

NBER WORKING PAPER SERIES

NEWS SHOCKS

Robert B. Barsky
Eric R. Sims

Working Paper 15312
<http://www.nber.org/papers/w15312>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2009

We are grateful to Rudi Bachmann, Susanto Basu, John Cochrane, Daniel Cooper, Lutz Kilian, Matthew Shapiro, and seminar participants at the University of Michigan for helpful comments and discussions. We thank John Fernald for providing us with his TFP data. All remaining errors are our own. Barsky acknowledges support from the Russell Sage Foundation as a visiting scholar, and Sims acknowledges the support of the Horace H. Rackham School of Graduate Studies at the University of Michigan. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Robert B. Barsky and Eric R. Sims. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

News Shocks

Robert B. Barsky and Eric R. Sims

NBER Working Paper No. 15312

September 2009

JEL No. E0,E00,E1,E10,E2,E20,E3,E30,E31,E32

ABSTRACT

We implement a new approach for the identification of "news shocks" about future technology. In a VAR featuring a measure of aggregate technology and several forward-looking variables, we identify the news shock as the shock orthogonal to technology innovations that best explains future variation in technology. In the data, news shocks account for the bulk of low frequency variation in technology. News shocks are positively correlated with consumption, stock price, and consumer confidence innovations, and negatively correlated with inflation innovations. The disinflationary nature of news shocks is consistent with the implications of sensibly modified versions of a New Keynesian model.

Robert B. Barsky

Department of Economics

University of Michigan

Ann Arbor, MI 48109-1220

and NBER

barsky@umich.edu

Eric R. Sims

University of Notre Dame

Department of Economics and Econometrics

723 Flanner Hall

Notre Dame, IN 46556

esims1@nd.edu

1 Introduction

Macroeconomists have devoted significant effort to the identification and study of technology shocks. The most commonly used empirical approach is structural vector autoregressions (VAR), frequently making use of long run restrictions (e.g. Shapiro and Watson (1988), Blanchard and Quah (1989), and Gali (1999)). Such identification leaves open the question of whether the resulting shocks affect technology on impact or are “news shocks” that point to future movements in technology while leaving current productivity largely unchanged. This distinction is critical because the two shocks have very different implications in most models, as detailed later in this paper and in Sims (2009).

News shocks have attracted growing interest from macroeconomists in recent years (Cochrane (1994b), Beaudry and Portier (2006), and Barsky and Sims (2008)). Much of this work has been theoretical (Beaudry and Portier (2004) and Jaimovich and Rebelo (2008)), with a focus on whether or not news about changes in future technology can be an important source of cyclical fluctuations. In comparison to the theoretical work in this area, there has been relatively little empirical work aimed at isolating these news shocks, and certainly no widely accepted method for identifying them.

This paper fills that void by proposing and implementing a generalized method for the identification of news shocks. In a vector autoregression (VAR) featuring a utilization adjusted measure of total factor productivity (hereafter “technology”) and several forward-looking variables, we identify the surprise technology shock as the innovation in technology. We then identify the news shock as the structural shock orthogonal to technology innovations that best explains future variation in technology. This identification strategy is an application of principal components. It identifies the news shock as the linear combination of reduced form innovations orthogonal to technology which maximizes the sum of contributions to technology’s forecast error variance over a finite horizon. This is a highly flexible empirical approach. It can be applied to systems estimated in levels or as stationary vector error correction (VECM) models, and on systems with a large number of variables without having to impose additional structure.

Cognizant of recent work questioning the ability of structural VARs to adequately identify economic shocks (e.g. Chari, Kehoe, and McGrattan (2008)), we provide simulation-based evidence that our empirical approach is likely to perform well in practice. We generate data from a New Keynesian model augmented with news shocks about future technology and apply our identification strategy to the simulated data. We find that our methodology applied to artificial data reliably identifies both news and surprise technology shocks as well as their dynamic implications for the variables of the model. In simulated samples of realistic

sizes, the estimated impulse responses to a news shock are roughly unbiased at all horizons, and the average correlation between true and identified shocks exceeds 0.85.

We focus on the implications of news shocks for long run growth and for forward-looking variables; Sims (2009) applies a similar methodology to study the implications of news shocks for the business cycle. We include in our benchmark VAR a quarterly version of the Basu, Fernald, and Kimball (2006) utilization-adjusted technology series, as well as measures of aggregate consumption, stock prices, consumer confidence, inflation, and interest rates. Beaudry and Portier (2006) document that surprise movements in stock prices are informative about future productivity movements, while Barsky and Sims (2008) reach similar conclusions for forward-looking measures of consumer confidence. Aggregate consumption should incorporate information about future fundamentals under the permanent income hypothesis, while inflation is a forward-looking jump variable in typical models with nominal frictions. The interest rate is included to allow the monetary authority to respond to news shocks as well as to check that the real interest rate implications of news shocks are consistent with the general equilibrium predictions of standard DSGE models.

In post-war US data, we find that news shocks are responsible for the bulk of low frequencies movements in productivity. In contrast, surprise innovations to measured technology appear largely transitory. Since information about new processes is typically available before any actual effect on productivity, this finding fits nicely with the idea that a narrow view of technology as the result of “inventions” is largely responsible for the trend, but that there are also a variety of real shocks that are difficult to pin down that behave similarly to the persistent but transitory productivity disturbances emphasized in the real business cycle literature (Kydland and Prescott (1982)). An historical simulation on the basis of our identified VAR shows that surprise technology shocks account for most of the short run variation in technology, while news shocks help to explain the productivity slowdown of the 1970s and ensuing speed up of the 1990s.

We find that favorable news shocks lead to increases on impact in both aggregate consumption and stock prices. Both of these series undershoot their long run responses; this undershooting is consistent with general equilibrium implications associated with increases in real interest rates. While news shocks account for large shares of the variation in aggregate consumption at most horizons, they only modestly contribute to the forecast error variance of stock prices at short horizons, explaining a larger share of stock price variation at lower frequencies. Indeed, there appear to be important movements in stock prices unrelated to technology shocks altogether. Our historical simulations show that news shocks can account for the general downward trend in stock prices from the 1960s through the early 1980s as well as the ensuing bull market from the early 1980s onwards. News shocks do not, however,

capture most of the short run cyclical fluctuations in stock prices evident in the data.

Consistent with the findings in Barsky and Sims (2008), favorable news shocks are positively correlated with surprise movements in forward-looking measures of consumer confidence. Rather strikingly, good news shocks are highly disinflationary, and explain a large share of the forecast error variance of inflation both on impact and at subsequent horizons. The historical simulations reveal that news shocks are capable of explaining most of the important movements in both consumer confidence and inflation over the sample period. In particular, news shocks explain well the coincident high inflation and low confidence of the 1970s and the reverse situation of the 1990s.

Our finding that news shocks are highly correlated with surprise movements in inflation is somewhat surprising. The strong correlation between news and inflation is potentially consistent with forward-looking models of price-setting, in which inflation is equal to a present discounted value of future real marginal costs. The prediction of the benchmark New Keynesian model augmented with a Taylor rule (1993), however, is actually for good news to be inflationary on impact, not disinflationary as we find in the data. In Section 4 we diagnose the reasons for this prediction of the model, and propose various modifications capable of making it better fit the data. We show that real wage rigidity of the type introduced by Blanchard and Gali (2007) is capable of making good news shocks disinflationary. In addition, we show that sensible variations on the Taylor rule – in particular ones in which the monetary authority responds to an activity measure different from the theoretical output gap – are also capable of generating disinflation. We then estimate a subset of parameters of the model with these proposed modifications. We use a minimum distance estimator to pick structural parameters to match the observed response of inflation to a news shock in the data. The parameterized model is capable of producing a disinflation in response to good news that is both quantitatively and qualitatively similar to what we estimate in the data.

The remainder of the paper is organized as follows. The next section lays out our empirical strategy in formal detail and provides simulation evidence that it is in fact capable of doing a good job. Section 3 presents our main results, while Section 4 rationalizes our finding that favorable news shocks are disinflationary in the context of the New Keynesian model with forward-looking price-setting. The final section concludes.

2 Empirical Strategy

We assume that aggregate technology is well-characterized as following a stochastic process driven by two shocks. The first is the traditional surprise technology shock of the real

business cycle literature, which impacts the level of technology in the same period in which agents see it. The second is the news shock, which is differentiated from the first in that agents observe the news shock in advance.

Letting A denote technology, this identifying assumption can be expressed in terms of the moving average representation:

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

$\varepsilon_{1,t}$ is the conventional surprise technology shock while $\varepsilon_{2,t}$ is the news shock. The only restriction on the moving representation is that $B_{12}(0) = 0$, so that news shocks have no contemporaneous effect on technology.¹

The following is an example process satisfying this assumption:

$$\ln A_t = A_{t-1} + g_{t-1} + \varepsilon_{1,t} \tag{1}$$

$$g_t = (1 - \kappa)\bar{g} + \kappa g_{t-1} + \varepsilon_{2,t} \tag{2}$$

Here log technology follows a random walk with drift, where the drift term itself follows a stationary AR(1) process. κ describes the persistence of the drift term and \bar{g} is the steady state growth rate. $\varepsilon_{1,t}$ is the conventional surprise technology shock. Given the timing assumption, $\varepsilon_{2,t}$ has no immediate impact on the level of technology but portends a period of sustained growth.

In a univariate context, it would not be possible to separately identify ε_1 and ε_2 . The identification of news shocks must come from surprise movements in variables other than technology. As such, estimation of a vector autoregression (VAR) seems sensible in this context. In a system featuring an empirical measure of aggregate technology and several forward-looking variables, we identify the surprise technology shock as the reduced-form innovation in technology. The news shock is then identified as the shock that best explains future movements in technology not accounted for by its own innovation. This identification follows directly from our assumption that two shocks characterize the stochastic process for technology. In practice, our identification strategy involves finding the linear combination of VAR innovations contemporaneously uncorrelated with technology innovations which maximally contributes to technology's future forecast error variance. This identification strategy is closely related to Francis, Owyang, and Roush's (2007) maximum forecast error variance

¹More generally, the shock to the level and the shock to the growth rate of technology may be correlated. If so, our orthogonalization assigns the common component to the surprise technology shock.

approach, which builds on Faust (1998) and Uhlig (2003, 2004). On the basis of simulations from a popular DSGE model, we show in subsection 2.2 that our approach is likely to perform well at identifying news shocks in practice.

