

An Estimated Dynamic, General Equilibrium Model for Monetary Policy Analysis* (Preliminary and Incomplete)

David Altig[†], Lawrence J. Christiano[‡], Martin Eichenbaum[§] and Jesper Linde[¶]

June 16, 2002

Abstract

We describe a strategy for fitting a multiple shock, dynamic general equilibrium model to the data. In this draft, we describe results for a version of the model with monetary policy shocks and with persistent shocks to technology. The model does well at reproducing the dynamic effects of these shocks.

*Christiano and Eichenbaum are grateful for the financial support of a National Science Foundation grant to the National Bureau of Economic Research. We are grateful for helpful comments from Susanto Basu, Valerie Ramey and Harald Uhlig. We particularly want to thank Eduard Pelz for excellent research assistance. This paper does not necessarily reflect the views of the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Cleveland, or the Riksbank.

[†]Federal Reserve Bank of Cleveland

[‡]Northwestern University, National Bureau of Economic Research, and Federal Reserve Banks of Chicago and Cleveland.

[§]Northwestern University, National Bureau of Economic Research, and Federal Reserve Bank of Chicago.

[¶]Riksbank

Our objective is to construct a dynamic general equilibrium model that matches the macroeconomic time series in detail. We want a model that can be used to address classic questions: what are the shocks that drive the business cycle? what is the optimal response of monetary policy to these shocks? is there a simple policy rule - like a Taylor rule - which comes close to supporting, as the unique equilibrium, the best outcome? is there a time consistency problem associated with this policy?

To develop credible answers questions like these requires a model that reproduces most of the variation in the data. The model we work with is a multishock version of the model estimated in Christiano, Eichenbaum and Evans (2002) (CEE). Our estimation strategy focuses on matching dynamic shock responses in the data to their counterparts in the model. The CEE model seems like a natural starting point, because the model has considerable success in matching impulse responses to a monetary policy shock.

Our model incorporates two types of disturbances: financial market shocks and non-financial market shocks. The former include a shock to household money demand, a shock to firm money demand, and a monetary policy shock. Non-financial market shocks include a persistent and a transitory shock to the technology for producing final goods, a shock to the technology for producing investment goods, a shock to the market power of firms and a shock to the market power of household labor suppliers. It has been argued that, in one way or another, each of these disturbances is important for the analysis of monetary policy questions. There is a large literature that stresses the importance of shocks to technology as impulses to the business cycle. There are also several papers which emphasize the role in business cycles of disturbances to household market power, which in our context are isomorphic to disturbances to preferences for leisure. Money demand shocks have been deemed to be so substantial, that many analysts of monetary policy find it convenient to abstract from money altogether. Finally, the importance of monetary policy shocks lies in the clues they leave regarding the structure of the economy.

To identify the effects of a monetary policy shock in the data, we adopt a recursiveness approach that has become standard in the literature. For the persistent technology shock, we follow Gali (1999) and Francis and Ramey (2001) in assuming that innovations to this shock are the only disturbance that affects the level of labor productivity in the long run. Of course, it is easy to challenge these identifying assumptions on a priori grounds.¹ But, this is always the case with identifying restrictions. In the end, the only convincing way to build confidence in such assumptions is to see if they help us to understand the data. One of the purposes of this project is to find out.

We propose a very different identification strategy for the other 6 shocks in the model.

¹For example, the presence of endogeneity in the evolution of technology in principle has the consequence that all shocks affect the level of productivity in the long run.

This strategy builds on the approach advocated in Uhlig (2001,2002) and also implemented in Canova (). We call this a ‘model-based identification strategy’. Under this strategy we characterize the space of possible impulse responses to the six shocks that are consistent with a Vector Autoregression (VAR). We then isolate the element in that space that matches most closely the impulse responses implied by our dynamic general equilibrium model. Although we describe our model based estimation methodology, we do not yet have results for it. We only report results for monetary policy and technology shocks.

We now briefly summarize our findings. Regarding policy shocks, our results are similar to what has been reported before. We find that an expansionary policy-induced shock to money growth delivers hump-shaped increases in consumption, investment, labor, and capital utilization. It produces a fall in the interest rate. After an initial drop, inflation slowly rises and peaks after nearly two years. Regarding technology, we find that a positive technology shock has an impact that students of real business cycles would recognize: output, capital utilization, employment and investment all rise after such a shock. To our surprise, we find that the model parameters which do best at jointly matching the estimated impulse responses to monetary policy and permanent technology shocks are roughly the same as the model parameters that produce the best match to monetary policy impulse responses alone.

In one respect, our results differ sharply from what is reported in the literature. While we find that employment rises after a positive shock to technology, Gali (1999), Francis and Ramey (2001), and others, report that employment falls after such a shock. We explore the reasons for the differences.

Our VAR-based estimates of the response to policy and permanent technology shocks are of independent interest. We find that policy shocks account for very little of the variation in the data, less than 5 percent. Regarding persistent technology shocks, we find that they account for nearly 50 percent of the variance in the data. When we take a closer look at the results, however, we find that the shocks do not go far towards explaining business cycle fluctuations. Instead, the technology shocks seem to be particularly useful for explaining lower frequency components of the data. A fuller assessment of whether technology shocks play an important role in the business cycle will be possible when we incorporate our two other technology shocks into the analysis.

