Expectations and the Business Cycle

by

Eric R. Sims

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2009

Doctoral Committee:

Professor Robert B. Barsky, Chair
Professor Lutz Kilian
Professor Hasan Nejat Seyhun
Professor Matthew D. Shapiro

Dedication

This dissertation is dedicated to my family – my parents Patty and Grant Sims, my brothers Ryan and Scott Sims, my grandfathers Jimmie Sims and Oscar Stewart, and my wife Jill.

Acknowledgements

I thank my committee members – Bob Barsky, Lutz Kilian, Nejat Seyhun, and Mattew Shapiro – for their support and guidance. I am particularly grateful to my chair, Bob Barsky, for several years of collaboration and friendship. Chapters II and IV of this dissertation are coauthored with him. His enthusiasm for economics is infectious. I consider him to be a valuable friend and mentor, and look forward to working with him more in the future. Special thanks are also due to Lutz Kilian, without whose countless hours of technical assistance and emotional support I would not be where I am today. Matthew Shapiro's extensive comments have greatly improved the quality of my work. I also acknowledge several other current and former members the Michigan economics faculty for their comments and assistance, including Rudi Bachmann, Susanto Basu, Chris House, Miles Kimball, and Serena Ng. In addition, former Michigan students Oli Coibion and Yuriy Gorodnichenko played valuable roles in my development as an economist.

I would not have reached this stage without the enduring support of my two closest friends in Michigan, Adam Cole and Daniel Cooper. Adam was my constant study companion throughout our first year in graduate school, and I am certain that I would not have passed my preliminary exams without his help. His piercing wit, encyclopedic knowledge, and ability to listen are greatly appreciated. Daniel has been my email and golf buddy for the last four years. He is always available to listen when I am down, and has provided valuable comments on my work. I hope that we have the opportunity to work together in the future.

Finally, I gratefully acknowledge the support of my immediate family. My parents, Patty and Grant, have set a wonderful example and have always provided me with encouragement. My two brothers, Ryan and Scott, constantly bring a smile to my face, providing me with welcome diversions from the every day drudgeries of research. Lastly, I thank my wife Jill for her unconditional love and dedication. Her positive outlook on life is a true inspiration.

Table of Contents

List of Figures

List of Tables

Chapter I

Introduction

Economists have long recognized the central role of expectations in the study of short run fluctuations. There has recently been a renewal of interest in the implications of expectations about changes in future fundamentals for the business cycle. This dissertation is broadly concerned with the identification and study of these so-called "news shocks" about future technological change. It proposes and implements a new approach for the empirical identification of these shocks. News shocks turn out to have important implications for a variety of forward-looking variables, such as aggregate consumption, stock prices, consumer confidence, and inflation. The effects of news shocks on these variables, as well as on measured technology itself, shed light on a variety of macroeconomic phenomena, including permanent versus transitory components of productivity, the interpretation of surprise movements in measured consumer confidence, the role for forward-looking models of price-setting, and the specification of the systematic component of monetary policy. Nevertheless, news shocks do not appear to be an important source of business cycles, inducing conditional comovement among aggregate variables at odds with the unconditional correlations in the data.

Chapters II and III propose and implement a new approach for the identification of news shocks. In the context of a vector autoregression (VAR) featuring a utilization-adjusted measure of total factor productivity (hereafter "technology") and several forward-looking variables, the news shock is identified as the structural shock orthogonal to technology innovations which 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 innovations which maximizes the sum of contributions to technology's forecast error variance over a finite horizon. Application of this empirical strategy to artificial data generated from a popular dynamic stochastic general equilibrium (DSGE) model confirms that it is likely to perform well in practice.

In US data news shocks account for the bulk of low frequency variation in productivity; in contrast surprise innovations in measured technology are quite transitory. Favorable news shocks are positively correlated with consumption, stock price, and consumer confidence innovations, and negatively correlated with inflation innovations. Consistent with general equilibrium implications of most models, they are also associated with higher real interest rates. News shocks 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.

Perhaps the most surprising result is the extent to which inflation innovations convey information about future productivity growth. This finding is potentially consistent with forward-looking models of price-setting. The prediction of the benchmark New Keynesian model with staggered price-setting and augmented with a Taylor rule, however, is actually for good news to be inflationary on impact, not disinflationary as found in the data. Chapter II diagnoses the reasons for this counterfactual prediction of the model and proposes various modifications capable of making it better fit the data. In particular, sensible variations on the Taylor rule – especially ones in which the monetary authority responds to an activity measure different from the theoretical output gap – are capable of generating disinflation in response to favorable news that is similar to what is estimated in the data.

Chapter III uses the same empirical strategy of Chapter II to study the business cycle implications of news shocks. It presents independent simulation evidence from a different DSGE model that the proposed empirical strategy is likely to reliably identify news shocks in practice. It also addresses the consequences of news shocks for VAR invertibility, and argues that any non-invertibilities resulting from the presence of news shocks are likely of limited practical importance. In post-war US data, a favorable news shock is associated with an increase in consumption and declines in output, hours of work, and investment on impact. After the impact effects, aggregate variables largely track predicted movements in technology. These findings are broadly consistent the theoretical predictions of a variety of standard DSGE models 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 the news shocks fail to account for output declines in four of the six most recent US recessions. These results suggest that news shocks about future productivity are not a dominant source of business cycles.

Chapter IV studies surprise movements in consumer confidence. Innovations to a variety of measures of confidence convey incremental information about economic activity far into the future. Motivated by economic theory, Chapter IV distinguishes between two competing hypotheses concerning the economic meaning of confidence. The first is that confidence innovations reflect shifts in sentiment unrelated to economic fundamentals; in particular, these "animal spirits" shocks represent overly optimistic or pessimistic expectations on the part of households. The alternative hypothesis is that confidence innovations reveal information about current and future economic fundamentals. In a calibrated New Keynesian model with shocks to current and expected productivity, the animal spirits shock behaves as a demand shock – a positive innovation leads to transitory increases in spending and inflation. In contrast, shocks to current and expected fundamentals are associated with permanent responses of real activity and are disinflationary given realistic specifications of monetary policy.

Which of these hypotheses is the better characterization of the data can be tested by comparing the conditional relationships between confidence and aggregate variables in the data with the predictions of the model. There is little apparent support for the animal spirits hypothesis, though it is not possible to completely rule out that animal spirits-like shocks do manifest themselves in measures of confidence. Surprise movements in confidence are associated with apparently permanent movements in real activity and are disinflationary. The implications of confidence innovations for aggregate variables are small at high frequencies, suggesting that the conditional relationships between confidence and aggregate variables largely reflect information about future economic prospects. These findings are broadly consistent with the results in both Chapters II and III, which simultaneously reveal that news shocks about future productivity are positively correlated with confidence innovations and negatively correlated with inflation innovations, but are not responsible for large high frequency fluctuations in economic activity.

Chapter V studies the role of Taylor-type nominal interest rate rules in the New Keynesian model of forward-looking price-setting. While not explicitly concerned with news shocks about future productivity, the interplay between nominal rigidity and the systematic component of monetary policy is central to understanding the disinflationary nature of news shocks documented elsewhere in the dissertation. The Taylor principle – roughly that central banks should raise real interest rates in response to increases in inflation – works quite differently in New Keynesian models relative to the "old" Keynesian models in which it was originally espoused. Whereas in the "old" models satisfaction of the Taylor principle is necessary to prevent nominal explosions, in the new models satisfaction of the Taylor principle is necessary to generate a unique equilibrium, conditional on ruling out nominal explosions *a priori*. This subtle difference has potentially important implications for both identification and estimation of the parameters of the Taylor rule.

This chapter first documents that, even in the New Keynesian model in which inflation is a forward-looking jump variable, the parameters of nominal interest rate rules are, in general, identified in both the region of determinacy and indeterminacy. Monte Carlo evidence documents that consistent estimates of these parameters may be obtained via instrumental variable regressions, though the finite sample properties of the estimators tend to be poor, especially for highly persistent real shocks. Identification of policy parameters fails when monetary policy is able to completely implement the flexible price equilibrium. The central bank may implement such a policy by augmenting its linear policy rule with a "stochastic intercept" which tracks changes in the Wicksellian natural rate of interest. The presence of a stochastic intercept of this nature leads to non-identification because it results in the set of valid instruments being null.

A testable implication of the stochastic intercept rule is that the conditional correlations between non-policy shocks and inflation should be zero. As documented in Chapter II, this testable implication clearly fails in the data. Shocks which permanently affect output are highly disinflationary on impact. This suggests that the stochastic intercept rule is a poor characterization of monetary policy. Shocks permanently affecting output are also associated with large predictable increases in spending in the data; predictable increases in spending necessitate higher real interest rates in most models (this is also a feature of the data, as documented in Chapter II). Conventional specifications of Taylor rules (without stochastic intercepts) have difficulty simultaneously generating disinflation and higher real interest rates, precisely owing to the Taylor principle which calls for lower real interest rates in responses to lower inflation. This finding indicates that modifications of conventional Taylor rules along the lines suggested in Chapter II are necessary in order to better fit the data.

Chapter VI offers concluding thoughts and ties the various themes of the dissertation together. It also speculates on possible avenues for future related research.

Chapter II

News Shocks

1 Introduction

Macroeconomists have devoted significant effort to the identification and study of technology shocks. The most commonly used empirical approach is the structural vector autoregression (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 Chapter and in Chapter III.

These so-called news shocks have attracted growing interest from macroeconomists in recent years (Cochrane (1994b), Beaudry and Portier (2006), and Chapter IV of this dissertation). 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 which 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, as well as 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 forward-looking variables; Chapter III of this dissertation 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 Chapter IV reaches 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 many 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. This finding fits nicely with the idea that the 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. 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 in response to favorable news shocks. 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 much 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 Chapter IV, 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 simple 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]

[1] More 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.

The following is an example process satisfying this assumption:

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

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

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 $\overline{g}$ 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 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. Without loss of generality, let an empirical measure of aggregate technology occupy the first position in $\mathbf{y}_t$. 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:

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

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

$$\mathbf{u}_t = \mathbf{A}_0 \boldsymbol{\varepsilon}_t \tag{4}$$

This implies the following structural moving average representation:

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

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, $\widetilde{\mathbf{A}}_0$ (e.g. a Choleski decomposition), the entire space of permissible impact matrices can be written as $\widetilde{\mathbf{A}}_0\mathbf{D}$, where $\mathbf{D}$ is a $k \times k$ orthonormal matrix ($\mathbf{D}\mathbf{D}' = \mathbf{I}$).

The h step ahead forecast error is:

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

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

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

The $\mathbf{e}_i$ denote selection vectors with one in the ith place and zeros elsewhere. The selection vectors inside the parentheses in the numerator pick out the jth column of $\mathbf{D}$, which we will denote by $\boldsymbol{\gamma}$. $\widetilde{\mathbf{A}}_0\boldsymbol{\gamma}$ is then a $k \times 1$ vector corresponding with the jth column of a possible orthogonalization. The selection vectors outside the parentheses in both numerator and denominator pick out the ith 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:

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

s.t.

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

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, $\widetilde{\mathbf{A}}_0$. 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 $\left(\mathbf{B}_{1,\tau}\widetilde{\mathbf{A}}_0\right)'\left(\mathbf{B}_{1,\tau}\widetilde{\mathbf{A}}_0\right)$ 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.

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 which leaves a large share of the variance of technology unexplained. As shown in Chapter III, 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 auxilliary assumptions about the other shocks. Third, we impose the stronger implication that news and surprise technology shocks account for variation in technology at all horizons, not just in the long run. As such, we explicitly 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 relative to what would obtain with a long run restriction. 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:

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

$$c_t = y_t \tag{7}$$

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

$$y_t = a_t + n_t \tag{9}$$

$$mc_t = w_t - p_t - a_t \tag{10}$$

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

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

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

These are the standard equations for the canonical New Keynesian model – see Woodford (2003) or Gali (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 $\overline{g} = 0.0025$. Technology (and thus output) grow at the annualized rate of 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 2.1 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 the model responses at all horizons. Similarly

[2] In 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. Similar results obtain under different assumptions about cointegration.

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. We should note that the small amount of spuriousness is not endimic to our approach. In particular, the combined long run/recursive restriction produces a similarly spurious shock when the data are generated without news shocks. When the data are generated with news shocks, our approach performs better than the long run identifying strategy (in the sense of lower mean-squared error in the impulse responses and a higher correlation between identified and true shocks).

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 Chapter III conducts simulation exercises for a similar empirical procedure on data generated from a real model with capital and reports similarly good simulation results. Chapter III 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. In practice, 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 or eliminates the impact of potential non-invertibilities. 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 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 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.[3] 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, obtained from Robert Shiller's website. 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.[4] For more on

[3] This series was constructed and given to us directly by John Fernald.

[4] The specific survey 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.

the confidence data, see Chapter IV.

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. Our results are also robust to a variety of different assumptions concerning the nature of cointegration. As suggested by a variety of information criteria, we estimate the system with four lags. In terms of the identification strategy outlined in the previous section, we set the truncation horizon at $H = 60$. In words then, the news shock is identified as the structural shock orthogonal to technology innovations which best explains technology movements over a fifteen year horizon.

Figure 2.2 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.[5] 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, quickly levelling off to a new permanently higher steady state. While the sharp predictable increase in stock prices following impact is consistent with the general equilibrium implications of higher real interest rates that we find in the data, it is nonetheless not possible to reject the hypothesis that the impulse response is a random walk. Consumer confidence rises strongly and significantly on impact in

[5] The 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.

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 Chapter IV 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 by far the largest response on impact.

Table 2.1 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 2.3 shows other impulse responses of interest from our benchmark estimation. 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.

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.[6]

Table 2.2 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 productivity growth over long horizons is predictable and that the low frequency component of productivity is largely unrelated to technology innovations are consistent with the conclusions in Rotemberg (2003).

The remaining two responses in Figure 2.3 depict 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, stock prices rise on impact and then revert, 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 2.4 we show several historical simulations from our benchmark system. The upper two figures plot the actual and simulated values of technology, with the

[6]The responses and variance decompositions of the other variables in the VAR to the surprise technology shock are small and are therefore ommitted. The surprise technology shock is associated with a small and transitory increase in consumption, a decline in real interest rates, and little significant movement in stock prices, inflation, or consumer confidence. For more on the business cycle relevance of surprise technology shocks, see Chapter III.

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 2.4 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 2.5 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.[7] 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.[8] News shocks account for a larger share of the innova-

[7] The 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 TFP. Nevertheless, Ball and Moffitt (2001) have argued that average labor productivity is a more exogenous measure of true technology than is total factor productivity.

[8] In a model with capital accumulation, it is to be expected that average labor productivity would

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

While some small quantitative discrepancies do exist, our qualitative results are very 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, different specifications of the truncation horizon in the optimization problem underlying identification, and different assumptions concerning cointegration (levels vs. VECM, etc.). 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, and positively correlated with consumption and stock price innovations.

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

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 2.5 are roughly consistent with this prediction.

more surprising, but is consistent with the evidence in Chapter IV. 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:

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

$(1-\theta)$ 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 2.6 replicates the theoretical responses of technology and inflation to a favorable news shock, using the calibration of the model described in Section 2.2. 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):

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

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 2.7 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 2.6). 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 New Keynesian 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 plentiful relative to the present, 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:[9]

[9] This 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).

$$r_t \approx \phi_y \left(y_t - y_t^f \right) + (\phi_\pi - 1)\, \pi_t \tag{16}$$

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 $\phi_\pi > 1$.[10] 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.[11]

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

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

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 2.8 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

[10] The 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 $\phi_\pi > 1$.

[11] One 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 $y_t \approx y_t^f$.

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:

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

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

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 2.9 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:

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

$\mathbf{M}(\Theta_1)$ 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 2.3 Figure 2.10 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 $\Theta_2 = \left(\rho,\, \phi_y,\; \phi_\pi,\, \delta\right)$; for the misperceptions model of policy we also estimate the parameter governing sluggishness in the perceived natural rate, $\Theta_3 = \left(\rho,\, \phi_y,\; \phi_\pi,\, \delta,\; \alpha\right)$.[12] 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.[13] 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

[12] We 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.

[13] The 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.

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 2.3. The estimated policy parameters governing monetary policy are in line with existing estimates, and the estimated values of real wage rigidity or sluggishness in the updating of the natural rate seem reasonable. Figure 2.11 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 2.12. 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.[14] One might nevertheless wonder how our conclusions would differ in a model with endogenous capital accumulation. In the extreme version without adjustment costs and perfect rental markets real wage rigidity ceases to be capable of alone generating disinflation in response to a good news shock. The reason for this is that real marginal costs tend to fall on impact in the model with capital following favorable news. The increase in inflation on impact is the result of higher marginal costs in the future, not high marginal costs today. Mitigating the effects of news shocks on marginal costs through real wage rigidity thus has only limited effects on the immediate response of inflation. Modifications of the policy rule continue to be capable of generating disinflation in response to good news, and a combination of several of these features remains able to roughly match the disinflationary nature of news shocks evident in the data.

[14] Indeed, 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.

5 Conclusion

In this paper we proposed a flexible VAR-based procedure for separately identifying surprise and news shocks about aggregate 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 which 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.

Table 2.1
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 2.2
Long Horizon Regressions

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

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 2.3

Parameter Estimates

$\widehat{\mathbf{\Theta}}_1$	κ	σ_2			
	0.89 (0.18) [0.66.0.98]	0.0035 (0.0036) [0.0018,0.0010]			
$\widehat{\mathbf{\Theta}}_2$	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]	
$\widehat{\mathbf{\Theta}}_3$	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 $\widehat{\mathbf{\Theta}}_1$ row are from the first stage estimates of the autoregressive process for technology growth. The estimates in the $\widehat{\mathbf{\Theta}}_2$ row are for other parameters of the baseline model with a standard Taylor rule and real wage stickiness. The estimates in the $\widehat{\mathbf{\Theta}}_3$ 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 2.1
Model and Monte Carlo Estimated Impulse Responses to News Shocks

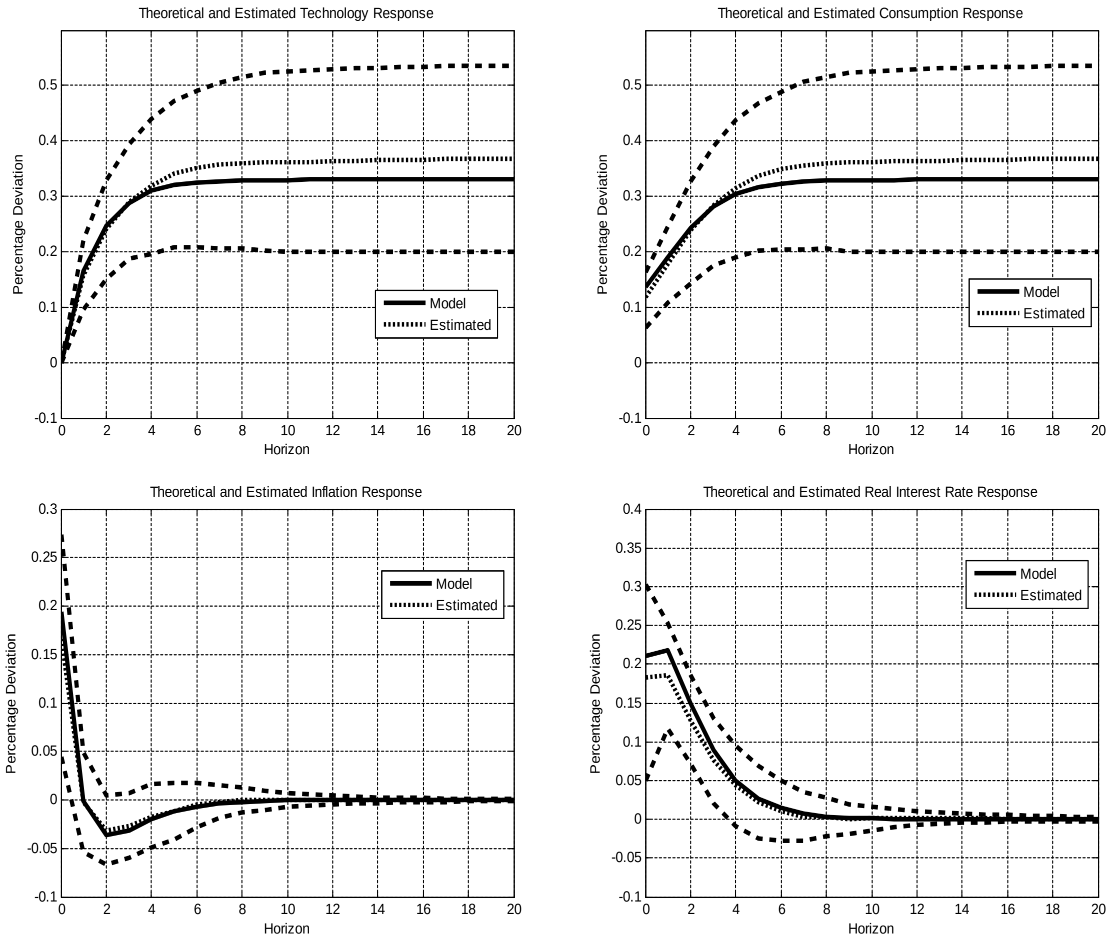

The black lines show the theoretical responses to a news shock from the model of Section 2.2. The solid blue line depicts the estimated responses averaged over the simulations, with the dashed blue lines showing the 10th and 90th percentiles of the distribution of estimated impulse responses.

Figure 2.2
Estimated Empirical Impulse Responses to a News Shock

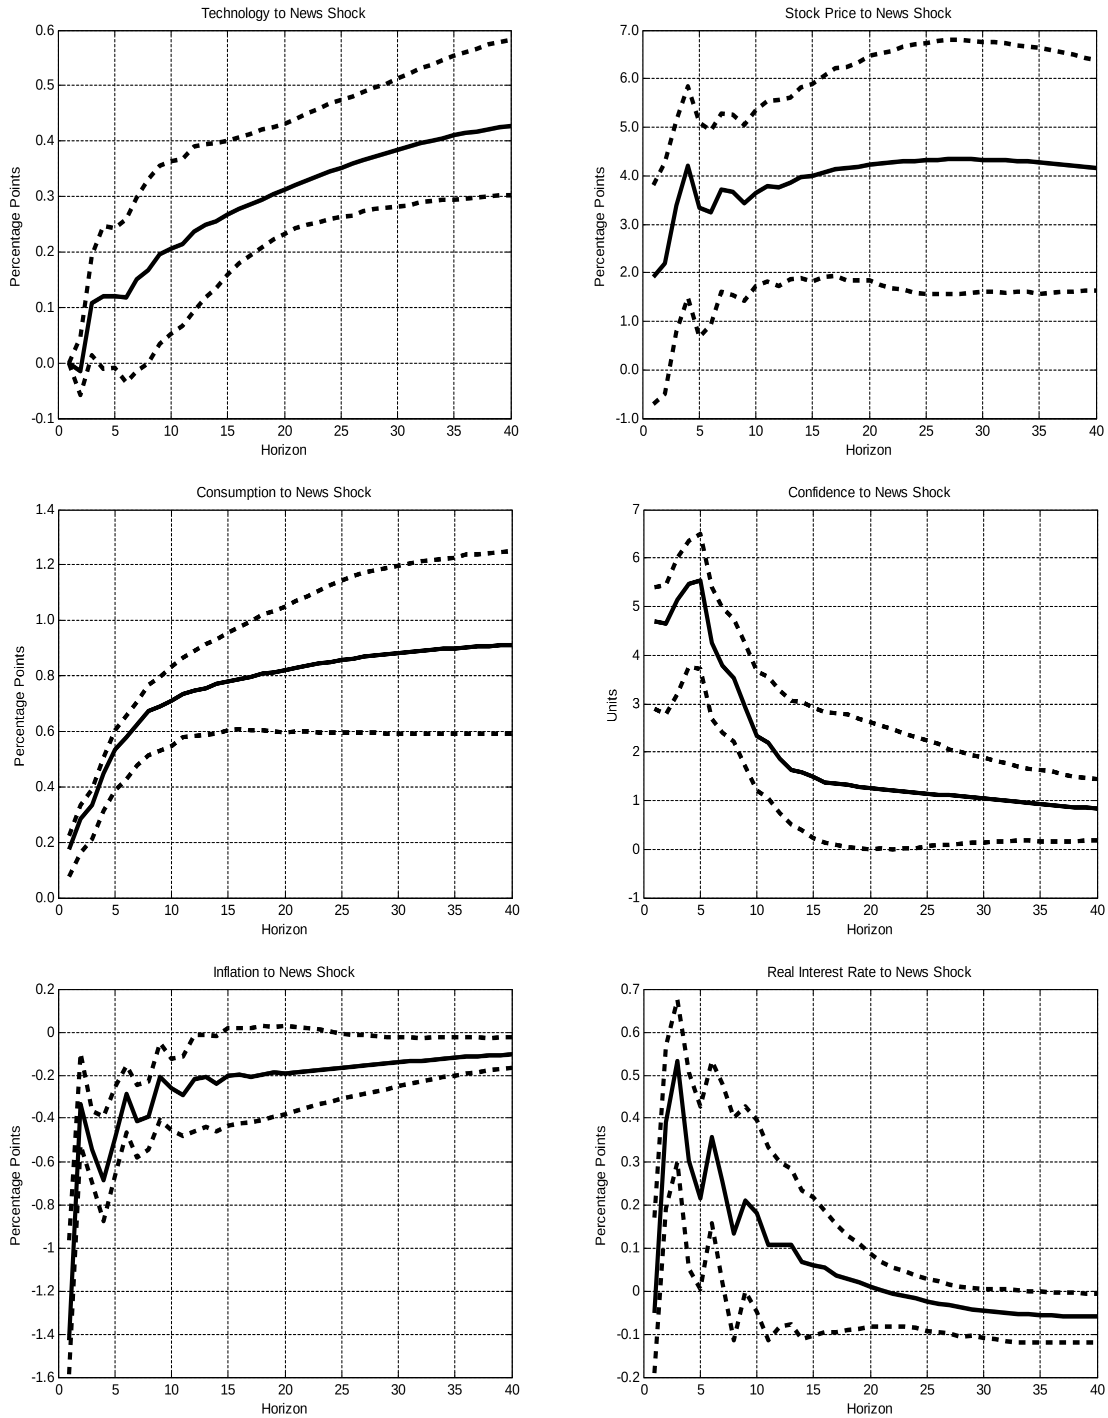

The dashed lines represent the 68 percent bias-corrected bootstrap confidence bands.

Figure 2.3
Other Estimated Empirical Impulse Responses

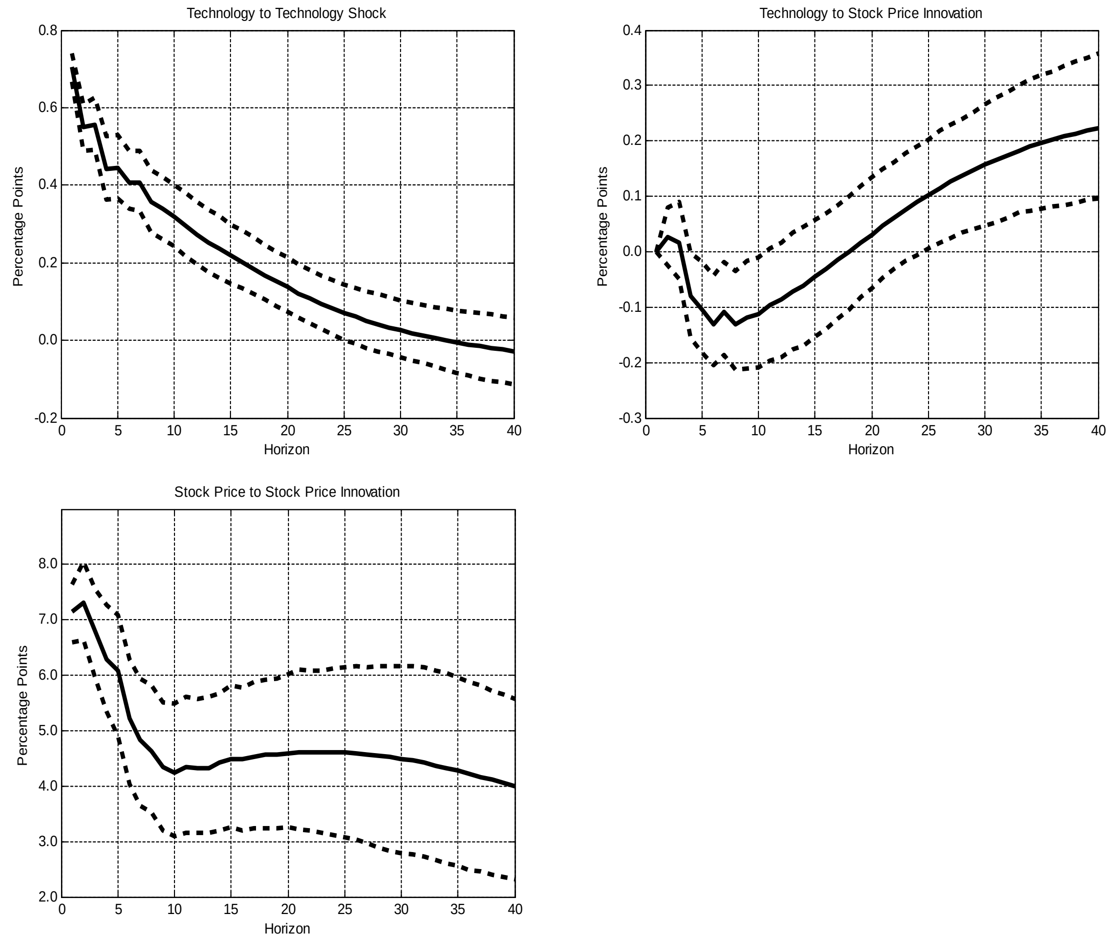

The dashed lines represent the 68 percent bias-corrected bootstrap confidence bands.

Figure 2.4
Historical Simulations

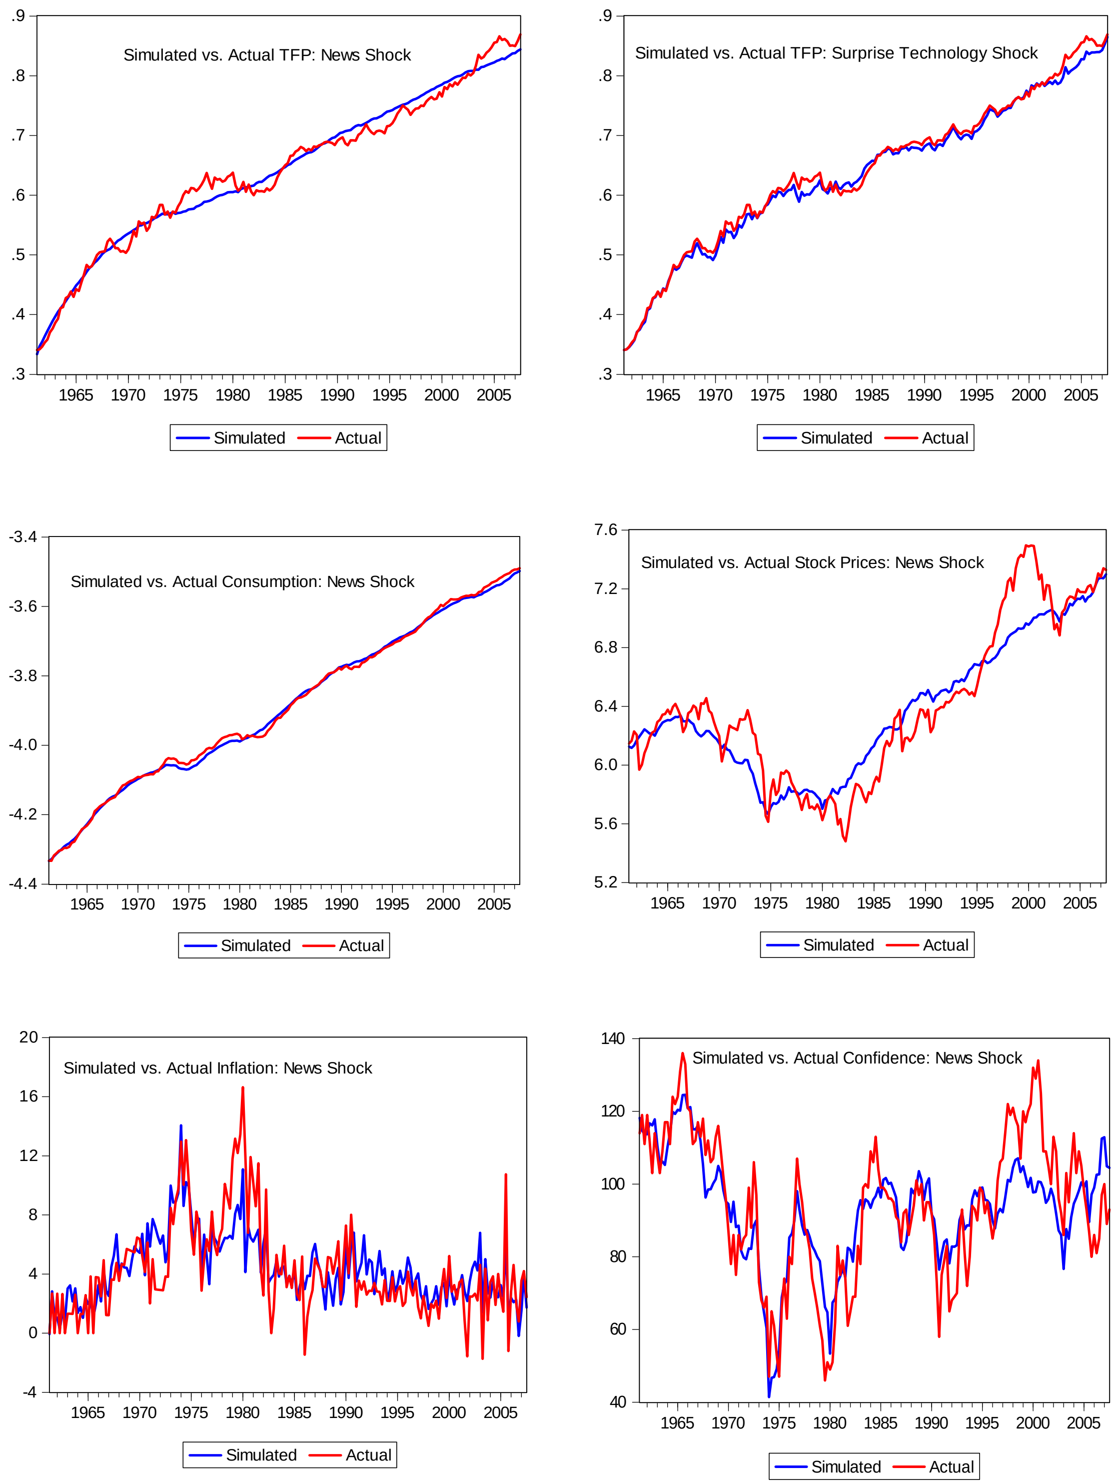

Figure 2.5
Estimated Empirical Impulse Responses to a News Shock
System with Average Labor Productivity

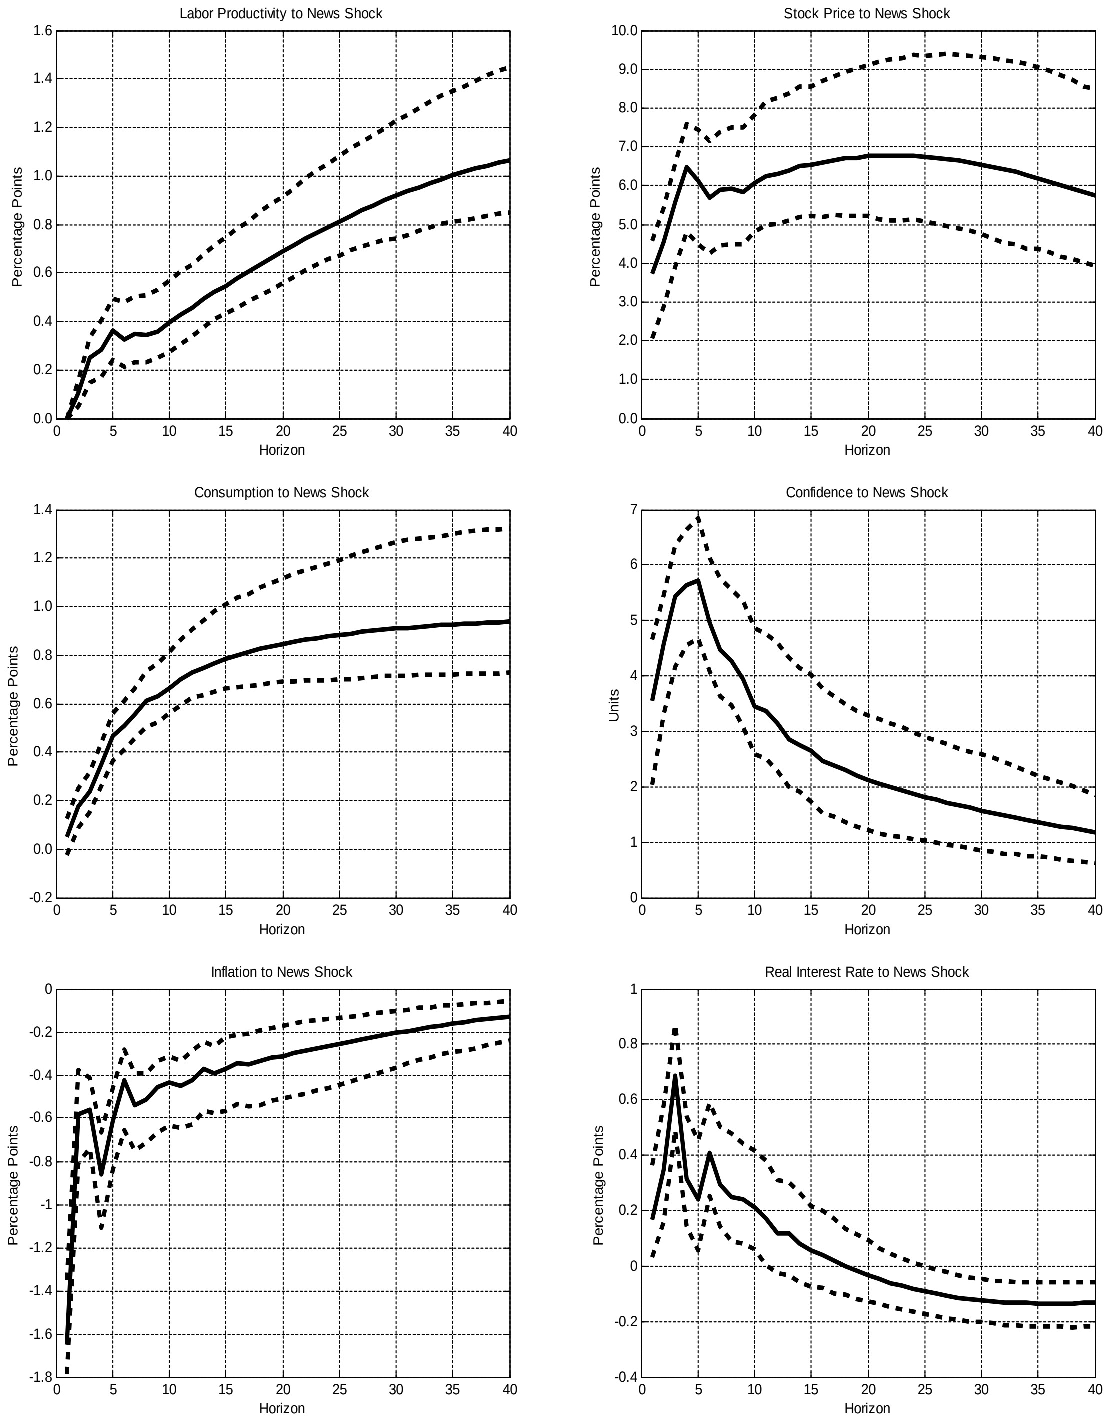

The dashed lines represent the 68 percent bias-corrected bootstrap confidence bands.