2.1 Identifying News Shocks

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

$$\mathbf{y}_t = \mathbf{B}(\mathbf{L})\mathbf{u}_t \quad (3)$$

Assume there exists a linear mapping between innovations and structural shocks:

$$\mathbf{u}_t = \mathbf{A}_0\boldsymbol{\varepsilon}_t \quad (4)$$

This implies the following structural moving average representation:

$$\mathbf{y}_t = \mathbf{C}(\mathbf{L})\boldsymbol{\varepsilon}_t \quad (5)$$

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 $\mathbf{A}_0\mathbf{A}_0' = \boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma}$ is the variance-covariance matrix of innovations, but it is not unique. For some arbitrary orthogonalization, $\tilde{\mathbf{A}}_0$ (e.g. a Choleski decomposition), the entire space of permissible impact matrices can be written as $\tilde{\mathbf{A}}_0\mathbf{D}$, where \mathbf{D} is a $k \times k$ orthonormal matrix ($\mathbf{D}\mathbf{D}' = \mathbf{I}$).

The h step ahead forecast error is:

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

The share of the forecast error variance of variable i attributable to structural shock j at horizon h is then:

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

The \mathbf{e}_i denote selection vectors with one in the i th place and zeros elsewhere. The selection vectors inside the parentheses in the numerator pick out the j th column of \mathbf{D} , which we will denote by $\boldsymbol{\gamma}$. $\tilde{\mathbf{A}}_0\boldsymbol{\gamma}$ is then a $k \times 1$ vector corresponding with the j th column of a

possible orthogonalization. The selection vectors outside the parentheses in both numerator and denominator pick out the i th row of the matrix of moving average coefficients, which we denote by $\mathbf{B}_{i,\tau}$.

Let technology occupy the first position in the system, and let the unanticipated shock be indexed by 1 and the news shock by 2. Our identifying assumption implies that these two shocks account for all variation in technology at all horizons:

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

We propose picking parts of the impact matrix to come as close as possible to making this expression hold. With the surprise shock identified as the innovation in technology, $\Omega_{1,1}(h)$ will be invariant at all h to alternative identifications of the other $k - 1$ 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 impact matrix 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 the impact matrix, this suggests choosing the second column of the impact matrix to solve the following optimization problem:

$$\begin{aligned} \boldsymbol{\gamma}^* = \arg \max \quad & \sum_{h=0}^H \Omega_{1,2}(h) = \frac{\sum_{\tau=0}^h \mathbf{B}_{i,\tau} \tilde{\mathbf{A}}_0 \boldsymbol{\gamma} \boldsymbol{\gamma}' \tilde{\mathbf{A}}_0' \mathbf{B}'_{i,\tau}}{\sum_{\tau=0}^h \mathbf{B}_{i,\tau} \boldsymbol{\Sigma} \mathbf{B}'_{i,\tau}} \\ & \text{s.t.} \end{aligned}$$

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

So as to ensure that the resulting identification belongs to the space of possible orthogonalizations of the reduced form, the problem is expressed in terms of choosing $\boldsymbol{\gamma}$ conditional on an arbitrary orthogonalization, $\tilde{\mathbf{A}}_0$. H is some finite truncation horizon. The first two constraints impose that the news shock has no contemporaneous effect on the level of technology. The third restriction (that $\boldsymbol{\gamma}$ have unit length) ensures that $\boldsymbol{\gamma}$ is a column vector belonging to an orthonormal matrix. Uhlig (2003) shows that this maximization problem can be rewritten as a quadratic form in which the non-zero portion of $\boldsymbol{\gamma}$ is the eigenvector

associated with the maximum eigenvalue of a weighted sum of the lower $(k-1) \times (k-1)$ submatrices of $(\mathbf{B}_{1,\tau}\tilde{\mathbf{A}}_0)'$ $(\mathbf{B}_{1,\tau}\tilde{\mathbf{A}}_0)$ 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.

The common assumption in the news shock literature is that a limited number of shocks lead to movements in aggregate technology. Our identification strategy is based solely on this assumption, and does not rely upon (potentially invalid) auxiliary assumptions about other shocks. Our approach is a partial identification strategy, only identifying the two technology shocks. As such, it can be conducted on a system with any number of variables without having to impose additional assumptions.

Our identification strategy is thus highly flexible, and encompasses the existing identifying assumptions in the empirical literature on news shocks. Beaudry and Portier (2006) and Beaudry, Dupaigne, and Portier (2008) propose identifying news shocks with the innovation in stock prices orthogonalized with respect to technology innovations. Were the conditions required for this identification to be valid satisfied, our approach would identify (asymptotically) exactly the same shock. Beaudry and Lucke (2009) propose using a combination of short and long run restrictions to identify news shocks. In particular, in systems featuring technology and stock prices, they use two long run restrictions to identify the two technology shocks, and differentiate the news shock from the surprise technology shock with an orthogonality restriction. This identification is identical to ours as the truncation horizon gets arbitrarily large (i.e. as $H \rightarrow \infty$). In practice the long run identification is problematic in that it identifies a news shock and a surprise technology shock that together leave a large share of the variance of technology unexplained. As shown in Sims (2009), the long run identification fails to account for as much as 40 percent of the variance of measured technology at business cycle frequencies.

Our approach has at least four advantages over previous work. First, we do not rely heavily upon stock prices as an information variable to help reveal movements in future technology. Indeed, we find that stock prices are fairly uninformative about future movements in technology relative to other forward-looking variables. Second, since ours is a partial identification strategy, we can include a large number of variables in the system without having to impose potentially invalid auxiliary assumptions about the other shocks. Third, we address the problem with existing work that the resulting shock leaves a large share of technology unexplained. Finally, our approach has better finite sample properties than the approach based on long run restrictions. Identification at frequency zero is based on sums of VAR coefficients, which are biased in finite samples. Summing up biased coefficients exacerbates the bias, and the resulting identification and estimation are often highly unreliable

(Faust and Leeper (1997)). Francis, Owyang, and Roush (2007) show that medium run identification similar to that proposed here performs better in finite samples than does long run identification.

2.2 Simulation Evidence

We now present simulation evidence which confirms that our proposed empirical strategy is indeed capable of doing a good job of identifying news shocks. We consider a simple New Keynesian model with exogenous price stickiness. The equilibrium conditions of the model log-linearized about the balanced growth path are:

$$E_t c_{t+1} = c_t + \sigma (i_t - E_t \pi_{t+1}) \quad (6)$$

$$c_t = y_t \quad (7)$$

$$\pi_t = \left(\frac{(1-\theta)(1-\theta\beta)}{\theta\beta} \right) mc_t + \beta E_t \pi_{t+1} \quad (8)$$

$$y_t = a_t + n_t \quad (9)$$

$$mc_t = w_t - p_t - a_t \quad (10)$$

$$\frac{1}{\eta} n_t = w_t - p_t - \frac{1}{\sigma} c_t + \psi_t \quad (11)$$

$$i_t = \rho i_{t-1} + (1-\rho) \left(\phi_y (y_t - y_t^f) + \phi_\pi (\pi_t - \pi^*) \right) + \varepsilon_{3,t} \quad (12)$$

$$\psi_t = \zeta \psi_{t-1} + \varepsilon_{4,t} \quad (13)$$

These are the standard equations of the New Keynesian model – see Woodford (2003) or Galí (2008) for a complete derivation. Equation (6) is the consumption Euler equation, with σ the elasticity of intertemporal substitution. Equation (7) reflects the accounting identity that, in the model without capital, all output must be consumed in equilibrium. Equation (8) is the conventional New Keynesian Phillips Curve, with θ describing the degree of exogenous price stickiness and β the subjective discount factor. Output is produced according to a constant returns to scale production function in technology and employment.

Let $a_t = \ln A_t$, and assume that it follows the stochastic process given in (1) and (2) above. Equation (10) defines real marginal cost as the (log) discrepancy between the real wage and the marginal product of labor. Equation (11) is the labor supply curve, with η the Frisch elasticity and ψ_t a stochastic preference parameter, which obeys equation (13). Equation (12) describes a partial adjustment nominal interest rate rule, with y_t^f corresponding to the level of output that would obtain in the absence of nominal rigidities.

We choose a baseline parameterization as follows: $\sigma = 1$, $\eta = 1$, $\beta = 0.99$, $\theta = 0.67$, $\rho = 0.75$, $\phi_y = 1$, $\phi_\pi = 1.5$, $\zeta = 0.6$, $\kappa = 0.5$, and $\bar{g} = 0.0025$. Technology (and thus output) grow at the annualized rate of one percent along the balanced growth; given the unit intertemporal elasticity of substitution, labor hours are stationary. We draw the four shocks from mean zero normal distributions with the following standard deviations: $\sigma_1 = 0.006$, $\sigma_2 = 0.00165$, $\sigma_3 = 0.001$, and $\sigma_4 = 0.001$. Given the calibration of κ , a one standard deviation news shock portends a level of technology that is one third of a percent higher along the new balanced growth path.

For this calibration of parameters, we simulate 2000 data sets with 200 observations each. For each simulation we estimate a four variable, unrestricted vector error correction model (VECM) in technology, consumption, inflation, and the interest rate with four lags.² Similar results obtain when the system is estimated as a VAR in levels. We identify the news shock by following the identification strategy outlined above, maximizing the variance share over a horizon of twenty quarters.

