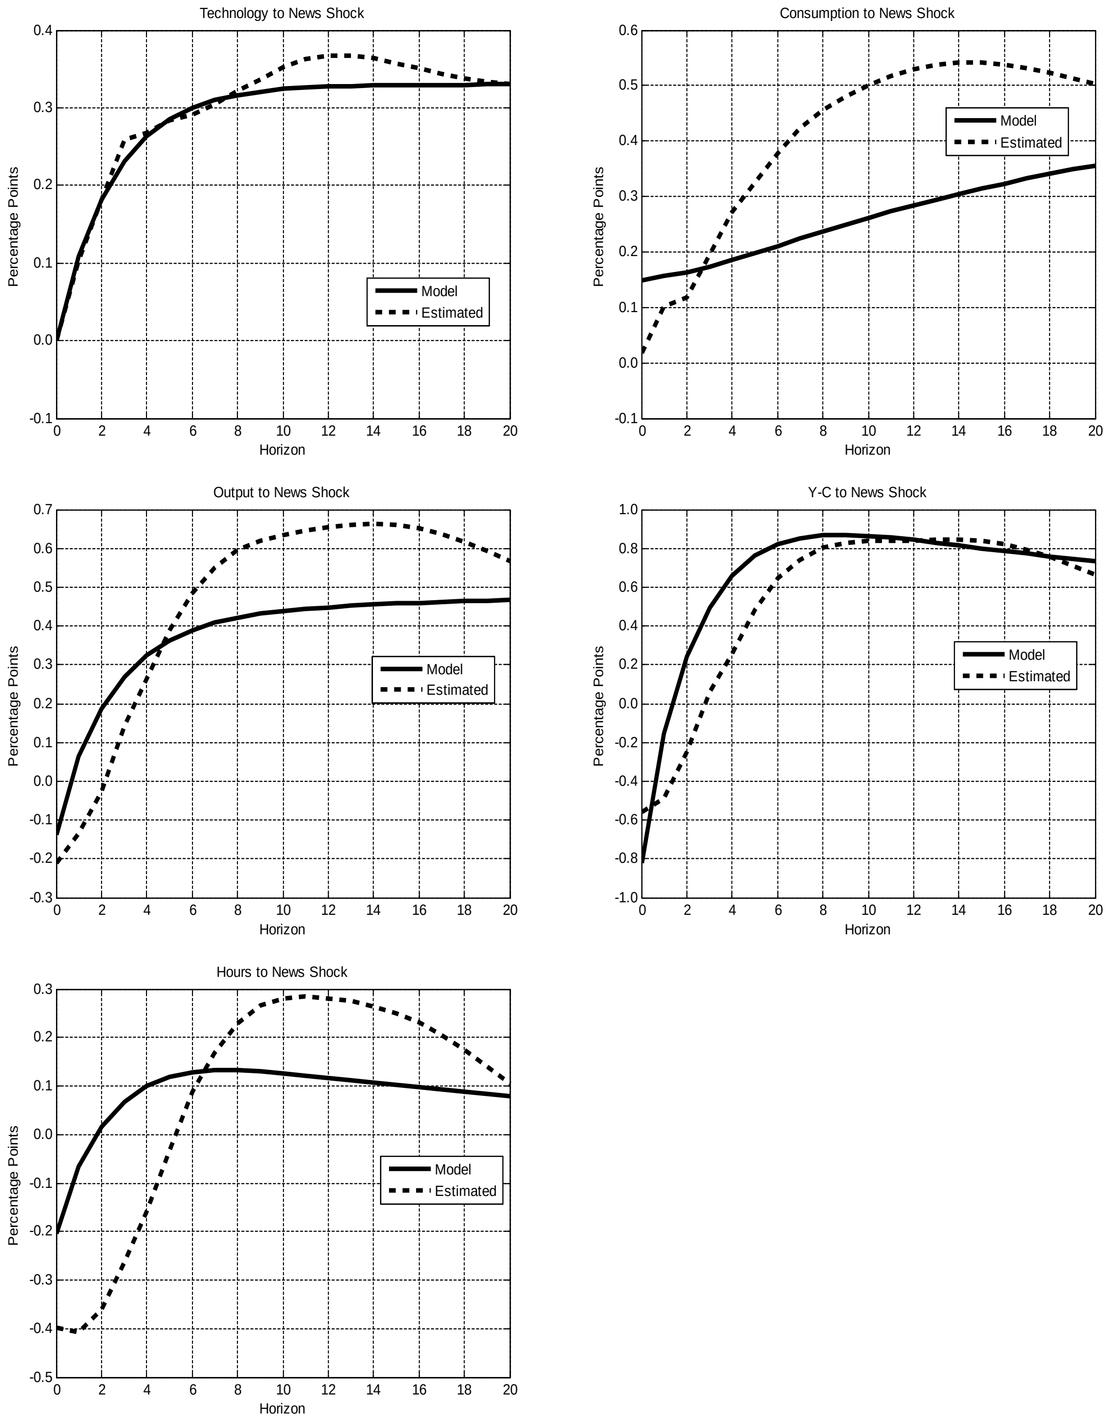

The solid line shows the model generated impulse responses from an RBC model with news shocks and a standard calibration. The dashed line shows the estimated response from the benchmark empirical VAR.