Figure 2.6
New Keynesian Model Responses to News Shock

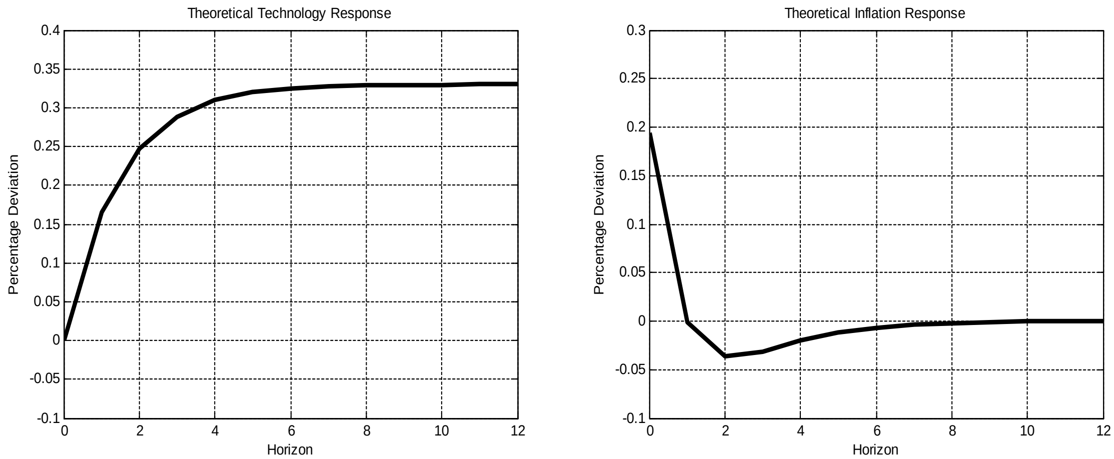

Figure 2.7
New Keynesian Model Inflation Response to News Shock
Real Wage Rigidity

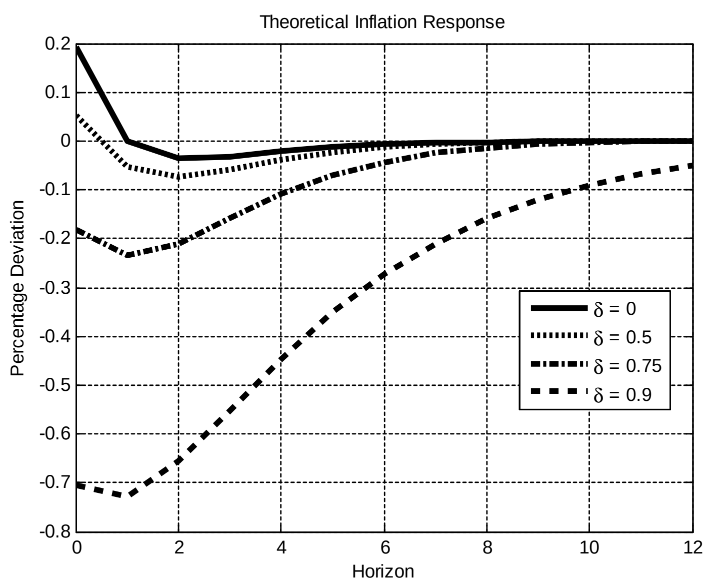

Figure 2.8
New Keynesian Model Inflation Response to News Shock
Growth Rate Policy Rule

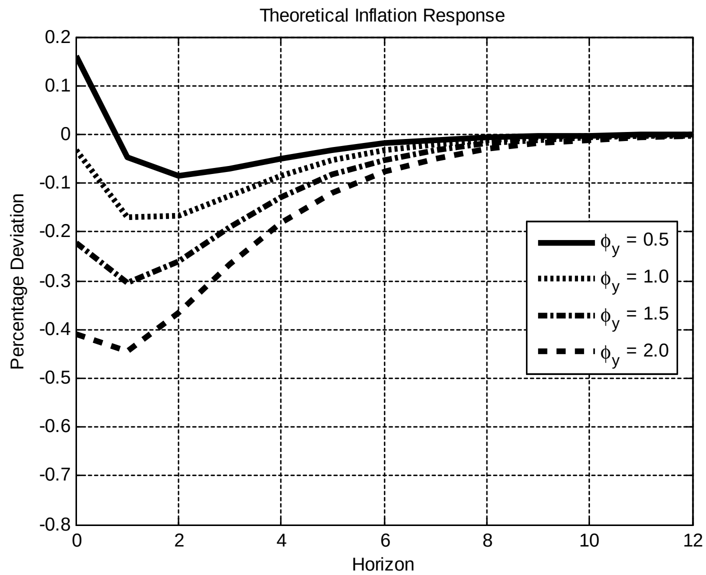

Figure 2.9
New Keynesian Model Inflation Response to News Shock
Incorrect Output Gap Rule

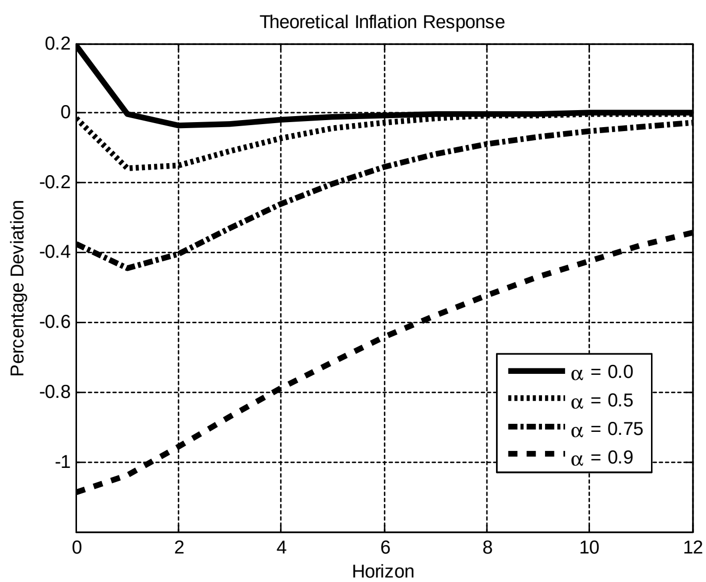

Figure 2.10
Technology Response to News Shock

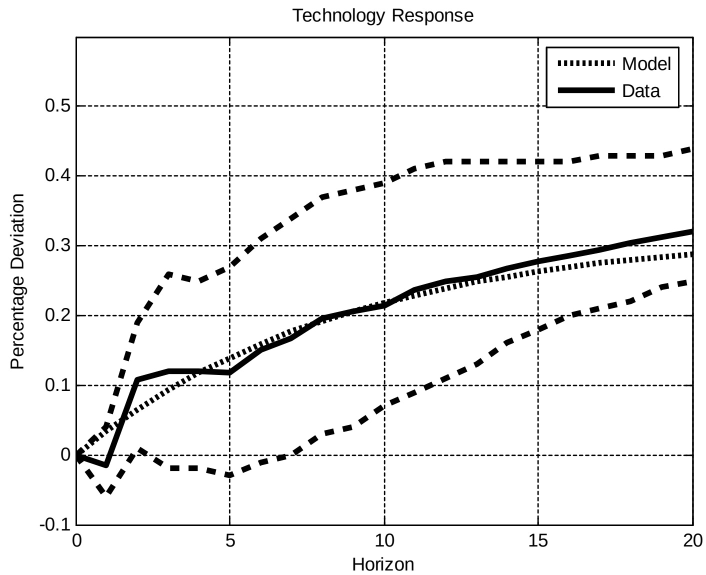

Figure 2.11
Inflation Response: Optimal Parameter Values
Sticky Real Wages, Conventional Taylor Rule

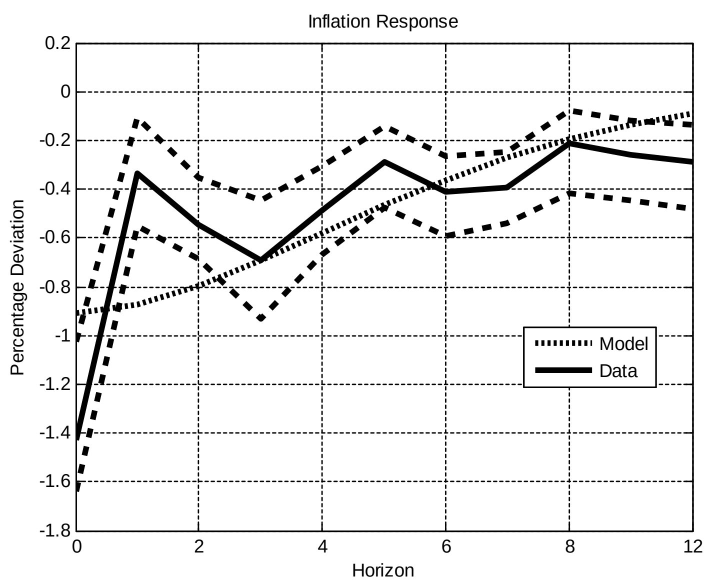

Figure 2.12
Inflation Response: Optimal Parameter Values
Sticky Real Wages, Misperception Taylor Rule

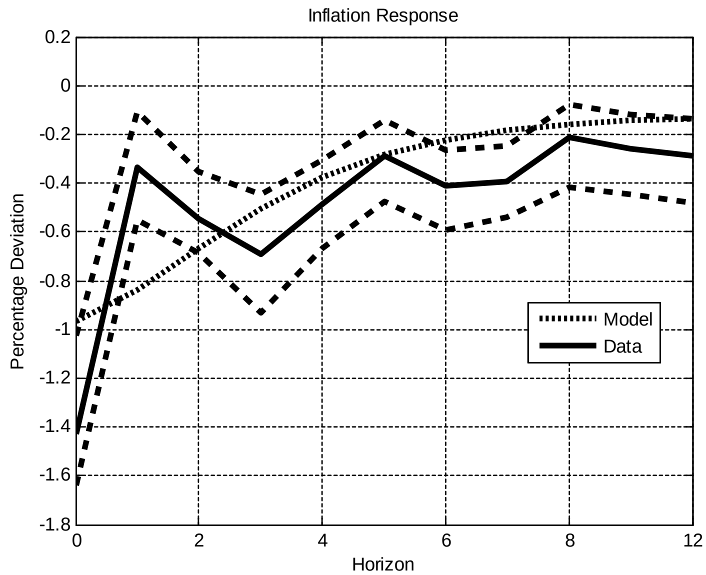

References

[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." Ch. IV, *Expectations and the Business Cycle*, Ph.D. Dissertation of Eric R. Sims, University of Michigan, 2009.

[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 and Jordi Gali. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit, and Banking* 39, 35-66, 2007.

[9] Blanchard, Olivier and Danny Quah. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79: 654-673, 1988.

[10] Calvo, Guillermo. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics* 12: 383-398, 1983.

[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] Cochrane, John. "Permanent and Transitory Components of GNP and Stock Prices." *Quarterly Journal of Economics* 109: 241-266, 1994a.

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

[14] Coibion, Olivier and Yuriy Gorodnichenko. "Strategic Interaction Among Price-Setters in an Estimated DSGE Model." University of California-Berkeley working paper, 2007.

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

[16] Faust, Jon and Eric Leeper. "When Do Long Run Identifying Restrictions Give Reliable Results?" *Journal of Business and Economic Statistics* 15: 345-353, 1998.

[17] Fernandez-Villaverde, Jesus and Juan Rubio-Ramirez. "How Structural are Structural Parameters?" *NBER Macroeconomics Annual* 22: 83-137, 2007.

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

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

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

[21] Gali, Jordi. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, 2008.

[22] Hamilton, James. *Time Series Analysis*. Princeton University Press, 1994.

[23] Ireland, Peter. "Technology Shocks and the New Keynesian Model." *Review of Economics and Statistics* 86: 923-936, 2004.

[24] Jaimovich, Nir and Sergio Rebelo. "Can News About the Future Drive the Business Cycle?" Northwestern University working paper, 2008.

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

[26] Orphanides, Athanasios. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics* 50: 983-1022, 2003.

[27] Orphanides, Athanasios. "Monetary Policy Rules and the Great Inflation." *American Economic Review* 92: 115-120, 2002.

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

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

[30] Sims, Eric. "Expectations Driven Business Cycles: An Empirical Evaluation." Ch. III, *Expectations and the Business Cycle*, Ph.D. Dissertation of Eric R. Sims, University of Michigan, 2009.

[31] Taylor, John. "Discretion vs. Policy Rules in Practice." *Carnegie-Rochester Series on Public Policy* 39: 195-214, 1993.

[32] Uhlig, Harald. "What Drives GNP?" EABCN working paper, 2003.

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

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

[35] Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press, 2003.

Chapter III

Expectations Driven Business Cycles: An Empirical Evaluation

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

[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.

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 explicitly dynamic dimension to the firm's profit maximization problem, 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 evidence 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 Chapters II and IV, 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 introduced in Chapter II 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 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 a simple DSGE model.

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. In contrast, surprise technology innovations appear quite transitory, and lead to large impact increases in both output and investment. 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 both news shocks and surprise technology shocks 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 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:

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

$\varepsilon_{1,t}$ 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.

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 Chapter II.[3] 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 innovation in technology. I then identify the news shock ($\varepsilon_{2,t}$) as the structural shock from the system which comes as close as possible to satisfying the identifying assumption that two shocks alone account for variation in technology. In particular, the news shock is identified as the structural shock which best explains future variation in technology not explained

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

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 which 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 the Appendix; see also the discussion in Chapter II. 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, by approach would (asymtotically) identify the same shock and impulse responses. As shown in Chapter II, there appear to be important movements in stock prices unrelated to technology shocks altogether.[4] 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 in technology at all horizons, not just in the long run. The shocks identified by these authors 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

[4]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.

costs) augmented with both news shocks and surprise technology shocks. The full description of the model and the details of the simulation exercise can be found in the Appendix. 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 3.1 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 3.2 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.

The average correlation between the identified news shock and the true news shock across 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. Importantly, 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.[5] The simulations are of similar high quality under alternative specifications of the model and over a variety of different parameterizations. Taken as a whole, the simulations suggest that my approach is capable of doing quite 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 especially 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 satisfied in the model as presented in the Appendix. The intuition for the failure of invertibility is that the news shock is both a shock and a state variable in the model – agents must keep track of its value for several periods until it loads onto the level of technology. Nevertheless, as stressed 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 remarkably 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.

[5] 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.

3 Empirical Evidence

In this section I present the main results of the paper. In particular, 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 standard 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.[6] As a robustness check, I will also present results using the more conventional measure of TFP which does not attempt to control for unobserved input variation.

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

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

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 technology 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 GDP and consumption data are from the BEA; the 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 Chapter IV.[7] The confidence data are available beginning in 1960; all other series begin in 1948.

3.2 Benchmark Results

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

[7] The specific survey 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 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.

correction model (VECM); I obtain very similar results when estimating the VAR in levels. Standard unit root tests overwhelmingly 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 3.3 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.[8] 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

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

the primary source 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.[9] 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. Strikingly, 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.

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 15 percent at high frequencies) is helpful in understanding the differences between my results and others in the literature. The news shock is associated with economically large declines in inflation and increases in consumer confidence. For further discussion of the responses of these "information variables" see Chapter II.

Table 3.1 depicts the share of the forecast error variance of the variables in the VAR attributable to the identified 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, confirming the notion that news shocks may indeed comprise a significant portion of the long run stochastic component of technology. The final row of the table shows the total contribution to technology's

[9] Similar results obtain when the interest rate is included directly in the system. See also the discussion in Chapter II.

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 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 3.4 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 correlation between the shock series and lagged output is essentially zero and is therefore omitted. 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 3.5 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.[10] 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

[10] 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.

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 is likely to be more informative about the business cycle relevance of news shocks than is a plot of the shocks alone. Figure 3.6 plots simulated and actual values of GDP 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 GDP 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 3.7 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.[11] 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

[11] 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 3.1. The surprise technology shock makes only small contributions to the forecast error variance of both consumption and hours.

Chapter II of this dissertation 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)).[12] My finding 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 "productivity shock" 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 do suggest that non-technology shocks are an important source of fluctuations. The final row of Table 3.1 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

In this subsection I consider several robustness checks on my benchmark result that favorable news about future technology leads to a rise in consumption and declines in output, hours, and investment on impact. The first is to use the more traditional Solow residual in place of the Basu, Fernald, and Kimball (2006) corrected technology measure. As described in detail above, this series is constructed similarly but without the correction for unobserved input variation. Without an attempt to control for such input variation, however, there are reasons to doubt the validity of the identifying assumption that the news shock is orthogonal to technology.

The estimated impulse responses are shown in Figure 3.8. The underlying system features all seven variables from the benchmark VAR, with the data running from

[12] 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 results presented here.

1960-2007. As before, I estimate the system as a VECM with a truncation horizon of $H = 50$. A favorable news shock still leads to negative impact comovement, but the responses differ in important and qualitative ways from the benchmark results. Whereas with the corrected measure the technology response to the identified news shock appears permanent, here the time path of TFP is quite transitory – rising sharply to a peak of some 0.5 percent higher than its pre-shock value before reverting. A natural explanation for this pattern of response is that the identification with uncorrected technology is picking up some purely cyclical features of the data. In spite of the reversion towards zero, the qualitative nature of the responses is the same with either measure of technology, with output, hours, and investment all falling on impact and consumption rising slightly following a favorable news shock.

I also estimate the VAR using the longer data set with observations from 1948 to 2007, which necessitates dropping the measure of consumer confidence from the system. These responses are omitted, but are qualitatively similar to my benchmark results. The main differences are that the both the impact increase in consumption and the impact decline in investment are larger. The response of technology itself is otherwise quite similar. Once again, the negative impact comovement is followed by the aggregate variables largely tracking movements in technology.

The main result that news shocks induce negative comovement is also 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. 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.

The original version 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 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 3.9 I show estimated impulse responses to a news shock from a smaller system. In particular, I omit consumer confidence, inflation, and investment 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.

Finally, I consider how sensitive my results are to the specification of the maximization problem underlying identification. Figure 3.10 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. The impact increase in consumption is larger here than in Figure 3.3, and 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. Indeed, as shown in Table 3.2, the news and combined technology shocks leave a large share of the variance of technology (as much as 40 percent) unexplained at business cycle frequencies when the news shock is identified with a long run restriction. This fact becomes important in understanding the differences between my conclusions and those of previous authors, and I discuss it in more depth in the next subsection.

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 truncation horizons.[13]

Not only do I reach different conclusions concerning the effects of news shocks on aggregate variables, there is a large quantitative and qualitative difference between 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.[14]

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

[14] 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

Indeed, Beaudry and Portier (2006) note that "growth beyond its [technology'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 technology 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 eight 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 3.2 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 3.1 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 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

data (available from the *American Economic Review* website) and obtain very similar results.

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.

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 3.11 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 the Appendix, 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 of the Appendix reduces to the standard real business cycle model. The remaining parameters of the model are calibrated as detailed in the Appendix.

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 3.7. 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 3.3. 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 3.3 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

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 $\mathbf{A}_0\mathbf{A}_0' = \mathbf{\Sigma}$ 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 is for all possible orthogonalizations is:

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

$\mathbf{B}_\tau\mathbf{A}_0$ 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:

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

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.

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

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. Uhlig (2003) shows that 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 $\boldsymbol{\varepsilon}_{2,t} = \boldsymbol{\varsigma}^{*\prime}\ \widetilde{\mathbf{A}}_0^{-1}\mathbf{u}_t$.

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 in the variables 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 (1) in the text.

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 the news shock by maximizing the variance share of technology over a horizon of twenty quarters. I follow the identification procedure outlined above and collect the estimated impulse responses and identified time series of news shocks for each simulation.

Table 3.1: 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	3.6	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 3.2: 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 3.3: 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 3.1

Model and Monte Carlo Estimated Impulse Responses to News Shocks

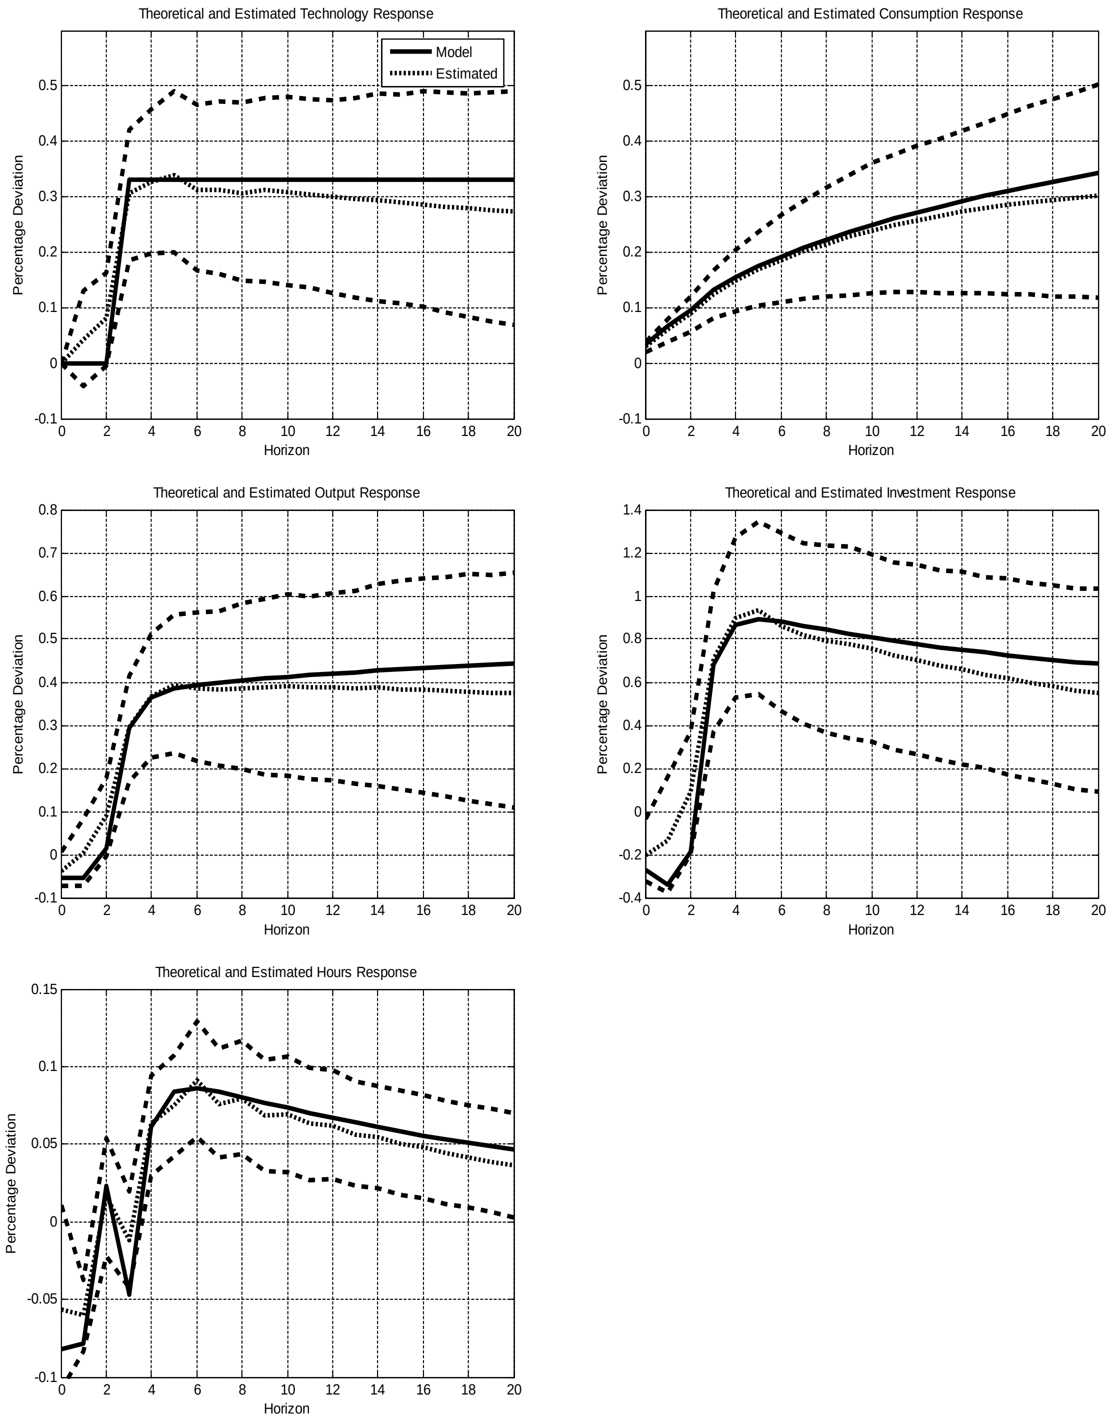

The solid lines show the theoretical impulse response to a news shock from the model of the Appendix, The dotted lines depict the average estimated impulse responses over 2000 Monte Carlo simultions, with the dashed lines representing the 10th and 90th percentiles of the distribution of estimated impulse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

Figure 3.2

Model and Monte Carlo Estimated Impulse Responses to Surprise Technology Shock

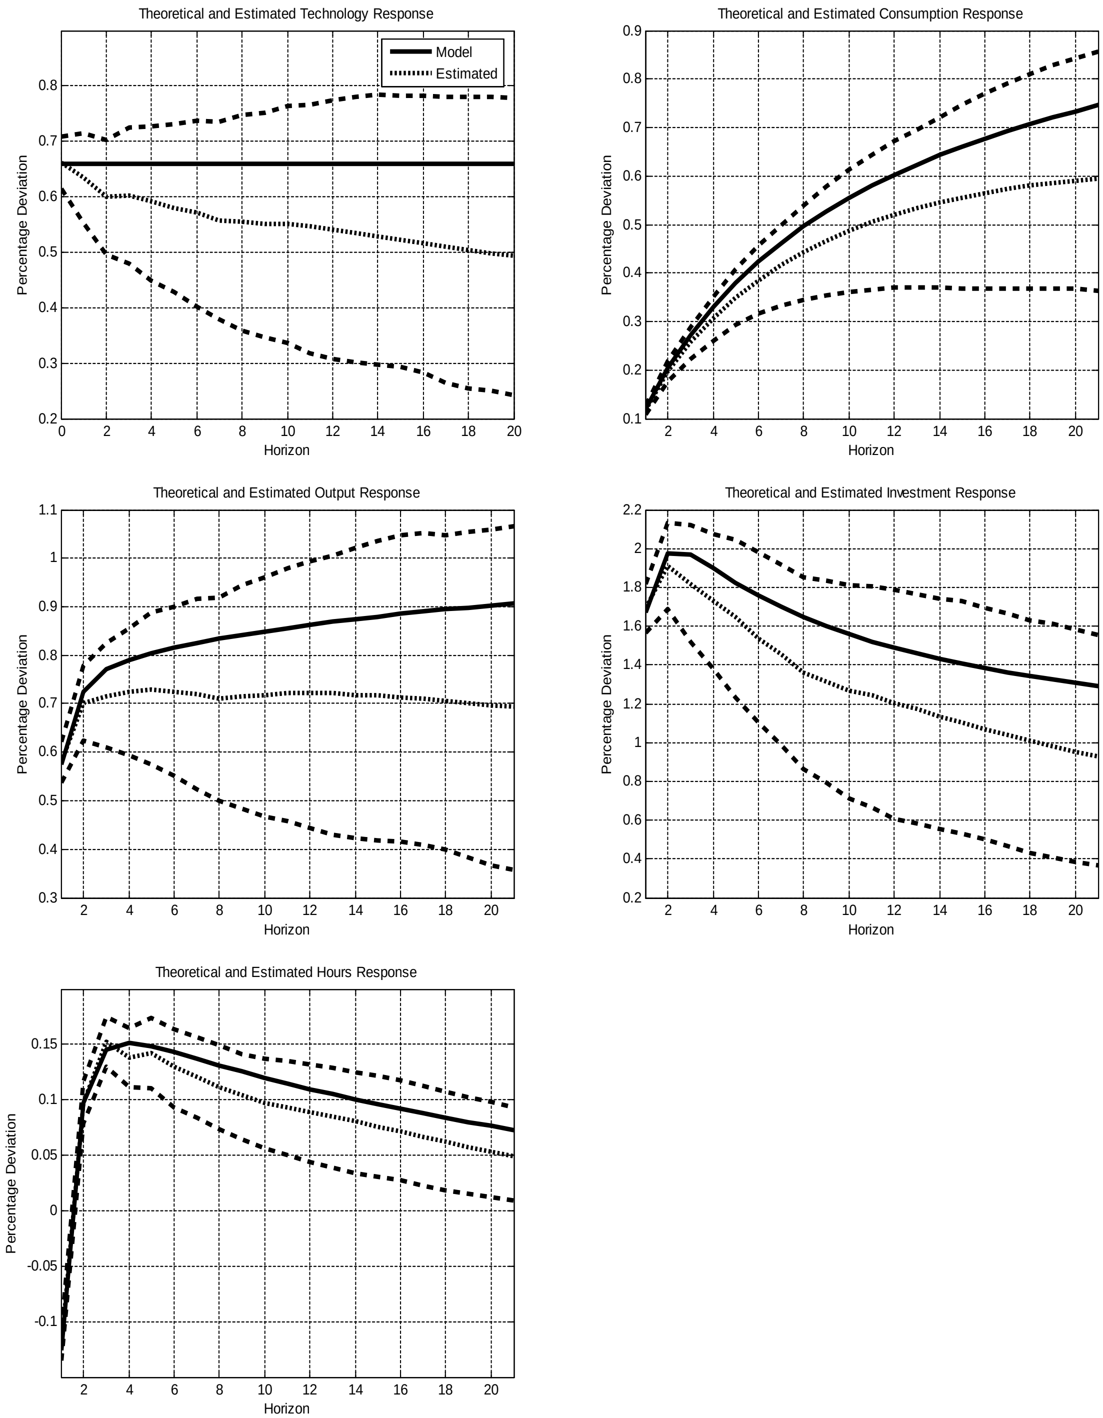

The solid lines show the theoretical impulse response to a surprise technology shock from the model of the Appendix, The dotted lines depict the average estimated impulse responses over 2000 Monte Carlo simultions, with the dashed lines representing the 10th and 90th percentiles of the distribution of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

Figure 3.3
Empirical Impulse Responses to News Shocks

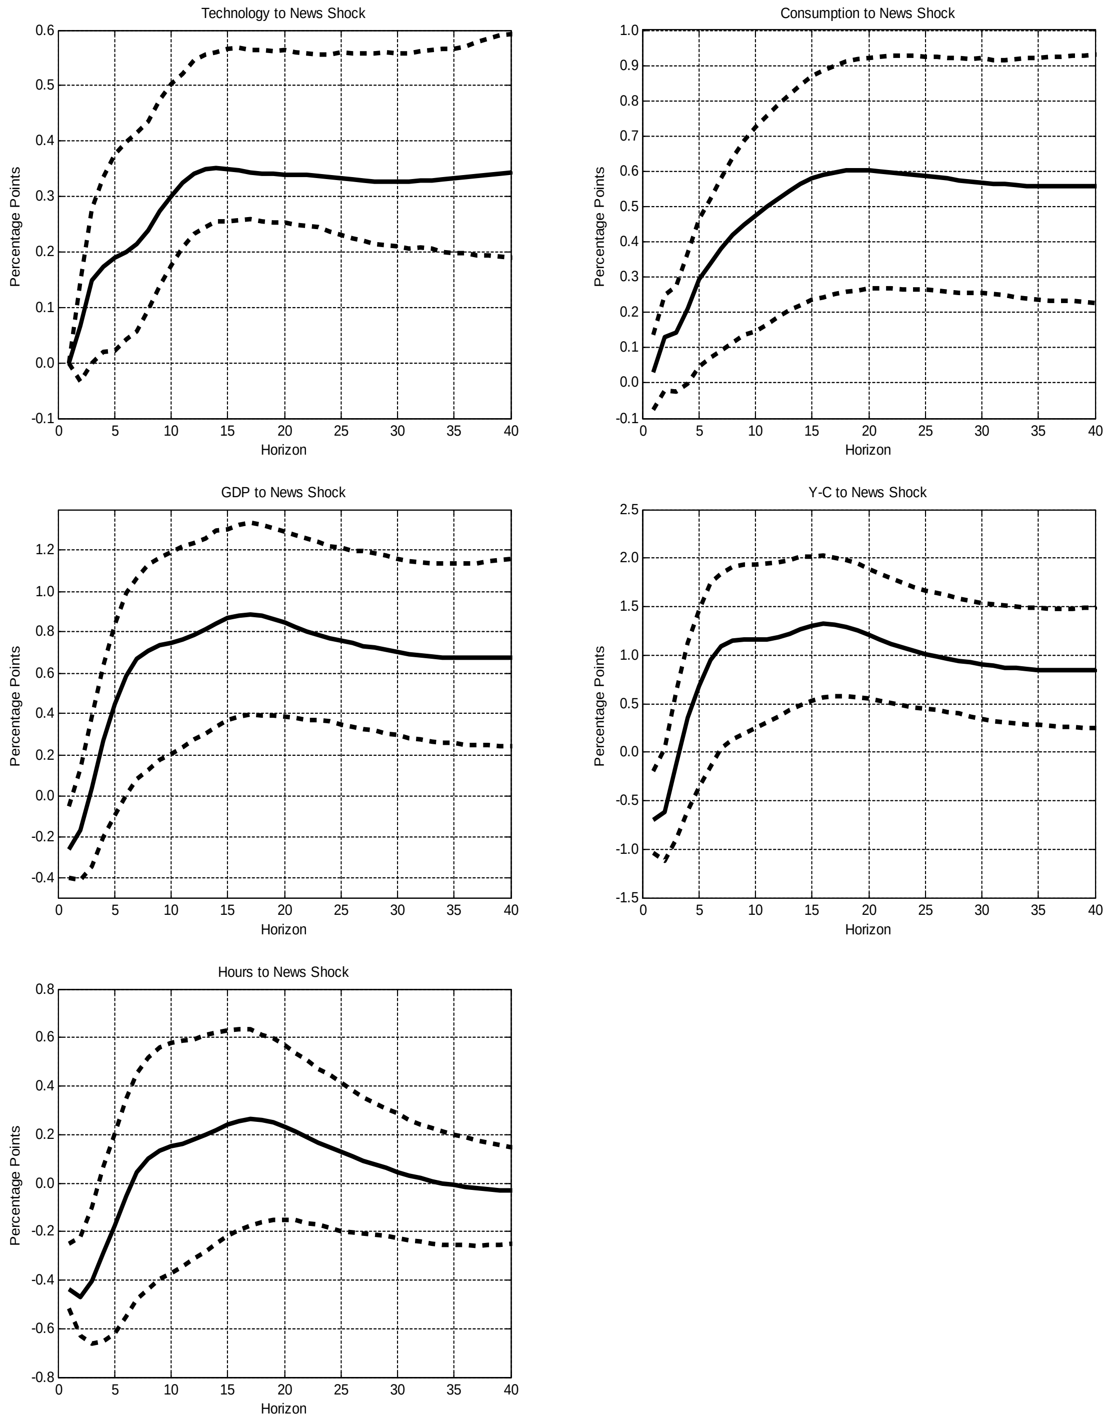

The dashed lines depict the 5th and 95th percentiles of the empirical distribution of impulse responses from a bias-corrected bootstrap procedure.

Figure 3.4
Cross Correlogram Between News Shock Series and HP Detrended Output

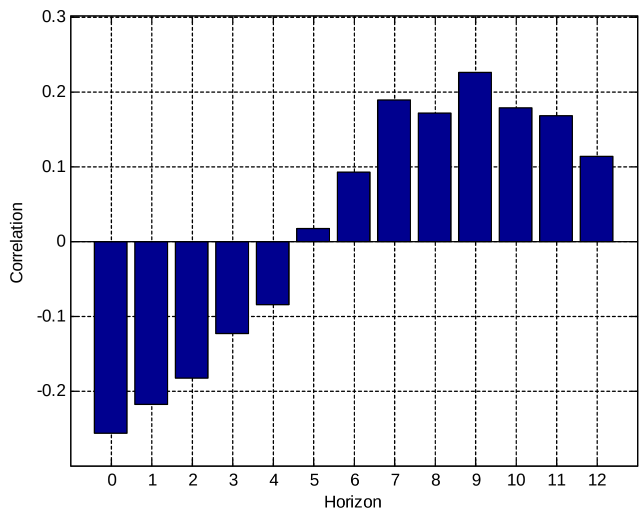

The above figure plots the correlation between the identified news shock series from the benchmark VAR with HP detrended output led over the specified horizon. As such, the number for horizon 0 is the contemporaneous correlation, the number for horizon 1 is the correlation between the shock series and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

Figure 3.5
Identified News Shock Time Series and US Recessions

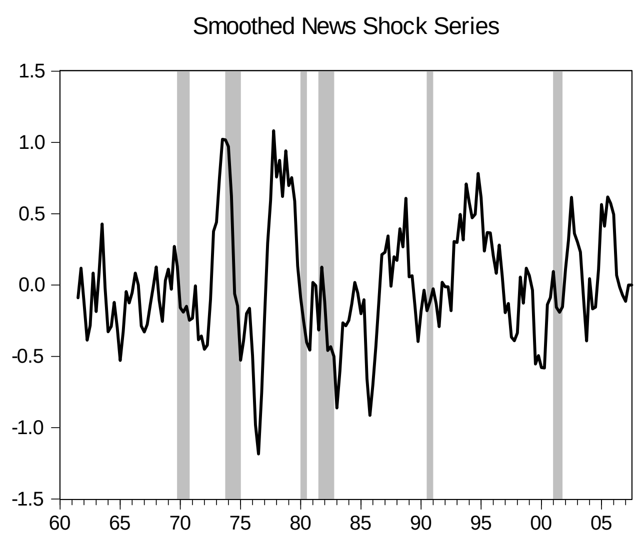

This figure plots the time series of identified news shocks from the benchmark VAR. So as to render the figure more readble, the plotted data is smoothed using a one year moving average.