Figure 1 depicts both theoretical and estimated impulse responses averaged over the simulations to a news shock. The theoretical responses from the calibrated model are in solid black, while the estimated responses averaged over the simulations are depicted by the dotted lines. The dashed lines depict the 10th and 90th percentiles of the distribution of estimated impulse responses. The real interest rate response in the simulations is imputed as the nominal interest rate response less the VAR forecast of one quarter ahead inflation. The interest rate and inflation responses are expressed at an annualized rate.

A cursory examination of the figure reveals that our empirical strategy is capable of performing well on model generated data. The estimated impulse responses to a news shock are roughly unbiased on impact and at subsequent horizons. There is some evidence of a slight upward bias in the estimated responses of technology and consumption at longer horizons, though it is very small. The estimated responses from the simulations capture well the dynamics implied by the model, and the distributional confidence bands contain

²In particular, we allow the matrix of cointegrating relations to be full rank, so that this is asymptotically equivalent to a VAR in levels with one more lag. This is an inefficient estimation procedure, as we know from the model that there is only one cointegrating relationship. Nevertheless, this is the conservative approach advocated by Hamilton (1994), and we will also apply it in the empirical section of the paper.

the model responses at all horizons. Similarly good results obtain when focusing on the surprise technology shock. The average correlation between the identified and true news shocks across the simulations is 0.83. The median correlation is 0.88, and the 10th and 90th percentile correlations are 0.67 and 0.94, respectively. As the sample size becomes arbitrarily large, the distributions of estimated responses collapse on the model responses and the correlation between true and identified shocks approaches one.

We also want to verify that we do not spuriously identify a news shock when no such shocks are present. When the data are generated without news shocks (i.e. with $\sigma_2 = 0$), our empirical procedure identifies a very small spurious news shock in the sense that, in finite samples, it identifies a positive long run response of technology (and consumption, given that they are cointegrated). Nevertheless, the estimated responses of interest rates and inflation (and consumption on impact and at high frequencies) to the non-existent news shock are unbiased. This small degree of spuriousness goes away as the simulated sample sizes become larger.

Alternative calibrations of the parameters of the model or slight differences in the empirical procedure (different truncation horizon, different lag lengths, VAR in levels instead of VECM, etc.) produce very similar results. The Appendix to Sims (2009) conducts simulation exercises for a similar empirical procedure on data generated from a real model with capital and reports similarly good simulation results. Sims (2009) also considers the role of any potential non-invertibilities (see Fernandez-Villaverde, Rubio-Ramirez, Sargent, and Watson (2007)) owing to the presence of news shocks and shows that these are likely of limited importance. Non-invertibilities arise when the variables included in the VAR fail to reveal the value of missing states. As stressed by Watson (1994), the inclusion of forward-looking variables mitigates the impact of potential non-invertibilities even if these variables do not fully reveal the missing state(s). Our simulation results, as well as the inclusion of a variety of additional forward-looking variables in our empirical VARs, suggest that one need not be overly concerned with non-invertibilities in this context.³

3 Empirical Results

Our empirical strategy requires a suitable measure of aggregate technology. The conventional Solow residual is not particularly appealing, as standard growth accounting techniques make

³Blanchard, L'Hullier, and Lorenzoni (2009) argue that the presence of news shocks observed with noise renders the system non-invertible, invalidating structural impulse response analysis. In their model it is not possible to separately identify the impulse responses to a noise disturbance, but structural VAR identification of the news shock from the perspective of the agents in the model continues to be capable of reliably identifying the model's structural impulse responses.

no attempt to control for unobserved input variation. Since our identification strategy requires orthogonalization with respect to technology, it is important that our measure of technology adequately control for unobserved input variation. To address this issue, we employ a quarterly version of the Basu, Fernald, and Kimball (2006) technology series.⁴ Their insight is to exploit the first order condition implying that firms should vary input intensity along all margins simultaneously. As such, they propose proxying for unobserved input variation with observed variation in hours per worker.

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 Solow residual is then $\Delta A = \Delta Y - \alpha\Delta K - (1 - \alpha)\Delta QH$, where α is capital's share. The utilization correction subtracts from this $\Delta U = \alpha\Delta Z + (1 - \alpha)\Delta E$, 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.

We measure consumption as the log of real consumption of non-durables and services. Similar results obtain when durable consumption is included. We convert this series to per capita by dividing by the civilian non-institutionalized population aged sixteen and over. Our results are insensitive to this transformation. Our measure of stock prices is the log of the real S&P 500 Index. The measure of inflation is the annualized percentage change in the CPI for all urban consumers. We use the three month Treasury Bill as our measure of the interest rate. The stock price, price index, and interest rate data are available at a monthly frequency. We convert to a quarterly frequency by taking the last monthly observation in 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.⁵ For more on the confidence data, see Barsky and Sims (2008).

We include the following variables in our benchmark system: the Basu, Fernald, and Kimball (2006) technology measure, stock prices, consumption, consumer confidence, inflation, and interest rates. The data begin in the first quarter of 1960 and end in the third quarter of 2007. We follow a conservative approach and estimate the system as an unrestricted vector error correction model (VECM); we obtain nearly identical results when estimating the system as a VAR in levels. As suggested by a variety of information criteria,

⁴This series was constructed and given to us directly by John Fernald.

⁵The question underlying the confidence data 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 and 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.

we estimate the system with four lags. In terms of the identification strategy outlined in the previous section, we set the truncation horizon at $H = 60$. The news shock is thus identified as the structural shock orthogonal to technology innovations that best explains technology movements over a fifteen year horizon.

Figure 2 shows the estimated impulse responses to a news shock. The dashed lines represent one standard error confidence bands, and are obtained from the bias-corrected bootstrap of Kilian (1998). Following a favorable news shock, technology grows smoothly for an extended period of time, with a long run response in the neighborhood of 0.5 percent. Consumption jumps up modestly on impact. After the impact effect, it grows rapidly for a number of quarters, reaching a new long run level of roughly 0.75 percent. The significant undershooting of consumption is consistent with the general equilibrium implications of higher real interest rates, which is broadly compatible with what we estimate in the data.⁶ The implied intertemporal elasticity of substitution from the estimated responses is 0.56, which is well within the range of other estimates in the literature.

Stock prices increase on impact in response to a favorable news shock, though this effect is statistically insignificant. Immediately after impact, they rise rather sharply over the next four to eight quarters, quickly levelling off to a new permanently higher steady state. The sharp predictable increase in stock prices following impact (though not statistically significant) is consistent with the general equilibrium implications of higher real interest rates that we find in the data.⁷ Consumer confidence rises strongly and significantly on impact in response to the favorable news. It rises further after impact before reverting to its initial value. This impulse response is consistent with the findings in Barsky and Sims (2008) that confidence innovations are prognostic of future productivity improvements. Perhaps the most striking impulse response is that of inflation. Following a good news shock, inflation jumps down sharply, and this effect is highly statistically significant. While the disinflation is statistically significant for a number of quarters after impact, it is not particularly persistent, with the largest response on impact.

⁶The real interest rate impulse response is imputed in the data as the nominal interest rate responses less the one quarter ahead VAR forecast of inflation, and is expressed at an annualized percentage rate. The point estimate of the impact response of the real interest is negative, though statistically insignificant, but is positive and significant at subsequent horizons. The calculation of the intertemporal elasticity is based on a regression of the consumption growth response on the non-annualized real interest rate response.

⁷There are specifications of our identification strategy in which the impact effect of the news shock on stock prices is negative (though also statistically insignificant). In particular, the impact effect on stock prices is smaller the smaller is the truncation horizon in the identification problem. The theoretical impact of favorable news on stock prices is ambiguous in most models when rates of return rise; an impact decline in stock prices is potentially consistent with the general equilibrium implications of rising real rates we find in the data. Regardless of the truncation horizon, the impact effect is always followed by positive growth in stock prices to a new higher steady state level.

Table 1 shows the forecast error variance decomposition for our benchmark estimation. The numbers in brackets are the one standard error bias-corrected bootstrap confidence intervals. The news shock explains a growing share of the variance of technology as the horizon increases; at a horizon of ten years, for example, news shocks explain more than half of the variation in technology. Our identified shock accounts for a modest, though non-negligible, share of the consumption innovation variance. The news shock quickly accounts for the bulk of the variance in consumption as the horizon grows. News shocks are only weakly correlated with stock price innovations on impact, but, similarly to consumption, account for a growing share of stock price movements at lower frequencies. The identified shock is positively and strongly correlated with consumer confidence innovations and explains a large share of movements in confidence at all horizons. News shocks explain a modest fraction of interest rate variations. Perhaps somewhat surprisingly, we find that news shocks account for the bulk of variation in inflation, explaining slightly more than 60 percent of its innovation variance.

Figure 3 shows impulse responses to the surprise technology shock. The upper left response shows the impulse response of technology to its own innovation. Strikingly, this response is quite transitory. In particular, technology jumps up roughly 0.7 percent on impact but begins to decline immediately, with the point estimate of the response roughly zero at horizons in excess of eight years. Technology's estimated response to its own innovation, in conjunction with the slowly-building response to the identified news shock, suggests that the bulk of the permanent component of productivity is attributable to news shocks.⁸ The surprise technology shock leads to small transitory increases in both consumption and stock prices; the reversion in these series is consistent with the equilibrium implications of lower real rates, which is what we find in the data. The surprise technology shock is associated with little important movement in consumer confidence, disinflation at high frequencies, and slightly higher inflation at longer horizons.

One narrow view of aggregate technology is that it reflects inventions and the development of new productive processes. It seems reasonable that this kind of technological progress is at least partly forecastable and thus known in advance. Implicit in the real business cycle literature, on the other hand, is the idea that there are also difficult to pin down real shocks which manifest themselves as transitory but persistent movements in measured technology. Our findings support the notion that the former is responsible for the trend, while the latter accounts for most of the high frequency variation in technology.