The following section lays out the model used in the analysis. After that, we discuss the estimation of impulse response functions. Finally, we report results.

1. A Dynamic, General Equilibrium Model

Following is a description of the model used in the analysis. The model builds on the one in Christiano, Eichenbaum and Evans (2001) (CEE). That model incorporates a single shock, a

disturbance to monetary policy. The model used in our analysis allows for 8 shocks, including a shock to monetary policy. The discussion of this section highlights the key features of the model, including the shocks, and explains the rationale for each. The 8 shocks in the model correspond to 3 financial market shocks and 5 non-financial market shocks. The latter include three shocks to technology: a permanent and a temporary shock to the aggregate goods-producing technology, and a transient shock to the productivity of investment. In addition, we include shocks to the market power of intermediate good firms and to the market power of household suppliers of differentiated labor services. The three financial market shocks include a monetary policy shock, a shock to household money demand and a shock to firm money demand. Monetary policy is endogenous, in that the control variable of the monetary authority - the aggregate stock of money - is permitted to respond to all shocks.

The shocks have been incorporated into our quantitative model, and in a later section we describe an econometric procedure for identifying and estimating the model and shock parameters jointly. However, at this time we have only estimated the version of the model with two shocks, the shock the technology of goods-producing firms and the shock to monetary policy. A later draft will incorporate results for all shocks.

In what follows we first describe the firm sector. We then describe the household sector and equilibrium.

1.1. Firms

Final goods are produced by competitive firms who use a continuum of intermediate goods along the lines of Dixit and Stiglitz:

$$Y_t = \left[\int_0^1 Y_{jt}^{\frac{1}{\lambda_{f,t}}} dj \right]^{\lambda_{f,t}},$$

where $\lambda_{f,t}$ is a stochastic process, and $\lambda_{f,t} \in [1, \infty)$. For estimation purposes, in this draft of the paper this shock is simply fixed at its mean value, λ_f . The price of the final good is P_t and the price of the i^{th} intermediate good is P_{it} . In the usual way, competition and profit maximization lead to the following relationship between these prices:

$$P_t = \left[\int_0^1 P_{jt}^{\frac{1}{1-\lambda_{f,t}}} dj \right]^{(1-\lambda_{f,t})}. \quad (1.1)$$

The shock, $\lambda_{f,t}$, shows up as a disturbance to the reduced form pricing equation of the model. Empirical analyses of inflation often find it important to include such a shock.²

²References for this to be added here.

Each intermediate good, $i \in (0, 1)$ is produced by a monopolist using the following production function:

$$Y_{it} = \begin{cases} \epsilon_t (z_t)^{1-\alpha} K_{it}^\alpha X_{it}^{1-\alpha} - z_t^* \phi & \text{if } (z_t)^{1-\alpha} K_{it}^\alpha X_{it}^{1-\alpha} \geq z_t^* \phi \\ 0, & \text{otherwise} \end{cases} \quad (1.2)$$

where z_t is a persistent shock to technology, ϵ_t is a stationary shock to technology, and K_{it} , X_{it} represent capital and labor services, respectively. In this draft of the paper, we set $\epsilon_t \equiv 1$. We assume

$$x_t = \log z_t - \log z_{t-1}.$$

where

$$x_t = (1 - \rho_x)x + \rho_x x_{t-1} + \varepsilon_{xt}, \quad x > 0. \quad (1.3)$$

We include a fixed cost in (1.2) to ensure that profits are not too big in equilibrium. We set the fixed cost parameter, ϕ , so that profits are zero along a nonstochastic steady state growth path. The fixed costs are modeled as growing with the exogenous variable, z_t^* :

$$z_t^* = z_t \Upsilon^{\left(\frac{\alpha}{1-\alpha}t\right)}, \quad \Upsilon > 1.$$

If fixed costs were not growing, then they would eventually become irrelevant. We specify that they grow at the same rate as z_t^* , which is the rate at which equilibrium output grows. Note that the growth of z_t^* exceeds that of z_t . This is because we have another source of growth in this economy, in addition to the upward drift in z_t . In particular, we posit a trend increase in the efficiency of investment. We discuss this further below.

Intermediate good producers are competitive in the market for capital and labor services, and so they take factor prices as given. Given our specification of technology in (1.2), marginal cost is the same for each firm, and independent of the scale of production. Let s_t denote the ratio of marginal cost to the aggregate price level. Then,

$$s_t = \left(\frac{1}{1-\alpha}\right)^{1-\alpha} \left(\frac{1}{\alpha}\right)^\alpha \left(\frac{R_t^k}{P_t}\right)^\alpha \left(\frac{W_t}{P_t} R_t^f\right)^{1-\alpha},$$

where R_t^k denotes the rental rate of capital and W_t is the wage rate, both denominated in currency units. The gross nominal rate of interest, R_t^f , appears here because intermediate good firms are assumed to have to borrow a fraction, ν_t , of their wage bill at the beginning of the period, and repay it at the end, when sales receipts come due. The gross nominal rate of interest at which they borrow is R_t , so that

$$R_t^f = \nu_t R_t + 1 - \nu_t.$$

In a later draft, ν_t will be treated as a stochastic process. For now, we suppose that $\nu_t \equiv 1$.