Figure 3.6
Historical Decomposition of Output per Capita in US Recessions

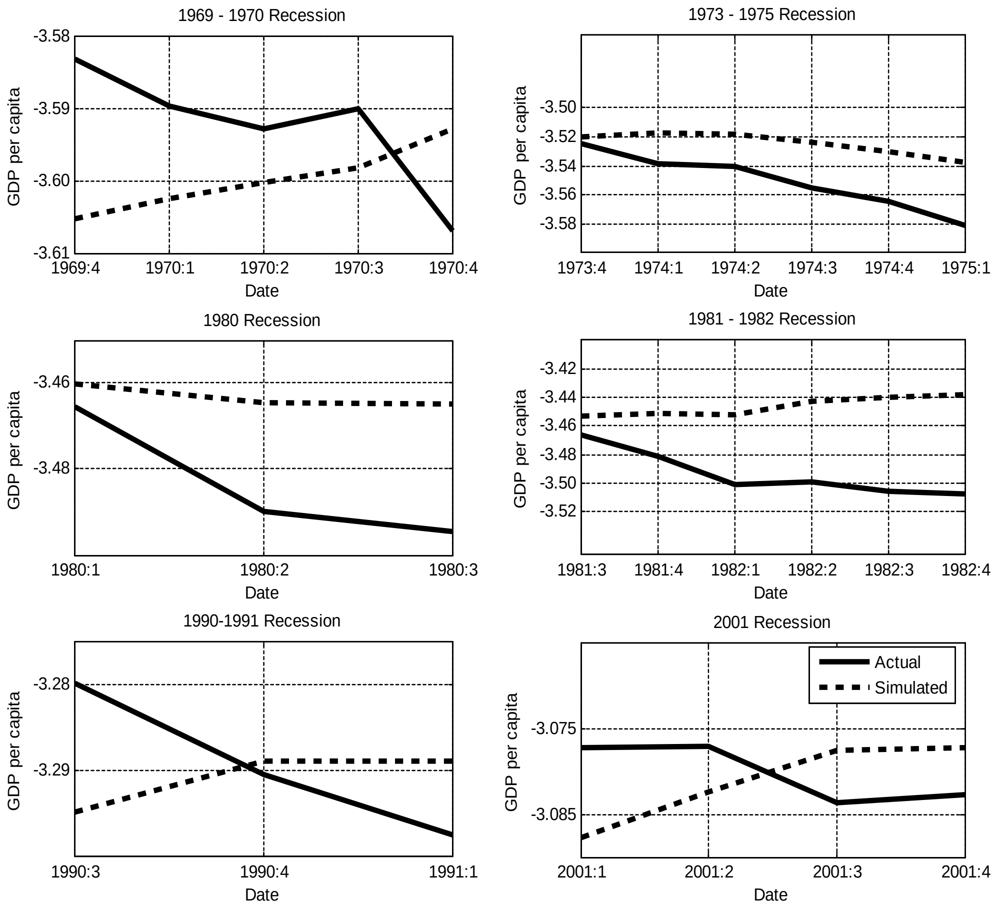

The above shows the actual (solid line) and simulated (dashed line) paths of GDP per capita during the six NBER dated US recessions since 1961.

Figure 3.7
Estimated Impulse Responses to Surprise Technology Shocks

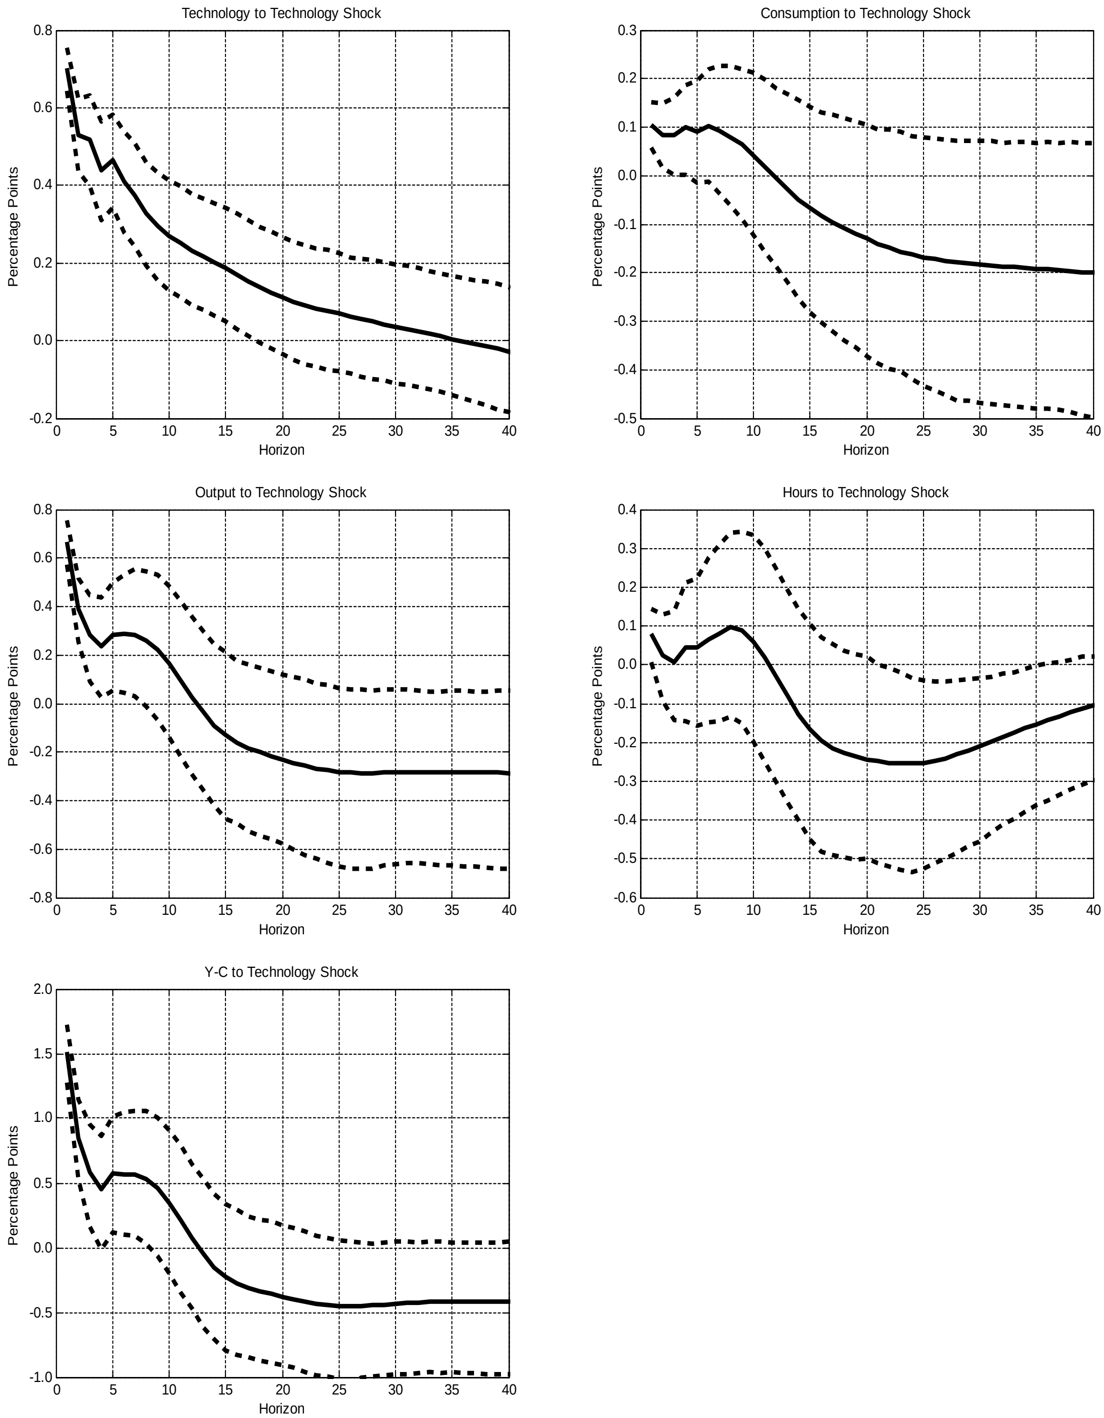

Figure 3.8
Estimated Impulse Responses to News Shocks: Uncorrected TFP

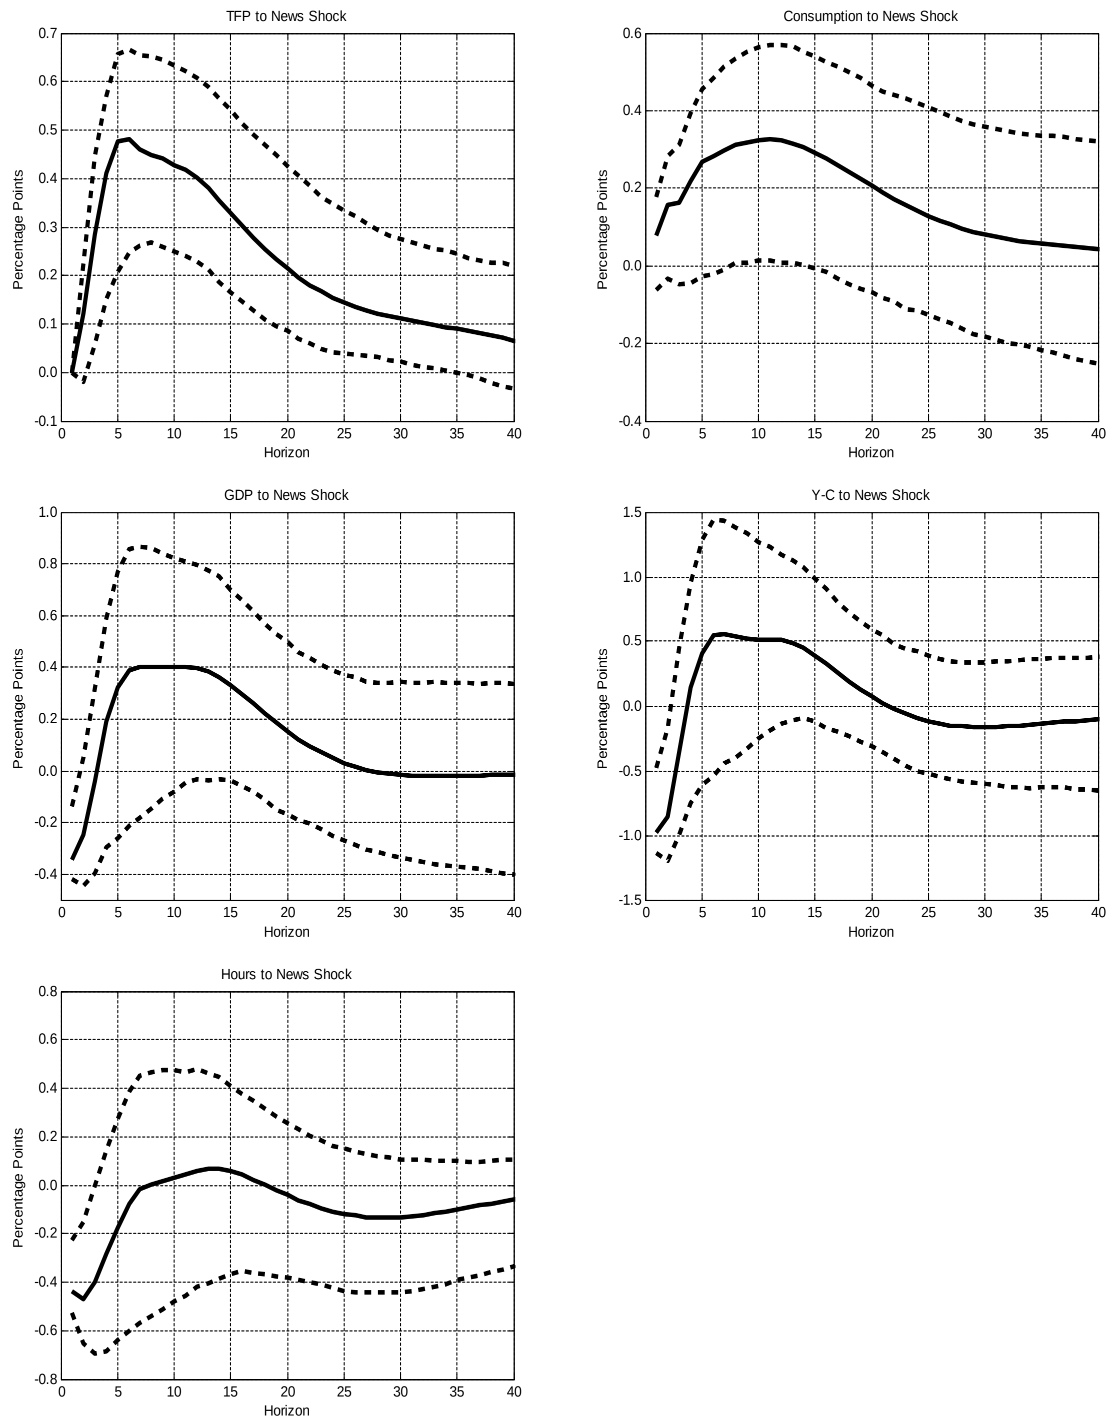

These are impulse responses from the benchmark VAR using the uncorrected measure of TFP.

Figure 3.9
Impulse Responses to News Shocks: Smaller System

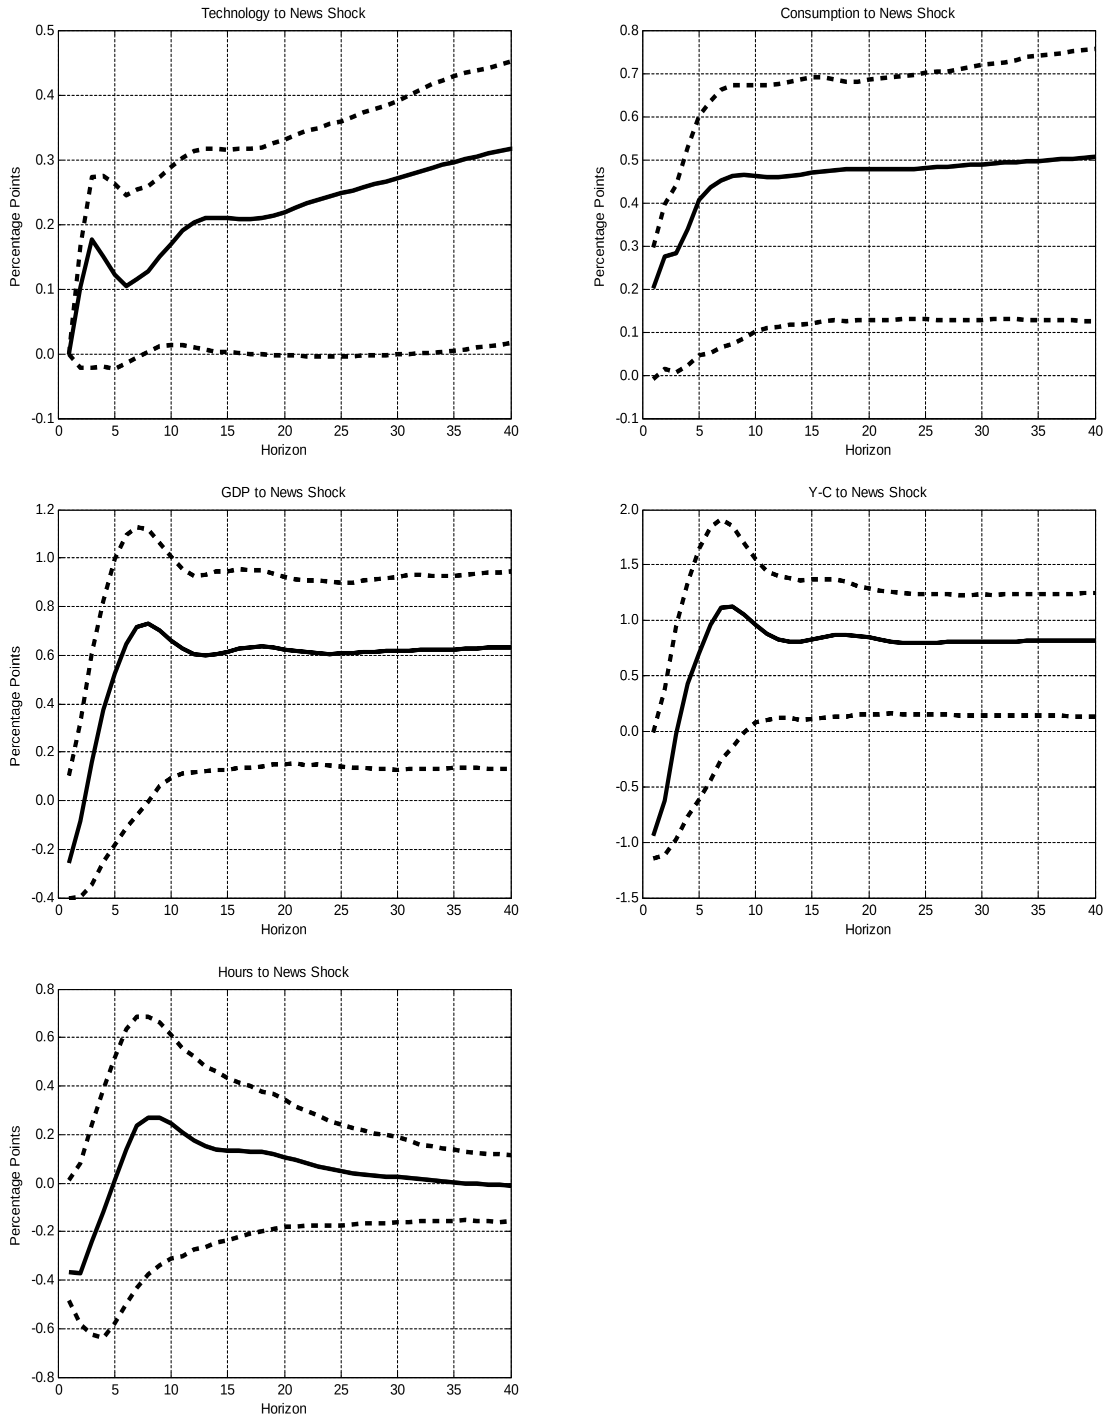

The above are responses from a VAR with technology, stock prices, consumption, output, and hours.

Figure 3.10
Impulse Responses to News Shocks
Long Run Identification

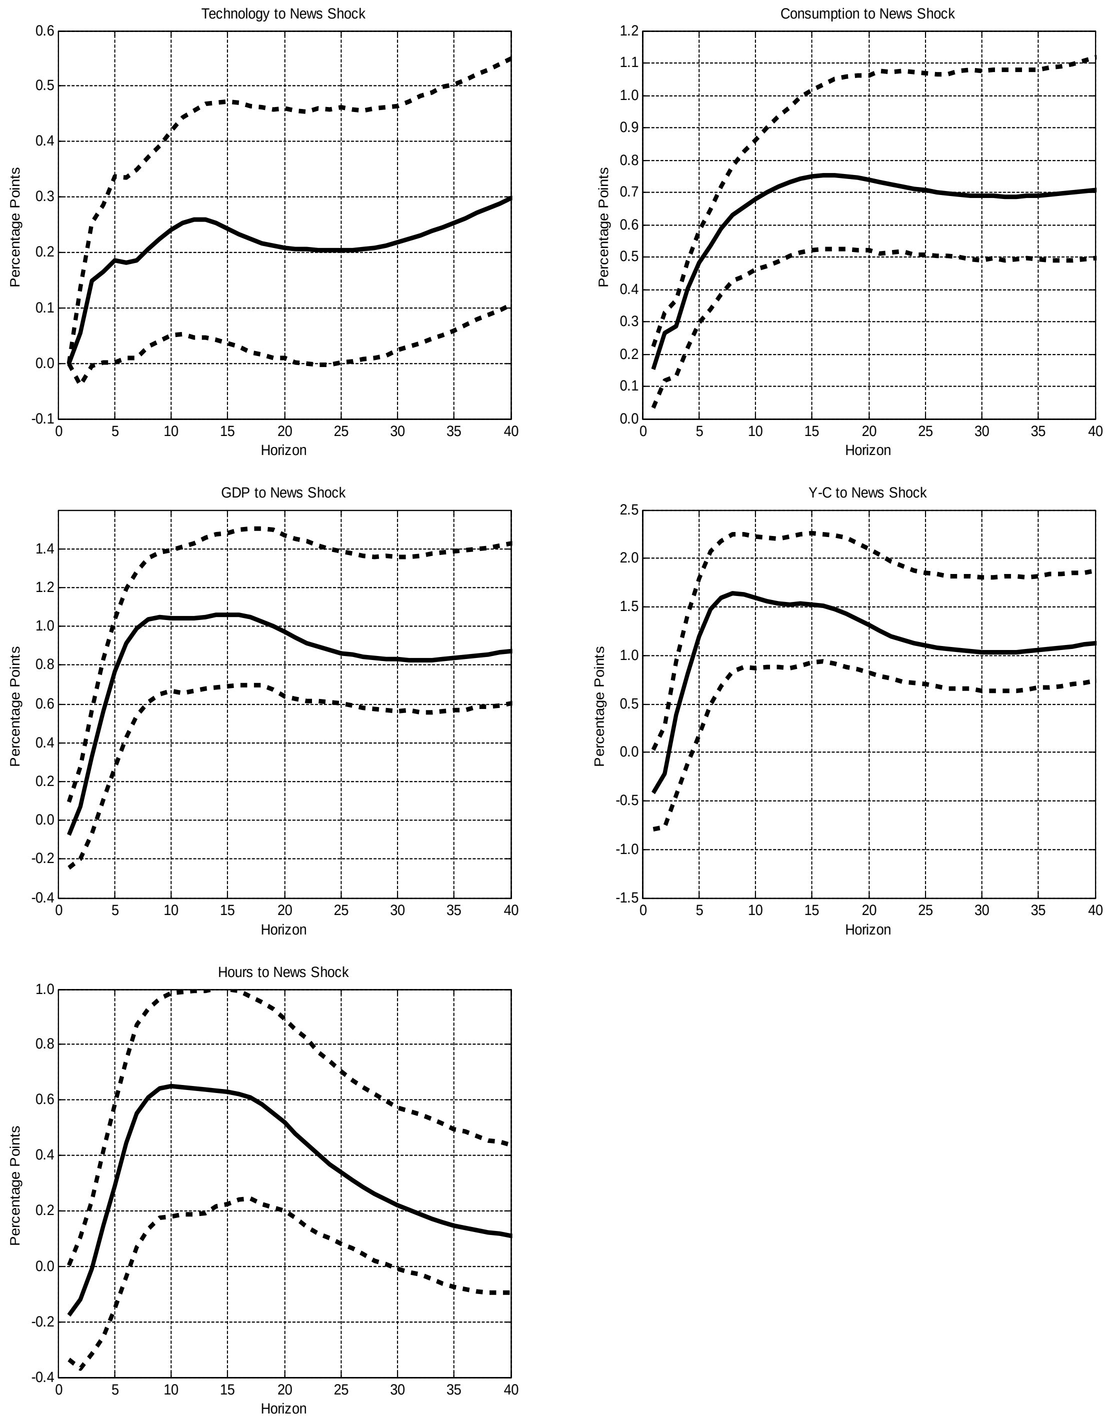

Figure 3.11

Estimated and Theoretical Responses to News Shocks

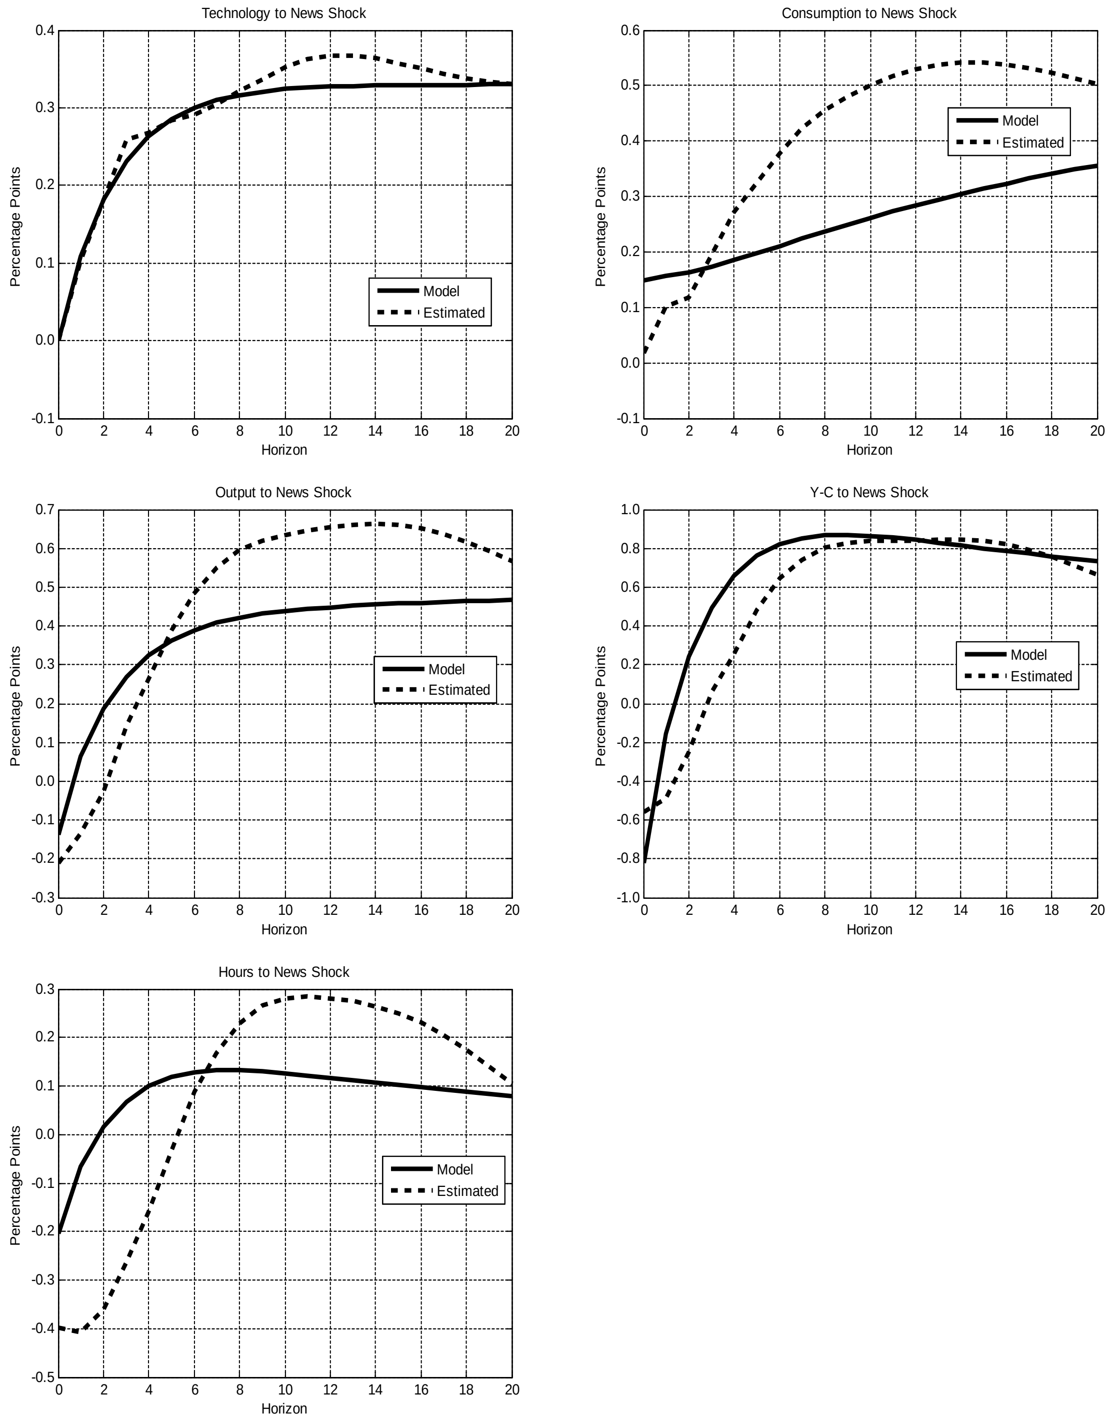

The solid line shows the model generated impulse responses from an RBC model with news shocks and a standard calibration. The dashed line shows the estimated response from the benchmark empirical VAR.

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." Chapter IV, *Expectations and the Business Cycle*, Ph.D. Dissertation of Eric R. Sims, University of Michigan, 2009.

[4] Barsky, Robert and Eric Sims. "News Shocks." Chapter II, *Expectations and the Business Cycle*, Ph.D. Dissertation of Eric R. Sims, 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] Lippi, Marco and Lucrezia Reichlin. "The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment." *American Economic Review* 83: 6445-652, 1993.

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

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

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

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

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

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

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

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

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

Chapter IV

Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence

1 Introduction

In the popular press and much of the business community it continues to be an article of faith that "consumer confidence" has an important role – both prognostic and causal – in macroeconomics. On the other hand, the stance of the rather limited academic literature on confidence is far more ambiguous. The judgments range from the conclusion that confidence measures have an important role both in prediction and understanding the cause of business cycles, to the view that they contain important information but have little role in the assignment of causality, to the verdict that they have no value even in forecasting.

There are, broadly speaking, two contrasting approaches to the role of confidence in macroeconomics. The first, which we will refer to as the "animal spirits" view, posits autonomous fluctuations in beliefs and consumption that in turn have causal effects on economic activity. In the proceedings of a symposium on the causes of the 1990-1991 recession, both Hall (1993) and Blanchard (1993) regard exogenous movements in consumption as a cause of business cycles.[1] Indeed, Blanchard proposes that the cause of the recession was a powerful, long-lasting negative consumption shock associated with an exogenous shift in pessimism that had a causal effect on consumption and overall aggregate demand. While not fully pursuing the idea in his

[1] In an interesting but almost forgotten early contribution, Hall (1986) – partially repudiating Hall (1978) – argues that an important fraction of the random walk in consumption comes not from the expectational surprise in the Euler equation but from a second disturbance that he has more recently referred to as "spontaneous consumption". In Hall (1993), this is interpreted as a shock to the taste for consumption relative to leisure.

brief paper, Blanchard proposes that one might be able to test this hypothesis on the basis of the observation that such an exogenous shift in pessimism ought to have only temporary effects on consumption.[2]

The second view of confidence – what we will call the "information view" – suggests that a relationship between innovations in measures of consumer confidence and subsequent macroeconomic activity arises because confidence measures contain fundamental information about the current and future states of the economy. For example, Cochrane (1994b) proposes that consumption surprises proxy for news that consumers receive about future productivity that does not otherwise show up in econometricians' information sets. His attempt to reconcile VAR evidence with theory closely anticipates the "news approach to business cycles" of Beaudry and Portier (2004, 2006). They analyze models where agents become aware of changes in future productivity orthogonal to current productivity, and argue that stock price innovations proxy for future technological improvement not reflected in current technology. The "information view" of confidence supposes that confidence innovations might contain similar information.

In Section 2 of the paper, we first show that unexplained innovations in several variables representing survey responses to forward-looking questions from the Michigan Survey of Consumers have powerful predictive implications for the future paths of macroeconomic variables. In particular, within the context of augmented consumption-income VARs, we show that unexplained innovations in the responses to several consumer confidence questions have significant, slowly building, and apparently permanent implications for output and consumption. Confidence is not highly Granger-caused by income or consumption, nor are its innovations highly correlated with innovations in those variables. Responses to little-used survey questions on "news heard" do help to somewhat explain confidence innovations, but with only a very modest incremental R^2. These observations point to the conclusion that these measures of consumer confidence are not merely noise, nor are they simply reflections of macroeconomic news reports or innovations in other variables with which they are correlated.

In Section 3 we attempt to distinguish the hypothesis that these impulse responses

[2] In some ways, a limiting case of animal spirits appears in the "sunspot" literature. Though pinned down only by extrinsic coordinating variables, expectations in the equilibria of these models are self-fulfilling, and thus not irrational (see Farmer (1999)). The existence of sunspot equilibria depend on strong increasing returns, supply externalities, or other mechanisms that are typically not accepted as empirically plausible. The notion of animal spirits in this paper does not encompass sunspots.

indicate a causal channel from sentiment to economic outcomes (the "animal spirits" view) from the alternative interpretation that the surprise confidence movements summarize information about economic prospects known to consumers (the "information" view). To provide a framework for distinguishing these alternative views of confidence, we present a highly stylized New Keynesian model with three kinds of shocks. The first shock is an immediate and unexpected improvement in productivity (a "level shock"). The second is a reflection of genuine news that productivity will grow more rapidly for a substantial period of time into the future (a "growth shock", also to be referred to as an "information shock" because it conveys information about future productivity that cannot be fully inferred from current productivity).[3] We only permit households to observe a noise-ridden signal of the information shock to technology. We interpret the noise innovation in the signal as an "animal spirits shock" as it is associated with erroneous consumer optimism or pessimism. This shock can be given alternative less structural interpretations, and in equilibrium its implications are similar to those of an exogenous innovation to the Euler equation. Regardless of the particular interpretation, a series of positive animal spirits shocks might capture the putative "irrational exuberance" of the 1920s or 1990s, while a predominance of negative shocks would usher in a period of excessive pessimism.

The model has clear implications for the response of the endogenous variables to each of the three shocks. "Animal spirits" shocks behave as aggregate demand shocks – they are associated with transitory increases in output that attenuate over time, and they produce both inflation and increases in real interest rates. "Information shocks" regarding future productivity and shocks to current productivity are followed by gradual movements in the macroeconomic variables that are not subsequently reversed. Both of these fundamental shocks are also associated with rising real interest rates. Thus, the model yields two primary criteria by which to distinguish animal spirits from fundamental shocks: positive animal spirits shocks are followed by transitory movements in real activity and increases in inflation, while favorable fundamental shocks may result in permanent movements in activity and may be either inflationary or deflationary.

In Section 4, we estimate an expanded VAR with the variables implied by the model augmented with a measure of confidence. As in the three variable systems of Section 2, the results show that confidence innovations are associated with little

[3] We employ the term "information" in the same way Cochrane (1994b) and Beaudry and Portier (2006) use the word "news". Lorenzoni (2008), somewhat confusingly, uses the term "news" to refer to noise in a public signal, which functions much like our animal spirits shock.

immediate response of real activity but prolonged growth in consumption, income, and measured productivity. There is no evidence of reversion in these variables – in particular, the point estimates suggest that income and consumption are higher by more than two-thirds of a percent in the long future in response to a confidence innovation, with the confidence bands associated with these impulse responses lying above zero at horizons in excess of ten years. Confidence innovations are associated with transitory increases in real interest rates and hours of work, and also lead to a large and persistent reduction in inflation. These empirical responses are not at all similar to the implications of animal spirits shocks in our model, nor are they particularly consistent with the theoretical responses to level shocks.

We next postulate a structural equation in which surprise movements in confidence are attributable to the signal agents receive about the growth rate and to the innovation in the current state of productivity. We estimate a subset of the structural parameters of the model via a modified version of simulated method of moments. We are able to resoundingly reject the hypothesis that animal spirits shocks (as specified in this paper) are an important source of the observed relationships between confidence innovations and macroeconomic variables. On the other hand, we do find convincing evidence in favor of the information interpretation of consumer confidence. The implications of confidence innovations for output and spending at short horizons are far too small for confidence to be primarily a reflection of changes in current fundamentals, yet the longer horizon implications are far too large and significant for confidence innovations to not be conveying information about fundamentals. Our results suggest that there are information shocks about future productivity not wholly reflected in current productivity, and that these shocks account for a significant fraction of the innovation in measured confidence.

2 Income, Consumption, and Confidence

We begin with the dynamics of income and consumption as implied by the bivariate vector autoregression discussed by Cochrane (1994a). In particular, we estimate a two variable system consisting of the log of real GDP and the log of real consumption of services plus non-durables, both in per capita terms after dividing by the civilian non-institutionalized population aged sixteen and over. The data are seasonally adjusted measures at a quarterly frequency from the first quarter of 1960 to the third quarter of 2007. The data strongly suggest that the variables are cointegrated, and the estimated cointegrating vector is sufficiently close to [1,-1] that we follow Cochrane and others in imposing it. While popular information criteria generally favor a small number of lags (one or two), we take a conservative stance and estimate the VAR with four lags.

Cochrane orthogonalizes the innovations so that consumption is ordered first. This ordering is implied by a simple permanent income model in which all information is immediately reflected in consumption.[4] However, the line of inquiry in his subsequent paper (Cochrane (1994b)) suggests a focus on the alternative ordering with income first; there the focus is on the information about future income embodied in consumption but not in current income. Figure 4.1 presents impulse responses under both orderings, with the solid line referring to the ordering with consumption first and the dashed line to the orthogonalization with income ordered first. The key feature of these impulse responses is that innovations in consumption – whether or not they are orthogonalized with respect to income – are associated with powerful and prolonged subsequent increases in income. At the shorter horizons, most of the movement in income is explained by its own innovation, but the "effects" of a consumption innovation build over time so that much or all of the permanent component of GDP appears to be captured by innovations in consumption. In short, results from this two variable VAR suggest that "consumption shocks" convey news about income many periods into the future.

As Cochrane (1994b) stresses, a natural explanation for the finding that consumption innovations predict much of future output is that agents have some advance knowledge about future income that they use when making consumption decisions. Forward-looking questions on surveys of consumer expectations and attitudes might potentially provide a direct measure of such information, and thus a direct test of Cochrane's hypothesis. Is much or most of the information embodied in consumption picked up by survey expectations of future output? Do the survey data indicate, on the other hand, that consumers receive a great deal of news that is not reflected in current consumption? We turn to these questions now, introducing some expectational measures from the Michigan Survey of Consumers and augmenting the bivariate consumption-income VARs with these variables.