⁸We do not impose that the long run response of technology to its own innovation is zero. Indeed, it is technically not – the point estimate of the response is roughly -0.1 percent at sufficiently long horizons. Likewise, the point estimates for the long run responses of both consumption and stock prices are slightly negative, though all are indistinguishable from zero in the both the statistical and economic senses.

Table 2 presents corroborating evidence for these conclusions from a series of long horizon regressions. In particular, the table shows the adjusted R^2 from several regressions of k step ahead technology growth on the current levels of the remaining variables in our benchmark system. While we are able to account for only about 3 percent of the one quarter ahead variation in technology growth, almost 25 percent of technology growth over a one year horizon is explicable by our forward-looking variables. This number rises to more than 50 percent at horizons in excess of five years. Our findings that a large fraction of productivity growth over long horizons is predictable and that the low frequency component of productivity is largely unrelated to technology innovations are similar to Rotemberg's (2003) model of smooth trends driven by slowly diffusing technical progress.

Figure 4 depicts the impulse responses of technology and stock prices to a stock price innovation orthogonalized with respect to the technology innovation. After a period of initial decline, technology grows slowly, with a positive long run response, though smaller in magnitude than technology's response to our identified news shock. This impulse response is nearly identical to the responses from the same identification in Beaudry and Portier (2006) and Beaudry, Dupaigne, and Portier (2008). The qualitative and quantitative discrepancies between technology's response to a news shock and its response to an orthogonalized stock price innovation are consistent with our finding that the news shock is only modestly correlated with stock price innovations. In response to its own innovation orthogonalized with respect to technology, the stock price rises on impact and then reverts, though levelling off to a new higher level in the long run. The estimated long run response is quantitatively similar in magnitude to the long run response of stock prices to the news shock. In conjunction with the estimated reversion to its own orthogonalized innovation at low horizons, this suggests that there is an important transitory component to stock prices. This finding is consistent with Cochrane's (1994a) conclusion that stock price innovations orthogonalized with respect to dividends are largely transitory.

In Figure 5 we show several historical simulations from our benchmark system. The upper two figures plot the actual and simulated values of technology, with the simulated values obtained using the estimated VAR coefficients assuming that news shocks or surprise technology shocks are the only stochastic disturbances in the system, respectively. News shocks appear to explain movements in technology over long horizons quite well, while the surprise technology shock accounts for almost all of the short run variation. In particular, the news shock simulation does a good job of accounting for the productivity slowdown in the 1970s and ensuing speedup in the 1990s. News shocks do not explain significant short run fluctuations in technology. These simulations are consistent with the findings from our impulse responses and variance decomposition that news shocks are the main driving force

behind low frequency movements in technology, while surprise technology shocks account for most of the high frequency variation.

The remaining plots in Figure 5 show the simulated and actual values of some of the other series in the benchmark system, assuming that news shocks are the only shock. Our identified news shock does an excellent job in accounting for historical movements in both inflation and consumer confidence. In particular, the news shock explains well the coincident high inflation and low confidence of the 1970s as well as the reverse situation in the 1990s. News shocks appear to do an exceptional job of explaining historical movements in consumption. Consistent with the results from the variance decomposition, news shocks do a good job accounting for low frequency movements in stock prices. In particular, the simulation does a good job at picking up the general downward trend in stock prices from the 1960s through the early 1980s as well as the bull market from the early 1980s onward. News shocks do not appear to account for the large cyclical variations in stock prices evident in the data.

Figure 6 shows estimated impulse responses to a favorable news shock from a system similar to our benchmark, but with average labor productivity in place of the utilization corrected technology measure.⁹ Our measure of labor productivity is output per hour in the non-farm business sector, and is obtained from the BLS. The estimation and identification of news shocks are otherwise the same as before. The results are qualitatively very similar to the results from the system with the corrected technology measure. Labor productivity grows smoothly and steadily in response to the news shock, with a long run response that is quantitatively somewhat larger than is the response of technology.¹⁰ News shocks account for a larger share of the innovation variance in stock prices in the system with labor productivity, and the impulse response of stock prices is quantitatively larger at all horizons. Consumption jumps up by less on impact in response to good news in the system with labor productivity, but otherwise follows a very similar dynamic path. Consumer confidence still rises on impact and at most horizons, though the response is somewhat smaller. As before, news shocks are highly disinflationary and are associated with higher real interest rates. News shocks continue to appear to account for a large share of the permanent component of productivity. The correlation between the news shock identified in this system with the shock from the

⁹The assumption that news shocks are contemporaneously orthogonal to the empirical measure of technology becomes apparently more precarious when using average labor productivity in place of technology. Nevertheless, Ball and Moffitt (2001) have argued that average labor productivity is a more exogenous measure of true technology than is total factor productivity.

¹⁰In a model with capital accumulation, it is to be expected that average labor productivity would respond more than true technology in the long run to a news shock of the same size. With a capital's share of one-third and stationary labor hours, a neoclassical model, for example, would predict a long run response of labor productivity 1.5 times that of true technology. The impulse responses in Figure 6 are roughly consistent with this prediction.

benchmark system with the utilization technology measure is also high at 0.86.

While some small quantitative discrepancies do exist, our qualitative results are robust to other sensible variations on our benchmark estimation. The general pattern of responses is similar when using the uncorrected Solow residual, though stock prices and consumption respond less in the long run and there is some evidence of reversion in the technology response to the news shock. Likewise, we obtain qualitatively similar results with different lag lengths and different specifications of the truncation horizon in the optimization problem underlying identification, as well as when the system is estimated as a VAR in levels as opposed to a VECM. We robustly find that favorable news shocks account for an important part of the permanent component of productivity, are strongly and negatively correlated with inflation innovations, positively correlated with consumer confidence innovations, positively correlated with consumption innovations, and are associated with increasing stock prices.

4 Inflation and News Shocks

Our main empirical findings can be summarized as follows. Shocks contemporaneously uncorrelated with technology innovations account of the bulk of productivity movements over long horizons, while technology innovations themselves are quite transitory. News shocks are associated with important fluctuations in aggregate consumption, stock prices, consumer confidence, and consumer price inflation. That forward-looking variables such as consumption or stock prices would incorporate news about future productive possibilities is not surprising. That a survey measure of consumer confidence would also accurately reflect information about the future may be more surprising, but is consistent with the evidence in Barsky and Sims (2008). That news shocks are so heavily incorporated into inflation innovations is the most intriguing and unexpected result, and we examine it in more detail in this section.

A natural framework for studying movements in inflation is the New Keynesian model with Calvo (1983) price-setting. This model offers a potential explanation for our empirical finding that favorable news about future productivity is highly disinflationary. Solving forward the New Keynesian Phillips Curve (see equation (8)), one sees that current inflation is equal to a present discounted value of expected future real marginal costs:

$$\pi_t = \frac{(1 - \theta)(1 - \theta\beta)}{\theta\beta} \sum_{j=0}^{\infty} \beta^j E_t m c_{t+j} \quad (14)$$

$(1 - \theta)$ is equal to the probability that firms will get to update their prices in any period, while β is the subjective discount factor. Other factors held constant, expected future

productivity improvements lower expected real marginal costs, and thus exert downward pressure on current inflation.

In general equilibrium, however, other factors are not held constant, and the prediction of the benchmark model as described in Section 2.2 is actually for good news shocks to be inflationary, not disinflationary. Figure 7 replicates the theoretical responses of technology and inflation to a favorable news shock, using the calibration of the model described above. In response to news that technology will grow more rapidly, inflation rises on impact before quickly reverting to zero in the model. There are at least two different but complementary ways of understanding why the model predicts that good news should be inflationary, and we propose and discuss different model features capable of overturning this prediction and more closely matching what we find in the data.

The first is to examine the behavior of real marginal cost in the model. From equation (10), one sees that the (log-deviation) of real marginal cost is equal to the log difference between the real wage and technology. Upon arrival of good news about the future, current productivity is unchanged. But the good news is a positive innovation to the lifetime wealth of households, and they therefore demand a higher real wage at any given level of employment. Put differently, the positive wealth effect from good news leads to an inward shift of the labor supply schedule, and there is thus a strong tendency for real wages to rise. Given no immediate change in productivity, higher real wages translate into higher real marginal costs, and thus upward pressure on prices.

One way to overturn the inflationary predictions of the model is thus to add some feature which mitigates the rise in real wages in anticipation of technological improvement. A simple way of doing this is to augment the model with exogenous real wage rigidity. We consider the specification in Blanchard and Gali (2007):

$$w_t - p_t = \delta (w_{t-1} - p_{t-1}) + (1 - \delta)mrs_t \quad (15)$$

Here mrs_t corresponds to the real wage which would obtain on the labor supply curve (given by equation (11) above), and δ is a measure of real wage rigidity. While this specification is obviously somewhat ad hoc, Blanchard and Gali (2007) show that it can be derived from explicit micro foundations. They also argue that the introduction of real wage rigidity improves the fit of the model along a number of other important dimensions.

High values of δ will dampen the extent to which favorable news shocks increase real marginal costs on impact, and thereby reduce the tendency of good news to be inflationary. Figure 8 shows the impulse response of inflation to a news shock for a variety of different values of δ (the response of technology is depicted in Figure 7). The remainder of the model is parameterized as described in Section 2.2. As expected, the impact increase in inflation is

strictly decreasing in the extent of real wage rigidity. For values of δ roughly in excess of 0.5 inflation falls on impact in response to good news. To achieve impact declines in inflation quantitatively similar to what we estimate in the data requires values of δ in excess of 0.9, which seems rather large. Nevertheless, it is clear that some real wage rigidity helps to improve the ability of the model to match the strongly disinflationary nature of news shocks evident in the data.