Intermediate good firms face price frictions using a modified version of the model in Calvo (1983). In particular, each period a random selection fraction of firms, $1 - \xi_p$, is permitted to reoptimize its price. The i^{th} firm among the ξ_p firms that do not reoptimize sets its price in the following way:

$$P_{it} = \pi_{t-1} P_{i,t-1},$$

where $\pi_t = P_t/P_{t-1}$ denotes the aggregate rate of inflation. Each intermediate good firm must satisfy its demand curve in each period. Optimizing firms discount future cash flows using the household's discount rate, $\beta \in (0, 1)$.³ This pricing behavior by firms, together with (1.1), leads to the following representation of inflation:

$$\hat{\pi}_t = \frac{1}{1 + \beta} \hat{\pi}_{t-1} + \frac{\beta}{1 + \beta} E_t \hat{\pi}_{t+1} + \frac{(1 - \beta \xi_p)(1 - \xi_p)}{(1 + \beta) \xi_p} E_t [\hat{s}_t + \hat{\lambda}_{f,t}],$$

where a ‘ $\hat{\cdot}$ ’ over a variable indicates percent deviation from steady state.⁴

1.2. Households

The j^{th} household discounts future consumption, C_t , labor, $h_{j,t}$, and real balances, Q_t/P_t , using the following preferences:

$$E_t^j \sum_{l=0}^{\infty} \beta^l \left[u(C_{t+l} - bC_{t+l-1}) - \frac{\psi_{L,t}}{2} (h_{j,t})^2 + \psi_{q,t} \frac{\left(\frac{Q_{t+l}}{z_{t+l}^* P_{t+l}} \right)^{1-\sigma_q}}{1 - \sigma_q} \right]$$

When $b > 0$, household preferences for consumption are characterized by habit persistence. This specification of preferences is standard in the monetary economics literature because it helps account for the hump-shaped response of consumption to monetary policy shocks. In addition, this specification has proved useful for understanding features of asset prices (see Boldrin, Christiano and Fisher (2001).) The terms, $\psi_{L,t}$ and $\psi_{q,t}$, represent stochastic shocks

³They actually do so using the Arrow-Debreu date and state-contingent prices. In equilibrium, these involve not just β , the household's discount factor, but also the marginal utility of consumption. However, with our linearization procedure (we use a standard procedure) the marginal utility of consumption drops out.

⁴In the case of $\hat{\lambda}_{f,t}$, this will be modeled as a zero mean, univariate time series process.

to preferences for leisure and real balances, respectively. In this draft of the paper, these variables are simply held constant.⁵

The j^{th} household is the only supplier of a differentiated labor service, h_{jt} . It sets its wage rate, W_{jt} , following a modified version of the setup in Erceg, Henderson, Levin (2000). This in turn follows the spirit of the price setting frictions in Calvo (1983). In each period, $1 - \xi_w$ households are randomly selected to reoptimize their wage. The j^{th} household among ξ_w who cannot reoptimize, set their wage according to

$$W_{jt} = \pi_{t-1} x_t W_{jt-1}.$$

Thus, non-optimizing households index their wage rate to the aggregate inflation rate, as do non-optimizing firms. In addition, non-optimizing households also add a technology growth factor to their wage. Households are required to be on their demand curve in each period. Demand for household labor derives from a competitive, representative ‘labor contractor’ who takes $h_{j,t}$, $j \in (0, 1)$, as input and produces aggregate, homogeneous labor services using the following production function:

$$X_t = \left[\int_0^1 (h_{j,t})^{\frac{1}{\lambda_w}} dj \right]^{\lambda_w}, \quad 1 \leq \lambda_w < \infty.$$

The labor contractor takes the price of labor services, W_t , as given, as well as the price of the j^{th} differentiated labor input.

The term, $\psi_{L,t}$, in the utility function is a disturbance to the preference for leisure. In the linearized solution to the model, λ_w and $\psi_{L,t}$ appear symmetrically, so we are free to interpret $\psi_{L,t}$ as a shift in the market power of workers. Various authors, including Shapiro and Watson (1988), Hall (1991), and Francis and Ramey (2001), have argued for the importance of these shocks as a source of business cycle fluctuations.

Note that we do not index C_t and Q_t in the utility function by j . In principle, different households would make different consumption and portfolio decisions because they differ in their labor market experiences. We rule out this sort of heterogeneity by the assumption that households have access to the appropriate insurance contracts.

Households own the physical stock of capital, \bar{K}_t . They make the investment decisions, I_t , which impact on the stock via the following capital accumulation technology:

$$\bar{K}_{t+1} = (1 - \delta)\bar{K}_t + \mu_{\Upsilon,t} \Upsilon^t [1 - S(I_t/I_{t-1})] I_t. \quad (1.4)$$

⁵From the point of interpreting $\psi_{L,t}$, it is interesting to note that this shock is observationally equivalent to a shock λ_w , a variable discussed below which measures the degree of labor market power that the household has.