The survey measure that we will make the most use of in this paper, which we call E5Y, summarizes responses to the following question: "Turning to economic conditions in the country as a whole, do you expect that over the next five years we will have mostly good times, or periods of widespread unemployment and depression,

[4] The consumption $\rightarrow$ income ordering also splits income fairly neatly into permanent and transitory components, as any innovation to income not reflected in consumption ought to be transitory under a partial equilibrium view of the PIH.

or what?" The variable is constructed as the percentage giving a favorable answer minus the percentage giving an unfavorable answer plus one hundred.[5] Our particular affinity for this question arises from the fact that it is aimed at gauging expectations over a relatively long horizon, and because of its specificity as to the relevant time frame.[6] However, its correlation with the response to a similar question specifying a horizon of only twelve months (a variable we call E12M) is 91 percent, and its correlation with another concerning expected changes in personal financial situation over the next twelve months is 85 percent. The correlation of E5Y with the overall expectations component of the Michigan index exceeds 95 percent. Our results in this section are essentially unchanged by the substitution of either of these alternative expectations variables. The alternative questions are described in more detail in the Appendix.

Figure 4.2 plots E5Y and E12M against time. Both series undergo repeated dramatic swings though (as we would expect) the twelve-month-ahead expectations are more volatile than the expectations over a five year horizon. Both variables are quite stationary. The cross-correlogram between E5Y and the conventional Hodrick-Prescott detrended GDP (not shown) indicates that the expectations are by no means a reflection of current output; the contemporaneous correlation between detrended GDP and E5Y is essentially zero. E5Y is negatively correlated with the output gap lagged several periods, and positively correlated with the gap several quarters ahead.

We begin by augmenting Cochrane's income-consumption VAR with E5Y. As before, the system is estimated allowing cointegration between consumption and income, with four lags of each variable. Because confidence measures are clearly stationary, E5Y cannot enter into the long run equilibrium relationship, and we once again impose that the cointegrating vector between consumption and income is [1, -1].[7] It is necessary to make some choices as to how to orthogonalize the innovations. It is important to understand that alternative orthogonalizations in this context are not to be thought of as minimum delay restrictions that delineate alternative structural

[5] Thus a value of 100 is a "neutral" position, while a value of 140 means that the fraction of responses reflecting optimism about the future exceeds the fraction reflecting pessimism by forty percentage points.

[6] Some might argue as well that this question gives the animal spirits hypothesis its "best shot". One argument is that individuals are likely to be more sober-minded in assessing family resources than in forming expectations about the national economy. Another is based on animal spirits models that focus on strategic complementary; in those models beliefs about the economic activities of other agents are central.

[7] Formally, the system features E5Y in levels, consumption and income in first differences, and the lagged (log) difference between consumption and income as an exogenous variable. Our results are virtually identical when estimating the full system in levels.

models; in almost any sensible model, innovations in the underlying structural shocks should affect all three variables instantaneously. Attempts to think about ordering should instead focus on "assigning" the common component of the information in innovations to one or another variable so as to provide upper and lower bounds for the amount of information content in each of the series.

To begin to assess the extent to which the "consumption shocks" in the bivariate VAR are in fact "information shocks" that are well captured by innovations in the survey expectations, we compute impulse responses with E5Y ordered first. As in Cochrane (1994a), income is ordered last, though our results from the augmented consumption-income VARs are largely invariant to the placement of income in the ordering. Figure 4.3 presents the impulse responses to E5Y and consumption innovations under this orthogonalization. The dashed lines represent 90 percent bias-corrected bootstrap confidence bands.[8] As in Cochrane's two variable system, consumption behaves roughly like a random walk in response to its own innovation. In response to a consumption innovation output jumps up on impact, follows a slight hump-shape, and levels off at roughly 0.4 percent higher than its pre-shock value. Though not shown, output displays a large and significant response to its own innovation that dissipates rather quickly. The part of the output innovation that is orthogonal to consumption predicts no significant movement in consumption at any horizon.

An innovation to E5Y has very small implications for both consumption and output on impact. The small impact effects are followed by slowly-building, statistically and economically significant, and apparently permanent responses of both consumption and output. In particular, a one standard deviation innovation to E5Y predicts levels of output and consumption that are roughly 0.7 percent higher forty quarters hence; further, the long run responses of both consumption and GDP to an E5Y innovation are both statistically significant at better than the 90 percent level. E5Y responds significantly neither to income nor consumption innovations; its own innovation accounts for more than 95 percent of its forecast error variance at all horizons under this ordering.

E5Y innovations thus clearly convey important information about the future time paths of real variables, with "effects" that show no tendency to attenuate even at long horizons. However, to what extent are innovations in E5Y simply reflective of

[8]In particular, we generate the confidence bands from the empirical distribution of impulse responses based on 2000 bootstrap draws using bias-corrected OLS slope coefficients as proposed by Kilian (1998).

information contained in consumption? To address this possibility, we re-order the variables in the system such that E5Y is orthogonalized with respect to consumption. As before, output is ordered last in the system. Figure 4.4 presents impulse responses with this particular ordering.

The qualitative features of the impulse responses are unaffected by the alternative orthogonalization. In particular, E5Y innovations orthogonal to consumption still predict slowly-building and permanent responses of both output and consumption. The point estimates are slightly smaller than in the case with E5Y ordered first, with a one standard deviation innovation to E5Y prognostic of long run increases in both consumption and output of slightly more than 0.5 percent (as opposed to 0.7 percent with E5Y ordered first). E5Y also responds significantly (in the statistical sense) to a consumption innovation, but the point estimate is small and the response is statistically significant only for a few quarters.

Figure 4.5 graphically depicts the variance decompositions of consumption, income, and E5Y under both orthogonalizations. Regardless of ordering, own innovations account for the bulk of the forecast error variance of output at short horizons and virtually none at longer horizons. Ordered first, E5Y innovations account for more than 60 percent of the forecast error variance of income and consumption at long horizons. Even after orthogonalization with respect to consumption, innovations to E5Y still account for more than 30 percent of the long horizon forecast error variance of both income and consumption. We can thus fairly easily reject the hypothesis that E5Y simply reflects information available in consumption. Rather, innovations in E5Y and in consumption each convey news about future output that is not subsumed in the other.

We now examine several variations on the three variable VAR using alternative measures of consumer confidence. First, we substitute the relative score from the question on the Michigan Survey concerning expected personal financial situation (PFE) in place of E5Y. This question gauges expectations analogously to E5Y and E12M, although it specifically asks for expectations concerning personal situations as opposed to aggregate expectations.[9] The second modification is to use the Index of

[9] Dominitz and Manksi (2004) express doubt that consumers can give meaningful responses to survey questions concerning aggregate as opposed to individual expectations, and they point to the higher volatility of responses to questions like E5Y versus questions like PFE as support. Given the structure of the questions, however, we would in fact expect aggregate questions to have greater volatility even if individuals are equally capable of answering both kinds of questions accurately. For example, even in severe recessions most people do not personally experience layoffs. The typical respondent who says that the national economy will exhibit "periods of widespread unemployment or depression" is predicting that a significant minority of others will experience layoffs while his or

Consumer Sentiment (ICS) in place of the purely forward-looking survey questions. While the ICS is the most reported measure of consumer confidence (both by the press and in the academic literature), it is an average of survey responses to both forward-looking and retrospective questions, and thus its interpretation is unclear. For a more detailed description of these alternative confidence measures and their statistical relationships with E5Y, the interested reader is referred to Table 4.1.

Figure 4.6 presents impulse responses to confidence innovations in our three variable system with three alternative measures of confidence: E5Y, PFE, and ICS. We order the confidence measure first in the system, impose cointegration between consumption and output, and employ a lag structure of four.[10] There is very little qualitative or quantitative difference between the results using E5Y or any of the other broad confidence measures. The seeming disparity between some of our results and others in the academic literature thus does not appear to be attributable to different measures of confidence.[11] Use of other alternative confidence measures – such as E12M or the expectations index of the Michigan Survey – and alternative measures of consumption and output (for example, durable goods consumption or private sector GDP) also produce very similar impulse responses.

In summary, innovations in expectational variables from the Michigan Survey of Consumers are powerful predictors of changes in output and future spending that last for the foreseeable future. This finding obtains regardless of whether or not the confidence innovations are orthogonalized with respect to current spending. In Section 3 we will argue, based on model with both shocks to information and animal spirits shocks, that the apparent permanence of the impulse responses of consumption and output to confidence shocks is more consistent with an information view of confidence than it is with an animal spirits interpretation.

Our finding that unexpected increases in confidence imply predictably higher subsequent consumption is somewhat related to the results of Carroll, Fuhrer, and Wilcox

her own income is stable by comparison.

[10] Alternative orderings with the confidence measure after consumption also produce quite similar results.

[11] Among papers in this literature that find a small role for consumer confidence measures in predicting the future time path of economic variables are Mishkin (1978), Leeper (1992), Mehra and Martin (2003), and Croushore (2005). Matsusaka and Sbordonne (1995) and Howrey (2001) report a much stronger prognostic role for confidence, while Ludvigson (2004) takes something of a middle ground. Souleles (2004) analyzes the micro data underlying aggregate confidence data used in the present paper. However, the most important difference between our results and the results in these papers is that by looking at impulse responses to confidence innovations many periods into the future, we are able to recover the longer run implications of confidence innovations that are in fact more powerful that are the short run business cycle "effects".

(1994), who focus on one-period-ahead consumption growth. These authors regard Granger causality from confidence to consumption growth partly as a failure of the PIH along the lines of short-term stickiness of consumption.[12] This focuses excessively on the short run and reflects a decidedly partial equilibrium approach. Our finding that consumption tracks predictable income increases over periods of several years suggests that the predictability of consumption growth is better thought of in terms of an endowment economy along the line of Lucas (1978), in which consumers may believe that income will be higher in the future, but can in the aggregate do little to increase current consumption in anticipation. One implication of this interpretation is that positive confidence innovations should be associated with increases in expected real rates of return. This implication will be explored in more detail in the next section, and we will see that, in addition to being an implication of a simple general equilibrium model, it also holds in the data.

In the augmented consumption-income VAR, E5Y and other overall confidence measures are roughly exogenous. With E5Y ordered first, more than 95 percent of the forecast error variance of confidence is explained by its own innovation at every horizon. Even when confidence is allowed to respond contemporaneously to consumption innovations, the fraction of the forecast error variance of confidence attributable to its own innovation always exceeds 85 percent.

What kinds of news might explain these surprise movements in consumer confidence? The Michigan Survey of Consumers, in addition to the questions already discussed, also asks respondents to report any recent “news heard” concerning the economy. It seems natural to include a brief investigation of the relationship between this reported economic news and responses to the survey questions concerning overall expectations of aggregate and individual economic conditions. For a complete description of the news heard questions, see the Appendix.

Respondents give answers to a question asking them to report favorable or unfavorable economic news, and their answers are tabulated into arbitrary, but generally well-defined, categories. Figure 4.7 presents spike plots for several of the more popular response categories across time. Most categories (such as trade deficit, government budget deficit, etc.) record very few responses in a typical quarter. Rather clearly, the most consistently popular concern news about prices and news about employment. Other responses stand out in particular time periods. Examples are a high incidence

[12] The proposed solution is that increases in confidence measures summarize information possessed by “rule of thumb consumers” whose consumption is excessively tied to current income. The authors do reject that this hypothesis is a complete explanation of the Granger causality from confidence to consumption.

of mention of "energy crisis" during periods of the 1970s and early 1990s as well as news heard concerning the stock market sporadically across the sample period, but most frequently during the 1990s.[13]

In Table 4.1 we present coefficient estimates from regressions of the E5Y innovations from the three variable VAR on selected categories of news. Most of the news heard categories have coefficients of the expected signs – an increase in the percentage of respondents reporting favorable news is positively correlated with the confidence innovation and vice versa. Favorable or unfavorable news about general prices and favorable news about the stock market are significant covariates with the E5Y innovation at the 10 percent level or better. News about employment and favorable news about the stock market have no significant correlation with the E5Y innovation. Unfavorable news about government policies also has a statistically significant coefficient at the 10 percent level. The adjusted R^2 from these regressions ranges from 0.10 to 0.15, suggesting that the bulk of E5Y innovations remain inexplicable from particular categories of news heard. Use of other more obscure categories of news heard produce insignificant coefficient estimates that frequently reduce the adjusted R^2 in the regressions. We also ran a specification that included the news heard variables in the income-consumption VARs directly. This produced impulse responses of consumption and income which were much weaker than when using the broader confidence measures.

Innovations to measures of consumer confidence evidently convey information about income many periods into the future, much of which is not reflected in current consumption or income innovations, and the surprise movements in the confidence measures are not attributable to tangible news. Some might find it surprising that the answers of largely naïve respondents to rather crude questions could be so informative. As emphasized in Cochrane (1994b), however, such expressions of surprise fail to recognize the role of information aggregation. As Cochrane puts it (see p. 350), "Ask a consumer about next year's GDP, and he will say 'I don't know.' But he may know that his factory is closing, and hence he is consuming less. This idiosyncratic shock is correlated with future GDP." Just as consumption data aggregate idiosyncratic information, consumer confidence data aggregate information from many sources and many individuals.[14]

[13]The data summarizing responses to the "news heard" questions do not have the statistical properties of "news" in the rational expectations sense. Rather, the data on news reports are highly serially correlated. This may be due to gradual diffusion of news reports along the lines of Carroll's (2003) epidemiological model, or it may reflect merely the wording of the question, which refers to news heard in the "last several months".

[14]Some who accept the notion that intangible news could be responsible for large movements

3 Information and Animal Spirits in a New Keynesian Model

The results of the previous section suggest that survey measures of consumer confidence ought to be taken seriously. The observation that unexpected movements in confidence appear to have permanent implications for output and consumption seems inconsistent with an interpretation in which confidence innovations represent autonomous fluctuations in sentiment (i.e. animal spirits), but perhaps consistent with the notion that confidence reflects households' information about current and/or future fundamentals. To subject these statements to further scrutiny requires reference to a theoretical model that contains both fundamental and animal spirits shocks.

In this section, we develop a simple New Keynesian general equilibrium model with three structural disturbances. The two fundamental shocks are a level shock and what we call an information shock. The level shock is an immediate and permanent innovation to the level of technology, while the information shock is a persistent but transitory innovation to the growth rate of technology. We call it an information shock because it portends of a permanent change in technology orthogonal to the present. We only allow households to observe a noise-ridden signal of the growth rate of technology, and interpret a pure noise innovation as an animal spirits shock, as it is associated with erroneous consumer optimism or pessimism.

We then develop the implications of each of the structural shocks for the endogenous variables of the model. The level and information shocks are associated with permanent movements in measures of real activity, while the animal spirits shock is associated with transitory increases in spending. Guided by the theoretical impulse responses of the model, we take up a more rigorous analysis of the meaning of consumer confidence innovations in Section 4.

3.1 Model

3.1.1 Households

Households have standard preferences over consumption and leisure, live forever, and are identical. They consume a final consumption good, c, and supply labor, n, to

in confidence might nevertheless be surprised at the volatility of responses to questions like E5Y. Our claim is not that all of the movements in measured confidence reflect genuine information, but rather that whatever relationship obtains between confidence and subsequent income or consumption is likely to reflect information. Our methodology does not unveil the meaning of innovations in measured confidence associated with words alone and not actions.

intermediate goods producers. Each period, they choose consumption, labor supply, and holdings of a riskless one period bond so as to maximize expected discounted lifetime utility subject to a nominal budget constraint:

$$\max_{c,b,n} \quad \sum_{t=0}^{\infty} \beta^t E_0 \left(\ln(c_t - \alpha c_{t-1}) - \frac{n_t^{1+1/\eta}}{1+1/\eta} \right)$$

s.t.

$$p_t c_t + b_t \leq w_t n_t + (1 + i_{t-1}) b_{t-1} + \Pi_t$$

β is a subjective discount factor; α is the degree of internal habit persistence; η is the Frisch labor supply elasticity; p is the price of the final consumption good; w is the nominal wage; b is a riskless one period bond paying nominal interest i; and Π denotes any lump sum profits or transfers households might receive.

The first order conditions characterizing the solution to the household's optimization problem are:

$$MU(c_t) = E_t \left(MU(c_{t+1})(1 + i_t) \frac{p_t}{p_{t+1}} \right) \tag{1}$$

$$n_t^{1/\eta} = MU(c_t) \frac{w_t}{p_t} \tag{2}$$

Equation (1) is the intertemporal consumption Euler equation and equation (2) is the labor supply condition. The marginal utility of consumption depends positively on lagged and led consumption and negatively on current consumption:

$$MU(c_t) = \frac{1}{c_t - \alpha c_{t-1}} - \alpha \beta E_t \left(\frac{1}{c_{t+1} - \alpha c_t} \right)$$

3.1.2 Final Goods

The final good is a CES aggregate of a continuum of intermediate goods, indexed by j along the unit interval:

$$y_t = \left(\int_0^1 (y_{j,t})^{\frac{\xi-1}{\xi}} dj \right)^{\frac{\xi}{\xi-1}}$$

The parameter ξ has the interpretation as the price elasticity of demand for intermediate goods, and is assumed to be greater than unity. Similarly, the price index for final goods is given by:

$$p_t = \left(\int_0^1 (p_{j,t})^{1-\xi} \, dj \right)^{\frac{1}{1-\xi}}$$

The model has neither capital nor a storage technology, so all final output must be consumed each period:[15]

$$y_t = c_t \tag{3}$$

3.1.3 Intermediate Goods

Intermediate goods are produced according to a linear production function:

$$y_{j,t} = A_t n_{j,t} \tag{4}$$

A denotes technology, which is common and freely available to all intermediate goods firms. It can be shown that profit maximization in the final goods sector implies a downward-sloping demand curve for each intermediate good:

$$y_{j,t} = \left(\frac{p_{j,t}}{p_t} \right)^{-\xi} y_t \tag{5}$$

We assume that intermediate goods firms are not freely able to adjust prices each period. In particular, following Calvo (1983), firms face a constant hazard of being able to adjust their price in any period equal to $1 - \theta$. Whenever a firm gets an opportunity to adjust its price, it solves the following maximization problem:

$$p_t^* = \arg\max \quad \sum_{t=0}^{\infty} (\theta\beta)^t \, E_0 \frac{MU(c_t)}{p_t} (p_{j,t} - mc_{j,t}) \left(\left(\frac{p_{j,t}}{p_t} \right)^{-\xi} y_t \right)$$

p^* is the firm's optimal reset price and mc denotes nominal marginal cost. Marginal cost can be found in the firm's static labor demand condition:

$$w_t = mc_{j,t} A_t \tag{6}$$

The optimal reset price will be a present discounted value of expected nominal marginal costs:

[15] As is commonplace in the sticky price literature, we abstract from the presence of capital in the model. While the central lessons we draw are unaffected by this simplifying assumption, the presence of capital does matter in an essential way for certain aspects of the model. We address some of these issues in the next section.

$$p_t^* = \frac{\xi}{\xi - 1} \left(\frac{\sum_{t=0}^{\infty} (\theta\beta)^t MU(c_t) p_t^{\xi-1} y_t mc_{j,t}}{\sum_{t=0}^{\infty} (\theta\beta)^t MU(c_t) p_t^{\xi-1} y_t} \right) \tag{7}$$

Because all firms face the same wage and technology, expected marginal costs will be the same across firms, implying that all firms with the ability to update their price will choose the same reset price. The aggregate price level will thus evolve according to:

$$p_t = \left(\theta p_{t-1}^{1-\xi} + (1-\theta) p_t^{*1-\xi} \right)^{\frac{1}{1-\xi}} \tag{8}$$

3.1.4 Technology

We assume that log technology ($a = \ln A$) evolves according to a random walk with drift:

$$a_t = g_{t-1} + a_{t-1} + u_t \tag{9}$$

The random variable u represents a level shock – a permanent and immediate innovation to the level of technology, while g is a drift term that is itself stochastic. We assume that g obeys a stationary autoregressive process:

$$g_t = (1-\kappa)\overline{g} + \kappa g_{t-1} + e_t \tag{10}$$

Where $\kappa < 1$ and $\overline{g}$ denotes the steady state growth rate. e (which is assumed orthogonal to u) is a growth shock – a persistent but stationary innovation to the growth rate of technology, heralding periods of above or below average growth. We call e an information shock because it portends changes in future levels of technology. It is simply a smooth version of the "news shocks" studied by Beaudry and Portier (2004) and Jaimovich and Rebelo (2006). Because of the assumed nominal rigidities in the model, there is an avenue here for output to expand upon the arrival of good news about the future and our model is not subject to the "bust" feature of neoclassical models in which agents receive advance signals about future technology.[16]

[16] By "bust" feature we are referring to the tendency of neoclassical models to yield output declines in response to good news of the future (see Beaudry and Portier (2004) or Jaimovich and Rebelo (2006)). This is because in these models output is completely supply determined, and, in the standard framework, there is no explicitly dynamic dimension to the firm's problem. As such, the wealth effect on the household side of the model usually induces a decrease in labor supply. Coupled with no change in labor demand, this results in a reduction of output in equilibrium. We do not have this problem because the price stickiness allows the "demand" effect of news about the future to work in the right direction. In particular, the increased desire to consume induces an (undesired)

3.1.5 Perceptions and Animal Spirits

While households observe the level of technology at each point in time, we assume that they never explicitly observe level shocks to technology, u, and observe only a noisy signal of growth rate shocks, e. The signal they receive is equal to:

$$s_t = e_t + v_t \tag{11}$$

v is a mean zero white noise disturbance uncorrelated with both growth and level shocks.

The setup described above implies that households imperfectly observe the drift term. We posit that they update their perceptions according to a simple linear filter:

$$g_t^p = \kappa(1 - \Omega_1)g_{t-1}^p + \kappa\Omega_1(a_t - a_{t-1}) + \Omega_2 s_t \tag{12}$$

κ is the autoregressive coefficient from equation (10), and the coefficients Ω_1 and Ω_2 are functions of the variances of the shocks in the economy. In particular:

$$\Omega_1 = \frac{\sigma_v^2}{\sigma_v^2 + \sigma_u^2} \quad \text{and} \quad \Omega_2 = \frac{\sigma_e^2}{\sigma_e^2 + \sigma_v^2}$$

To see why these coefficients look the way they do, it is helpful to consider a couple of extreme cases. If $\sigma_v^2 = 0$ (i.e. there is no noise in the signal concerning the growth rate shock) then $\Omega_2 = 1$, $\Omega_1 = 0$, and the perceived drift term is equal to the truth at all times. If $\sigma_v^2 = 0$ (i.e. there are no shocks to the current level of technology), then $\Omega_1 = 1$. In this case agents will be uncertain about the growth rate between today and tomorrow due to the noise in the signal, but the realization of technology tomorrow will reveal perfectly to them today's actual growth rate shock, so that there will be no endogenous persistence of a false signal for more than one period. Intermediate cases are more interesting. As the variance of the noise term in the signal grows, Ω_2 becomes smaller and Ω_1 gets bigger – people will place little weight on a very noisy signal but will place a lot of weight on the realization of actual technology growth relative to their previous period's perception in updating their current belief. As σ_u^2 gets bigger, Ω_1 becomes smaller, so household perceptions about the technology drift term will be more persistent. Intuitively, a very high variance of level technology shocks means that a realization of technology growth different from what was expected is less likely to mean that the original perception of the persistent growth rate was wrong, and

---

reduction in firm markups, which leads to an increase in labor demand, thus allowing employment and output to expand in anticipation of the realization of good news.

more likely that there was simply an offsetting level shock.

While we assume that households observe the drift term in technology with imprecision, we allow firms to view both level and information shocks without noise. Although it seems both intuitive and realistic that firms have superior information relative to individuals, this setup is incompatible with the usual structure in which firms are owned by households. To avoid this complication, we can simply assume that management is separated from ownership, with managers risk neutral agents with the sole objective to maximize profits.

The disparate information to which households and firms are privy presents an additional complication. Even though households are unable to immediately differentiate between legitimate news about the drift term and pure noise, the equilibrium effects of noise and genuine information shocks on the endogenous variables of the model will be different, owing to the fact there is a shock to the supply side of the model in the case of a true growth rate shock, whereas there is not in response to a noise shock to the households' signal. Therefore, the linear filter given by equation (11) is not the optimal filter when firms have better information than households. In particular, the optimal filter would include information revealed to households through wages, interest rates, and prices, whose equilibrium behavior would reveal to them the underlying nature of the signal. We simply assume this complication away. The filtering specification in (11), though not fully optimal, is both intuitive and simple, and household perceptions converge to the truth in the long run.[17]

We will interpret a noisy innovation to the households' signal of the drift term as an animal spirits shock. A positive v means that households erroneously believe that the future will be better. Given this belief, they will desire to consume more immediately. Because firms do not share this belief, there is no shock on the supply side of the model. In this way, our animal spirits shock is a pure demand shock, and is similar to the kind of shock studied in Lorenzoni (2008).[18] The animal spirits shock will play a role in equilibrium nearly identical to a preference shock manifesting itself

[17]Under a fully optimal filter, there would be no endogenous persistence of the noise shock on households' perceptions of the drift term. This is because the general equilibrium behavior of the endogenous variables would immediately reveal to households the true nature of the signal. A more complicated version of this model which would preserve the endogenous persistence of these noise shocks under a fully optimal filter would introduce an additional shock, unobservable to households, into the model (e.g. a markup shock or a monetary policy shock).

[18]Conceptually, the only fundamental difference between Lorenzoni's specification and ours is that in his paper agents receive a noisy signal about the current level of technology, whereas in our framework the noisy signal concerns future levels of technology. His specification gives animal spirits a better shot at inducing significant fluctuations, and we address this possibility in further detail below.

as an exogenous innovation in the consumption Euler equation. As such, one could give this disturbance an alternative, less structural interpretation as a taste or rate of time preference shock.[19]

3.1.6 Monetary Policy Rule

We close the model with a nominal interest rate rule. In particular, we postulate that the central bank sets nominal interest rates according to a partial adjustment mechanism where the interest rate in any period is equal to a convex combination of the lagged interest rate and the central bank's target rate. The target rate is adjusted in response to deviations of output growth and inflation from constant targets.

$$i_t = \rho i_{t-1} + (1-\rho)\left(\phi_y\left(y_t - y_{t-1} - \Delta y^*\right) + \phi_\pi\left(\pi_t - \pi^*\right)\right) \tag{13}$$

The parameter ρ captures the degree of interest rate smoothing. We abstract from the presence of monetary disturbances, so we show no error term.

Our specification of the policy rule differs slightly from the ubiquitous Taylor rule (1993) in which the nominal rate is adjusted in response to inflation and the output gap. We prefer our specification for two reasons. First, the informational requirements imposed on the central bank are much more reasonable when assuming that it responds to output growth relative to its long term trend as opposed to its "potential", which is itself time-varying. Secondly, a rule such as this is capable of matching certain features of the data which a standard Taylor rule is not. In particular, we know from other work that information shocks about future productivity appear to be strongly deflationary (Chapter II). In general equilibrium, predictable increases in output and consumption must be associated with rising real interest rates. It is extremely difficult to simultaneously generate a large increase in real interest rates and a large disinflation under a standard Taylor rule. A rule in which the bank reacts to output growth as opposed to the gap is capable of matching the data along this dimension, and is thus the one which we adopt here.

A policy rule similar to (13) is also not without precedent in the literature. In particular, a number of recent papers make use of very similar rules – for example, Coibion and Gorodnichenko (2007), Fernandez-Villaverde and Rubio-Ramirez (2007), and Ireland (2004). In particular, our exclusion of a theoretical output gap from the policy rule is consistent with Ireland's (2004) finding that the coefficient on the gap

[19] Some might prefer to think of an animal spirits shock this way in the first place. Irrational exuberance, for instance, could be interpreted as an emphasis on the enjoyment of current consumption at the implicit expense of future consumption.

in an estimated rule does not differ significantly from zero. Orphanides (2003) argues that a rule responding to output growth provides at least as good a description of actual US monetary policy over the last thirty years as does the more standard formulation in which the central bank responds to an output gap. We will discuss the implications of alternative monetary policy rules for our results below.

3.2 Theoretical Impulse Response to Shocks

We solve the model by log-linearizing the equations above about the non-stochastic balanced growth path. Solving the model requires picking values of the structural parameters. We assume the following: $\beta = 0.995$ (with the interpretation of the unit of time as one quarter, this corresponds to an annual discount rate of roughly two percent), $\alpha = 0.5$, $\eta = 1.0$, $\theta = 0.66$ (meaning that firms get to update their prices on average once every three quarters), $\rho = 0.75$, $\phi_y = 2.5$, $\phi_\pi = 4.5$, $\kappa = 0.85$, $\sigma_u = 1$, $\sigma_e = 0.125$, and $\sigma_v = 0.125$. Because of the assumed high degree of persistence to information shocks ($\kappa = 0.85$), it is necessary that the standard deviation of shocks to the drift term be small relative to that of level shocks in order to generate data in which actual productivity growth is approximately white noise, which appears to be the case in the US. We calibrate the variance of the animal spirits shock so that agents place a fairly high weight (in this case $\Omega_2 = 0.5$) on the signal in updating their perceptions of the drift term. The choice of the habit persistence term is similar to the estimates in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2003). The calibration of the labor supply elasticity is in the middle of the range of estimates from micro studies (which are typically small) and those in the business cycle literature (which are usually much higher than unity), and is equal to the central point estimate in Kimball and Shapiro (2003). Our calibration of the parameters of the monetary policy reaction function is in line with empirical estimates from similarly specified rules (Coibion and Gorodnichenko (2007), Ireland (2004), and Paciello (2008)).

Figure 4.8 shows the theoretical responses of output, technology, hours, inflation, and real interest rates to the level and information shocks. The sizes of the shocks are chosen so that each leads to an ultimate increase in technology of one percent. For output, technology, and hours, the figures show the percentage response relative to the initial non-stochastic steady state. For inflation and the real interest rate, the figures show the annualized percentage point response (for example, a response of inflation of -0.2 means that inflation falls from, say, 4.0 percent to 3.8 percent at an annualized rate, and similarly for the real interest rate).

By construction, the level shock leads to an immediate jump in technology that is expected to remain forever at the new higher level, whereas the information shock is orthogonal to current technology but portends a sustained period of smooth growth. In response to both the level and information shocks, output jumps on impact and is expected to rise towards its new steady state value. Quite naturally, the impact jump in output is smaller for the information shock than for the level shock. Relative to the perfect information version of the model (where the variance of the animal spirits shock is zero), output overshoots in response to a level technology shock and undershoots in response to an information shock. The intuition for these effects is clear. When agents receive a signal that productivity growth will be higher, they place some weight on the possibility that the signal is purely noise, and so they react less than if they knew the shock with certainty. Likewise, when the level of technology jumps up unexpectedly, households place some weight on the possibility that there was an unseen growth rate shock at some point in the past that was buried in noise. They therefore place some weight on the possibility that productivity growth will be higher in the near future, and so react more than if they knew that it were only a one time level shock.

Employment rises on impact in response to the growth rate shock, while it falls on impact following the level shock. The fall in hours in response to the level shock is the well-known "contractionary technology shock" due to the undesired increase in firm markups following immediate technological improvement.[20] The response of hours to any shock in the model is constrained to be transitory because household preferences are in the class of preferences described by King, Plosser, and Rebelo (1988) consistent with balanced growth. Both kinds of shocks are associated with a rising real interest rate, which is to be expected, as both shocks make the future plentiful relative to the present.

Both the level and information shocks are disinflationary in the model, with the magnitude of the disinflation smaller but more persistent for the information shock than the level shock. A different behavior of prices could be rationalized under a different policy rule. In a standard Taylor rule where rates are adjusted in response to inflation and the deviation of output from the theoretical gap, for example, inflation would be almost completely stabilized in response to both shocks and would in fact rise slightly in response to the information shock. An exogenous time path for the money supply with a quantity type money demand equation, on the other hand, would

[20] See, for example, Basu, Fernald, and Kimball (2006).

produce a behavior of prices quite similar to what is shown here. As such, there is no robust implication of the model for prices in response to the two fundamental shocks.

Figure 4.9 shows the responses to the animal spirits shock in the model. The size of shock is chosen so that it is the same as the information shock (i.e. both shocks raise the signal by an amount prognostic of an ultimate increase in the level of technology of one percent). By construction, the shock never has any effect on the actual level of technology. The animal spirits shock is differentiated from the level or information shocks in that it is associated with a transitory response of output and rising prices. All three kinds of shocks raise the real interest rate.

For the given calibration of the parameters of the model, the impact of animal spirits on output is small (the maximal output effect is roughly 0.06 percent), though the effects are fairly persistent. The reason for the weak response of output is straightforward – this response is essentially the "aggregate demand" effect of an information shock (i.e. the increase in output coming from the consumer side of the model following a good signal). The response to the true information shock combines this demand effect with the "aggregate supply" effect, which is positive given the forward-looking nature of the Phillips Curve. The impact effect of animal spirits is thus bounded from above by the impact effect of a true information shock, which is itself modest.

The animal spirits shock also leads to a fairly significant increase in inflation. In the New Keynesian model, inflation is equal to a present discounted value of real marginal costs. Real marginal costs should be weakly higher at all horizons following an animal spirits shock, and inflation should thus rise. The intuition for this effect is straightforward. Real marginal cost in the model is equal to the log difference between real wages and technology. Following a positive animal spirits disturbance, households feel wealthier and thus demand a higher real wage for a given level of employment. Since there is never any effect on actual or perceived technology on the firm side of the model, the wealth effect on the household side dictates that real marginal costs are always weakly higher following a noise innovation to the household signal. As such, the model has the implication that animal spirits shocks (as specified) are inflationary.

There are alternative calibrations of the model's parameters which yield more significant responses to the animal spirits shock. In particular, the impact effects of animal spirits are larger for very low values of κ. The intuition for this is straightforward. When κ is small, most of the expected improvement in productivity occurs sooner as opposed to later, resulting in a larger innovation to perceived permanent income, which in turn leads to a larger effect on overall aggregate demand. We will allow the data to inform us on the value of this parameter in the next section.

4 Expanded Analysis and a Structural Model

4.1 Reduced Form Empirical VAR

To begin to assess the relative importance of the structural shocks of the theoretical model in the determination of consumer confidence innovations, we first estimate a VAR with E5Y, annualized CPI inflation, the three month Treasury Bill rate, the BLS measure of aggregate per capita hours in the non-farm business sector, real non-durables plus services consumption per capita, and real GDP per capita.[21] Aside from the fact that we include separate measures of output and consumption, the variables in this empirical VAR coincide with those in the theoretical model of Section 3.[22] The real interest rate is implicitly defined as the nominal three month Treasury Bill rate less the VAR forecast of one quarter ahead inflation. As in the empirical VARs of Section II, the data are quarterly from 1960:01 – 2007:03, we choose a lag order of four, and we impose cointegration between output and consumption.[23] Labor hours, E5Y, the interest rate, and inflation enter in the VAR in levels.[24]

Figure 4.10 presents the impulse responses of the variables in the empirical VAR to an E5Y innovation (ordered first in a block recursive system). As before, the dashed lines represent 90 percent confidence bands from a bias-corrected bootstrap procedure. Consumption and output both jump slightly on impact in response to an E5Y innovation, but thereafter continually rise, with no tendency to attenuate.

[21] Earlier versions of this paper reported results with the civilian unemployment rate in place of hours, and expressed some concern about reverting impulse responses of consumption and income to E5Y innovations. This reversion, which apparently depends on a marginally significant coefficient implying that higher unemployment is associated with higher consumer confidence, largely disappears as one increases the lag length. Nevertheless, that reversion is nonetheless statistically indistinguishable from a permanent response, and, at any rate, the responses of output and consumption to E5Y innovations with unemployment in the VAR (even with very low lag lengths) are far too persistent to be taken as positive evidence in favor of an important animal spirits component.

[22] The theoretical model without capital does not distinguish between consumption and output. We include consumption as well as GDP in the empirical VARs of this section so as to facilitate comparison with the results of Section 2. The empirical results are unaffected by using either consumption or output in isolation.

[23] As before, we impose that the cointegrating vector between consumption and output is [1,-1]. The imposed lag order of four is somewhat higher than the choices of a variety of widely accepted information criteria (which, on average, favor two lags). Alternative lag structures produce nearly identical results. A VAR with all variables entering in levels yields nearly identical impulse responses.

[24] There is a large debate over whether labor hours are $I(1)$ or $I(0)$ (see, for example, Christiano, Eichenbaum, and Vigfusson (2004)). Our results are qualitatively similar whether hours enter the VAR in levels, first differences, or as deviations from a trend.

The point estimates suggest that a one standard deviation to E5Y is prognostic of consumption and output that are higher in the long run by roughly 0.67 percent. Even at a horizon of forty quarters, these responses are statistically different from zero at better than the 90 percent level. The E5Y innovation is associated with transitory and significant increases in both real interest rates and hours of work, with both responses following hump-shaped patterns. Inflation falls by roughly one quarter of a percentage point on impact and is persistently below its initial value for a number of quarters. Though there is no significant impact effect of E5Y on measured labor productivity (imputed within the VAR as the output response less the hours response), an E5Y innovation is prognostic of a permanent increase in productivity of more than two thirds of a percent, with the long run response statistically different from zero at horizons in excess of forty quarters.

A cursory comparison of the responses in Figures 4.8 and 4.10 reveals that the empirical responses to a confidence innovation look similar to the theoretical responses to what we have deemed an information shock in our model. In particular, a positive innovation to E5Y is associated with a prolonged and permanent increase in real activity, a transitory rise in both real rates and hours of work, and a strong and persistent disinflation. These are roughly the qualitative predictions of the model in response to a favorable information shock. It therefore seems natural to associate innovations in consumer confidence with information shocks – in particular, persistent shocks to productivity growth. Does this observation mean that there is no role for animal spirits, and no noise in measured confidence? In the next subsection we specify and estimate a variance components model of confidence innovations that allows us to address these questions.

4.2 A Model of Consumer Confidence

In the context of the theoretical model of the previous section, we assume that a measure of consumer confidence follows a stationary autoregressive process, with its innovation a linear combination of the unexpected change in the current state of the economy and the signal concerning the persistent growth term:

$$CC_t = \delta CC_{t-1} + \lambda_1 \left(a_t - E_{t-1} a_t\right) + \lambda_2 s_t \tag{14}$$

The conditional relationships between confidence innovations and the other variables of the model will depend both on the λs and the variances of the structural shocks. For instance, $\lambda_1 > 0$ and $\lambda_2 = 0$ would mean that that the confidence innovation is purely a reflection of the change in the current state of the economy, while $\lambda_2 > 0$ would mean that innovations to confidence at least partially reflect signals households receive about the future. If λ_2 is relatively large and the signal, s, is not very noisy, the relationships between confidence innovations and macroeconomic variables in the model will look similar to the theoretical responses to an information shock. On the other hand, if λ_2 is large and the signal is quite noisy, then confidence innovations may generate patterns similar to the theoretical responses to an animal spirits shock.