We next consider the role of monetary policy. Because favorable news shocks make the future output high relative to its present level, the strong tendency is for real interest rates to rise in general equilibrium. Under conventional specifications of interest rate rules along the lines of Taylor (1993), it is extremely difficult to simultaneously generate higher real interest rates and lower inflation. To see this, note the linearized Fisher relationship between real and nominal rates: $r_t = i_t - E_t\pi_{t+1}$. Using the approximation that $i_t \approx i_{t-1}$ and $\pi_t \approx E_t\pi_{t+1}$, one can simplify the policy rule (12) to:¹¹

$$r_t \approx \phi_y \left(y_t - y_t^f \right) + (\phi_\pi - 1) \pi_t \quad (16)$$

Absent monetary policy disturbances, the current real interest rate depends positively on the gap between the actual and flexible price equilibrium level of output and positively on current inflation, assuming that the so-called Taylor principle is satisfied with $\phi_\pi > 1$.¹² In the standard model with a policy rule of this form, movements in the output gap are extremely small. In other words, the Taylor type rule comes very close to restoring the flexible price equilibrium with $y_t \approx y_t^f$. Simplifying further with this approximation, one sees that real interest rates and inflation must, to a first order approximation, commove positively in the absence of policy disturbances.¹³

This discussion suggests that another way to reverse the inflationary predicts of the model is to alter the specification of the monetary policy rule. We entertain what we consider to be two sensible variations on the rule which are capable of better fitting the data. The first

¹¹This approximation is very good for conventional parameterizations of the New Keynesian model. It results from the fact that the nominal interest rate is a state variable for $\rho > 0$, and thus its current value will be close to its lagged value, while inflation is a jump variable, and thus its current value will be close to its expected value next period (for a sufficiently high discount factor).

¹²The actual condition required for determinacy of a rational expectations equilibrium in the New Keynesian model is $\phi_\pi + \frac{1-\beta}{\xi}\phi_y > 1$, where ξ is slope of the Phillips Curve expressed in terms of the output gap. See Woodford (2003) for a full derivation. For values of the discount factor sufficiently close to 1, it is easy to see that the condition for determinacy is still approximately that $\phi_\pi > 1$.

¹³One might wonder how this conclusion is consistent with the results above that real wage rigidity, in the context of the New Keynesian model with a conventional Taylor rule, can simultaneously generate disinflation and higher real interest rates. As stressed by Blanchard and Gali (2007), the presence of real wage rigidity breaks what they term the “divine coincidence”. The fluctuations in the output gap become large with sufficient real wage rigidity, invalidating the approximation that $y_t \approx y_t^f$.

is to suppose that the policy rule reacts not to the output gap, but rather to output growth. Formally:

$$i_t = \rho i_{t-1} + (1 - \rho) (\phi_y (y_t - y_{t-1} - \Delta y^*) + \phi_\pi (\pi_t - \pi^*)) + \varepsilon_{3,t} \quad (17)$$

Rules of this sort in which the central bank reacts to output growth relative to its long term trend as opposed to an output gap have been gaining traction in the literature – for example, see Coibion and Gorodnichenko (2007), Fernandez-Villaverde and Rubio-Ramirez (2007), and Ireland (2004). Orphanides (2003) argues that such a rule fits the data better than the traditional gap specification.

Figure 9 shows theoretical responses of inflation to a news shock from the benchmark model with policy rule (17) for different values of ϕ_y . The impact increase in inflation is decreasing in ϕ_y , and is indeed negative for values of this parameter above a modest cutoff. The intuition for why the growth rate rule can produce disinflation in response to favorable news shocks is straightforward. Output must grow faster than normal for an extended period of time in order to reach its new higher steady state value. Positive output growth exerts upward pressure on nominal (and thus real) interest rates in the policy rule, reducing the need for inflation to rise to generate rising real rates. Put differently, in the growth rate rule the monetary authority follows a policy that is too contractionary relative to the baseline Taylor rule, thereby allowing for the possibility of disinflation following good news shocks.

Our second proposed modification of the policy rule is one in which the monetary authority does respond to an output gap, but that this gap does not correspond to the theoretical gap between the actual and flexible price equilibrium levels (i.e. the “natural rate”) of output. In particular, we propose a rule of the form:

$$i_t = \rho i_{t-1} + (1 - \rho) (\phi_y (y_t - y_t^p) + \phi_\pi (\pi_t - \pi^*)) + \varepsilon_{3,t} \quad (18)$$

$$y_t^p = \alpha y_{t-1}^p + (1 - \alpha) y_t^f \quad (19)$$

Above y_t^p denotes the authority’s perceived natural rate of output. We assume that the current perceived natural rate is a convex combination of the previous period’s perception and the current true natural rate. This specification captures nicely the idea that the monetary authority may react cautiously and therefore sluggishly to the variety of real disturbances reflected in y_t^f . The flexible price equilibrium level of output, y_t^f , is not directly observable, and is indeed a highly complex function of shocks and deep structural parameters. As such, assuming that the central bank responds to some activity measure other than the theoretical

gap seems fairly innocuous.

Figure 10 shows impulse responses of inflation to a news shock from the benchmark parameterization of the model with a policy rule given by (18)-(19) for different values of α . For sufficiently high values of α inflation falls on impact in response to good news. Similarly to the growth rate specification, for high values of α the monetary authority follows too contractionary a policy relative to the standard Taylor rule. In particular, for high degrees of sluggishness, the central bank perceives a large positive output gap for a number of periods into the future and reacts accordingly, when in fact no such gap materializes. This action raises real interest rates more than would happen in a model with flexible prices, thereby choking off aggregate demand and exerting disinflationary pressures. Such a scenario is similar to one explanation for the high inflation of the 1970s – that the US Fed failed to recognize an adverse natural rate shift and therefore followed too loose a monetary policy (Orphanides (2002)).

We next consider the above modifications to the standard New Keynesian model simultaneously. In particular, we estimate several of the parameters of the modified model to investigate whether it is capable of quantitatively matching the estimated empirical response of inflation to a news shock. Our estimation proceeds in two steps. In the first step, we pick the persistence (κ) and standard deviation of the news shock (σ_{ε_2}) to match the estimated empirical response of technology to a news shock. Formally, the estimated parameter vector $\Theta_1 = (\kappa, \sigma_{\varepsilon_2})$ is the solution to the following optimization problem:

$$\Theta_1^* = \arg \min \quad (\mathbf{M}(\Theta_1) - \mathbf{M}^*)' \mathbf{W} (\mathbf{M}(\Theta_1) - \mathbf{M}^*)$$

$\mathbf{M}(\Theta_1)$ is a $(K \times 1)$ stacked vector of the impulse response of technology to a news shock up to horizon K for a particular draw of the parameters. \mathbf{M}^* is the stacked vector of the empirically estimated impulse response of technology to a news shock from our benchmark estimation in Section 3. \mathbf{W} is a diagonal weighting matrix, with elements equal to the inverse of the standard error of the estimated impulse response. We set $K = 20$, fitting the model and estimated impulse responses of technology over a five year horizon. The estimated parameters and standard errors are in the first row of Table 3. Figure 11 shows the model and estimated response of technology to a news shock for these parameter values, along with the empirical confidence bands. The resulting fit is quite good.

In the second step we estimate other parameters of the model to match the estimated empirical response of inflation to a news shock. For the conventional gap specification of monetary policy we estimate the parameter vector $\Theta_2 = (\rho, \phi_y, \phi_\pi, \delta)$; for the misperceptions model of policy we also estimate the parameter governing sluggishness in the perceived

natural rate, $\Theta_3 = (\rho, \phi_y, \phi_\pi, \delta, \alpha)$.¹⁴ The remaining parameters of the model are calibrated as in Section 2.2.

We estimate the parameters in two steps because the inflation impulse response in the model is a function of both κ and σ_{ε_2} – in particular, inflation will in general respond more on impact the less persistent is the news shock.¹⁵ Our goal is to see whether or not the model is capable of matching the inflation response to a news shock *given* the response of technology. If we proceeded in one step, the estimated values of κ and σ_{ε_2} would be chosen not only to match the empirical response of technology to a news shock but also the inflation response. Θ_2 and Θ_3 are otherwise estimated analogously to Θ_1 . In particular, these parameters are chosen to minimize the weighted squared distance between the model and empirical inflation response to a news shock, taking as given the estimated values of κ and σ_{ε_2} from the first stage. As before, the weighting matrix is diagonal with elements equal to the inverse of the estimated standard errors of the inflation impulse response.

The estimated parameters and standard errors are in the second and third rows of Table 3. Figure 12 shows the model and estimated impulse responses of inflation to a news shock using the estimated parameters, assuming a conventional Taylor rule specification. The model does a good job at capturing the dynamic response of inflation to a news shock, though it is unable to fully match the large impact decline. The better-fitting version of the model is that with both real wage rigidity and the misperceptions model of monetary policy. The estimated and model impulse responses are shown in Figure 13. This version of the model produces a slightly better overall fit. The model still has some difficulty fully matching the estimated impact decline in inflation, though the impact effect in the model is within one standard error of the estimated response in the data. Further, the model does a good job at matching the qualitative nature of the dynamics following a news shock.

We have thus far only considered the simple New Keynesian model without capital. For the purposes of elucidating the basic mechanisms at work this simplification is justified.¹⁶ One might nevertheless wonder how our conclusions would differ in a model with endogenous capital accumulation. The addition of capital to the basic model does not significantly alter the effects of news shocks on inflation, nor does it qualitatively impact the effects of the various “fixes” we have proposed. The main role of the presence of capital is to alter the effects of news on the intertemporal allocation of consumption and savings. In the benchmark

¹⁴We do not report estimates for the growth rate specification of monetary policy, as these yield a similar fit with the conventional policy rule augmented with real wage rigidity.

¹⁵The reason for this is evident upon inspection of the Phillips Curve solved forward (14). For a given long run movement in technology, the present discounted value of changes in expected real marginal cost will be larger the sooner most of the productivity improvement occurs.

¹⁶Indeed, Woodford (2003) has argued that the simple model without capital serves as a good approximation to a more elaborate model with sufficient investment adjustment costs.

model with capital favorable news is inflationary, and the variety of alternative specifications we have proposed continue to be capable of making favorable news disinflationary.