The term in square brackets reflects the presence of costs of adjusting the flow of investment. We suppose that S and its derivative are zero along a steady state growth path for the economy. The second derivative of this function in steady state, $S'' > 0$, is a parameter that we estimate. We place adjustment costs on the change of investment, rather than, say, the level, to enable the model to account for the hump-shaped response of investment to a monetary policy shock.

The terms multiplying the square brackets in (1.4) represent an exogenous process governing the evolution of the efficiency of investment. There is a positive trend in this term, since $\Upsilon > 1$. This term gives rise to a trend fall in the relative price of capital goods, $P_{k',t}$, in our model economy. It captures the trend increase in the efficiency of investment that Greenwood, Hercowitz and Krusell (1998) argue is a key engine of growth for the US economy. The other term, $\mu_{\Upsilon,t}$, is a stationary stochastic disturbance to the efficiency of investment. Greenwood, Hercowitz and Krusell (1998a) argue that this is an important source of business cycle fluctuations in the US. In this draft of the paper, we set $\mu_{\Upsilon,t} \equiv 1$. In a later draft, when we estimate $\mu_{\Upsilon,t}$, we will be able to evaluate the Greenwood, Hercowitz and Krusell (1998a) claim.

Households control the amount of capital services supplied to the capital services market by choosing the utilization rate of capital. In particular, capital services are determined according to:

$$K_t = u_t \bar{K}_t.$$

To ensure that u_t is finite, we suppose that the household faces convex costs, in terms of final goods, of increasing utilization, in the form:

$$a(u_t) \Upsilon^{-t} \bar{K}_t.$$

We suppose that $a = 0$ along a steady state growth path (when $u_t = 1$) and we set a' in steady state to the scaled, real rental rate of capital. A free parameter for estimation is $a''/a' > 0$, where a'' is the second derivative of a , evaluated at $u_t = 1$. Note that for a given rate of utilization, u_t , and stock of capital, \bar{K}_t , the cost of utilization falls over time. This is to ensure that the model has a balanced growth steady state, one in which hours worked, capital utilization and various ‘great ratios’ are constant.⁶ If the term were not present, the fact that the growth rate of capital is relatively rapid would imply that utilization costs would grow too fast in steady state to be consistent with $u_t = 1$.

The presence of variable capital utilization in the model, by causing the supply of capital services to be elastic, helps damp the response of marginal costs to a monetary policy shock.

⁶By the great ratios, we mean the ratio of the consumption good value of capital to output and the ratio of consumption to output.

This in turn is key for the model's ability to account for the inertial response of inflation to a monetary policy shock. The assumption that utilization costs are denominated in goods helps assure that capital utilization rises after a positive monetary shock. In several computational experiments in which utilization costs take the form of increased depreciation of physical capital, we have found that capital utilization has a tendency to drop after a positive monetary policy shock. This is because a positive monetary shock leads to a rise in physical investment which, via the adjustment costs, leads to a rise in the price of physical capital. With capital more expensive, households find it desirable to reduce their utilization of capital.

The household has a portfolio decision. At the beginning of the period, it is in possession of the economy-wide stock of high-powered money, M_t . It splits this between deposits with a financial intermediary and Q_t . The deposits at the financial intermediary are combined with a money injection from the central bank, and loaned on to firms who need the funds to finance their wage bill. The interest received by the financial intermediary on its loans is transferred to households at the end of the period. Households are willing to forego interest earnings to hold Q_t , because Q_t generates services that are captured in the last term in square brackets in the utility function. The exogenous shifter, z_t^* , in the utility function guarantees that, in a steady state growth path, the ratio of Q_t/P_t to output is constant.

1.3. Monetary Authority

We adopt the following specification of monetary policy:

$$\hat{\mu}_t = \hat{\mu} + \hat{\mu}_{p,t} + \hat{\mu}_{x,t}, \quad (1.5)$$

where $\hat{\mu}_t$ represents the growth rate of high powered money, M_t . We model $\hat{\mu}_{p,t}$ and $\hat{\mu}_{x,t}$ as follows:

$$\begin{aligned} \hat{\mu}_{p,t} &= \rho_{\mu_p} \hat{\mu}_{p,t-1} + \varepsilon_{\mu_p,t} \\ \hat{\mu}_{x,t} &= \rho_{\mu_x} \hat{\mu}_{x,t-1} + c_{\mu_x} \varepsilon_{x,t} \end{aligned} \quad (1.6)$$

Here, $\varepsilon_{\mu_p,t}$ represents a shock to monetary policy and we suppose that the response of money growth to this is characterized as a scalar first order autoregression. The term, $\hat{\mu}_{x,t}$, captures the response of monetary policy to an innovation in technology, $\varepsilon_{x,t}$. The contemporaneous response is governed by the parameter, c_{μ_x} . The dynamic response of $\hat{\mu}_{x,t}$ to $\varepsilon_{x,t}$ is characterized by a first order autoregression. Initially, we worked with more elaborate parameterizations of $\hat{\mu}_{p,t}$ and $\hat{\mu}_{x,t}$. However, we found that the simple representations in (1.6) are adequate in practice.