As written, consumer confidence responds only to structural shocks in the economy. One should not take this description too literally. No one would wish to maintain that all of the variation in the observed confidence data reflects genuine information (or even changes in beliefs, however formed) – there is always sampling error, misunderstandings on the part of respondents, etc. As such, any realistic specification of confidence should include some kind of measurement error. We abstract from the presence of pure measurement noise here simply because, using our empirical method detailed below, it is not possible to separately identify its empirical properties.

We estimate the parameters of the confidence equation using a modified version of the simulated method of moments. In particular, we would like to know what parameter configuration is most likely to generate data yielding impulse responses similar to what we see in the actual data. As such, our SMM estimator tries to match impulse responses to confidence innovations from simulated data from the model to the impulse responses from the actual data. This approach is similar to that in Christiano, Eichenbaum, and Evans (2005).[25]

As our main focus is only on the parameters directly influencing consumer confidence, we first calibrate many of the other parameters of the model. In particular, we set the preference parameters, the parameter governing price stickiness, and the policy coefficients as in the calibration of the previous section, and we set the autoregressive coefficient in the confidence equation at 0.8.[26] We normalize the variance of level technology shocks to be unity. The remaining parameters to be estimated are given by the vector $\mathbf{\Theta} = [\lambda_1 \quad \lambda_2 \quad \kappa \quad \sigma_e \quad \sigma_v]'$.

For a given guess of the parameter vector $\mathbf{\Theta}$, we simulate a data set of length τT,

[25] Our approach is similar to that of CEE in that we are choosing parameter values so as to match impulse responses as opposed to unconditional moments. CEE differ slightly from us in that their objective function is to match theoretical impulse responses from their model to those in the data, whereas we match impulse responses from VARs estimated on simulated data from our model.

[26] We experimented with many different values for the other parameters of the model. Alternative calibrations of these parameters have little noticeable impact on our results. $\delta = 0.8$ is approximately the estimated autoregressive coefficient for E5Y in the data.

where T is the length of the actual data set and $\tau = 10$ (the shocks are drawn from normal distributions). After discarding the first 100 observations from the simulated data set (so as to limit the influence of arbitrary starting values), we estimate a five variable VAR similar to the one in subsection (a) with simulated confidence, output growth, inflation, hours of work, and the interest rate and compute impulse responses to the confidence innovation (ordered first in the VAR).[27] Letting $\mathbf{m}(\mathbf{\Theta})$ denote the stacked vector of impulse responses for the given guess of $\mathbf{\Theta}$ and $\mathbf{m}^*$ the corresponding stacked vector of impulse responses from the data, we iterate on our guess of $\mathbf{\Theta}$ so as to minimize the following:[28]

$$\mathbf{\Theta}^* = \arg\min \quad (\mathbf{m}(\mathbf{\Theta}) - \mathbf{m}^*)' \, \mathbf{W} \, (\mathbf{m}(\mathbf{\Theta}) - \mathbf{m}^*)$$

As in Christiano, Eichenbaum, and Evans (2005), we set $\mathbf{W} = \mathbf{V}^{-1}$, where $\mathbf{V}$ is a diagonal matrix with elements equal to the variances of the impulse responses from the data. The optimal parameter vector is then that which minimizes the weighted sum of squared deviations between the estimated impulse responses on model simulated data and the corresponding responses from the actual data. This choice of weighting matrix places more weight on those responses which are most precisely estimated in the data.

The estimates of the parameters of interest and corresponding confidence bands are in Table 4.2. These coefficients are of the expected signs, and with the exception of the coefficient on the unexpected change in the current state of the economy, are all different from zero at conventional levels of significance. In order to provide some interpretation to these quantitative estimates, the total variance of the structural confidence innovation is seen to be:

[27] As laid out above, our model only has three structural shocks, meaning that any combination of three or more simulated series from the model would be perfectly collinear – i.e. the model suffers from "stochastic singularity". So as to be able to estimate a VAR with more than three variables, we introduce three additional shocks into the model for the purposes of estimation. In particular, we introduce a shock to the monetary policy rule, a "cost-push" shock in the Phillips Curve, and a preference shock in the Euler equation. The model E5Y innovation is unrelated to these disturbances, so they should not (in large enough samples) impact the conditional correlations between confidence and the other variables. Rather, the only role of these shocks is to ensure that five variables in the model are not perfectly collinear. The estimated VAR here is identical to the empirical VAR of subsection (a), except that the model does not differentiate between income and consumption.

[28] The impulse responses making up our objective function include the impact effect on confidence itself, the impulse response of the level of output over ten years, and the responses of hours, interest rates, and inflation over five years. We also include the autocorrelation of productivity growth in the objective function. The observed first order correlation of measured labor productivity growth in our sample is roughly 0.045.

$$var\left(e_{cc}\right)=\lambda_1^2 var\left(a_t-E_{t-1}a_{t-1}\right)+\lambda_2^2\left(\sigma_e^2+\sigma_v^2\right)$$

Given the low estimate of λ_1, we see that the unexpected change in the current state of the economy evidently accounts for less than one half of one percent of the innovation variance of confidence in the model. Information and animal spirits shocks each account for roughly one half of the innovation variance in measured consumer confidence, with the noise disturbance mattering slightly more.

Figure 4.11 presents the average impulse responses of output, inflation, hours, the real interest rate, and confidence to a confidence innovation from simulated data using these parameter estimates. In particular, we simulated 2000 sets of data with 200 observations each based on these parameters, with the structural shocks drawn from a normal distribution. For each simulated data set we estimated a five variable VAR with confidence, output growth, hours, interest rates, and inflation, with four lags of each variable. For output, interest rates, and inflation, these responses look similar to those from the empirical VAR depicted in Figure 4.10. In the simulations, the average E5Y innovation is associated with a small impact effect on output followed by a sustained period of growth, a significant and persistent disinflation, and persistently high real interest rates.

The dimension along which the model is least successful in matching the empirical impulse responses is in the response of hours. In the data confidence innovations are associated with a small but reasonably persistent increase in hours. While the model produces data roughly matching the impact response of hours to an E5Y innovation, it fails to match the persistence. After the small positive impact effect and a few quarters of being above trend, the response of hours is slightly negative for a number of quarters. The intuition for this response is reasonably straightforward. A significant portion of the E5Y innovation is accounted for by true information shocks, which begin to exert a contractionary effect on employment in the model once the technological improvement starts to take hold.

Two seemingly contradictory conclusions emerge from our structural estimation results. On the one hand, animal spirits disturbances seem to account for an important portion of the structural confidence innovation. On the other hand, the model responses to a confidence innovation look very much like the theoretical responses to an information shock, and nothing at all like the responses to an animal spirits disturbance. The resolution of this apparent contradiction is that information shocks have implications for the other variables of the model which simply dwarf those of animal spirits shocks. As such, the conditional correlations between confidence innovations and the other variables of the model are dominated by the information shock, even though animal spirits shocks account for a significant component of the structural confidence innovation.

Nevertheless, our results do allow us to reject the hypothesis that animal spirits account for a significant portion of the observed relationship between consumer confidence and macroeconomic variables. Given the estimated persistence of information shocks, the implications of an animal spirits shock for the variables of the model are very small (see the discussion earlier or the theoretical impulse responses in Figure 4.9). As such, the animal spirits shock is difficult to differentiate from pure measurement noise in data generated from the model. The wide confidence bands on the estimate of the standard deviation of animal spirits shocks shown above confirm this and seem to suggest that this parameter is probably poorly identified. Forcing the variance of animal spirits shocks to zero and re-estimating the model leads to little noticeable difference in the estimates of the other parameters or in the overall fit of the model. Eliminating information shocks from the model and re-estimating, however, leads to a much poorer overall fit.

While confidence innovations evidently reflect both information and animal spirits (which are in practice difficult to differentiate from pure measurement error), our results suggest that the relationships between confidence and macroeconomic variables are largely driven by information about future fundamentals. The impulse responses in Figure 4.12 make this point perhaps even more clear. These responses are from a bivariate VAR featuring the growth rate of a utilization-corrected measure of aggregate total factor productivity (TFP) and E5Y.[29] This exercise is similar to the stock price-TFP VARs in Beaudry and Portier (2006). Confidence is ordered second, so that the structuralized innovation in E5Y is contemporaneously orthogonal to TFP. Two observations stand out. First, the confidence innovation orthogonal to TFP predicts a permanent increase in TFP of roughly 0.7 percent, with this effect highly significant even at very long horizons. In quantitative terms, this long run response of TFP is about the same magnitude as TFP's response to its own innovation, which looks very much like a pure random walk. The TFP response to E5Y is both smooth and prolonged, and looks similar to the theoretical response to a growth rate shock discussed in the previous section. Secondly, consumer confidence does not respond significantly (either statistically or economically) to the TFP innovation. Both of these findings corroborate our estimation results above, which did not make explicit

[29]We are grateful to John Fernald for providing us with this measure. The responses show the level response of TFP, which is simply the cumulated growth rate response.

use of any TFP measure. In particular, there appear to be information shocks about the future which may account for significant component of long run productivity. These information shocks appear to be reflected in consumer confidence innovations, which are evidently unrelated to contemporaneous productivity shocks.

4.3 Information Shocks and Business Cycle Fluctuations

We close the body of the paper with a brief comment on the differences and similarities between our model and those of Beaudry and Portier (2004) and Jaimovich and Rebelo (2006). One difference is that their models are laid out in neoclassical settings, whereas we assume the presence of nominal rigidities. The nominal rigidities turn out to be important, for they introduce an explicitly forward-looking dimension into the firm's profit maximization problem which provides an avenue by which output can expand in response to good news about future productivity. The other primary difference is that we abstract from the presence of capital. As both of the above papers make clear, it is difficult to generate a simultaneous increase in consumption, output, and investment in response to favorable news about future productivity. In a pure neoclassical model, there is no explicitly dynamic dimension to the firm's problem. As such, the wealth effect of higher future productivity usually induces a reduction in labor supply and an increase in consumption, the combined implications of which are reduced output and investment.

As our goal has been to study the meaning of surprise movements in consumer confidence – and not the requisite model structures needed to deliver broad-based co-movement following information shocks – we have made the deliberate decision to sidestep the issue of co-movement by abstracting from endogenous capital accumulation altogether. While not without loss of generality, this decision has enabled us to elucidate the apparent necessity of shocks to future fundamentals orthogonal to the present in generating confidence data consistent with what we see in the world. That we have found an apparent empirical counterpart to the kinds of information shocks studied by other authors suggests that further study of more sophisticated models with these kinds of shocks – as well as the model features which will produce positive co-movement – is likely to be a fruitful avenue for future research.

5 Conclusion

While many in the popular press and business community regard measures of consumer confidence as essential in understanding the evolution of the aggregate economy, economists have devoted little attention to the economic interpretation of variation in measured confidence. Most of the scant academic literature focuses on the extent to which confidence measures help to improve forecasts of spending and output over relatively short horizons. While related to that line of research, this paper goes further in attempting to ascertain the underlying meaning of surprise movements in confidence.

We began our inquiry with an analysis of simple consumption-income VARs augmented with forward-looking measures of confidence. As noted by Cochrane (1994a), innovations to consumption are powerful predictors of subsequent movements in income. We demonstrated that measures of consumer confidence play a role similar to that of consumption innovations in that they foretell important movements in future output. Even after orthogonalization with respect to consumption, confidence innovations remain prognostic of significant movements in output and spending, especially at longer horizons.

We then turned more formally to the question of what economic concept underlies surprise movements in confidence, beginning with two polar hypotheses. The first – which we deemed "animal spirits" – posits that surprise movements in measured confidence proxy for exogenous changes in sentiment, which in turn have causal effects on aggregate demand. Such an interpretation of confidence was given by Blanchard (1993) in a paper on the causes of the 1990-1991 recession. The second hypothesis – the "information view" – supposes that there exists no causal relationship from confidence to economic activity, but rather that measured confidence reflects aggregated information individuals possess regarding present and future economic fundamentals.

We developed a New Keynesian model incorporating both animal spirits and fundamental shocks. The animal spirits shock is a manifestation of overly optimistic or pessimistic perceptions on the part of households, and leads them to desire more or less consumption than is optimal under perfect information. The two fundamental shocks in the model are a current level shock and an anticipated growth rate shock to productivity. In general equilibrium, the animal spirits disturbance plays the role of an aggregate demand shock – it is associated with transitory movements in spending and with higher inflation. Both fundamental shocks, on the other hand, are likely to be disinflationary and are associated with movements in spending that are not subsequently reversed. The information shock is distinguished by a small initial response of output followed by a prolonged period of growth.

We estimated an empirical VAR including the variables in the model as well as a measure of consumer confidence. Income and consumption appear to respond permanently to a confidence innovation, with a positive innovation to confidence associated with income, consumption, and labor productivity that are appreciably higher in the long run. The implications of a confidence innovation for output, spending, and productivity are much larger at longer horizons than at shorter ones, and positive confidence innovations are associated with a strong and persistent disinflation. In light of the theoretical model, the impulse responses from the empirical VAR provide essentially no support for animal spirits and point strongly to the information interpretation of confidence.

After positing a structural equation for consumer confidence, we then turned to a more formal estimation of the parameters of the model. We can resoundingly reject the hypothesis that animal spirits shocks (as specified in this paper) can account for the bulk of the relationships between consumer confidence and macroeconomic variables. If ever one hoped to find empirical support for animal spirits like shocks, surely it would be found in survey responses of seemingly naïve consumers. That we are unable to find compelling evidence in support of the animal spirits hypothesis thus casts expectations-driven theories of demand shocks such as Lorenzoni (2008) into serious doubt. On the other hand, we do find convincing evidence in favor of the information interpretation of consumer confidence. The implications of confidence innovations for output and spending at short horizons are far too small for confidence to be primarily a reflection of changes in current fundamentals, yet the longer horizon implications are far too large and significant for confidence innovations to not be conveying information about fundamentals. Putting the two together, it would appear as though confidence innovations are likely conveying information about future fundamentals, and in particular long run productivity. A bivariate TFP-confidence VAR seems to lend credence to this conclusion.

A recent line of research studies the extent to which news about future fundamentals can drive the business cycle (Beaudry and Portier (2004, 2006) and Jaimovich and Rebelo (2006)). Our results provide empirical support for the notion that agents do receive advance signals about future fundamentals, but they do not yet indicate that such information shocks play a pivotal role in short run fluctuations. Our ongoing research builds on the results of this paper and further addresses the business cycle implications of information shocks.

6 Appendix

This Appendix details the survey questions underlying the confidence data used in this paper.

Questions:

E5Y: 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?

E12M: Now turning to business conditions in the country as a whole – do you think that during the next twelve months we'll have good times financially or bad times or what?

PFE: Now looking ahead – do you think that a year from now you (and your family living there) will be better off financially, worse off, or just about the same as now?

News Heard: During the last few months, have your heard of any favorable or unfavorable changes in business conditions?

Answer Choices and Variable Construction: For most questions (including E5Y, E12M, and PFE), individuals are given three answer choices that amount to "favorable", "neutral" or "don't know", and "unfavorable". The "relative score" – the variable we use in this paper – is then constructed as the percentage giving a favorable response less the percentage giving an unfavorable response plus one hundred.

> Thus, a relative score of 100 indicates that an equal number of people gave a favorable response as an unfavorable response. If 30 percent of respondents give a favorable response and 20 percent given an unfavorable response, with the remaining 50 percent either "neutral" or "don't know", then the relative score will be 110 (i.e. 30 – 20 + 100).
>
> If, out of 100 people, 1 person switches from an unfavorable response to a neutral response, the index score will go up by 1. If that person switches from unfavorable to favorable, the index score goes up by 2. If someone leaves the state of "neutral" to either "favorable" or "unfavorable" the index score moves up or down by 1.

The Index of Consumer Expectations (ICE) is constructed based on the relative scores for PFE, E12M, and E5Y as follows:

$$ICE = \frac{PFE + E12M + E5Y}{4.1134} + 2.0$$

The Index of Consumer Sentiment (ICS) is constructed based on the relative scores for the PFE, E12M, and E5Y, plus two other questions. The first we'll call PFP and is similar to PFE, except that it asks respondents to make a comparison of their current financial situation relative to one year ago. The second we'll call DUR and it asks respondents whether or not it is currently a good time to buy "large household items" (i.e. durable goods). The ICS is constructed as:

$$ICS = \frac{PFE + E12M + E5Y + DUR + PFP}{6.7558} + 2.0$$

Table 4.1

Regressions of Confidence Innovations on News Heard Categories

News Heard Category	Coefficient		
Favorable Employment	0.248**	0.113	0.140
	(0.13)	(0.13)	(0.13)
Favorable Prices	1.001**	0.889*	1.005*
	(0.51)	(0.51)	(0.58)
Unfavorable Employment	-0.064	-0.071	0.035
	(0.05)	(0.05)	(0.06)
Unfavorable Prices	-0.363***	-0.342***	-0.312***
	(0.13)	(0.13)	(0.15)
Favorable Stocks		0.915**	0.845**
		(0.38)	(0.38)
Unfavorable Stocks		-0.235	-0.259
		(0.16)	(0.17)
Favorable Government			0.342
			(0.53)
Unfavorable Government			-0.604**
			(0.24)
Favorable Credit			-0.342
			(0.27)
Unfavorable Credit			0.124
			(0.19)
Energy Crisis			-0.393*
			(0.22)
Adj. R^2	0.10	0.12	0.15

The above are coefficient estimates from a regression of the reduced form innovation in E5Y obtained from the three variable system described in Section II on the percentage of respondents reporting having heard either favorable or unfavorable news concerning employment, prices, or stock prices. The sample period is 1961:1 - 2007:3. OLS standard errors are in parentheses.

Table 4.2
Structural Parameter Estimates

Parameter	Estimate	90 percent confidence interval
λ_1	0.15	[-1.13, 1.38]
λ_2	29.11	[7.44, 57.91]
κ	0.76	[0.48, 0.91]
σ_e	0.17	[0.01, 0.09]
σ_v	0.21	[0.001, 1.00]

Innovation variance in consumer confidence:

Due to unexpected change in current state:	0 percent
Due to information shocks:	41 percent
Due to animal spirits shocks:	59 percent

The above are estimates of the parameters of the model augmented with a structural specification of consumer confidence, as described in Section IV. The confidence intervals are computed as the 5th and 95th percentiles of a Monte Carlo simulated distribution.

Figure 4.1

Impulse Responses in Cochrane's Bivariate VAR

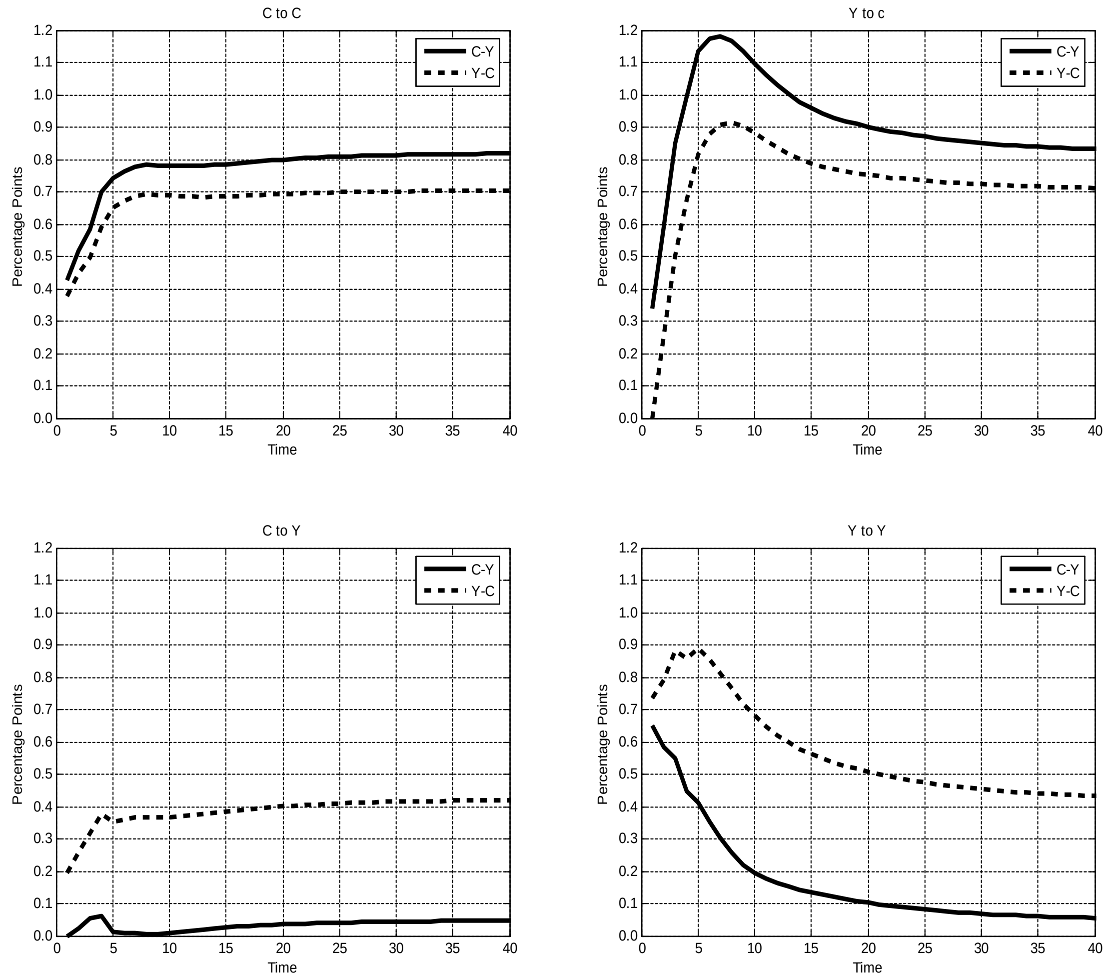

These are impulse responses from the bivariate consumption-income VAR discussed in Cochrane (1994a). The solid line shows responses from the orthogonalization with consumption ordered first (Cochrane's principle interpretation), while the dashed line refers to an orthogonalization with income ordered first.

Figure 4.2
E5Y and E12M

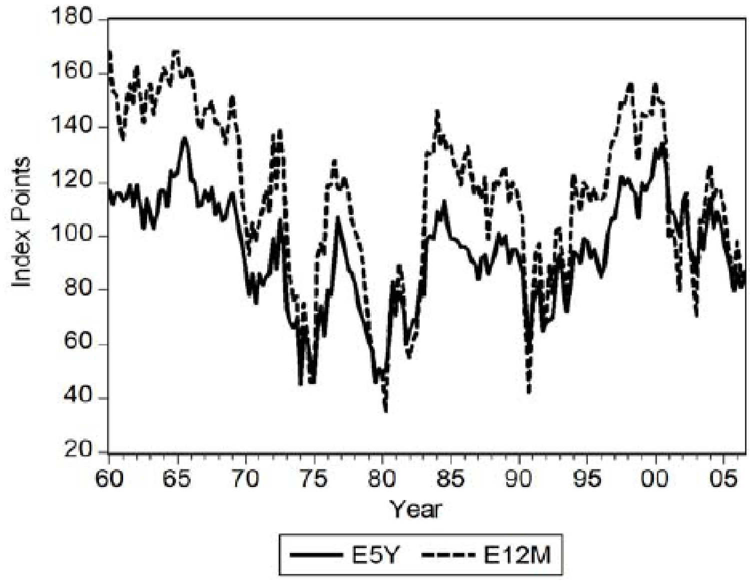

Figure 4.3

Impulse Responses: Ordering $E5Y \to C \to Y$

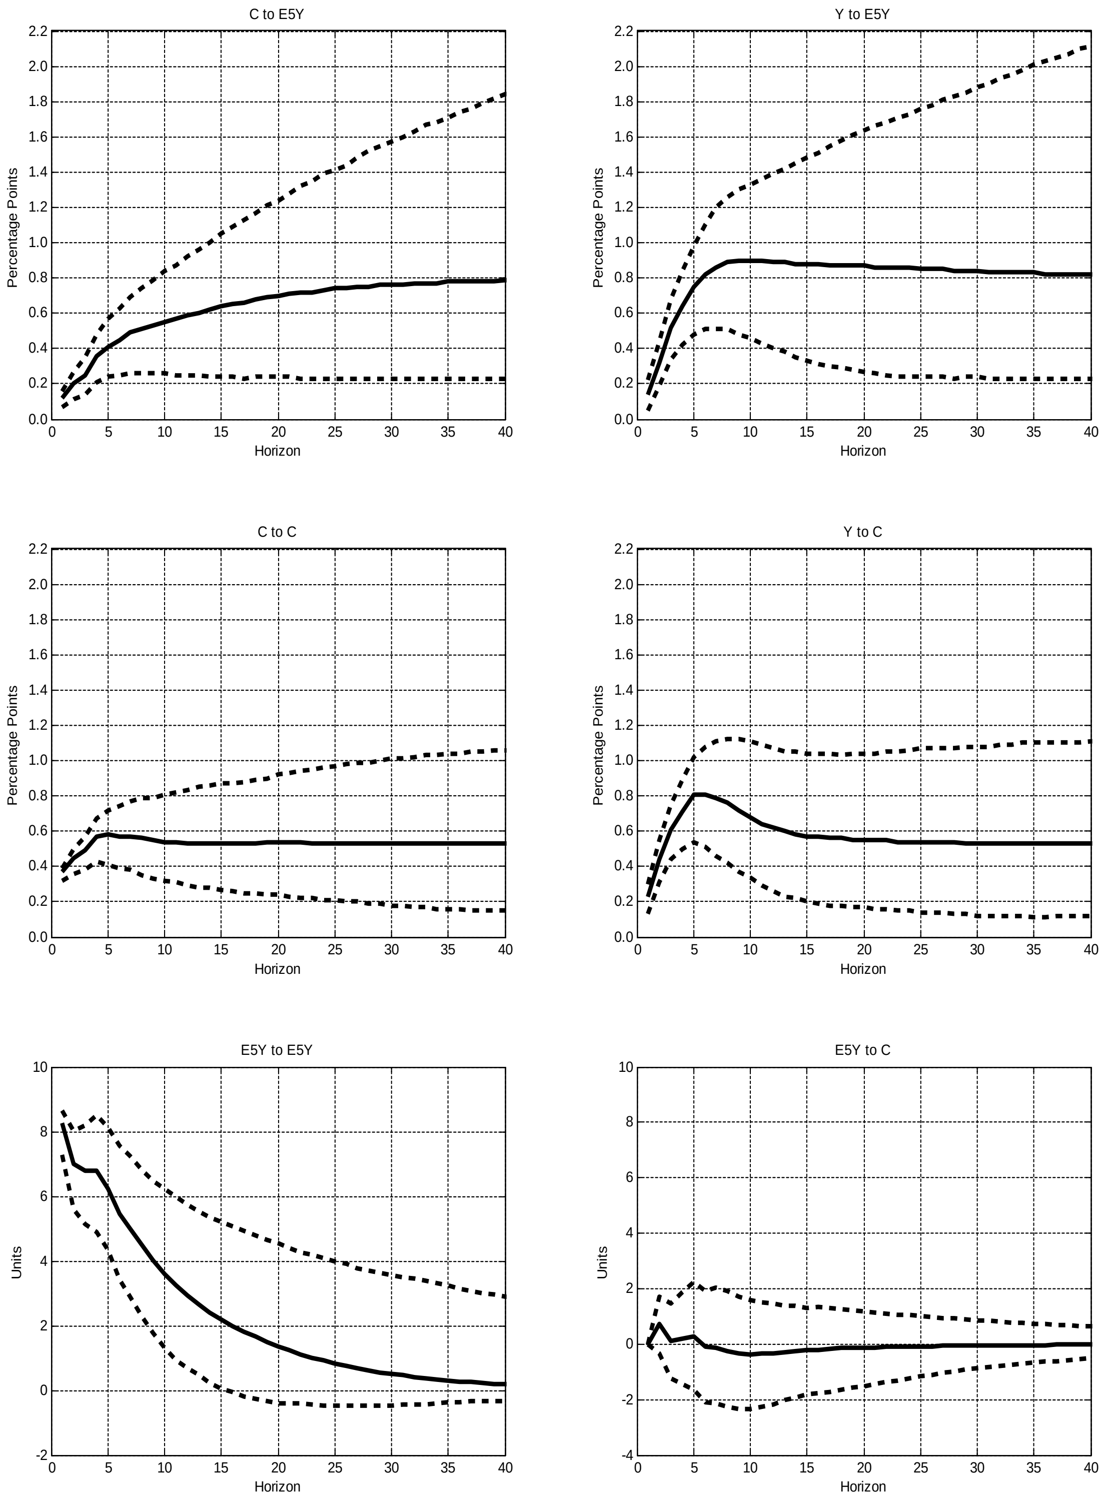

These are impulse responses from the three variable VARs described in Section 2. We omit the responses to output innovations. Dashed lines represent 90 percent confidence bands.

Figure 4.4

Impulse Responses: Ordering $C \rightarrow E5Y \rightarrow Y$

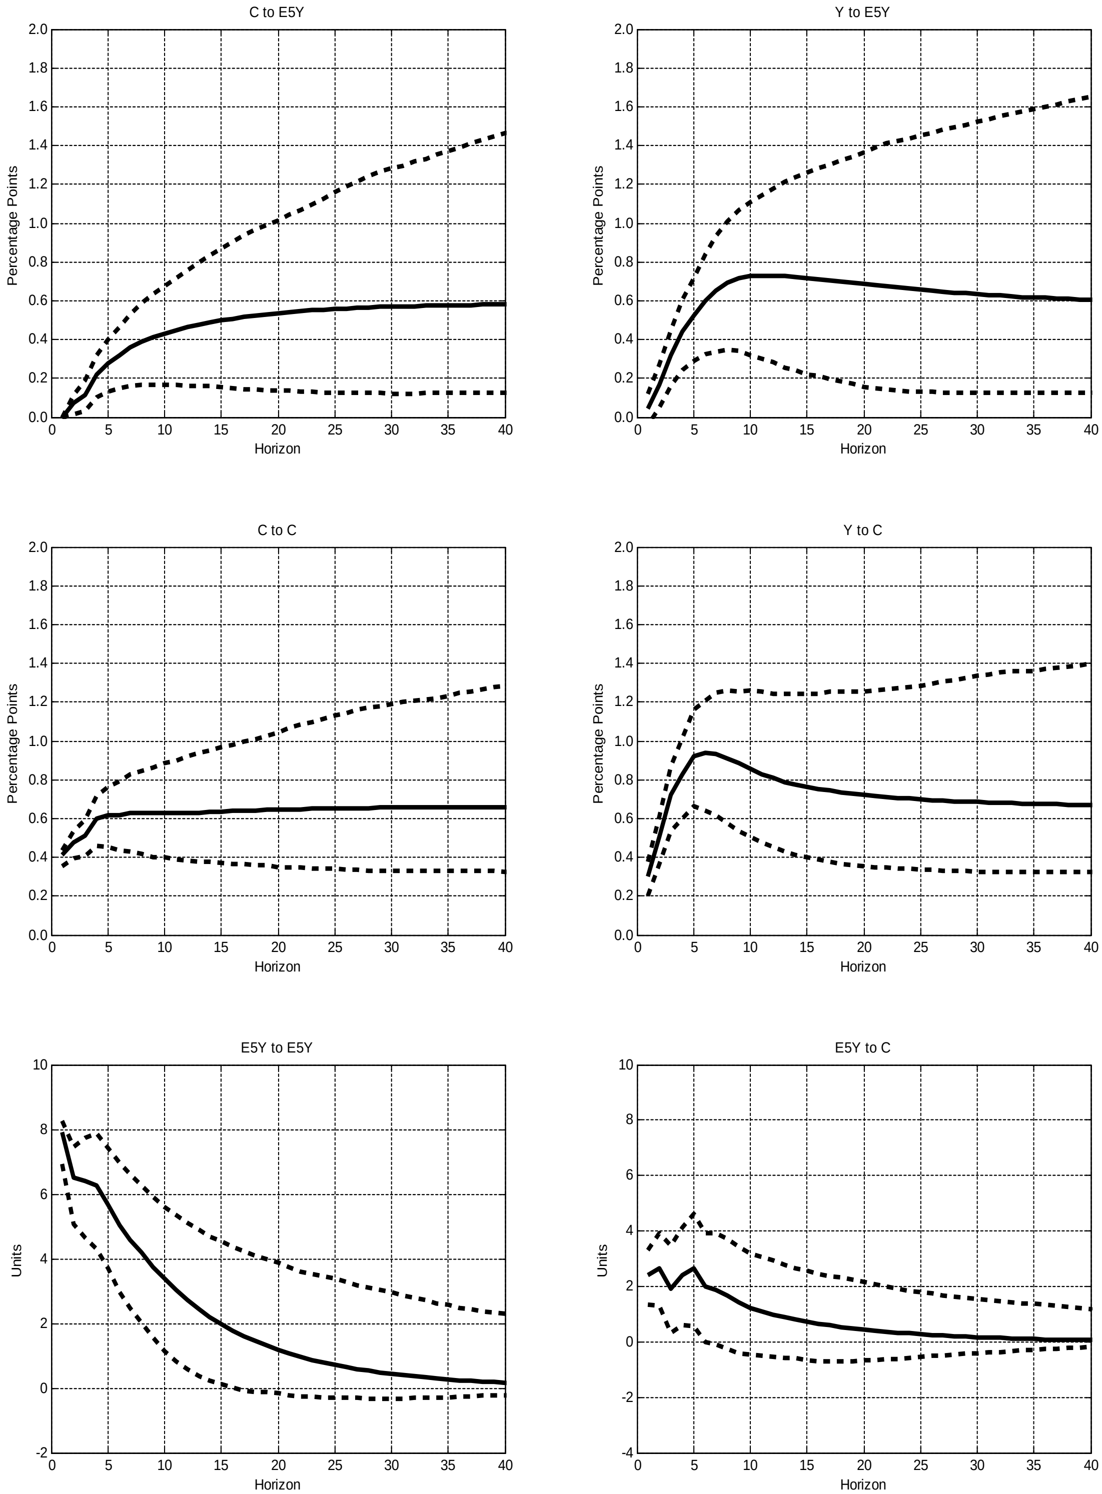

These are impulse responses from the three variable VARs described in Section 2. We omit the responses to output innovations. Dashed lines represent 90 percent confidence bands.

Figure 4.5
Forecast Error Variance Decomposition

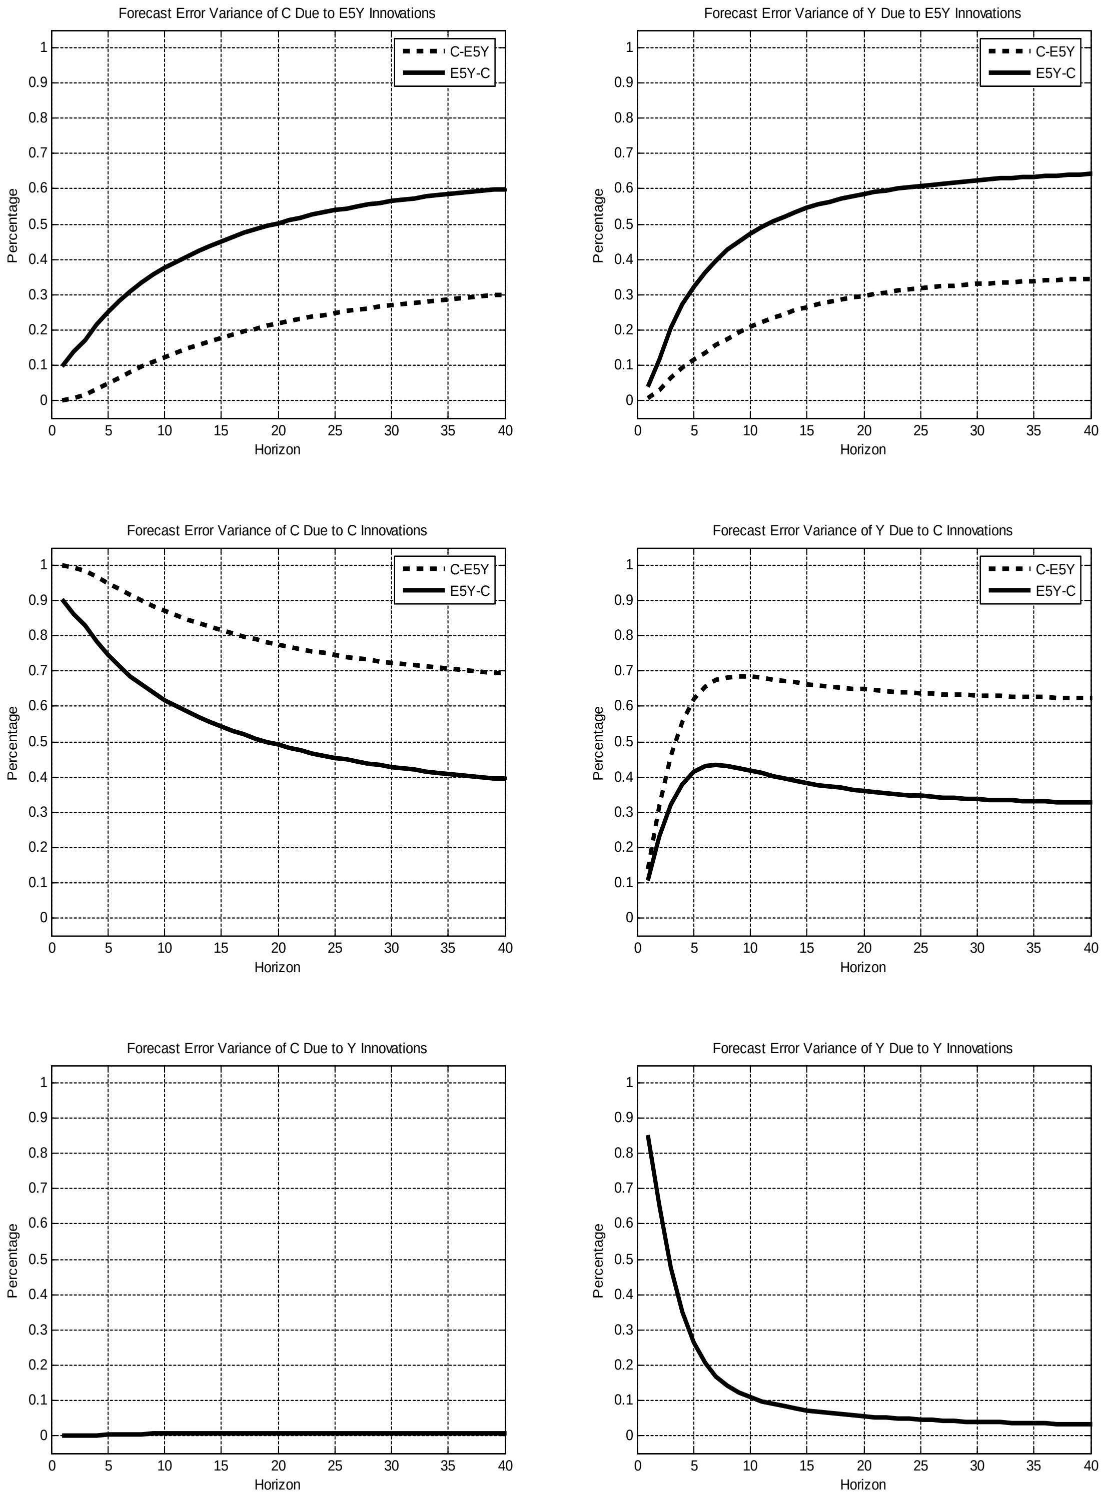

Figure 4.5 (Cont.)