5 Conclusion

In this paper we proposed a flexible VAR-based procedure for separately identifying surprise technology shocks and news shocks about future technology. We identify the surprise technology shock as the innovation in a measure of technology and the news shock by applying principal components to the VAR innovations, identifying this shock as the structural shock orthogonal to technology that best explains future variation in technology. We showed through simulation of DSGE models that this approach is likely to perform well in practice, and argued that it represents an important improvement over existing proposed identification strategies found in the literature.

In post-war US data we find that news shocks are responsible for the bulk of low frequencies movements in productivity. In contrast, surprise innovations to technology appear largely transitory. Favorable news shocks are positively correlated with innovations to consumption, stock prices, and consumer confidence, and negatively correlated with inflation innovations. News shocks do a good job at accounting for movements in consumption at all horizons, and for stock prices at lower frequencies. News shocks explain a large share of the forecast error variance of both confidence and inflation at all horizons, and historical decompositions reveal that news shocks do an excellent job at accounting for historical movements in both of these series.

Perhaps the most surprising empirical result is that news shocks are so strongly (negatively) correlated with inflation innovations. While forward-looking models of price-setting suggest that inflation should incorporate news about future productive possibilities, the prediction of the benchmark New Keynesian model is actually for good news to be inflationary, not disinflationary as in the data. We proposed a variety of sensible modifications of the model capable of better fitting the data, and showed that these versions of the model are in fact capable of generating an impulse response of inflation to a news shock that is similar to what we estimate in the data. Though the fit is imperfect, we view the ability of the basic forward-looking model of price-setting to generate disinflation in response to good news about future productivity as something of a success.

References

- [1] Ball, Laurence and Robert Moffitt. “Productivity Growth and the Phillips Curve.” *The Roaring Nineties: Can Full Employment Be Sustained?* Alan Krueger and Robert Solow (editors), 2001.
- [2] Barsky, Robert and Eric Sims. “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence.” Working paper, University of Michigan, 2008.
- [3] Basu, Susanto, John Fernald, and Miles Kimball. “Are Technology Improvements Contractionary?” *American Economic Review* 96: 1418-1448, 2006.
- [4] Beaudry, Paul and Bernd Lucke. “Letting Different Views About Business Cycles Compete.” University of British Columbia working paper, 2009.
- [5] Beaudry, Paul and Franck Portier. “An Exploration Into Pigou’s Theory of Cycles.” *Journal of Monetary Economics* 51: 1183-1216, 2004.
- [6] Beaudry, Paul and Franck Portier. “News, Stock Prices, and Economic Fluctuations.” *American Economic Review* 96: 1293-1307, 2006.
- [7] Beaudry, Paul, Martial Dupaigne, and Franck Portier. “The International Propagation of News Shocks.” University of British Columbia working paper, 2008.
- [8] Blanchard, Olivier, Jean-Paul L’Hullier, and Guido Lorenzoni (2009). “News, Noise, and Fluctuations: An Empirical Exploration.” MIT working paper, 2009.
- [9] Blanchard, Olivier and Jordi Gali. “Real Wage Rigidities and the New Keynesian Model.” *Journal of Money, Credit, and Banking* 39, 35-66, 2007.
- [10] Blanchard, Olivier and Danny Quah. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review* 79: 654-673, 1988.
- [11] Calvo, Guillermo. “Staggered Prices in a Utility Maximizing Framework.” *Journal of Monetary Economics* 12: 383-398, 1983.
- [12] 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.
- [13] Cochrane, John. “Permanent and Transitory Components of GNP and Stock Prices.” *Quarterly Journal of Economics* 109: 241-266, 1994a.

- [14] Cochrane, John. “Shocks.” *Carnegie-Rochester Conference Series on Public Policy* 41: 295-364, 1994b.
- [15] Coibion, Olivier and Yuriy Gorodnichenko. “Strategic Interaction Among Price-Setters in an Estimated DSGE Model.” University of California-Berkeley working paper, 2007.
- [16] Faust, Jon. “The Robustness of Identified VAR Conclusions About Money.” *Carnegie-Rochester Conference Series on Public Policy* 49: 207-244, 1998.
- [17] Faust, Jon and Eric Leeper. “When Do Long Run Identifying Restrictions Give Reliable Results?” *Journal of Business and Economic Statistics* 15: 345-353, 1998.
- [18] Fernandez-Villaverde, Jesus and Juan Rubio-Ramirez. “How Structural are Structural Parameters?” *NBER Macroeconomics Annual* 22: 83-137, 2007.
- [19] Fernandez-Villaverde, Jesus, Juan Rubio-Ramirez, Thomas Sargent, and Mark Watson. “ABCs (and Ds) for Understanding VARs.” *American Economic Review* 97: 1021-1026.
- [20] Francis, Neville, Michael Owyang, and Jennifer Roush. “A Flexible Finite Horizon Identification of Technology Shocks.” Federal Reserve Bank of St. Louis working paper, 2007.
- [21] Gali, Jordi. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review* 89: 249-271, 1999.
- [22] Gali, Jordi. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, 2008.
- [23] Hamilton, James. *Time Series Analysis*. Princeton University Press, 1994.
- [24] Ireland, Peter. “Technology Shocks and the New Keynesian Model.” *Review of Economics and Statistics* 86: 923-936, 2004.
- [25] Jaimovich, Nir and Sergio Rebelo. “Can News About the Future Drive the Business Cycle?” Northwestern University working paper, 2008.
- [26] Kilian, Lutz. “Small Sample Confidence Intervals for Impulse Response Functions.” *Review of Economics and Statistics* 80: 218-230, 1998.
- [27] Kydland, Finn and Edward Prescott. “Time to Build and Economic Fluctuations.” *Econometrica* 50: 1345-1370, 1982.

- [28] Orphanides, Athanasios. “Historical Monetary Policy Analysis and the Taylor Rule.” *Journal of Monetary Economics* 50: 983-1022, 2003.
- [29] Orphanides, Athanasios. “Monetary Policy Rules and the Great Inflation.” *American Economic Review* 92: 115-120, 2002.
- [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, Eric. “Expectations Driven Business Cycles: An Empirical Evaluation.” Working paper, University of Michigan, 2009.
- [33] Taylor, John. “Discretion vs. Policy Rules in Practice.” *Carnegie-Rochester Series on Public Policy* 39: 195-214, 1993.
- [34] Uhlig, Harald. “What Drives GNP?” EABCN working paper, 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.
- [37] Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, 2003.

Table 1
 Fraction of Forecast Error Variance Explained by News Shock

	$h = 0$	$h = 4$	$h = 8$	$h = 16$	$h = 24$	$h = 40$
Tech.	0.0	1.9	4.4	14.8	28.3	51.4
	[0.0,0.0]	[0.7,7.2]	[1.3,18.0]	[6.2,37.0]	[19.1,49.0]	[41.0,65.5]
Stock Price	7.1	18.0	23.5	33.9	37.7	41.5
	[1.3,29.7]	[4.3,41.3]	[6.4,47.0]	[10.2,57.8]	[12.0,62.9]	[12.3,68.1]
Consumption	21.7	57.7	81.4	91.7	91.6	87.3
	[4.8,35.0]	[26.9,68.2]	[49.4,85.3]	[64.3,92.6]	[66.1,94.0]	[59.8,93.3]
Inflation	63.9	53.6	55.1	46.7	43.8	43.3
	[28.1,73.3]	[28.7,58.6]	[29.4,59.9]	[27.3,55.2]	[25.6,53.8]	[25.0,53.8]
Confidence	39.3	57.1	66.7	62.1	57.3	54.6
	[15.0,49.2]	[27.2,65.1]	[35.6,72.2]	[32.9,69.6]	[29.8,65.9]	[27.6,64.9]
Interest Rate	18.5	11.7	8.7	10.7	13.5	18.6
	[2.9,33.6]	[3.1,31.0]	[3.3,29.7]	[6.8,29.8]	[11.1,30.8]	[13.6,37.7]

The numbers in brackets are the 68 percent bias-corrected bootstrap confidence intervals.

Table 2
 Long Horizon Regressions

$$a_{t+k} - a_t = \alpha + \sum_{i=1}^N \beta_i x_{i,t} + e_t$$

Horizon	Adjusted R^2
$k = 1$	0.034
$k = 4$	0.235
$k = 8$	0.357
$k = 16$	0.491
$k = 40$	0.512