In the discussion above, we have described 6 additional shocks: a shock to household money demand, $\psi_{q,t}$, a shock to firm money demand, ν_t , a shock to household preferences for leisure (or, equivalently, to their degree of labor market power), $\psi_{L,t}$, a stationary shock to technology, ϵ_t , an investment-specific technology shock, $\mu_{\Upsilon,t}$, and a shock to intermediate good firm market power, λ_{ft} . For now, these shocks are held constant. When they are non-trivial stochastic processes, we will add six additional terms to the representation of monetary policy, (1.5), one corresponding to the monetary policy response to each shock.

1.4. Timing, Market Clearing and Equilibrium

We adopt the following timing specification in the model. At the beginning of the period, the non-financial market shocks are realized. Then, prices and wages are set and households make their consumption, investment and capital utilization decisions. After this, the financial market shocks are realized. Then, households make their portfolio decision, goods and labor markets meet and clear, and production investment and consumption occur.

Clearing in the goods market requires:

$$C_t + I_t \leq Y_t - a(u_t)\Upsilon^{-t}\bar{K}_t,$$

where Y_t is the output of final goods. The measure of final goods and services that we compare with aggregate output in the data is $C_t + I_t$. Clearing in the money market requires:

$$W_t X_t = M_t - Q_t + (1 + \hat{\mu}_t)M_t.$$

The demand for funds appears on the left, and the supply appears on the right.

We adopt a standard sequence-of-markets equilibrium concept. The equilibrium prices and quantities in the model can be represented as follows:

$$\begin{aligned} C_t &= c_t z_t^* \\ I_t &= i_t z_t^* \\ Y_t &= y_t z_t^* \\ \bar{K}_{t+1} &= \bar{k}_{t+1} z_t^* \Upsilon^t \\ R_t^k &= P_t \Upsilon^{-t} r_t^k \\ P_{k',t} &= \Upsilon^{-t} p_{k',t} \\ W_t &= P_t z_t^* w_t. \end{aligned}$$

Here, lower case variables to the right of the equality are covariance stationary and converge to constant steady state values when all shocks are held at their unconditional mean values.

Also, $P_{k',t}$ is the price of \bar{K}_{t+1} at time t , in consumption goods units. According to these expressions, consumption, investment, output and the real wage grow at the rate of growth of z_t^* . The value, in consumption units, of the physical stock of capital also grows at the rate of growth of z_t^* . However, its relative price falls over time and the growth rate of the physical quantity of capital is greater than the growth rate of z_t^* . These balanced growth properties of our model are just the properties of Solow's model of investment specific technical change, recently emphasized by Greenwood, Hercowitz and Krusell (1998). An interesting feature of these properties is the logarithm of the growing variables are a linear combination of a unit root process, z_t , and a deterministic time trend, $\log(\Upsilon)t$. In practice, the literature emphasizes the possibility of one type of process, or the other, but not both.

For numerical analysis, we approximate the model's solution by linearizing the first order conditions and identities that characterize equilibrium about the non-stochastic steady state values of the scaled variables. We apply standard solution methods to the resulting linear system (see Christiano (2002).)

2. Estimation of Impulse Response Functions

We first briefly describe the data. We then describe how we go about estimating impulse responses to shocks using Vector Autoregressions. We report findings for the dynamic effects of monetary policy shocks and for a permanent shock to technology. The dynamic effects to a monetary policy shock are similar to what has been reported before in Christiano, Eichenbaum and Evans (1999, 2001). This in itself is a notable finding, because of its implications for robustness. Although the basic recursiveness assumption on monetary policy is used in all these settings, other details about the estimation vary substantially. These include the list of variables used in the analysis, the estimation period and whether the data are assumed to be trend stationary or difference stationary. Through all these applications, the basic qualitative nature of the results is always the same.

Turning to the analysis of the consequence of the permanent shock to technology, our results are in some respects surprising. In particular, we find that the response to a technology shock corresponds roughly to what a student of real business cycles would expect: hours worked, investment, consumption and output all increase. This finding is surprising because it conflicts with a recent literature which argues that hours worked actually fall after a positive shock to technology. We devote some space to reconciling our results with those in this literature. Our preliminary results suggest the possibility that the findings of this literature are consistent with the hypothesis that they are an artifact of over differencing the data. Whether this is the most plausible hypothesis is something that we are currently studying.