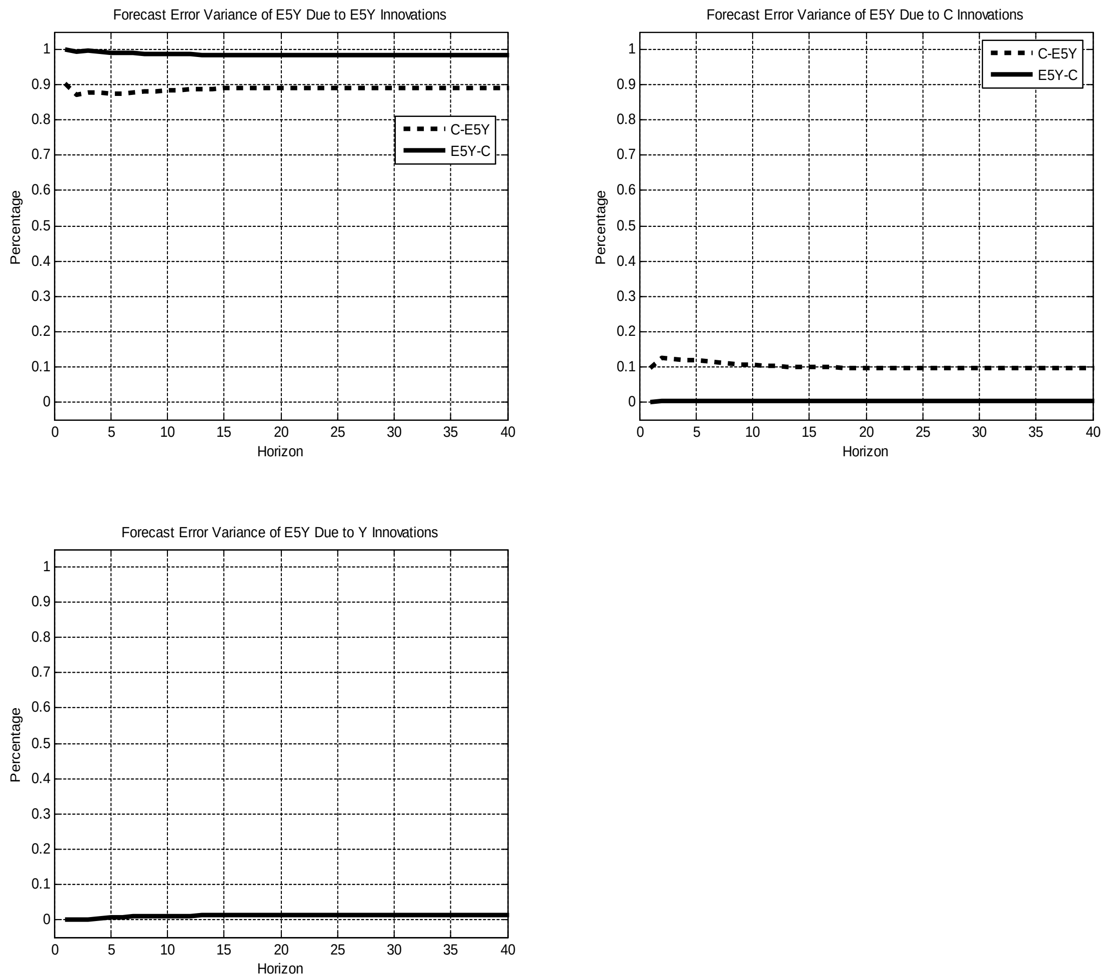

Figure 4.6

Impulse Responses with Alternative Confidence Measures

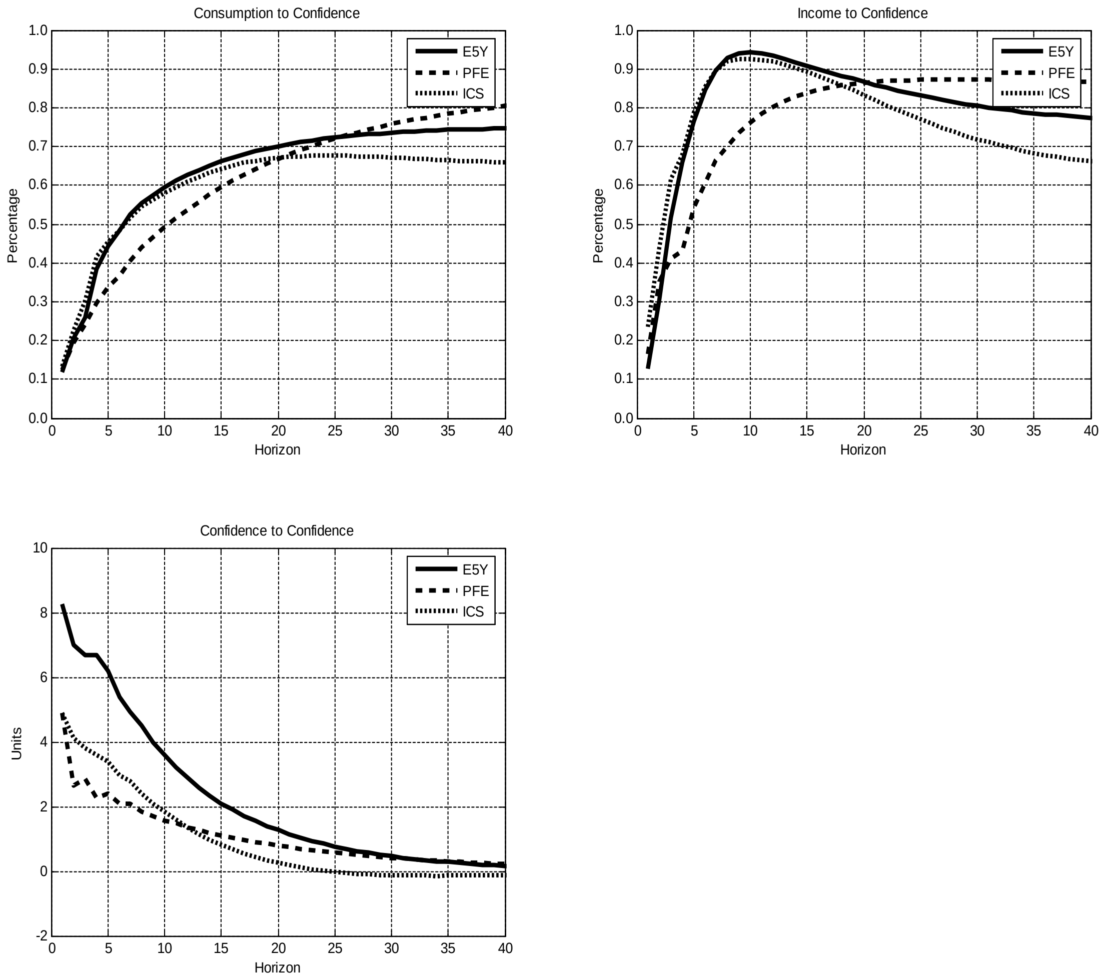

Figure 4.7
Spike Plots of News Heard Categories

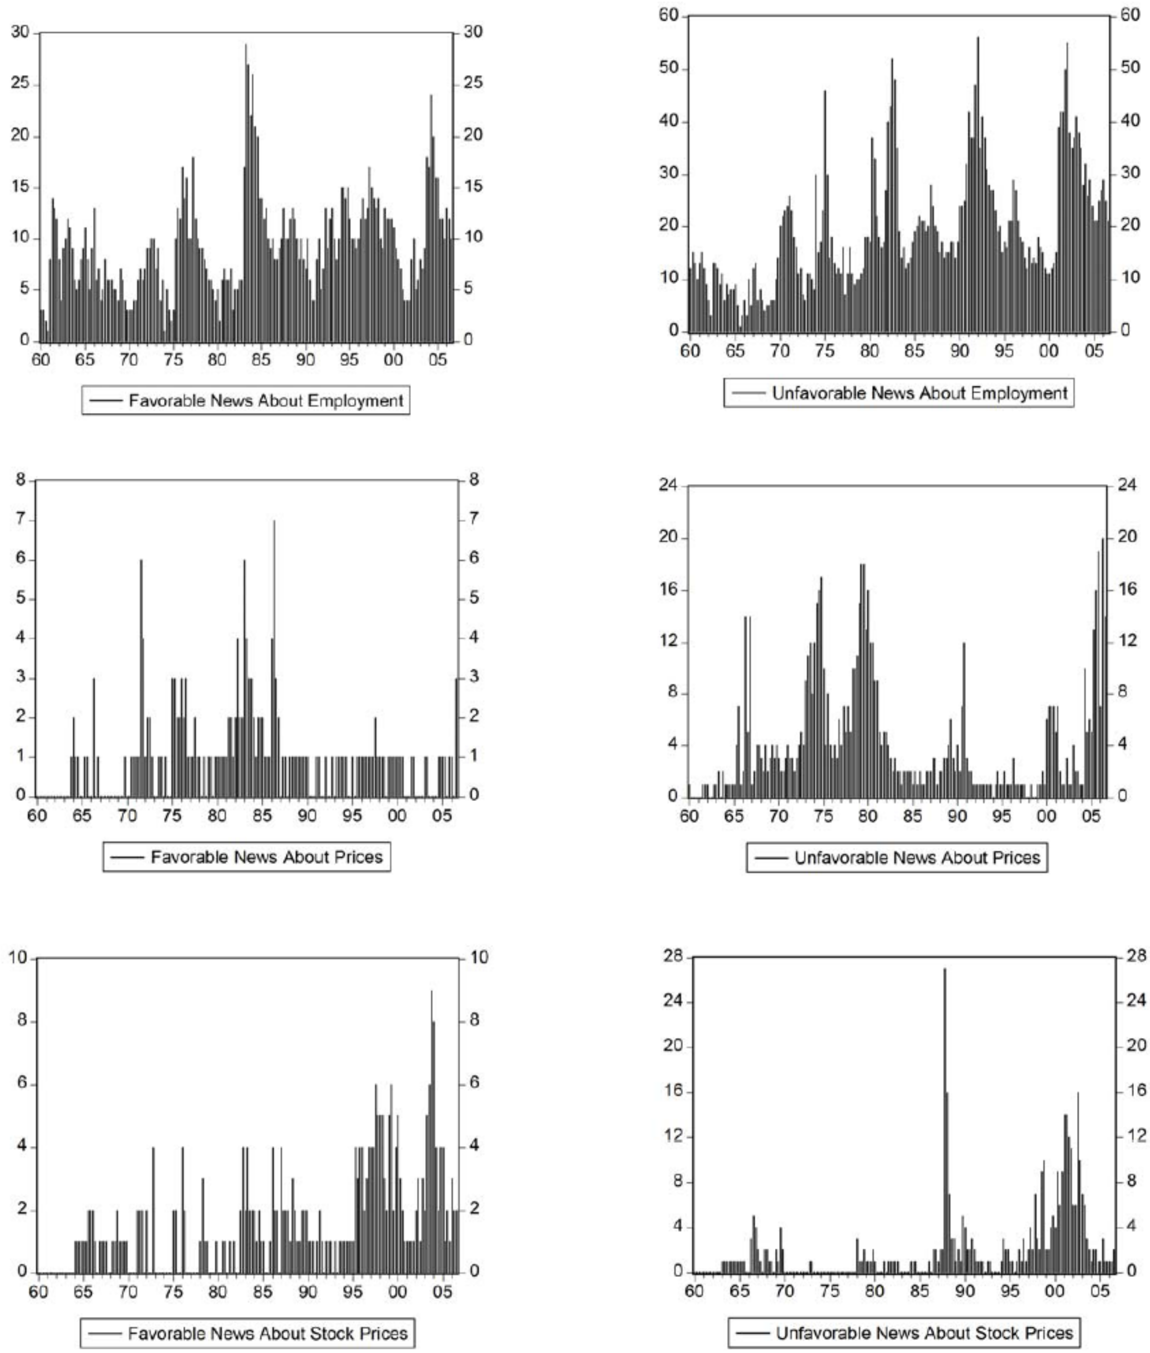

Figure 4.8

Theoretical Responses to Level and Information Shocks in Model

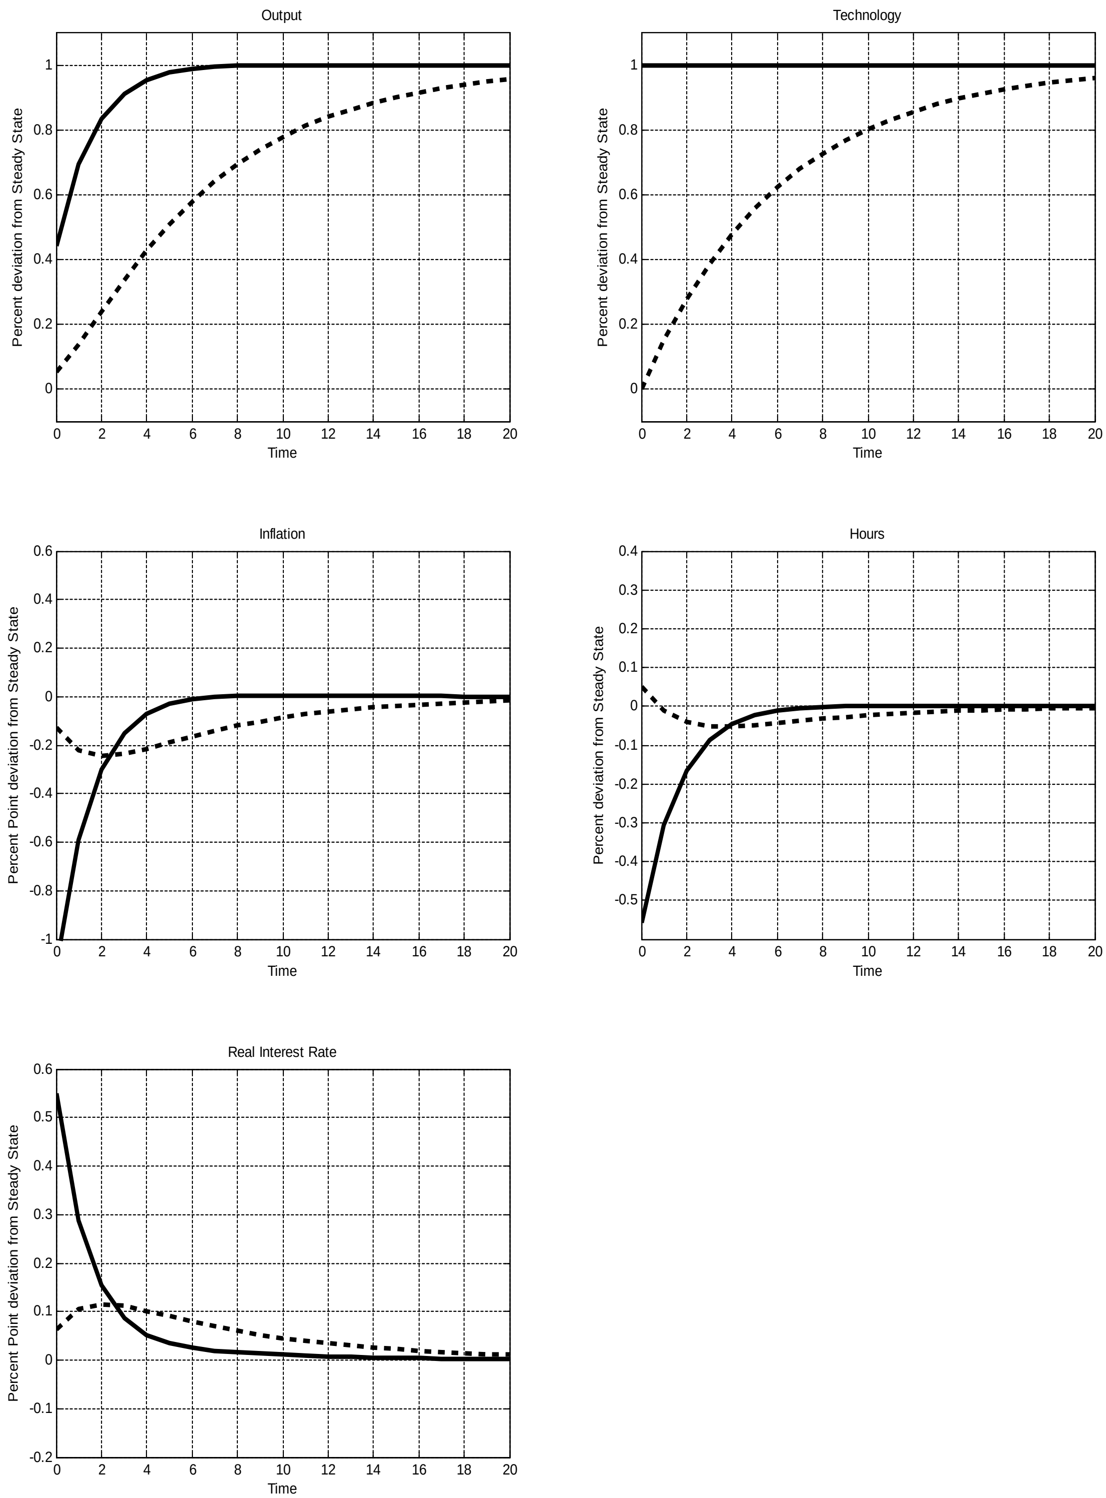

The above are theoretical impulse responses from the model using the calibration noted in Section 3.

Figure 4.9
Theoretical Responses to Animal Spirits Shock

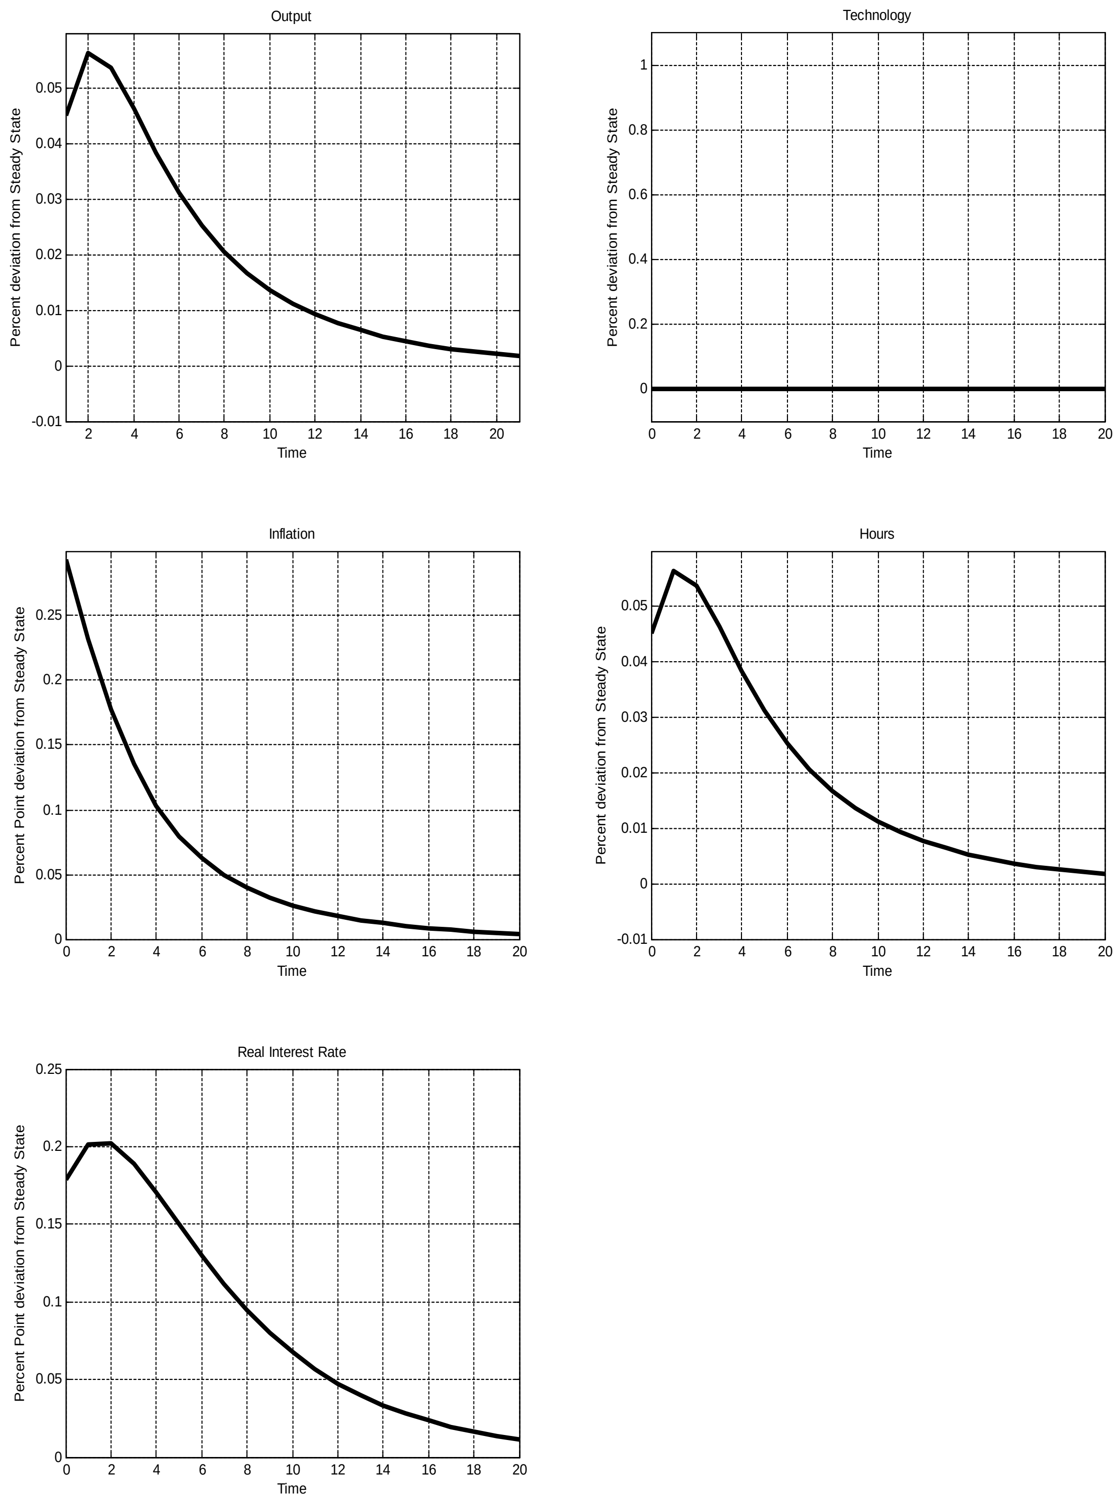

Figure 4.10

Responses to Confidence Innovation in Expanded VAR

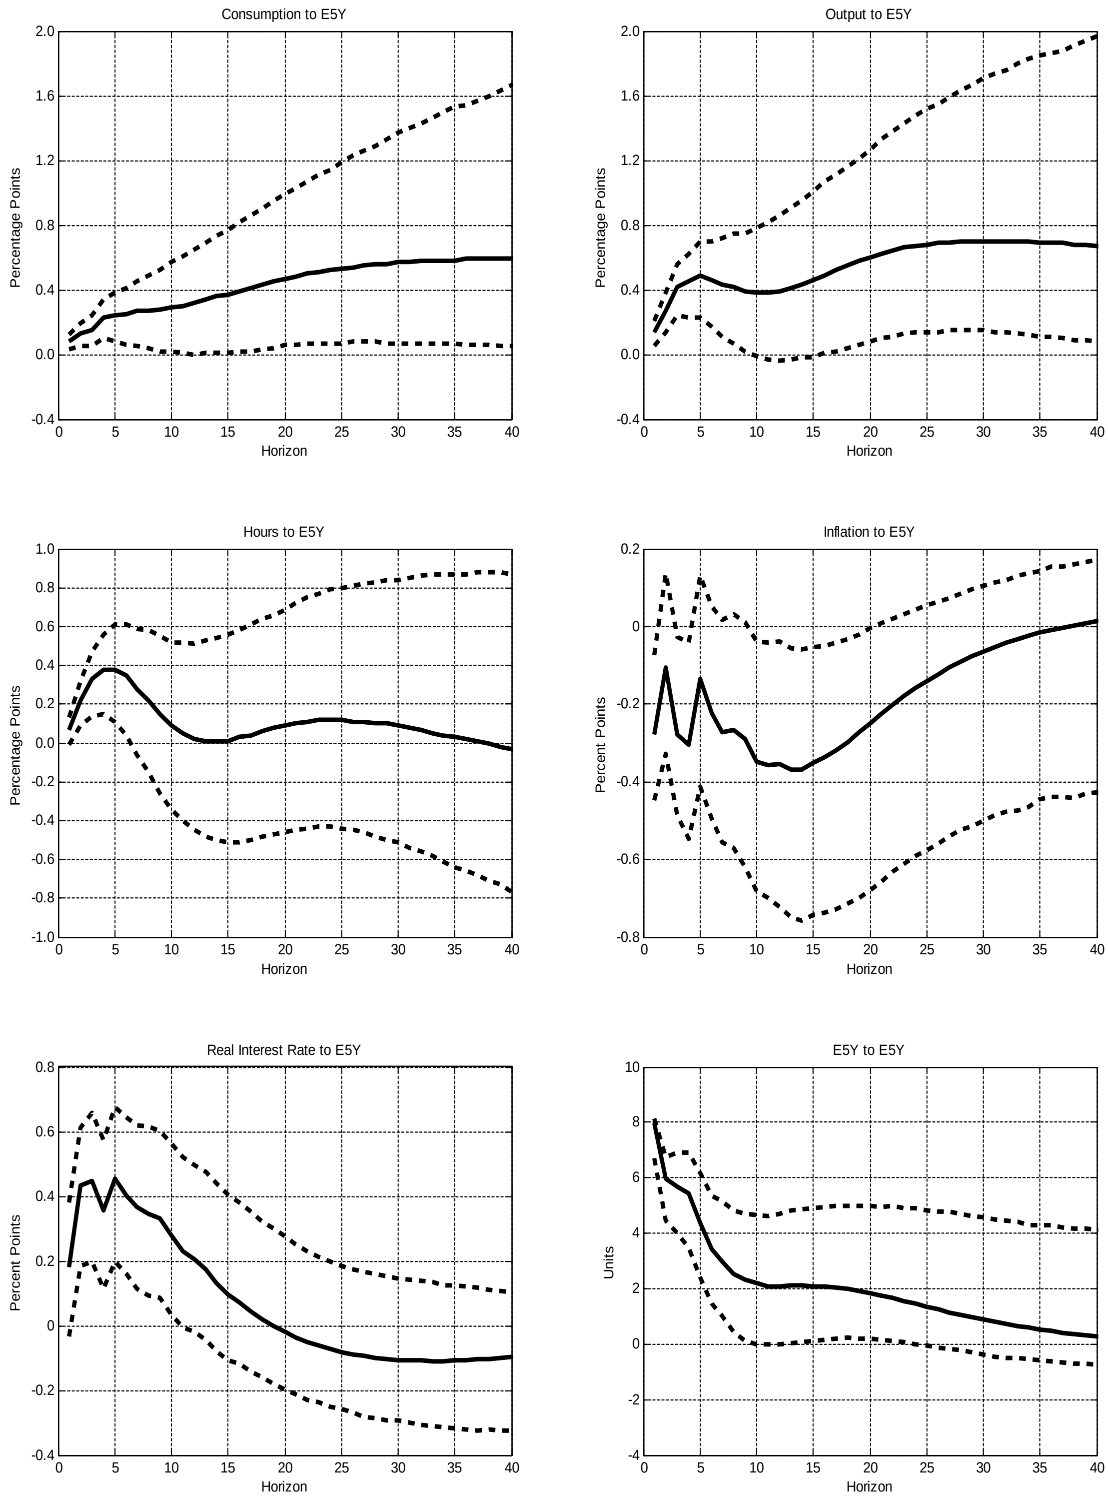

These are empirical responses from the six variable VAR described in Section 4. The dashed lines represent 90 percent bootstrap after bootstrap confidence bands.

Figure 4.11
Average Responses from Simulated Data Using Parameter Estimates

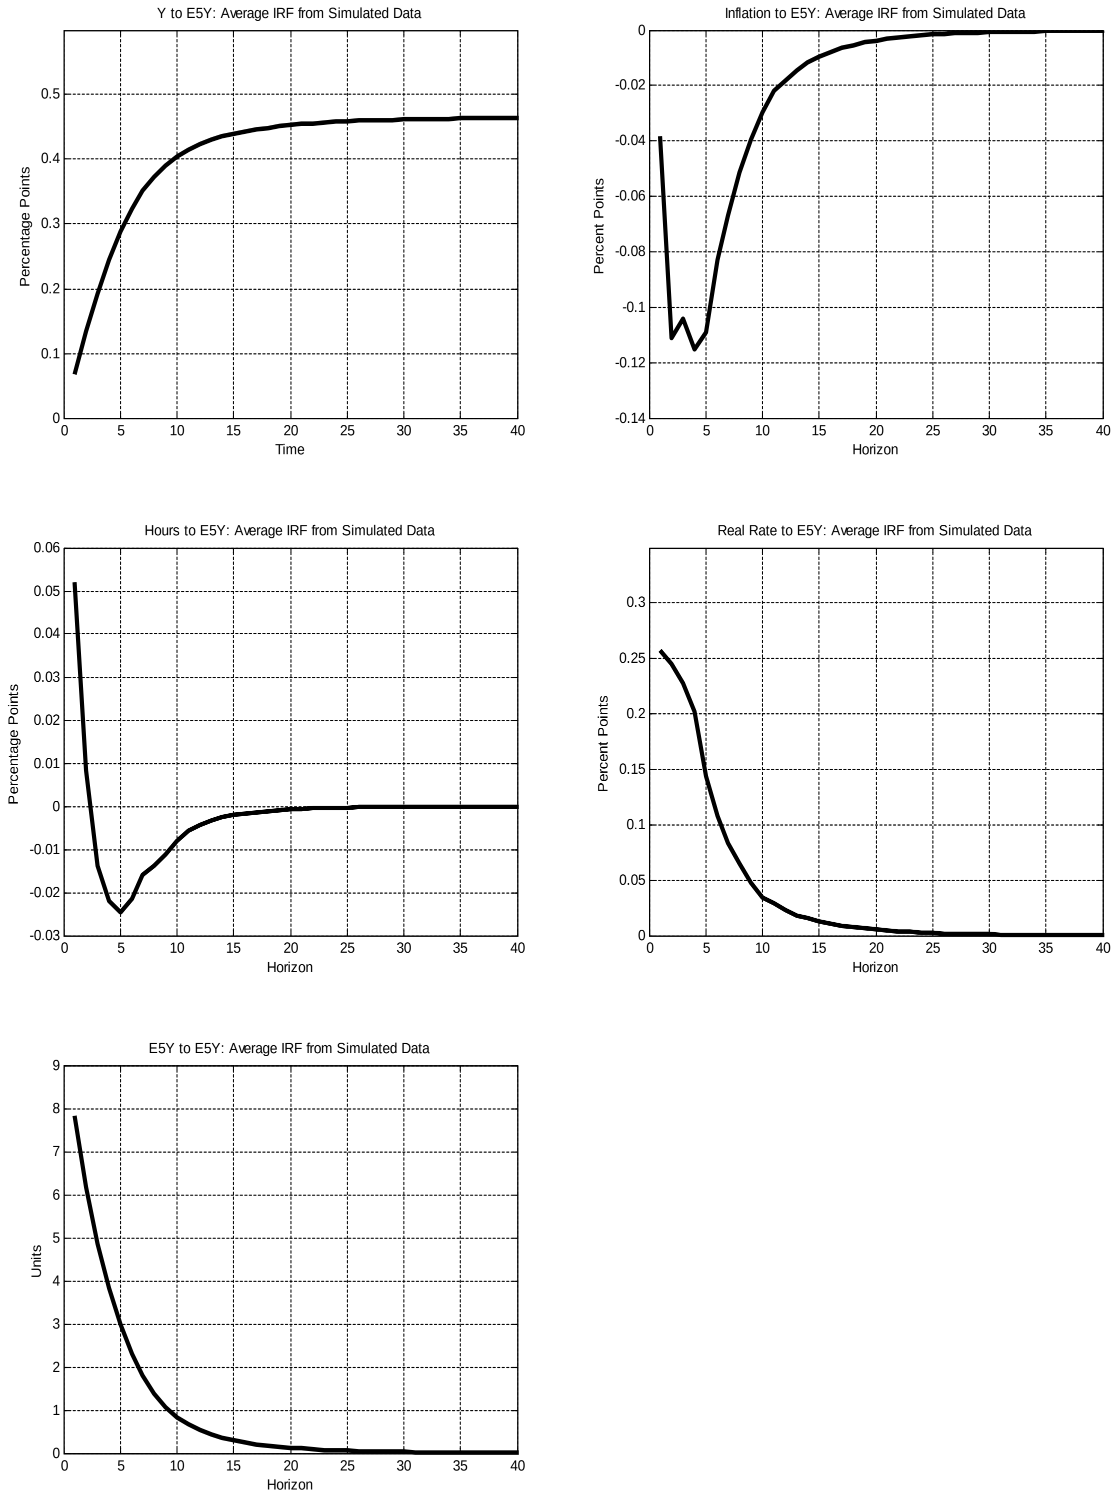

These figures show average responses from 2000 simulations with 200 observations each using the parameter estimates.

Figure 4.12
Impulse Responses from TFP-E5Y VAR

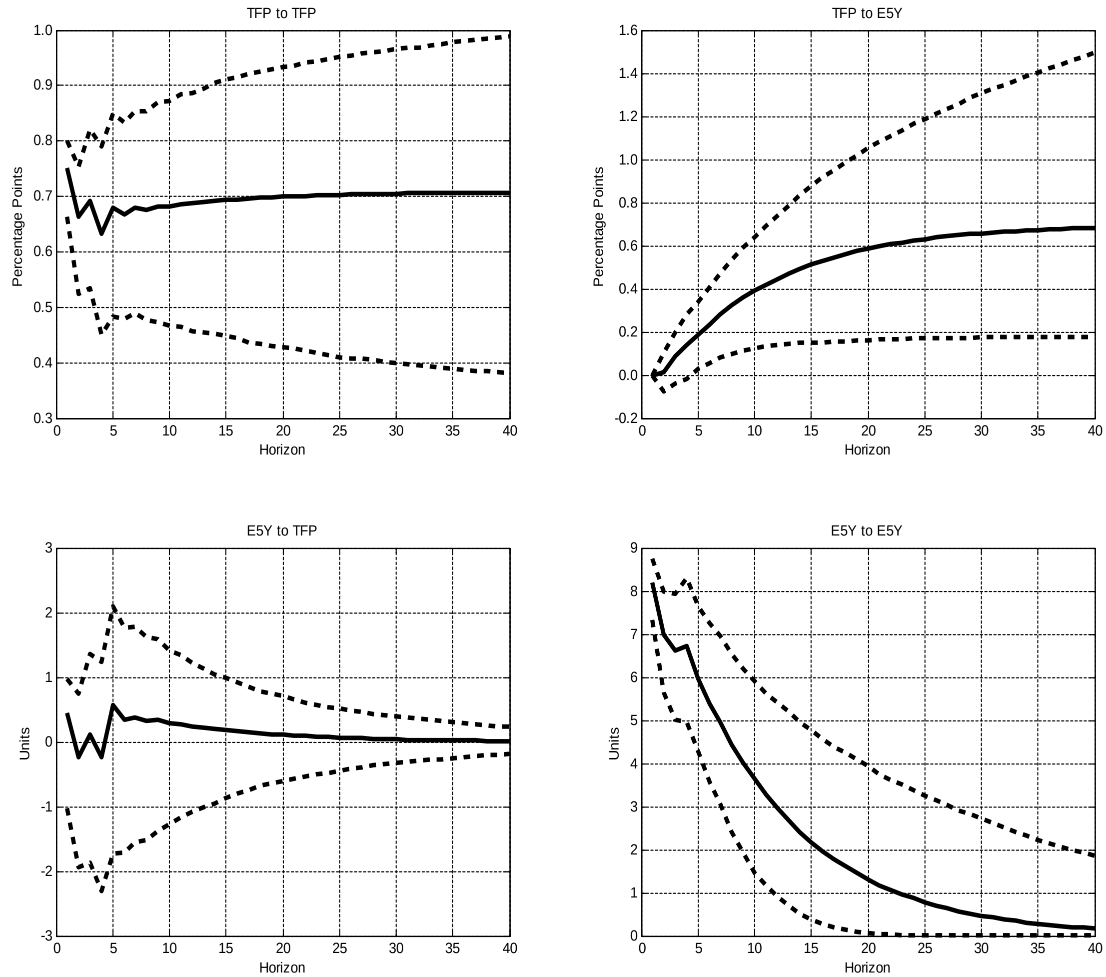

These are impulse responses from a bivariate VAR featuring a utilization corrected measure of total factor productivity and E5Y. The innovations are orthogonalized so that the structuralized E5Y innovation has no contemporaneous effect on the level of TFP. The dashed lines are 90 percent confidence bands from a bias-corrected bootstrap.

References

[1] Barsky, Robert and Eric Sims. "News Shocks." Chapter II, *Expectations and the Business Cycle*, Ph.D. Dissertation of Eric R. Sims, University of Michigan, 2009.

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

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

[4] Beaudry, Paul and Franck Portier. "Stock Prices, News, and Economic Fluctuations." *American Economic Review*, September 2006, 1293-1307.