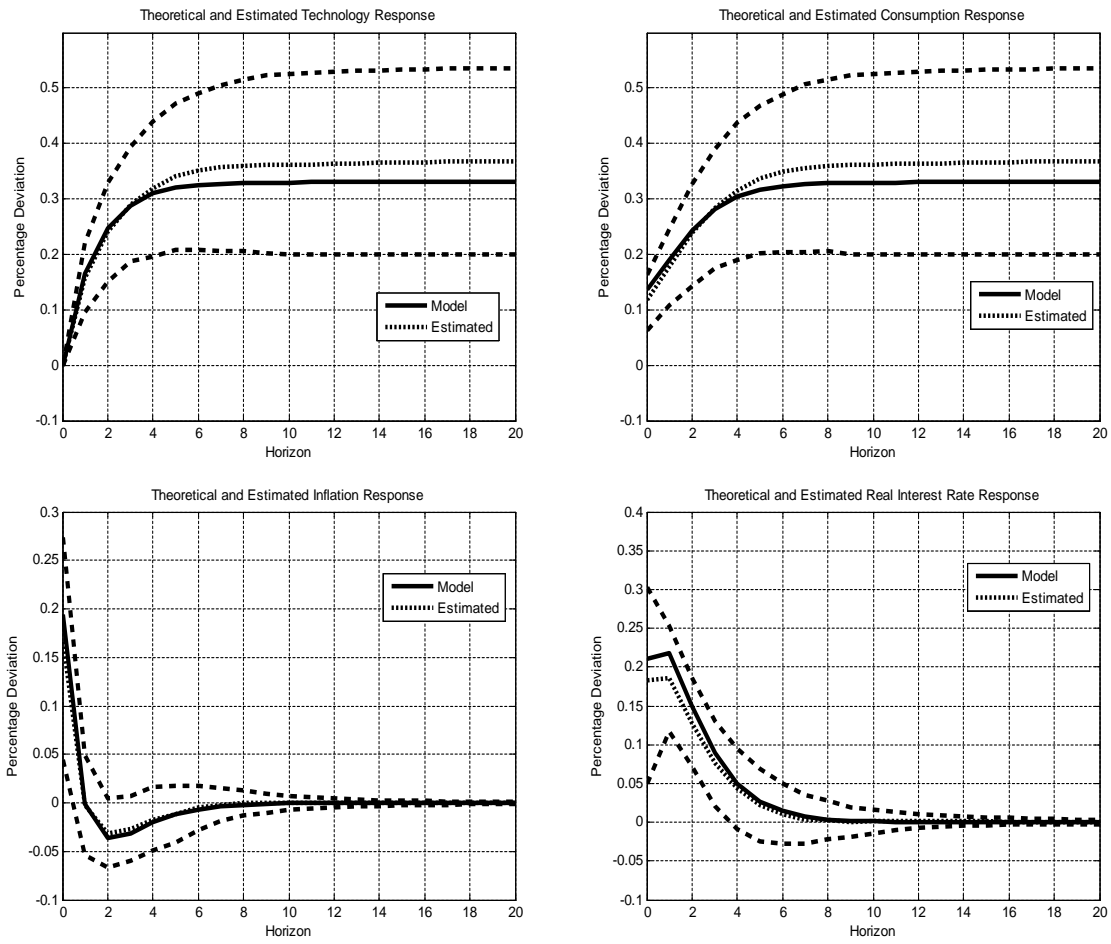
These are results from long horizon regressions of technology growth on the current levels of consumption, stock prices, consumer confidence, inflation, and the interest rate.

Table 3
Parameter Estimates

$\widehat{\Theta}_1$	κ	σ_2			
	0.89 (0.18) [0.66,0.98]	0.0035 (0.0036) [0.0018,0.0010]			
$\widehat{\Theta}_2$	ρ	ϕ_y	ϕ_π	δ	
	0.97 (0.18) [0.59,0.99]	1.24 (0.30) [1.08,1.51]	1.80 (0.24) [1.25,1.87]	0.91 (0.09) [0.87,0.94]	
$\widehat{\Theta}_3$	ρ	ϕ_y	ϕ_π	δ	α
	0.97 (0.11) [0.83,0.99]	1.49 (0.46) [0.92,1.96]	1.61 (0.29) [1.25,2.01]	0.70 (0.25) [0.10,0.79]	0.82 (0.15) [0.71,0.98]

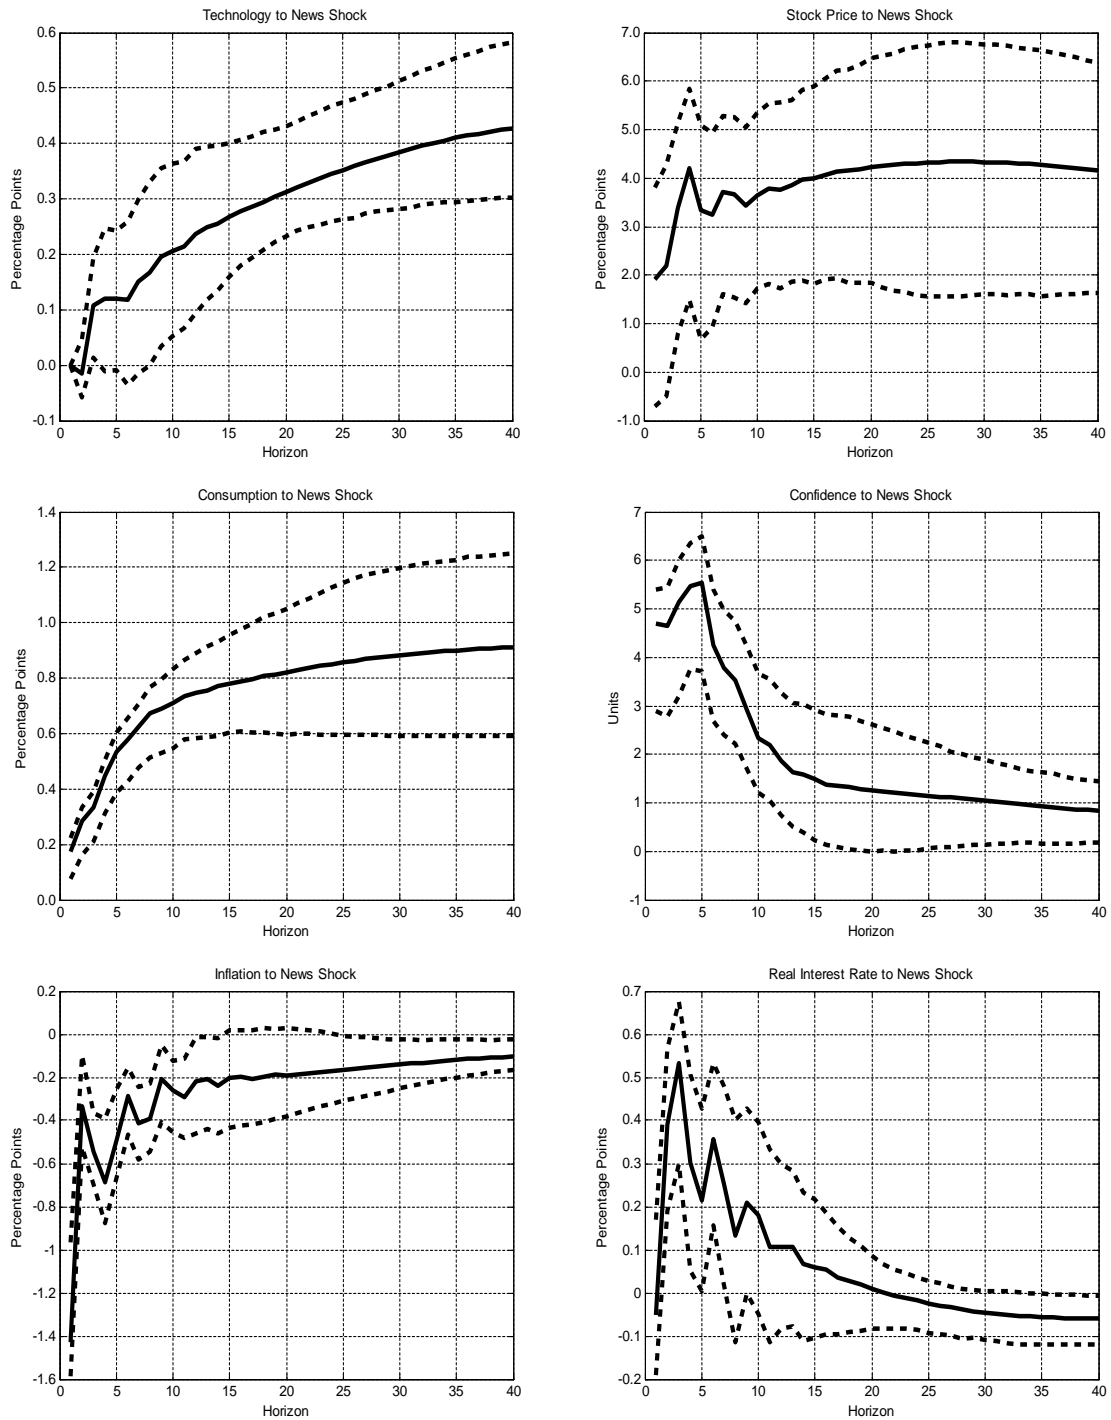
This table presents parameter estimates from the estimation of Section 2.4. The estimates in the $\widehat{\Theta}_1$ row are from the first stage estimates of the autoregressive process for technology growth. The estimates in the $\widehat{\Theta}_2$ row are for other parameters of the baseline model with a standard Taylor rule and real wage stickiness. The estimates in the $\widehat{\Theta}_3$ row are for the model with both real wage stickiness and the misperceived output gap Taylor rule. The bootstrap standard errors are in parentheses, and the numbers in brackets are the one standard error bootstrap confidence bands.

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



The black lines show the theoretical responses to a news shock from the model of Section 2.2. The solid blue line depicts the estimated responses averaged over the simulations, with the dashed blue lines showing the 10th and 90th percentiles of the distribution of estimated impulse responses.

Figure 2
 Estimated Empirical Impulse Responses to a News Shock



The dashed lines represent the 68 percent bias-corrected bootstrap confidence bands.