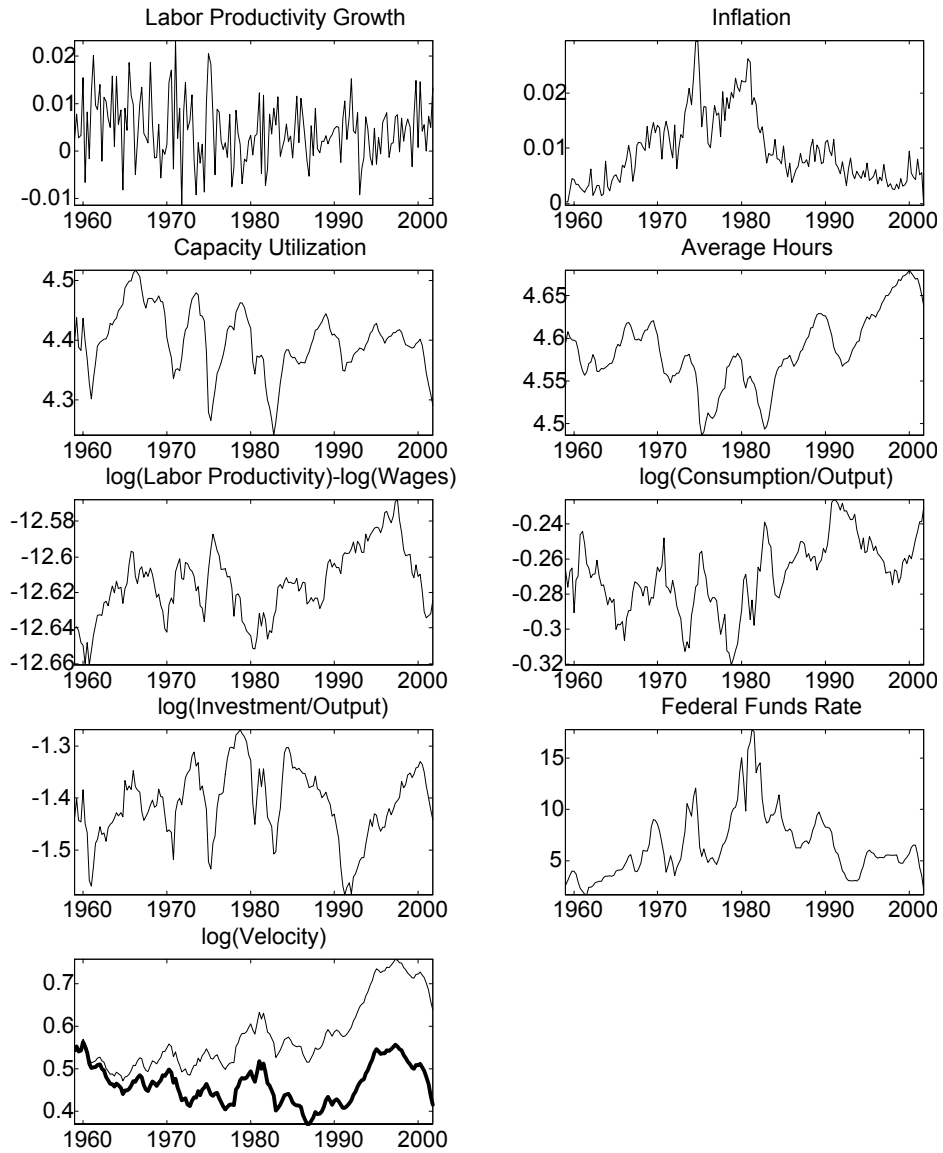
2.1. Data

The data used in the analysis were taken from the DRI Basic Economics Database.⁷ In our analysis, we require that productivity growth, the interest rate, inflation, the log consumption to output ratio, $\log(c/y)$, the log investment to output ratio, $\log(i/y)$, log capacity utilization, log per capita hours worked, the log of the productivity to real wage ratio ($\log(y/h)-\log(w)$) and the log of *M2* velocity all be stationary. These variables are graphed in Figure 1. Thin lines indicate the raw data. For the most part, the data appear roughly consistent with our stationarity assumption. One exception is velocity which rises very sharply in the 1990s. We detrended these data prior to analysis using a linear trend. The detrended data are indicated

⁷The data were taken from <http://economics.dri-wefa.com/webabstract/index.htm>. Nominal gross output is measured by *GDP*, real gross output is measured by *GDPQ* (real, chain-weighted output). Nominal investment is *GCD* (household consumption of durables) plus *GPI* (gross private domestic investment). Nominal consumption is measured by *GCN* (nondurables) plus *GCS* (services) plus *GGE* (government expenditures), money is measured by *FM2*. These variables were converted into per capita terms by *P16*, a measure of the US population over age 16. A measure of the aggregate price index was obtained from the ratio of nominal to real output, *GDP/GDPQ*. Capacity utilization is measured by *IPXMCA* the manufacturing industry's capacity index (there is a measure for total industry, *IPX*, but it only starts in 1967). The interest rate is measured by the federal funds rate, *FYFF*. Hours worked is measured by *LBMNU* (Nonfarm business hours). Hours were converted to per capita terms using our population measure. Nominal wages are measured by *LBCPU*, (nominal hourly non-farm business compensation). This was converted to real terms by dividing by the aggregate price index.

by the thick line. The other data were used without further transformation.

Figure 1
Data Used In VAR



2.2. Impulse Response Functions: How We Compute Them

We adopt standard strategies for identifying monetary policy and technology shocks. To identify monetary policy shocks, we adopt the recursive method pursued in CEE. To identify innovations to technology, we adopt the strategy in Gali (2001), Gali, Lopez-Salido, and Valles (2002) and Francis and Ramey (2001). In particular, we suppose - as is true in our model - that innovations to technology are the only shock that affects the level of labor productivity in the long run.⁸ To identify the remaining six shocks, we developed a method of identification which we call ‘model-based’. We use the restrictions implied by the model itself to do identification, using a method that is inspired by the strategy pursued recently by Uhlig (2001). Our approach differs from Uhlig’s in that he imposes a priori sign and shape restrictions, while we impose the restrictions of the model.