[5] Blanchard, Olivier J. "Consumption and the Recession of 1990-1991." *American Economic Review*, May 1993, 270-274.

[6] Calvo, Guillermo. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics*, September 1983, 383-398.

[7] Carroll, Christopher D. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics*, February 2003, 269-298.

[8] Carroll, Christopher D., Jeffrey C. Fuhrer and David W. Wilcox. "Does Consumer Sentiment Forecast Household Spending? If so, Why?" *American Economic Review*, December 1994, 1397-1408.

[9] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, February 2005, 1-45.

[10] Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfussion. "The Responses of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology." *Journal of the European Economic Association*, May 2004, 381-395.

[11] Cochrane, John H. "Permanent and Transitory Components of GNP and Stock Prices." *Quarterly Journal of Economics*, February 1994a, 241-266.

[12] Cochrane, John H. "Shocks." *Carnegie-Rochester Conference Series on Public Policy*, December 1994b, 295-364.

[13] Coibion, Olivier and Yuri Gorodnichenko. "Strategic Interaction Among Price-Setters in an Estimated DSGE Model." Working paper, 2007.

[14] Croushore, Dean. "Do Consumer Confidence Indexes Help Forecast Consumer Spending in Real Time?" *North American Journal of Economics and Finance*, December 2005, 435-450.

[15] Dominitz, Jeff and Charles Manksi. "How Should We Measure Consumer Confidence?" *Journal of Economic Perspectives*, Spring 2004, 51-66.

[16] Farmer, Roger. *The Macroeconomics of Self-Fulfilling Prophecies*. Cambridge, MA: MIT Press. 1999.

[17] Fernandez-Villaverde, Jesus, and Juan F. Rubio-Ramirez. "How Structural are Structural Parameters?" Forthcoming, *NBER Macroeconomics Annual 2007*.

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

[19] Gali, Jordi and Mark Gertler. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics*, October 1999, 195-222.

[20] Gorodnichenko, Yuri and Matthew D. Shapiro. "Monetary Policy When Potential Output is Uncertain: Understanding the Growth Gamble of the 1990s." *Journal of Monetary Economics*, May 2007, 1132-1162.

[21] Hall, Robert E. "Macro Theory and the Recession of 1990-1991." *American Economic Review*, May 1993, 275-279.

[22] Hall, Robert E. "Stochastic Implications of the Life Cycle Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, December 1978, 971-987.

[23] Hall, Robert E. "The Role of Consumption in Economic Fluctuations." *The American Business Cycle: Continuity and Change*, edited by Robert J. Gordon. Chicago: UCP, 1986, pp. 237-255 and 265-266.

[24] Howrey, E. Phillip. "The Predictive Power of the Index of Consumer Sentiment." *Brookings Papers on Economic Activity*, January 2001, 175-207.

[25] Ireland, Peter N. "Technology Shocks and the New Keynesian Model." *Review of Economics and Statistics*, April 2004, 923-936.

[26] Jaimovich, Nir and Sergio Rebelo. "Can News About the Future Drive the Business Cycle?" Stanford University Working Paper, August 2006.

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

[28] Kimball, Miles S. and Matthew D. Shapiro. "Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?" University of Michigan Working Paper, May 2003.

[29] King, Robert G., Charles I. Plosser, and Sergio Rebelo. "Production, Growth, and Business Cycles I: The Basic Neoclassical Model." *Journal of Monetary Economics*, 21 (2-3), 195-232.

[30] Leeper, Eric M. "Consumer Attitudes: King for a Day." *Federal Reserve Bank of Atlanta Economic Review*, July-August 1992, 1-15.

[31] Lorenzoni, Guido. "A Theory of Demand Shocks." MIT Working Paper, January 2008.

[32] Lucas, Robert E. "Asset Prices in an Exchange Economy." *Econometrica*, November 1978, 1429-1445.

[33] Ludvigson, Sidney. "Consumer Confidence and Consumer Spending." *Journal of Economic Perspectives*, Spring 2004, 29-50.

[34] Matsusaka, John G. and Argia M. Sbordonne. "Consumer Confidence and Economic Fluctuations." *Economic Inquiry*, April 1995, 296-318.

[35] Mehra, Yash P and Elliot W. Martin. "Why Does Consumer Sentiment Predict Household Spending?" *Federal Reserve Bank of Richmond Economic Quarterly*, Fall 2003, 51-67.

[36] Mishkin, Frederic S. "Consumer Sentiment and Spending on Durable Goods." *Brookings Papers on Economic Activity*, 1, 1978, 217-231.

[37] Orphanides, Athanasios. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics*, July 2003, 983-1022.

[38] Paciello, Luigi. "The Response of Prices to Technology and Monetary Policy Shocks Under Rational Inattention." Northwestern University working paper, February 2008.

[39] Ruge-Murcia, Francisco J. "Methods to Estimate Dynamic Stochastic General Equilibrium Models." *Journal of Economic Dynamics and Control*, August 2007, 2599-2636.

[40] Smets, Rafael and Frank Wouters. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, September 2003, 1123-1175.

[41] Souleles, Nicholas S. "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys." *Journal of Money, Credit, and Banking*, February 2004, 39-72.

[42] Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ: Princeton University Press, 2003.

Chapter V

Identification and Estimation of Interest Rate Rules in the New Keynesian Model

1 Introduction

One of the most celebrated papers in economics is by Taylor (1993), in which he shows that a simple linear rule relating the Fed Funds Rate to the deviation of GDP from trend and the deviation of inflation from target characterizes US monetary policy well. Taylor argues that it is important that the central bank commit to raising nominal rates more than one for one with inflation. This so-called "Taylor Principle" essentially calls for the Fed to raise real rates in response to inflation, thereby eliminating the possibility of self-fulfilling inflations.

While the Taylor Principle was first espoused in the context of an "old" Keynesian model, researchers over the last fifteen years have adopted Taylor's nominal interest rate rule (i.e. "Taylor Rules") in the context of "new" Keynesian models. As Cochrane (2007a and 2007b) eloquently points out, in many respects the "new" Keynesian models are substantially different from their predecessors, and the "old" Keynesian intuition does not always carry over into the "new" models. Nevertheless, something very similar to the Taylor Principle obtains in the "new" model. Although the exact conditions are somewhat more complicated, a unique and non-explosive equilibrium in the New Keynesian model still requires that the central bank respond aggressively enough to movements in inflation and/or output.

Whether or not the Fed's nominal interest rate target has always responded sufficiently enough to inflation is potentially a key issue in understanding the moderation of the US economy since the 1970s. In a provocative paper, Clarida, Gali, and Gertler (2000, hereafter CGG) develop a New Keynesian model with a policy rule based on

Taylor (1993) and derive conditions for the determinacy of a rational expectations equilibrium. They then estimate regressions based on their model's policy rule and argue that the Fed failed to satisfy the modified Taylor Principle in the 1970s but did so afterwards. If their empirical findings are robust, this suggests that policy itself might have been an important source of instability in the 1970s. Not only would the economy be subject to indeterminate "sunspot" equilibria were the Taylor Principle not satisfied, the economy's responses to fundamental shocks might also have been affected in important and welfare-reducing ways.

In a recent paper, Cochrane (2007b) argues that CGG's interpretation of macroeconomic history is flawed. In particular, he asserts that the parameters of the nominal interest rate rule are not identified in the New Keynesian model. His essential argument is that satisfaction of the Taylor Principle in the model is tantamount to a threat which is never actually carried out in equilibrium. As such, he argues, data generated from the model can never reveal the values of the policy parameters in the rule which lend credibility to the threat.

The present paper delves deeper into this issue of identification, and considers the conditions under which Cochrane's claim of non-identification holds and the conditions under which it does not. To foreshadow my conclusions, I show that Cochrane's non-identification result is not a generic implication of the model, but rather obtains only under a particular (and likely unrealistic) specification of the central bank's reaction function. For more standard specifications of policy, I show that the central bank's key policy parameters are in fact identified and may be estimated by conventional means.

The key ingredient allowing identification of the central bank's policy parameters is for non-policy shocks to lead to equilibrium movements in inflation and the output gap. In the standard New Keynesian model with a conventional Taylor Rule, non-policy shocks do affect inflation and the gap through the Phillips Curve. Thus, the New Keynesian model generates movements in inflation orthogonal to structural policy errors. Further, these movements in inflation are monotonically related to the parameters of the Taylor Rule. As such, those parameters are in fact identified. I demonstrate that a standard linear regression, with proper instrumental variables, will in fact consistently estimate the central bank's policy parameters. Nevertheless, I document that the small sample properties of single equation estimates of the Taylor rule are likely rather poor. In particular, the further the central bank is from the Taylor principle cutoff, the more imprecise are the estimates.

Cochrane obtains his non-identification result in a framework in which the policy rule contains a “stochastic intercept” term which shifts endogenously in response to non-policy shocks in such a way as to completely stabilize the output gap and inflation. In particular, he assumes that the “stochastic intercept” moves one for one with fluctuations in the Wicksellian natural rate of interest, which Woodford (2003) defines as the real interest rate which would obtain in the absence of nominal frictions.In such a world, the only observable movements in inflation will are due to policy shocks, and hence the model yields no exogenous variation in inflation or the output gap off of which to identify the central bank’s policy parameters.

Such a specification of the policy rule is potentially appealing from a normative perspective. As shown in Blanchard and Gali (2007), stabilization of the output gap (up to first order approximation) maximizes welfare. Implementation of such a rule requires that the central bank know the natural rate of interest, which is not directly observable, but is rather a theoretical construct that is a potentially highly complex function of several structural shocks and parameters. As such, it is unlikely that any central bank would be able to implement such a rule with much precision. This fact can lead to a reinterpretation of policy shocks not as “helicopter drops of money”, but rather as errors in the observation of the natural rate of interest.

Cochrane argues that it is restrictive to make the common assumption of a constant intercept in the nominal interest rate rule. While it is true that the stochastic intercept rule welfare dominates the more standard specification of a Taylor type rule, the relative welfare losses from a constant intercept rule are small. I show that the standard deviation of the welfare relevant output gap in a constant intercept specification of the rule is only slightly higher than it is under the stochastic intercept rule for plausible calibrations of the parameters of the model. As such, it takes only a very slight reduction in the variance of policy shocks for the constant intercept rule to welfare dominate the stochastic intercept rule. Since the constant intercept rule does not require the central bank to react to anything which is not directly observable, it seems likely that such a rule would be associated with a significantly lower variance of policy shocks, and thus with higher welfare.

I present empirical evidence that the Fed does not, in fact, follow a stochastic intercept rule. There are testable implications of the stochastic intercept rule which do not depend on the exact parameter values. In particular, such a rule has the stark prediction that the only source of variation in inflation are policy shocks/errors. This prediction can be tested by examining the conditional relationships between inflation and non-policy shocks. I identify a “supply” shock from a bivariate VAR using a long run restriction. The shock explaining the unit root in output accounts for well more than half of the innovation variance in inflation. This empirical finding is incompatible with a nominal interest rate rule featuring a stochastic intercept which adjusts one for one with fluctuations in the Wicksellian natural rate of interest.

While not consistent with the predictions of a stochastic intercept rule, the pattern of responses from the bivariate VAR to "supply" shocks are also not consistent with more standard specifications of the interest rate rule. The impulse responses from the bivariate VAR suggest that there is both a large disinflation and a significant predictable increase in output following a positive "supply" shock. The predictable increase in output would have to be associated with an increase in real interest rates to be consistent with the Euler equation. The standard New Keynesian model with the common Taylor type nominal interest rate rule is simply not capable of delivering a substantial increase in real rates and disinflation at the same time. This result suggests that modifications to the New Keynesian/Taylor rule framework are necessary in order to better fit the data.

2 The New Keynesian Model with a Taylor Rule

The canonical New Keynesian model is firmly rooted in the pillars of modern macroeconomics: explicit micro foundations, agent optimization, rational expectations, and market clearing. It differs from its neoclassical counterparts in that it introduces nominal frictions into the model, thereby inducing non-trivial distortions from the first best. These frictions serve two purposes. First, they allow for real effects of monetary disturbances. Second, these frictions fundamentally alter the economy's response to real shocks. Modern research in monetary economics focuses on the second point, and in particular the question of how to better design the systematic component of policy so as to minimize the distortions associated with the economy's response to non-policy shocks in a world with nominal frictions.

I present an extremely stylized version of the basic New Keynesian model. Other than the monetary policy rule, the two primary equations of the model are the Euler/IS equation and the Phillips Curve. The log-linearized "IS" equation is given by:

$$E_t y_{t+1} - y_t = \sigma(i_t - E_t \pi_{t+1}) \tag{1}$$

Reflecting the accounting identity that all output must be consumed, the equation relates expected consumption growth to the *ex ante* real interest rate, with σ denoting the intertemporal elasticity of substitution.[1] For simplicity, I assume a constant discount factor, no preference shocks, and no government consumption.

The New Keynesian Phillips Curve is derived under the assumption of Calvo (1983) style price-setting and relates inflation to the output gap and expected future inflation:

$$\pi_t = \kappa(y_t - y_t^f) + \beta E_t \pi_{t+1} \tag{2}$$

κ is a reduced form parameter reflecting the degree of exogenous price stickiness, the rate of time preference, and the output elasticity of real marginal cost.[2] β is a subjective discount factor, and y_t^f refers to the flexible price equilibrium level of output (i.e. that level of output that would obtain in the absence of price stickiness).

One can think of the flexible price equilibrium level of output as either being exogenous or as endogenously determined given exogenous shock processes. For the sake of simplicity, I choose the former, and model the flexible price level of output as obeying an autoregressive process:

$$y_t^f = \gamma y_{t-1}^f + \varepsilon_t \tag{3}$$

In this simple environment, one can think of the flexible price equilibrium level of output as being driven by fluctuations in technology, though in a more complicated setting it would reflect any of several forces that would affect output in a standard RBC model – government purchase, preference, or tax shocks, to name but a few. It is also helpful to define another theoretical construct which I will refer to as the Wicksellian natural rate of interest (following Woodford (2003)):

$$r_t^f = \frac{1}{\sigma}(E_t y_{t+1}^f - y_t^f) \tag{4}$$

This is the real interest rate which would obtain in the absence of nominal rigidities, and is found by solving the Euler equation for the real interest rate consistent with output being at the flexible price equilibrium level.

[1] The New Keynesian model usually abstracts from capital and investment. This abstraction simplifies matters a great deal, and turns out to be largely irrelevant to the issues studied in the current paper.

[2] Formally, $\kappa = \frac{(1-\theta)(1-\theta\beta)}{\theta}(1/\sigma + 1/\eta)$, where θ is the exogenous probability of a firm being able to change its price in any period and η is the Frisch labor supply elasticity. The term $(1/\sigma + 1/\eta)$ is the elasticity of the output gap with respect to real marginal cost. I abstract from features which would introduce real rigidities (also known as strategic complementarities). Such features would affect the slope of the Phillips Curve, but would not affect the analysis concerning the identification of parameters from the monetary policy rule.

The full solution to the model requires the introduction of an additional equation describing monetary policy. Early variants of the model included a money demand function and specified an exogenous time path for the money supply. Most recent research closes the model with a nominal interest rate rule. In the spirit of Taylor (1993), one might suppose that the central bank sets nominal interest rates according to a rule similar to:

$$i_t = i^* + \phi_y(y_t - y_t^f) + \phi_\pi(\pi_t - \pi^*) + v_t \tag{5}$$

Taylor argues that such a rule is both a good description of historical Fed policy as well as a good normative prescription for how the Fed ought to conduct policy. i^* is the Fed's target nominal rate when both output and inflation are at their targets. Both response coefficients on the gap and inflation are assumed to be non-negative. v_t is some exogenous disturbance. One may think of this disturbance either as an exogenous shift in the stance of policy or, perhaps more realistically, as reflecting misperceptions of current inflation and/or the gap or the influence of some omitted variable(s).

The positive fact that nominal interest rates are highly persistent and the normative observation that banks may find it desirable to smooth rates over time has led researchers to consider modifications of (5) in which there is an explicit smoothing parameter. One simple way to do this is to assume that v is autocorrelated:

$$v_t = \alpha v_{t-1} + u_t \tag{6}$$

Another is to include a lagged interest rate term on the right hand side of (5), leading to a reinterpretation of the policy rule as one of partial adjustment:

$$i_t = \rho i_{t-1} + (1-\rho)i^* + (1-\rho)\left(\phi_y(y_t - y_t^f) + \phi_\pi(\pi_t - \pi^*)\right) + v_t \tag{7}$$

While both (5) with an autocorrelated error term and (7) with a white noise error term will produce persistent effects of policy shocks, they are not the same. In particular, (7) assumes that the bank desires to smooth rates in response to all shocks, whereas in (5) rates are only explicitly smoothed in response to policy shocks. The partial adjustment specification is the more popular of the two in the literature, though Rudebusch (2002) argues that the standard Taylor rule with an autocorrelated disturbance is more consistent with the data.

The full linear system of equations satisfies the Markov property and features two jump variables and two state variables.[3] The determinacy and boundedness of the equilibrium of the model depend on the eigenvalues of the transition matrix. A unique and non-explosive rational expectations equilibrium requires one unstable eigenvalue for each jump variable. As shown by Woodford (2003) and others, a necessary condition for the uniqueness of the equilibrium is that the parameters of the nominal interest rate rule satisfy:[4]

$$\phi_\pi + \frac{1-\beta}{\kappa}\phi_y > 1$$

This is slightly more complicated than the Taylor principle as originally espoused, which simply calls for the coefficient on inflation in the policy rule to exceed unity. That being said, with the discount factor close to one, the condition requisite for the uniqueness of equilibrium is still approximately that the coefficient on inflation in the rule be greater than one.

While cosmetically similar, the underlying economics behind this condition are quite different than what Taylor (1993) had in mind. The Taylor principle was originally cast in the context of an "old" Keynesian model, which differs from the "new" model in that inflation is a state, rather than a jump, variable. In such a world, a coefficient on inflation in the policy rule in excess of unity is necessary to head off nominal explosions. In particular, the central bank must raise real interest rates whenever inflation increases to prevent inflation from accelerating. In the "new" Keynesian model, however, we typically rule out nominal explosions by assumption.[5] There, the satisfaction of the modified Taylor principle is necessary not to rule out explosions, but rather to ensure the uniqueness of the equilibrium.

As Cochrane (2007a and 2007b) stresses quite eloquently, though apparently similar, the Taylor principle in the old and new variants of the Keynesian model is quite different, and applying the "old" logic to the new model can be misleading. He offers an intriguing interpretation of satisfaction of the Taylor principle in the New Keyne-

[3] The jump variables are output and inflation. If one takes the rule to be given by (5), then the nominal interest rate can be eliminated, with v being a state variable. Under a partial adjustment rule like (7) the nominal interest rate is a state variable. The other state variable is the flexible price equilibrium level of output.

[4] Note that this condition holds for both (5) and the partial adjustment rule given by (7), since I wrote the $(1-\rho)$ term as multiplying the response coefficients to inflation and the output gap in (7).

[5] In his companion paper, Cochrane (2007a) criticizes the assumption of ruling out nominal explosions, arguing that there is nothing in economic theory which justifies this (i.e. there is no transversality condition for inflation). Allowing for non-bounded solutions would require analyzing the non-linear system of equations. The log-linearization as shown here is only valid in the neighborhood of the zero inflation steady state.

sian model. In particular, he argues that parameterizations of the policy parameters in the region of determinacy are tantamount to an off-the equilibrium path threat to "blow up the world". Since we rule out nominal explosions by assumption, we never see this threat actually carried out. Thus, Cochrane argues, data generated by the model can never reveal the parameters of the policy rule lending credibility to this threat, and concludes that these parameters are thus not identified.

3 Identification and Estimation

Cochrane's finding of non-identification of policy parameters is not a generic implication of the satisfaction of the Taylor principle in the New Keynesian model. While it is true that we never see the world "blow up" in the model, we do observe equilibrium fluctuations of the endogenous variables in response to exogenous shocks, and for policy rules like (5) or (7), these equilibrium fluctuations do reveal information about the underlying parameters of the rule.

In order to see this, it is helpful to first construct theoretical impulse responses to the two exogenous shocks of the model. Doing so requires picking values of the structural parameters. I set the discount factor, β, to 0.99, implying an annualized steady state discount rate of roughly four percent. I set the elasticity of intertemporal substitution, σ, to 1, corresponding to the common case of log utility over consumption. The slope of the Phillips Curve is set to 0.3.[6] I assume that the autocorrelation coefficient of the flexible price equilibrium level of output, given by γ in (3) above, is 0.95. So as to focus in on the primary coefficient governing determinacy in the Taylor rule, I set the coefficient on the gap, ϕ_y, in the rule equal to zero.[7] To highlight the role of different parameterizations in the equilibrium fluctuations of the model, I consider two values of the parameter on inflation in the region of determinacy, 1.25 and 1.5. I consider both Taylor's original specification of the rule given by equation (5) and the partial adjustment specification, (7). In the specification given by (5), I assume that the error term is autocorrelated as given by (6), with $\alpha = 0.8$. Under the partial adjustment specification in (7), I assume that the smoothing parameter, ρ, is also 0.8.

Figure 5.1 shows responses to a one percent technology shock and twenty-five

[6] This calibration corresponds with a Calvo parameter of 2/3, implying that firms get to change their prices on average once every three quarters, and a Frisch elasticity of labor supply of roughly 1. See Footnote 2.

[7] This is done for simplicity and does not affect any of my conclusions. Most empirical evidence suggests that the Fed does not respond strongly to gap measures anyway.

basis point policy shock under specification (5) for both parameterizations. The first panel shows that both inflation and the output gap fall in response to the technology shock. The fall in both variables is monotonically decreasing in the policy rule parameter – that is, for higher values of ϕ_π, the downward jumps in both inflation and the gap in response to a favorable technology shock are smaller. The second panel shows the response to the monetary policy shock. Again, the impact effects on both inflation and the output gap from the shock are decreasing in the size of the policy rule coefficient. Figure 5.2 repeats the same exercise for the partial adjustment description of monetary policy as given by specification (7). Here again, we see that the impact jumps in both inflation and the gap are smaller for larger values of ϕ_π, though the discrepancy in the jumps for the two calibrations is smaller than under specification (5).

We thus see that both policy and non-policy shocks affect inflation and the output gap, and further that the effects of shocks on these variables depend on the values of the parameters in the policy rule. As such, the policy rule parameters ought, in principle, to be identified in the model. As noted in the Introduction, the key for identification is that non-policy shocks affect both inflation and the gap. If the only observed variation in these two variables was due to policy shocks, we would not be able to recover the values of the policy parameters. While it is true that the impact responses of inflation and the gap following a policy shock are dependent on the policy parameters, the size of the responses also depends on the size of the policy shock. In the second panels of Figure 5.1, I could have simply altered the size of the policy shock to force the responses of both inflation and the gap to lie everywhere on top of one another for both $\phi_\pi = 1.25$ and $\phi_\pi = 1.50$. Without knowledge of the variance of policy shocks, we can thus not identify the policy parameters off of variation in the variables due to policy shocks. Identification must come from the interaction of non-policy shocks with inflation and the gap.

One can go about estimating the policy rule parameters either by single or multiple equation methods. Simple inspection of (5) and (7) reveals that OLS will not produce consistent estimates of the parameters of the rule. This is because both inflation and the gap are jump variables, and are thus contemporaneously correlated with the structural error term in both specifications. As such, consistent estimation of the policy rule parameters requires the use of valid instrumental variables.

The set of permissible instruments depends on the specification of the policy rule. The most commonly used instruments in empirical papers include lagged values of inflation, the gap, and nominal interest rates. While potentially valid under the partial adjustment specification, lagged endogenous variables are not acceptable instruments under (5). To see this, we can lag the specification one period and write it with a white noise error term:

$$i_t = (1-\alpha)i^* + \alpha i_{t-1} - \alpha\phi_y(y_{t-1} - y^f_{t-1}) - \alpha\phi_\pi(\pi_{t-1} - \pi^*) + \phi_y(y_t - y^f_t) + \phi_\pi(\pi_t - \pi^*) + u_t$$

Because u_t above is white noise, the lagged values of inflation, the interest rate, and output are econometrically exogenous. The current values of inflation and output still respond to the white noise policy innovation, and thus we still need instruments. Once lagged values of the endogenous variables will not work, as they already (implicitly, at least) appear in the rule. Because of the Markovian structure of the model, twice or more lagged endogenous variables are not permissible instruments either. This is because, conditional on the first lag of the endogenous variables, twice or more lagged variables convey no information about the current state of the system.

The only permissible instruments under Taylor's original specification of the policy rule with an autocorrelated error term are thus current and lagged values of the flexible price equilibrium level of output itself (or shocks to it). If the policy error is not autocorrelated, then lagged values of the endogenous variables are valid instruments, provided there is some persistence to natural rate shocks.[8] This presents a potential complication – both to central bankers and to economists trying to understand central bank behavior – as the flexible price equilibrium level of output is not directly observable. In reality, however, there are many factors affecting the economy's natural rate – among them government spending, tax rates, natural disasters, etc. – which are observable and which may serve as valid instruments for both inflation and the output gap. For the purposes of this paper, I simply assume that I can observe the flexible price equilibrium level of output.

To assess the ability of a simple linear regression to estimate the policy rule parameters, I simulate a data set of 500 observations using the calibration described above under the rule given by (5) with $\phi_\pi = 1.5$. The standard deviation of the technology shock is normalized to be unity and I set the standard deviation of the policy shock to be 0.5 (both innovations are drawn from a standard normal distribution). I discard the first 100 observations so as to limit the influence of arbitrary initial conditions. I then estimate a regression of the nominal interest rate on current inflation and the lags of inflation and the nominal interest rate. I include the lags so as to force the error term to be white noise, as discussed above, and instrument for current inflation

[8] If neither policy nor natural rate shocks have any persistence, then the correlation between current and lagged endogenous variables will be zero.

with the current level of the flexible price equilibrium level of output. I repeat this process 1000 times.

The mean estimate of ϕ_π from the 1000 Monte Carlo simulations is 1.56, which is very close to the true value of 1.5. The mean value of α comes out to be 0.81, almost exactly equal to the true value. Figure 5.3 depicts the histogram of estimates of across the Monte Carlo simulations. While approximately unbiased, the implied standard errors of the estimate are rather large, with the 95 percent empirical confidence bands given by [1.19, 2.36]. Nevertheless, only one percent of the estimates lie below the critical value of unity. One can easily verify that, letting the size of the sample become arbitrarily large, the empirical distribution of point estimates collapses on the true value. The precision of the estimates (as measured by the confidence bands) is decreasing in the size of the true coefficients. As the policy rule coefficients become large, the amount of exogenous variation in inflation and the gap becomes smaller (see the theoretical impulse responses in Figure 5.1), and thus estimation of the policy parameters becomes more difficult.

I next consider a similar Monte Carlo exercise when policy is governed by the partial adjustment mechanism. Here the set of permissible instruments is different. Because we now interpret the error term to be white noise, lagged values of inflation and the output gap are valid instruments, because lags of these variables do not otherwise appear on the right hand side of the rule. Lagged values of the interest rate are not permissible, nor are twice or more lagged values of inflation or the gap once the first lags are included in the instrument set. Current and lagged values of the flexible price equilibrium level of output remain valid. There is a potential incongruity between valid instruments in the model and the instruments used in actual estimation in a number of empirical papers. CGG (2000), for example, assume a policy rule nearly identical to (7).[9] Their instrument set includes multiple lags of the nominal rate, inflation, and an empirical measure of the gap, as well as lags of other variables which do not appear in their model. As noted above, however, multiple lags of the endogenous variables are not permissible instruments due to the Markovian structure of the model.

I examine the properties of the IV estimator using both lagged inflation and the current flexible price equilibrium level of output as instruments in isolation and together. The mean estimates and empirical distributions of estimates are virtually

[9]CGG assume that the Fed reacts to expected inflation, not current inflation as presented here. There is no explicit policy shock in their rule, but there is a structural error term in their regression due to forecast errors. The assumption that the Fed reacts to expected inflation as opposed to current has little bearing on any of my results.

identical using the different instrument sets. For all three instrument sets, the mean estimate of ϕ_π across the Monte Carlo simulations is about 1.65, while the average point estimate of the smoothing parameter comes out to be 0.79. Figure 5.4 depicts the empirical distribution of estimates of ϕ_π using either once lagged inflation or the current flexible price equilibrium level of output as instruments in isolation. Here the small sample bias appears to be more significant than under Taylor's original specification of the rule. The dispersion of the estimates also appears to be greater under rule (7) than rule (5). Indeed, the empirical 95 percent confidence bands come out to be roughly [0.74, 3.60] for either instrument, which is significantly wider than before.

The intuition for the less precise estimates under the partial adjustment rule is made obvious by close inspection of the theoretical impulse responses to the shocks shown in Figures 5.1 and 5.2. The difference in jumps of inflation and the gap in response to the technology shock for the different calibrations of the policy parameter is significantly smaller under the partial adjustment rule relative to the original Taylor rule, thus likely accounting for the poor estimates under rule (7).[10] Even with a sample size of 400 (which is much larger than the samples to which researchers typically have access), while we can reject that it is greater than zero, we cannot reject that ϕ_π is greater than unity at any sensible level of statistical significance. As with the original Taylor Rule, as I let the sample size become arbitrarily large, the distribution of estimates collapses on the true value.

As noted by Lubik and Schorfheide (2004), full system based estimates of the policy parameters are more efficient than single equation IV estimates. The drawback of multiple equation based estimates is that they are more susceptible to misspecifications elsewhere in the model. Systems based estimation is also complicated by the fact that, as written, the model suffers from stochastic singularity. Because there are more variables than shocks, and more parameters than variables, identification of the complete parameter vector of the model is extremely difficult, if not impossible. Though I do not fully pursue this type of estimation in this paper, I have verified that system based GMM estimates of the policy parameters are consistent once I calibrate some of the other parameters of the model.

In spite of the fact that the small sample properties of the single equation IV estimators are not particularly good, it is clear that one can, in principle, consistently

[10] One can see that the impact jumps in both inflation and the gap with $\phi_\pi = 1.25$ are roughly double those with $\phi_\pi = 1.50$ under Taylor's original specification (Figure 1(a)). Under the partial adjustment rule (Figure 1(b)), the difference in jumps is closer to a factor of 1.5.

estimate the policy parameters from data generated from the baseline New Keynesian model with a standard description of the central bank's nominal interest rate rule. How, then, does Cochrane (2007b) obtain his non-identification result? He assumes a policy rule that looks nearly identical to (5), but with one small (and important) difference – there is a time subscript on the target nominal rate:

$$i_t = i_t^* + \phi_y(y_t - y_t^f) + \phi_\pi(\pi_t - \pi^*) + v_t \tag{8}$$

Cochrane refers to the time-varying target rate as a "stochastic intercept". This is an unfortunate and potentially misleading use of terminology. After all, one could interpret innovations to v as a time-varying intercept as well, so it is not immediately clear that (8) is conceptually any different than (5). Rather than allowing i^* to vary randomly, however, Cochrane assumes that it tracks fluctuations in the Wicksellian natural rate of interest, as defined in (4) above:

$$i_t^* = r_t^f$$

It is easy to see that, if the target rate varies one for one with the Wicksellian natural rate of interest, then a no gap/no inflation outcome is a potential equilibrium of the model absent policy shocks. The modified Taylor principle necessary for the uniqueness of the equilibrium is unaffected by the presence of the stochastic intercept. Hence, so long as the policy rule parameters are chosen so as to satisfy the modified Taylor principle, the equilibrium of the model will be unique, and both inflation and the output gap will be completely stabilized in response to any non-policy shock. Note that it is not, in general, possible for a partial adjustment rule such as (7) to implement the no gap/no inflation equilibrium, even with the stochastic intercept term tracking fluctuations in the Wicksellian natural rate.[11]

Figure 5.5 shows theoretical impulse responses to the two exogenous shocks of the model under the same calibration as above but with the stochastic intercept rule. As before, I show responses under two values of the policy parameter on inflation, 1.25 and 1.5. As noted, the non-policy shock induces no equilibrium movement in

[11] Following King (2000), Cochrane usually write the rule in terms of deviations from the Fed's chosen equilibrium, e.g. the term on the RHS of the rule would be $\pi_t - \pi_t^*$, where π_t^* denotes the central bank's desired level of inflation at that particular moment in time. Under a stochastic intercept rule, the central bank solely determines the equilibrium level of inflation. Thus, the rule allows the Fed to insure that the actual level of inflation is equal to its desired level, so that $\pi_t = \pi_t^*$ at all times, and the policy parameter is not identified. Under the more standard rules such as (5) or (7), the central bank is not able to always implement its desired level of inflation in the face of external shocks.

either inflation or the gap for either value of the policy rule parameter. Because the stochastic intercept is unaffected by a policy shock, the responses to a policy shock are identical to what is shown in Figure 5.1. It is thus obvious that the policy rule parameters do appear in the equilibrium dynamics of data generated from the model, even under the stochastic intercept rule.

Nevertheless, the policy parameters are not identified. While the model does produce variation in the variables of the model which is influenced by those parameters, this variation is not exogenous. In particular, the lack of identification results because the assumptions on the policy rule ensure that non-policy shocks have no effect on either inflation or the output gap. Thus, the model yields no permissible instruments and hence no way of consistently estimating the policy parameters in a linear regression.

There is also no hope of estimating these parameters through systems based estimation. In particular, the likelihood function is not single-peaked for this specification of the model. Intuitively, the volatilities of inflation and the gap are both influenced by the variance of policy shocks and the magnitude of the policy rule parameters. The model is capable of matching given volatilities either with a relatively high value of the variance of policy shocks and a low value of the policy rule parameters, or with a low variance of the shock and high coefficients in the rule. Without knowledge of the variance of policy shocks, it is simply not possible to recover the true value of the policy rule parameters under this specification of the rule.

Because the modified Taylor principle itself is unaffected by the presence of the stochastic intercept, it is clear that non-identification is not a generic implication of the satisfaction of the Taylor principle in the New Keynesian model, but is rather the result of this particular specification of the policy rule. Cochrane's assertion of non-identification does not apply to existing empirical work, since most or all existing empirical papers assume interest rate rules like (5) or (7).

All of the above results assume that the policy rule parameters are such that the economy is in the region of determinacy – in other words, I have assumed throughout that the modified Taylor principle is always satisfied and therefore that the equilibrium of the economy is unique. I close this section with a brief discussion about estimation and identification when the Taylor principle is not satisfied.

Although it may at first seem non-intuitive, policy coefficients in the region of indeterminacy ought to be better identified (and more precisely estimated) than those yielding uniqueness. In the simple setup above, with and $\phi_y = 0$ and $\phi_\pi < 1$, the policy function mapping states into output and inflation would not be unique. In other words, for a given state of the world, there are an infinite number of combinations of output and inflation consistent with the equations of the model holding and with non-explosion. Pinning down the actual equilibrium from this set of possible equilibria requires the introduction of an extrinsic coordinating variable – a "sunspot" – which is otherwise completely independent of the rest of the model. For a clear description of how to incorporate sunspots into the model, the interested reader is referred to the working paper version of CGG (available as NBER WP # 6442, 1998).

The mapping between the sunspot realizations and the jump variables (output and inflation) of the model is completely arbitrary, so it is difficult to say much about the actual behavior of the economy in the region of indeterminacy. We can, however, make a few general observations. First, the sunspot will, in general, lead to fluctuations in both inflation and output which are unrelated to the two fundamental shocks. Second, the presence of sunspots will alter the response of output and inflation to the fundamental shocks. As it pertains to identification and estimation, the first observation is important. In particular, the sunspot introduces an additional source of exogenous variation to both inflation and the output gap. Under policy rules like (5) or (7), this source of additional variation has no bearing on whether or not the policy rule parameters are identified. Because it is an additional source of exogenous variation in both inflation and the gap, however, the sunspot itself is now also in the set of permissible instruments (the other valid instruments remain the same as in the discussion above). Even if the sunspot is unobservable, the additional source of exogenous variation in both inflation and the gap is likely to improve the precision of IV estimates of the policy rule coefficients.

In the region of determinacy, Cochrane's specification of the rule given by (8) eliminates all exogenous variation in both inflation and the gap, thus leading to his non-identification result. Once the coefficients place the economy in the region of indeterminacy, however, the policy shock is no longer the only source of variation in these variables. In particular, the non-fundamental sunspot shock will, in general, affect both inflation and the gap on impact. Furthermore, fundamental non-policy shocks (which manifest themselves as fluctuations in the Wicksellian natural rate of interest), will also, in general, affect both inflation and the output gap. Thus, Cochrane's non-identification result only applies under rule (8) when the Taylor principle is satisfied. If the Taylor principle is not satisfied, there will exist exogenous variation in inflation and the gap off which to identify the policy parameters. Thus, these parameters may, as a matter of principle, be consistently estimated.

I repeat the Monte Carlo exercise from above for the system in the region of indeterminacy assuming that I can observe the sunspot realizations. I assume that the policy rule is given by (8) and that $\phi_y = 0$ and $\phi_\pi = 0.9$. The rest of the parameters are calibrated as above, and I draw sunspot realizations from a standard normal distribution. For simplicity, I assume that the mapping between sunspot realizations and the jump variables is such that the innovation in current inflation is equal to the sunspot realization.[12] I run a regression of the nominal interest rate on its own lag, the lag of inflation, and current inflation, which forces the structural error term to be white noise. I do not include the Wicksellian natural rate of interest in the regression. While this leads to a composite error term which is not white noise, it does not affect the consistency of the estimates of the policy rule coefficients, since the sunspot realization which serves as my instrument is uncorrelated with the natural rate of interest.

Figure 5.6 depicts the histogram of estimates of ϕ_π under this policy rule in the region of indeterminacy. The estimates are approximately unbiased, with the mean estimate of α equal to 0.79 and the mean estimate of ϕ_π equal to the truth at 0.90. The empirical confidence bands are much tighter than those in the region of determinacy, with 95 percent of the estimates of lying in the region [0.80, 1.00]. For the kind of question posed by CGG (2000) – did the Fed move from a "passive" policy in the 1970s to a more "active" one in the 1980s – the appropriate null hypothesis is that the policy parameters are in the region of indeterminacy. The Taylor rule parameters are always identified under this null hypothesis, even under the stochastic intercept rule.

4 Welfare Evaluation and Empirical Evidence

The previous section established that it is not a generic implication of the New Keynesian model that the policy rule parameters are not identified in the region of determinacy. Rather, non-identification only obtains in the case in which the central bank's rule features a stochastic intercept term which adjusts one for one with movements in the Wicksellian natural rate of interest. In this section, I address both a normative and a positive question related to this kind of policy rule. First, should the central bank attempt to adhere to a stochastic intercept rule? Second, do central banks try to follow such a rule? My answer to the former is probably not, and almost certainly not to the latter.