Figure 3
Impulse Responses to Surprise Technology Shock




Figure 4

Impulse Responses to Stock Price Innovation Orthogonalized with Respect to Technology

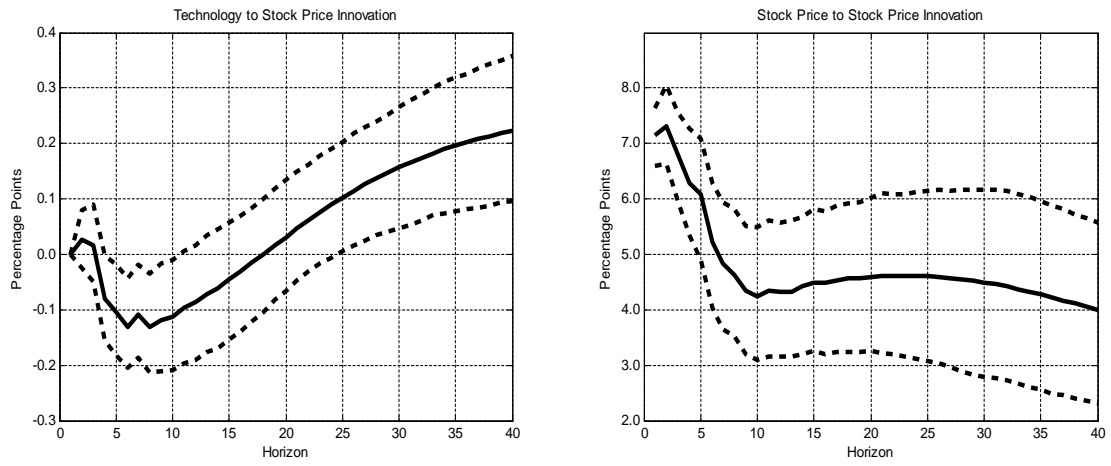


Figure 5
Historical Simulations






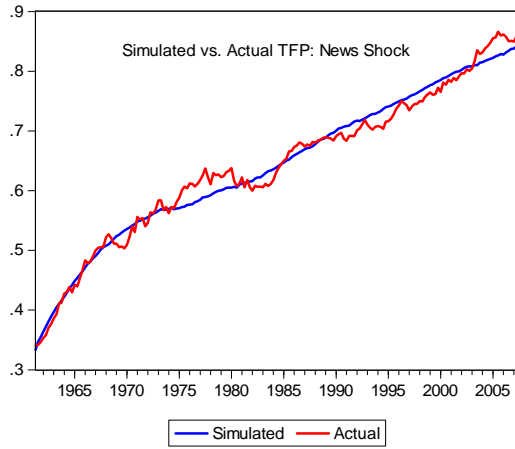
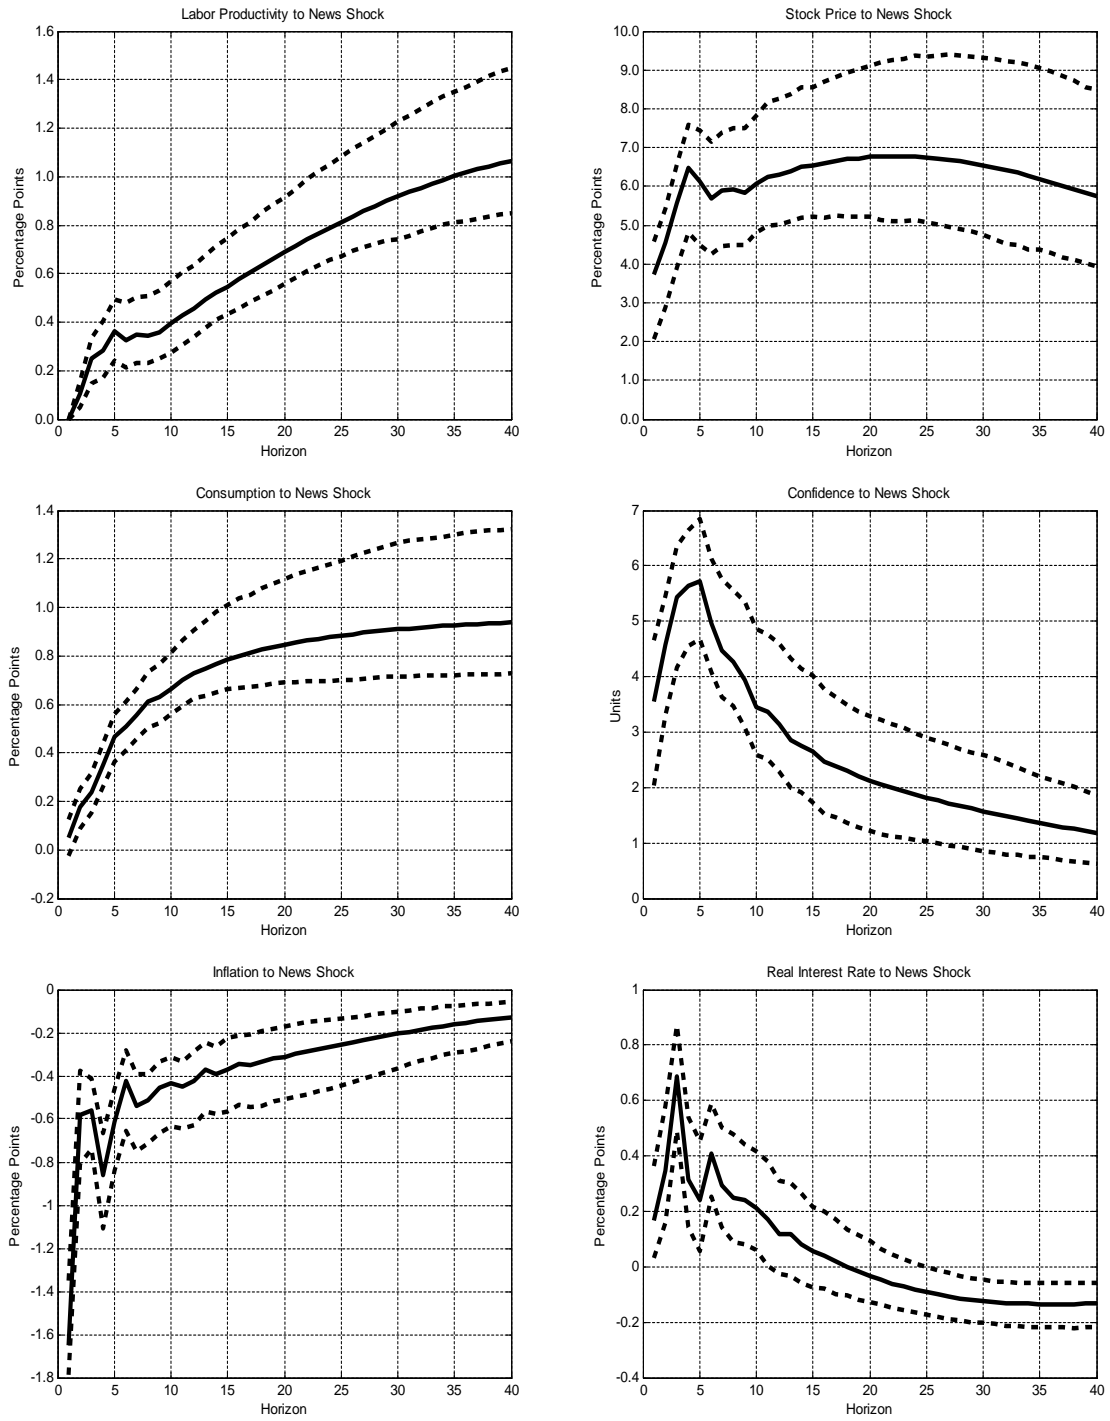


Figure 6
 Estimated Empirical Impulse Responses to a News Shock
 System with Average Labor Productivity



The dashed lines represent the 68 percent bias-corrected bootstrap confidence bands.

Figure 7
New Keynesian Model Responses to News Shock




Figure 8
New Keynesian Model Inflation Response to News Shock
Real Wage Rigidity




Figure 9
 New Keynesian Model Inflation Response to News Shock
 Growth Rate Policy Rule




Figure 10
 New Keynesian Model Inflation Response to News Shock
 Incorrect Output Gap Rule

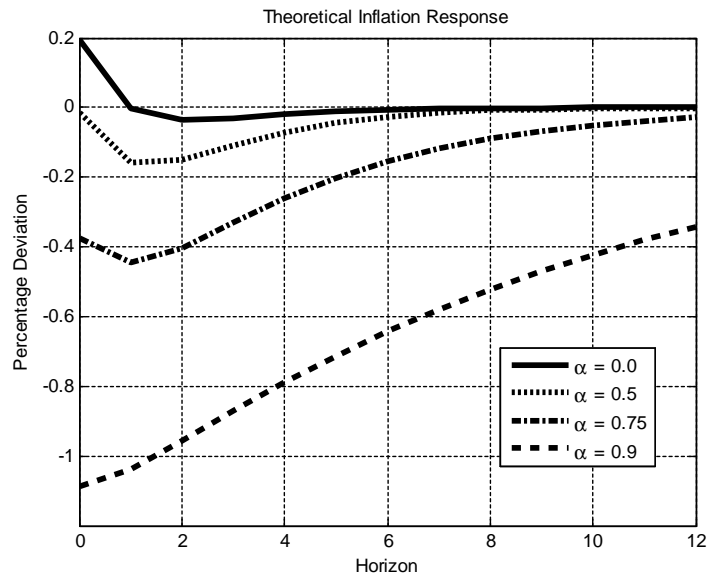


Figure 11
Technology Response to News Shock

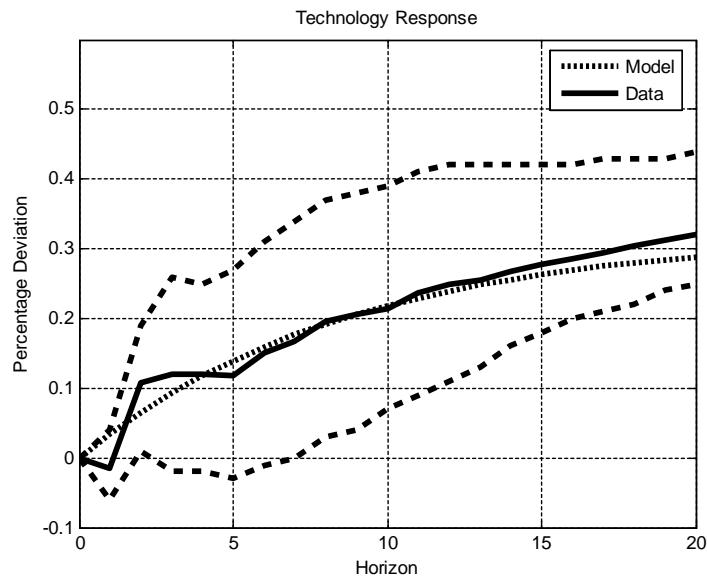


Figure 12
Inflation Response: Optimal Parameter Values
Sticky Real Wages, Conventional Taylor Rule




Figure 13
Inflation Response: Optimal Parameter Values
Sticky Real Wages, Misperception Taylor Rule