We now discuss the calculations of the impulse response functions using the data just described. Consider the following *reduced form* vector autoregression:

$$\begin{aligned} Y_t &= \alpha + B(L)Y_{t-1} + u_t, \\ Eu_t u_t' &= V \end{aligned} \tag{2.1}$$

The ‘fundamental’ economic shocks, e_t , are related to u_t by the following relation:

$$u_t = Ce_t, \quad Ee_t e_t' = I.$$

To obtain the dynamic response function to, say, the i^{th} fundamental shock, e_{it} , we need $B(L)$ and the i^{th} column of C , C_i , and we simulate:

$$Y_t = B(L)Y_{t-1} + C_i e_{it}. \tag{2.2}$$

This section discusses how we compute $B(L)$ and C_i for the shocks we wish to identify.

⁸It is of course easy to imagine models in which *all* shocks have a permanent impact on productivity. For example, an endogenous growth model in which shocks lead to a transitory change in the rate of growth of technology has such a property. Any set of identification assumptions can be challenged on a priori grounds, and ours are no exception. Ultimately, a defense of any particular set of identification assumptions is determined by how far one can go with them in explaining empirical observations. Considerably more experience is needed before we can say with confidence what sort of identification assumptions are useful for understanding business cycle observations. This paper is part of a broader research program involving many other researchers that attempts to build the necessary experience.

In the analysis, Y_t is defined as follows:

$$\underbrace{Y_t}_{9 \times 1} = \begin{pmatrix} \Delta \ln(GDP_t/\text{Hours}_t) \\ \Delta \ln(GDP \text{ deflator}_t) \\ \text{capacity utilization}_t \\ \ln(GDP_t/\text{Hours}_t) - \ln(W_t/P_t) \\ \ln(\text{Hours}_t) \\ \ln(C_t/GDP_t) \\ \ln(I_t/GDP_t) \\ \text{Federal Funds Rate}_t \\ \ln(GDP \text{ deflator}_t) + \ln(GDP_t) - \ln(M2_t) \end{pmatrix}$$

$$= \begin{pmatrix} \underbrace{\Delta y_t}_{1 \times 1} \\ \underbrace{Y_{1t}}_{6 \times 1} \\ \underbrace{R_t}_{1 \times 1} \\ \underbrace{Y_{2t}}_{1 \times 1} \end{pmatrix}.$$

We partition e_t conformably with the partitioning of Y_t :

$$e_t = \begin{pmatrix} \underbrace{\varepsilon^z}_{1 \times 1} \\ \underbrace{e_{1t}}_{6 \times 1} \\ \underbrace{\varepsilon_t}_{1 \times 1} \\ \underbrace{e_{2t}}_{1 \times 1} \end{pmatrix}.$$

An alternative representation of our system is given by the *structural form*:

$$A_0 Y_t = A(L) Y_{t-1} + e_t. \quad (2.3)$$

The parameters of the reduced form are related to those of the structural form by:

$$C = A_0^{-1}, \quad B(L) = A_0^{-1} A(L). \quad (2.4)$$

We obtain impulse responses by first estimating the parameters of the structural form, then mapping these into the reduced form, and finally simulating (2.2). We specify the VAR to have four lags, so that $A(L) = A_1 + A_2L + A_3L^2 + A_4L^3$. Our data are quarterly and cover the period is 1959Q1 - 2001Q4 (the estimation period drops the first 4 quarters, to accommodate the 4 lags).

The following two subsections consider first, identification of the monetary policy and technology shocks, and then identification of the other shocks.

2.2.1. Restrictions on Monetary Policy and Technology Shocks

We assume policy makers manipulate the monetary instruments under their control in order to ensure that the following interest rate targeting rule is satisfied:

$$R_t = f(\Omega_t) + \varepsilon_t, \quad (2.5)$$

where ε_t is the monetary policy shock. We interpret this as a kind of ‘reduced form’ Taylor rule. Conventional representations of the Taylor rule include a smaller set of variables than we do. Typically, these ‘structural representations’ of the Taylor rule include expected future inflation and the output gap. We interpret our (2.5) as a convolution of the structural representation of the Taylor rule with the (linear) functions which relate the variables in the structural Taylor rule to the variables in our VAR. By representing the Taylor rule in this way, we sidestep difficult and controversial questions, such as how it is that the monetary authorities actually compute the output gap. To ensure identification of the monetary policy shock, we assume f is linear, Ω_t contains $Y_{t-1}, Y_{t-2}, Y_{t-3}, Y_{t-4}$ and the only date t variables in Ω_t are the ones above R_t in Y_t . Finally, we assume that ε_t is orthogonal with Ω_t . It is easy to verify (see, e.g., Christiano, Eichenbaum and Evans (1999)) that these identifying assumptions correspond to the following restrictions on A_0 :