[12]This assumption is made only for simplicity. My results are unaffected for more complicated mappings between sunspot realizations and the jump variables.

There are two distortions reflected in the equilibrium level of output in the New Keynesian model as described above. One comes from the assumed price stickiness, which potentially leads to discrepancies between the actual and flexible price equilibrium levels of output. The second distortion is the monopolistic competition which gives rise to price-setting in the first place. Optimal policy should target the Pareto optimal level of output that would obtain in the absence of both of these distortions. As shown in Blanchard and Gali (2007), in the standard New Keynesian model, the first best and flexible price equilibrium levels of output differ by a constant, and thus stabilization of the output gap (the gap between actual and flexible price output) is equivalent to maximization of the welfare relevant gap (the gap between actual and first best output). As shown above, the stochastic intercept rule stabilizes both inflation and the output gap, and thus is the welfare maximizing policy rule in the basic model.

While the stochastic intercept rule maximizes welfare in the model, it may or may not be optimal once real world considerations are taken into account. In particular, such a rule requires the central bank to observe the Wicksellian natural rate of interest in real time. This natural rate is in fact not observable and is a potentially complex function of several underlying shock processes. If the central bank can only observe the natural rate with noise, then it is possible that the resulting increase in the variance of policy shocks could result in lower welfare relative to the more standard constant intercept specification.

I compare welfare under both the standard Taylor rule (5) and the stochastic intercept rule (8). Because it is generally not possible to implement the no gap, no inflation equilibrium under a partial adjustment rule, a welfare comparison between the stochastic intercept rule and a rule similar to (7) is unnatural. I use the standard deviation of the output gap as a welfare metric. The first two columns of Table 5.1 show the model standard deviations of the gap under both the stochastic intercept and standard Taylor rules for different values of γ, which is the persistence of the flexible price equilibrium. The rest of the model parameters are calibrated as above. In particular, the policy rule parameters are set at $\phi_y = 0$ and $\phi_\pi = 1.5$.

Because the stochastic intercept rule stabilizes the gap in response to non-policy shocks, the standard deviation of the gap under such a rule is invariant to different values of γ. For high values of γ, the standard deviation of the gap under the standard Taylor rule is almost identical to the stochastic intercept rule (standard deviation of 0.6352 vs. 0.6331 when $\gamma = 0.95$). As γ decreases the welfare differences between the two rules increase. The intuition for this is straightforward upon inspection of the definition of the Wicksellian natural rate (equation (4)). For high values of γ, fluctuations in the Wicksellian natural rate are small, and thus a constant intercept rule comes close to replicating the stochastic intercept rule.[13] As real shocks become less persistent, the natural rate of interest fluctuates more, and the constant intercept rule is further from the stochastic intercept rule. That being said, even when real shocks are almost white noise ($\gamma \to 0$), the standard deviation of the gap under the standard Taylor rule is only about 50 percent larger than under the stochastic intercept rule.

I next make the realistic assumption that the central bank can only observe a noisy signal of the Wicksellian natural rate. In particular:

$$r_t^* = r_t^f + n_t \tag{9}$$

r_t^* denotes the central bank's observed natural rate and n_t is the noise in their observation. I assume that n_t is uncorrelated with all other shocks in the economy. I allow misperceptions of the natural rate to be persistent by modeling the noise term as a stationary AR(1):

$$n_t = \delta n_{t-1} + \eta_t \tag{10}$$

With this specification of the observed natural rate, the stochastic intercept rule can now be written as:

$$i_t = r_t^f + \phi_\pi(\pi_t - \pi^*) + v_t + n_t \tag{11}$$

This is identical to (8), but now with a composite error term reflecting exogenous shifts in policy (v_t) and misperceptions of the natural rate of interest (n_t).

The last three columns of Table 5.1 show the value of the standard deviations of noise shocks (std(η)) for different persistence parameters of noise for which the standard Taylor rule with a constant intercept welfare dominates the stochastic intercept rule. For highly persistent real shocks, one can see that it only takes a very small amount of noise in the central bank's observation of the natural rate for the constant intercept rule to be optimal. To get an idea for how large the standard deviation of noise shocks must be, Table 5.2 shows the standard deviation of the Wicksellian natural rate under different calibrations of γ. For highly persistent real shocks, one

[13] In the model without capital, if real shocks obey an exact random walk ($\gamma = 1$), then the Wicksellian natural rate of interest is constant and the constant intercept rule is equivalent to a stochastic intercept rule.

can see that the required standard deviation of noise disturbances can be quite a bit smaller than the standard deviation of the natural rate itself. As the persistence of real shocks rises, the variance of noise shocks required for the constant intercept rule to be welfare dominant also rises. If one believes that real economic shocks are highly persistent (which they appear to be in the US), then it is clear that it takes only a small amount of imprecision in the central bank's observation of the Wicksellian natural rate for a standard Taylor rule to result in higher welfare than a stochastic intercept rule.

There are additional reasons why a stochastic intercept rule may be either infeasible or undesirable. For one thing, central banks cannot implement such a rule in times in which the natural rate of interest is negative, given the zero lower bound on the nominal interest rate. Secondly, as stressed repeatedly throughout Woodford (2003), it is not so much the actual coefficients and structure of the policy rule which matter, but rather that the agents in the economy know the structure of the rule and believe that the central bank will remain committed to it. For the same reasons that the parameters cannot be identified by an econometrician, there is no mechanism by which agents in the economy could ever learn the values of the response coefficients under a perfect stochastic intercept rule. Even if the central bank were to publicly announce values of ϕ_π and ϕ_y, there would be no observable action lending credibility to the announcement, and thus no reason for households to believe them. Lastly, there are many realistic modifications of the baseline model in which the stochastic intercept rule is not necessarily optimal. Such modifications frequently involve a "cost push" shock in the Phillips Curve, which makes it impossible for the central bank to simultaneously stabilize inflation and the gap. More generally, there are several realistic features in which the distance between the flexible price and first best equilibrium levels of output is not constant; such features include real wage stickiness or time-varying monopoly power (see Blanchard and Gali (2007)). In such instances, optimal policy would not even attempt to stabilize the actual output gap or inflation in response to external shocks.

I now turn to a positive question: does the Fed attempt to follow a stochastic intercept rule? It is not possible to directly estimate the policy rule under the null hypothesis of the stochastic intercept term. There are, however, testable implications of the model with such a rule which do not directly depend on the values of the policy parameters (as long as the economy is in the region of determinacy). In particular, the model with the stochastic intercept rule has the sharp prediction that the only source of variation in inflation is the policy shock. Real shocks (by which I mean shocks which affect the flexible price equilibrium of output) have no impact on inflation in the model augmented with a stochastic intercept rule.

It is thus possible to test whether or not the Fed obeys a stochastic intercept rule by examining the conditional relationships between real shocks and inflation. To identify real shocks I estimate a bivariate VAR featuring real GDP and inflation and orthogonalize the innovations into "demand" and "supply" through a long run restriction that "demand" shocks have no long run effect on the level of output. I use the terms "demand" and "supply" only begrudgingly so as to facilitate comparison with the literature. In reality, the identified "supply" shock is a compilation of anything which can have a long run effect on output, while the "demand" shock reflects forces only having a temporary effect. This VAR is similar to the one in Gali, Lopez-Salido, and Valles (2003), with the exception that their VAR features output per hour. My identifying assumption is that real shocks may have a permanent effect on output, but monetary policy shocks may not. Thus, the structural policy shock from the model above is subsumed in the identified "demand" shock while other factors imparting a unit root on output are reflected in the "supply" shock.[14]

I estimate the VAR using quarterly data on real GDP growth and CPI inflation from 1960-2006 with four lags of each variable.[15] The impulse responses of both the level of real GDP and inflation to the identified "demand" and "supply" shocks are show in Figure 5.7. The dashed lines are 90 percent confidence bands from the bias-corrected bootstrap procedure of Kilian (1998). The shock identified as having a permanent effect on output explains the bulk of the inflation innovation variance. In particular, a shock raising output in the long run is associated with a reduction in inflation (at an annualized rate) of more than one and a half percent on impact. This effect is statistically significant for roughly six quarters. By comparison, the "demand" shock induces an increase in inflation, though the impact effect is not statistically significant. After a few quarters the demand shock explains inflation more significantly. These results are clearly at odds with the predictions of a stochastic intercept rule, under which there should be no relationship between the identified "supply" shock and inflation. The point estimates are somewhat smaller and the statistical significance is weaker in the post-Volcker part of the sample, though

[14] In the baseline model of the previous section there is no unit root in output. In a version of the model in which there are shocks inducing a unit root in output it is still the case that monetary policy shocks have no long run effect on output.

[15] I arrive at a quarterly inflation measure by calculating the percentage change in the seasonally adjusted CPI from the last month of each quarter. Output enters the VAR in first differences and I show the cumulated level response in the figures.

the same qualitative pattern emerges. That the response of inflation to “supply” shocks is smaller in the latter half of the sample is consistent with larger response coefficients in the policy rule.[16]

While clearly not consistent with a stochastic intercept rule, the impulse responses from the empirical VAR are also not consistent with a more standard Taylor type rule. The impulse responses suggest that there is a significant predictable increase in output following a “supply” shock. To be consistent with the Euler equation, there must also be an increase in real interest rates following such a shock. However, it is extremely difficult for a standard Taylor rule with coefficients in the region of determinacy to produce a simultaneous decrease in inflation and increase in real rates in response to non-policy shocks. With $\beta \approx 1$ and θ not too close to zero, the Phillips Curve (2) suggests $\pi_t \approx E_t\pi_{t+1}$. Using this approximation, subtracting $E_t\pi_{t+1}$ from both sides of (5), and ignoring the gap term, we see that:[17]

$$r_t = i_t - E_t\pi_{t+1} \approx (\phi_\pi - 1)\pi_t$$

With $\phi_\pi > 1$, an increase in real interest rates absent a policy disturbance requires an increase in inflation.[18] I should stress that this is an approximate result – there are values of the parameters in which there is a simultaneous increase in real rates and decrease in inflation, but this only occurs when the increase in real interest rates is small. For the amount of predicted increase in output shown in Figure 5.7, however, the increase in real interest rates would have to be rather large. In that case, it is almost certainly true that the model with a Taylor rule would require an increase in inflation, not a decrease as we see in the data. These empirical results thus indicate not only that the Fed likely does not adhere to a stochastic intercept rule, but that the more commonplace specifications of interest rate rules may also be misspecified.

[16] Gali, Lopez-Salido, and Valles (2003) find an insignificant response of inflation to identified technology shocks in the latter half of the sample, and intrepret this as evidence in favor of the Fed following a stochastic intercept rule in that time period. The point estimates even in that sample are far too large for such an interpretation. It is extremely difficult to ascertain statistical significance on the basis of roughly 80 quarterly observations with relatively persistent data series. While not necessarily statistically significant, the identified response of inflation to “supply” shocks remains economically significant in the post-Volcker part of the sample.

[17] Ignoring the influence of the gap is relatively innocuous here. With this same approximation, fluctuations in the gap are always very small and tend to co-move positively with inflation anyway.

[18] The exact same condition also obtains under the partial adjustment specification. Assuming that both $i_t \approx i_{t-1}$ and $\pi_t \approx E_t\pi_{t+1}$ and simplifying yields the same approximate relationship between real interest rates and inflation under (7) as (5).

5 Conclusion

A Taylor type nominal interest rate rule has become ubiquitous in the monetary economics literature and is almost universally accepted by macroeconomists as both a good description and prescription for the conduct of monetary policy. Likewise, the New Keynesian model with explicit micro foundations and optimizing behavior has become one of the workhorse models for analyzing short run fluctuations. As such, fully understanding the interplay between interest rate rules and the New Keynesian model – and whether or not data generated from the model are informative about the structure of the policy rule – is an important task.

The received wisdom in macroeconomics is that the Fed helped to stabilization inflation and the US economy by switching to a more active monetary policy under Paul Volcker in the early 1980s. This conclusion is largely based on regression estimates of interest rate rules, which generally find response coefficients on inflation and other variables that were too low in the 1970s and much higher thereafter. In a recent paper, Cochrane (2007b) has challenged this conclusion, arguing that the policy rule coefficients are not identified in the New Keynesian model. Whereas in the "old" Keynesian model sufficiently large policy rule coefficients are necessary to prevent explosive dynamics, in the New Keynesian model the policy rule coefficients must be sufficiently large to render the equilibrium unique. Satisfaction of the Taylor principle in the New Keynesian model imparts an unstable root into the dynamic system of equations. Cochrane argues that, since New Keynesian modelers rule out explosive behavior *a priori*, the equilibrium dynamics of the model cannot reveal any information about the unstable root. In other words, he argues, policy rule parameters in the region of determinacy are not identified.

I demonstrated that non-identification is not a generic implication of the model, but rather results from a particular (and unrealistic) assumption on the policy rule itself. For a standard specification of the interest rate rule – similar either to Taylor's original specification or the more common partial adjustment specification – the policy rule parameters are in fact identified and may be estimated using standard techniques. The key for identification is for real shocks, by which I mean shocks which would affect output in a flexible price model, to influence inflation and the output gap. This condition is satisfied in the standard New Keynesian model, and thus the model produces exogenous variation in inflation and the output gap off which one may consistently estimate the policy rule parameters.

Cochrane's non-identification result requires that the interest rate rule feature a "stochastic intercept" which tracks fluctuations in the Wicksellian natural rate of interest. Provided that the response coefficients on inflation and the gap satisfy the modified Taylor principle, such a rule serves to completely stabilize inflation and the output gap in response to any non-policy shock. By eliminating any exogenous variation in inflation and the gap, such a rule renders the policy rule parameters unidentified. Since the gap between the actual and flexible price levels of output is the appropriate welfare metric, such a rule is optimal provided that the central bank can in fact observe the Wicksellian natural rate in real time. Nevertheless, the welfare losses from a more standard constant intercept specification of the policy rule are small, and I showed that it requires only a small amount of noise in the bank's observation of the Wicksellian natural rate for the stochastic intercept rule to actually result in lower welfare.

While it is not possible to directly estimate the policy rule under the null of a stochastic intercept, there are testable implications of such a rule which do not rely on the particular values of the rule's response coefficients. In particular, the stochastic intercept rule leads to the stark prediction that the only source of variation in inflation are monetary policy shocks. Since policy shocks can only have temporary effects on output in the model, a direct test of the stochastic intercept rule is to see whether or not things permanently influencing output and spending affect inflation. The results from a structural bivariate VAR suggest that they do. In particular, I find that shocks permanently affecting output account for the bulk of the innovation variance in inflation. I interpret this finding as a rejection of the stochastic intercept rule. Without the stochastic intercept term, there is no inherent identification issue with respect to Taylor type rules in the New Keynesian model.

Nevertheless, the results from the bivariate VAR also pose challenges to a standard Taylor rule with a constant intercept, and suggest that either the policy rule or the model itself may be inconsistent with the data. The "supply" shock identified from the bivariate VAR is associated both with a large reduction in inflation and large predictable increases in output. For this to be consistent with the Euler equation, such permanent supply shocks must lead to higher real interest rates. It is extremely difficult for a standard Taylor rule satisfying the Taylor principle to simultaneously generate large declines in inflation and large increases in real interest rates. As this is exactly the pattern implied by the data, this suggests that there may be a more general misspecification issue regarding the interest rate rule or the New Keynesian model itself.

Table 5.1

Standard Deviations of Output Gaps Under Different Policy Rules

	Stochastic Intercept	**Standard Taylor**	**Noise Std.**		
	std$(y-y^f)$	std$(y-y^f)$	std(η)		
			$\delta=0.0$	$\delta=0.4$	$\delta=0.8$
$\gamma=0.95$	0.6331	0.6352	0.0770	0.0567	0.0407
$\gamma=0.75$	0.6331	0.7053	0.4654	0.3430	0.2455
$\gamma=0.50$	0.6331	0.8062	0.7474	0.5508	0.3942
$\gamma=0.10$	0.6331	0.8996	0.9570	0.7052	0.9136

The table shows the anaylitcal standard deviation of the output gap under the stochastic intercept and standard Taylor rules for different values of the persistence of the flexible price equilibrium level of output. The last three columns show the required standard deviation of noise in the observation of the natural rate of interest for different persistence terms for the standard Taylor rule to welfare dominate the stochastic intercept rule.

Table 5.2

Standard Deviation of Wicksellian Natural Rate

	std(r_t^f)
$\gamma=0.95$	0.1601
$\gamma=0.75$	0.3780
$\gamma=0.50$	0.5774
$\gamma=0.10$	0.9041

This table shows the standard deviation of the Wicksellian natural rate for different values of the persistence of the natural rate of output.

Figure 5.1

Theoretical Responses to Shocks under Rule (5)

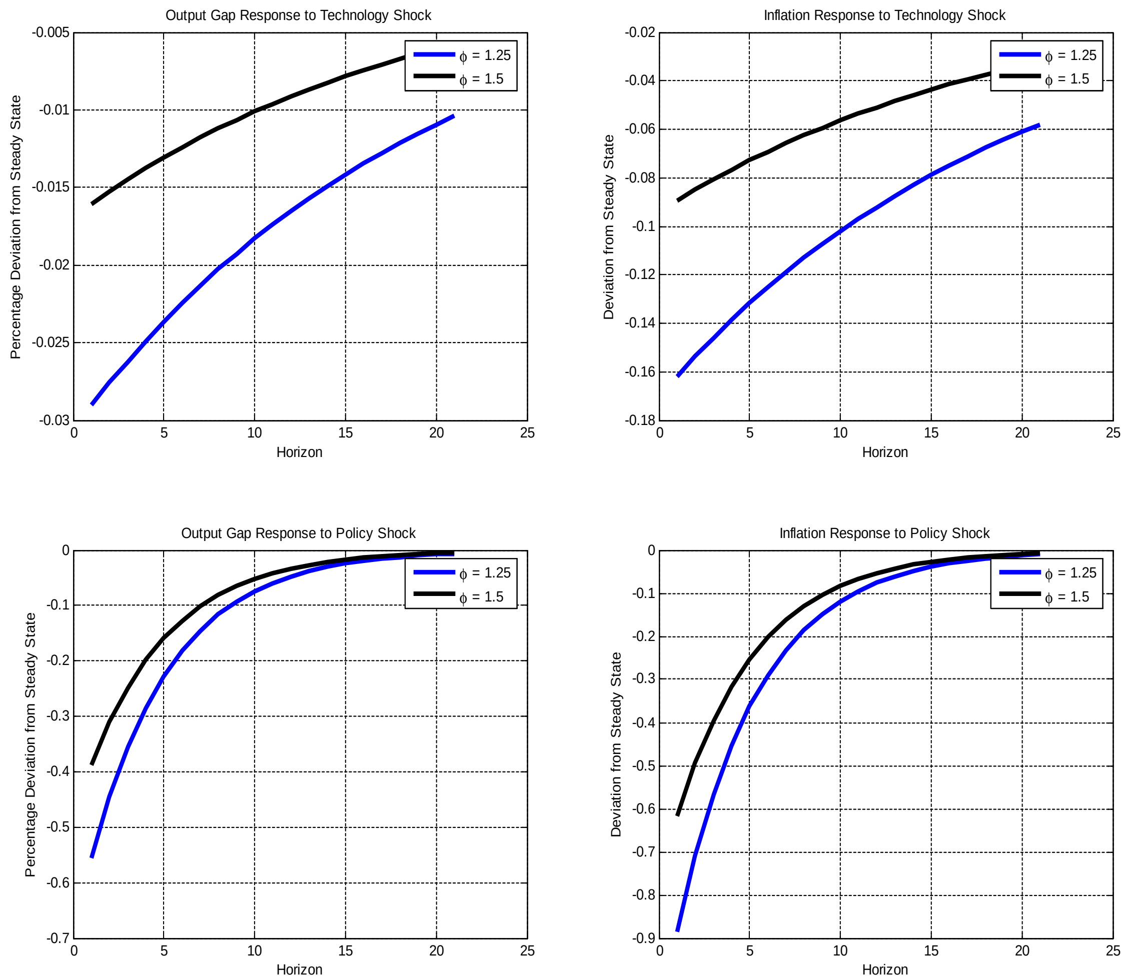

Figure 5.2
Theoretical Resposnes to Shocks under Rule (7)

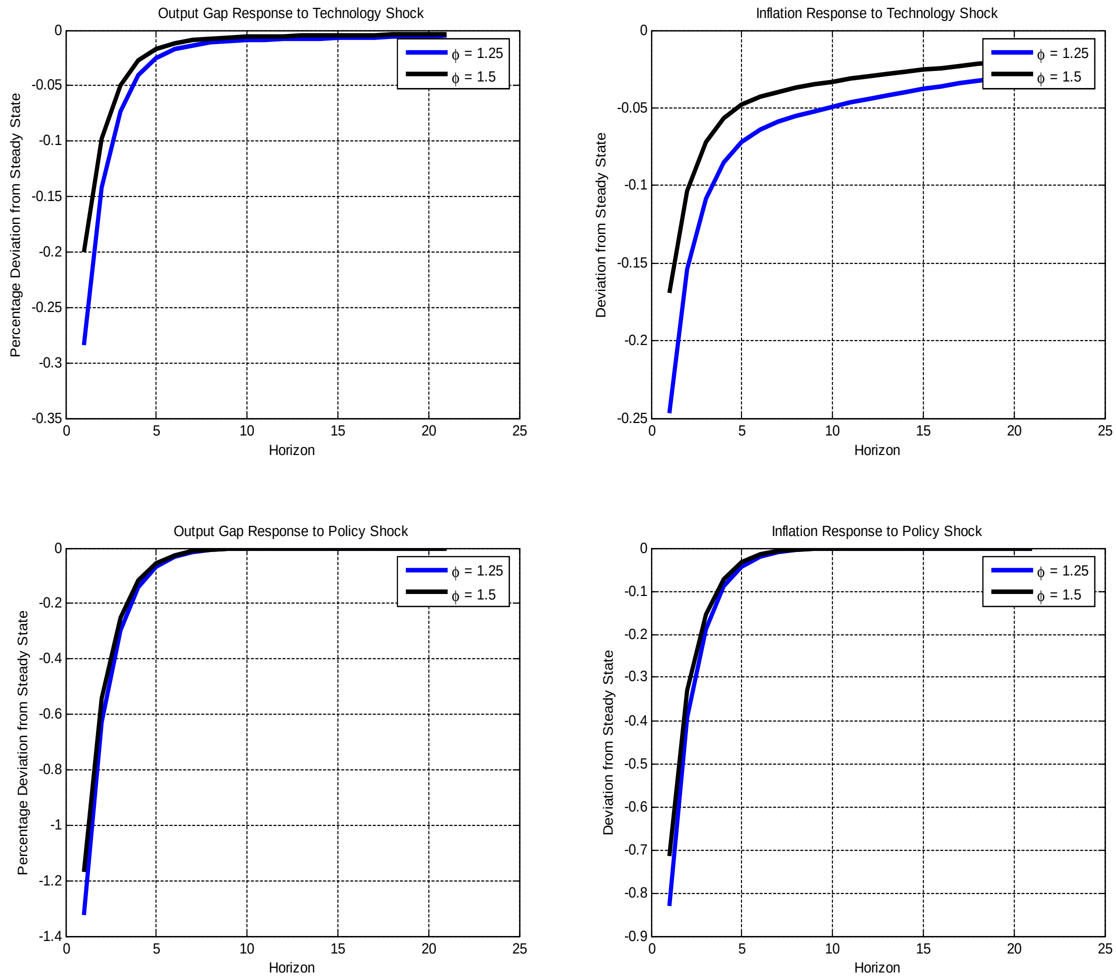

Figure 5.3

Histogram of IV Estimates of ϕ_π under Rule (5)

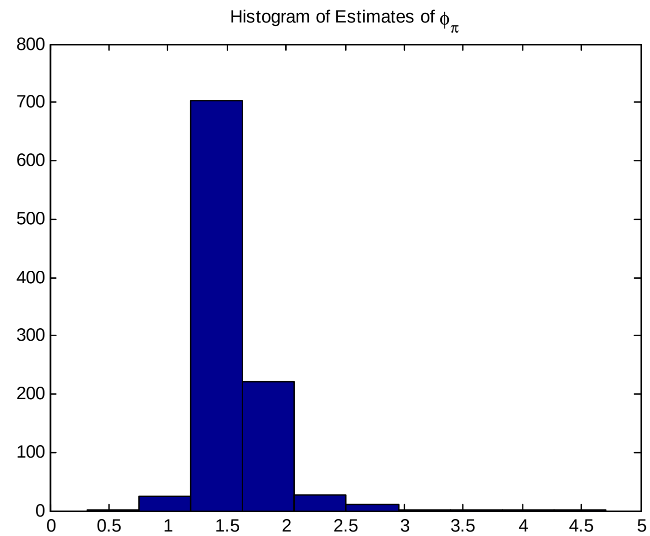

Figure 5.4

Histogram of IV Estimates of ϕ_π under Rule (7)

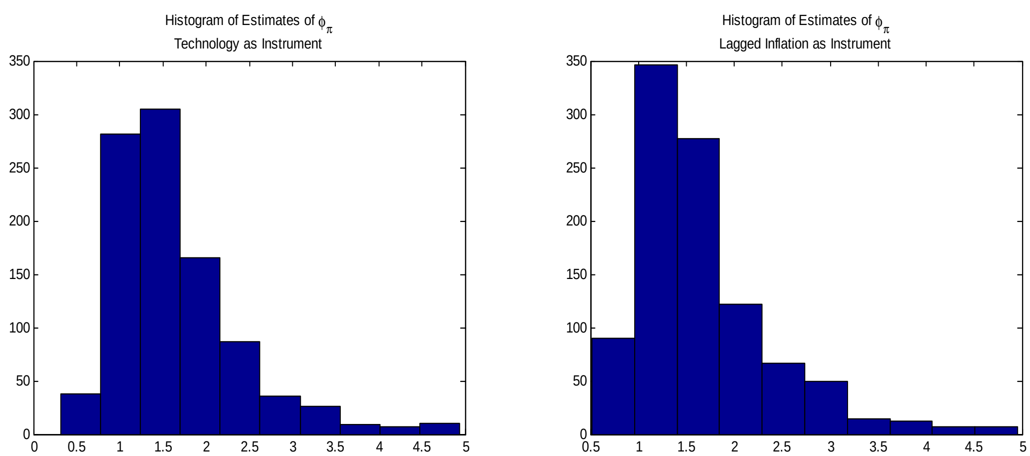

In both figures the true value of ϕ_π is 1.5.

Figure 5.5

Theoretical Impulse Responses to Shocks under Stochastic Intercept Rule (8)

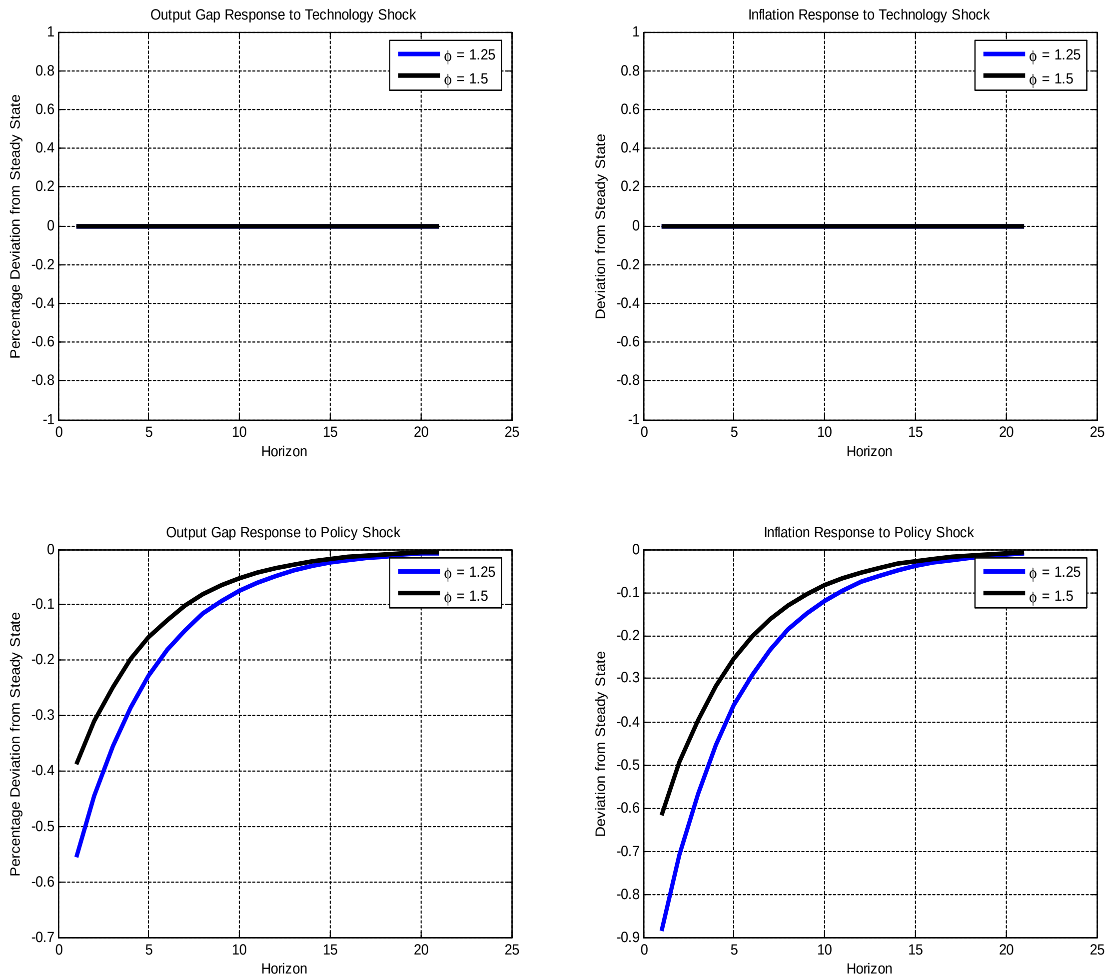

Figure 5.6
Histogram of Estimates of ϕ_π under Stochastic Intercept Rule (8)
In Region of Indeterminacy

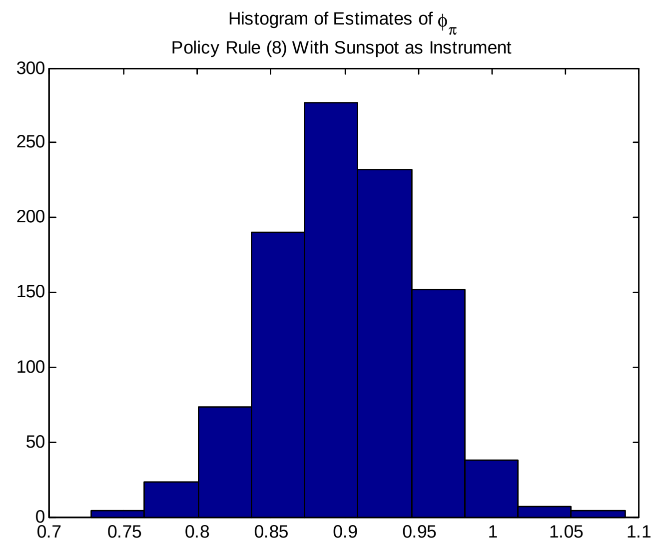

The true value of ϕ_π above is 0.9.

Figure 5.7
Impulse Responses of Output and Inflation from Bivariate VAR

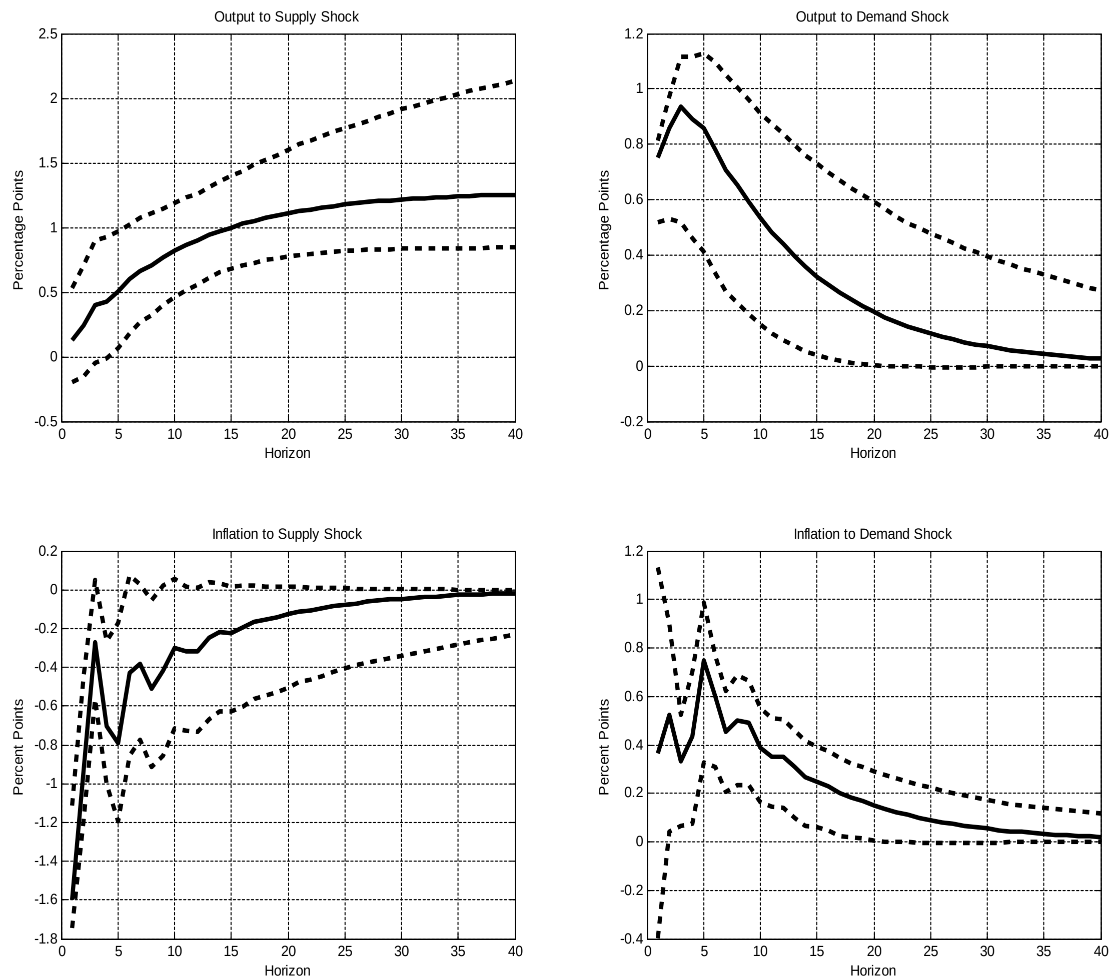

The above are impulse responses from a bivariate VAR featuring real GDP growth and inflation. The shocks are identified by a long run restriction that only supply shocks may lead to a long run response of output. The dashed lines are 90 percent bias-corrected bootstrap confidence bands.

References

[1] Calvo, Guillermo (1983). "Staggered Prices in a Utility Maximizing Framework." Journal of Monetary Economics, 12(3), pp. 383-398.

[2] Clarida, Richard, Jordi Gali, and Mark Gertler (2000). "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." Quarterly Journal of Economics, 105:1, pp. 147-180.

[3] Cochrane, John (2007a). "Inflation Determination with Taylor Rules: A Critical Review." University of Chicago Graduate School of Business working paper.

[4] Cochrane, John (2007b). "Identification with Taylor Rules: A Critical Review." University of Chicago Graduate School of Business working paper.

[5] Gali, Jordi (1999). "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" American Economic Review, pp. 249-271.

[6] Gali, Jordi and Olivier Blanchard (2007). "Real Wage Rigidities and the New Keynesian Model." Journal of Money, Credit, and Banking, 39, pp. 35-56.

[7] Gali, Jordi, J. David Lopez-Salido, and Javier Valles (2003). "Technology Shocks and Monetary Policy: Assessing the Fed's Performance." Journal of Monetary Economics, 50, pp. 723-743.

[8] Kilian, Lutz (1998). "Small Sample Confidence Intervals for Impulse Response Functions." Review of Economics and Statistics, 80, pp. 218-230.

[9] King, Robert (2000). "The New IS-LM Model: Language, Logic, and Limits." Federal Reserve Bank of Richmond Economic Quarterly, 86, pp. 45-103.

[10] Lubik, Thomas and Frank Schorfheide (2004). "Testing for Indeterminacy: An Application to US Monetary Policy." American Economic Review, pp. 190-217.

[11] Rudebusch, Glenn (2002). "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia." Journal of Monetary Economics, 49, pp. 1161-1187.

[12] Taylor, John B (1993). "Discretion vs. Policy Rules in Practice." Carnegie-Rochester Series on Public Policy, North Holland, 39, pp. 195-214.

[13] Woodford, Michael (2003). Interest and Prices. Princeton: Princeton University Press.

Chapter VI

Conclusion

There has been a recent resurgence of interest in “news shocks” about changes in future productivity. This dissertation has provided an in-depth study of these news shocks. It proposed and implemented a new VAR-based approach for the empirical identification of these shocks, and showed that this approach is likely to perform very well in practice. While news shocks have interesting implications for a variety of forward-looking variables, they do not appear to be an important source of business cycles.

The results presented in this dissertation touch on a number of different literatures in macroeconomics. Chapters II and III make a methodological contribution to the literature on VAR-based identification of economic shocks, and more broadly show that structural VAR techniques often perform quite reliably in practice. Chapter II documents that news shocks have important implications for a variety of forward-looking variables, the most surprising of which is inflation. In contrast to recent empirical results, Chapter III argues that news shocks about future productivity are not a dominant source of business cycles, showing that news shocks induce conditional comovement among aggregate variables strongly at odds with the salient fact of aggregate comovement in the data. Chapter IV contributes to the literature on the meaning of surprise movements in consumer confidence, finding that confidence innovations are surprisingly informative about movements in future fundamentals, while finding little evidence in support of an important animal spirits interpretation. Chapter V fits into the literature on estimation of monetary policy rules.

There are several different avenues for future related research. The proposed empirical methodology could be applied to the identification of news shocks in other plausibly exogenous series, such as oil prices, fiscal policy, etc.. The results in Chapters II and V call for further work on the specification of monetary policy rules. One particular finding of interest in Chapters II and III is that stock price innovations are only weakly explained by news shocks. While news shocks appear to account for important low frequency movements in stock prices, stock price innovations appear to have a roughly transitory component. Understanding what this transitory component is – perhaps investor animal spirits – is an important task for future work.