$$A_0 = \begin{bmatrix} A_0^{1,1} & A_0^{1,2} & 0 & 0 \\ 1 \times 1 & 1 \times 6 & 1 \times 1 & 1 \times 1 \\ A_0^{2,1} & A_0^{2,2} & 0 & 0 \\ 6 \times 1 & 6 \times 6 & 6 \times 1 & 6 \times 1 \\ A_0^{3,1} & A_0^{3,2} & A_0^{3,3} & 0 \\ 1 \times 1 & 1 \times 6 & 1 \times 1 & 1 \times 1 \\ A_0^{4,1} & A_0^{4,2} & A_0^{4,3} & A_0^{4,4} \\ 1 \times 1 & 1 \times 6 & 1 \times 1 & 1 \times 1 \end{bmatrix}. \quad (2.6)$$

To understand this, consider first the second to last row of A_0 . This row corresponds to the monetary policy rule, (2.5), and the zero in this row reflects that the monetary authority does not look at the last variable in Y_t . Now consider the first 7 rows. The right two columns reflect our assumption that a monetary policy shock has no contemporaneous impact on Δy_t

or Y_{1t} . The two sets of zeros reflect the two distinct channels by which this impact could occur. The second to last column of zeros reflects that the interest rate cannot enter directly into the first set of 7 equations. The second reflects that the interest rate cannot also enter indirectly, via its contemporaneous impact on the last variable.

The assumption that only the technology shock has a non-zero impact on the level of output at infinity implies that the matrix

$$A_0 - A(1), \quad (2.7)$$

has all zeros in its first row, except the 1,1 element, which could potentially be non-zero. To see this, note that the impact of the vector of shocks on the level of y_t at $t = \infty$ corresponds to the first row of $[A_0 - A(1)]^{-1}$. Thus, our long-run restriction is that only the first element in the first row may be non-zero, while the others are zero. But, this is true if, and only if, the same restriction is satisfied by (2.7).

It is useful to write out the equations explicitly, taking into account the restrictions implied by our assumptions about long-run effects, and by our assumptions about the effects of a monetary policy shock:

$$\Delta y_t = a_{\Delta y}(L)\Delta y_{t-1} + \tilde{a}_1(L)\Delta Y_{1t} + \tilde{a}_R(L)\Delta R_{t-1} + \tilde{a}_2(L)\Delta Y_{2,t-1} + \frac{\varepsilon_t^z}{A_0^{1,1}},$$

where $\Delta = (1 - L)$, and the polynomial lag operators correspond to the relevant entries of the first row of $A_0 - A(L)L$, scaled by $A_0^{1,1}$. Note that among the right hand variables in this expression, the only one whose current value appears here is ΔY_{1t} . This fact rules out ordinary least squares as a strategy for obtaining a consistent estimate of the coefficients in this equation, because we expect ε_t^z to be correlated with ΔY_{1t} . An instrumental variables method can be constructed based on the insight that lagged variables are correlated with ΔY_{1t} , but not with ε_t^z . Suppose that an initial consistent estimate of the coefficients have been obtained in this way. The coefficients in the first row of the structural form can then be obtained by scaling the instrumental variables estimates up by $A_0^{1,1}$, where $A_0^{1,1}$ is estimated as the (positive) square root of the variance of the fitted disturbances in the instrumental variables relation.

The next set of 6 equations can be written like this:

$$A_0^{2,1}\Delta y_t + A_0^{2,2}Y_{1t} = b(L)Y_{t-1} + e_{1t} \quad (2.8)$$

The following equation is just the policy rule:

$$R_t + \frac{A_0^{3,1}}{A_0^{3,3}}\Delta y_t + \frac{A_0^{3,2}}{A_0^{3,3}}Y_{1t} = c(L)Y_{t-1} + \frac{\varepsilon_t}{A_0^{3,3}}.$$

Consistent estimates of the parameters in this expression may be obtained by ordinary least squares with R_t as the dependent variable, by our assumption that ε_t is not correlated with Δy_t and Y_{1t} . The parameters of the 8th row of the structural form are obtained by scaling the estimates up by $A_0^{3,3}$, where $A_0^{3,3}$ is estimated as the positive square root of the variance of the fitted residuals. Finally, according to the last equation:

$$Y_{2t} + \frac{A_0^{4,1}}{A_0^{4,4}} \Delta y_t + \frac{A_0^{4,2}}{A_0^{4,4}} Y_{1t} + \frac{A_0^{4,2}}{A_0^{4,4}} R_t = d(L)Y_{t-1} + \frac{e_{2t}}{A_0^{4,4}}.$$

The coefficients in this relation can be estimated by ordinary least squares. This is because e_{2t} is not correlated with the other contemporaneous variables in this relation. This reflects that Y_{2t} does not enter any of the other equations. The parameter, $A_0^{4,4}$, can be estimated as the square root of the estimated variance of the disturbances in this relation. The parameters in the last row of the structural form are then estimated suitably scaling up by $A_0^{4,4}$.

The previous argument establishes that the 1st, 8th and last rows of A_0 are identified. The block of 6 rows in the middle are not identified. To see this, let w denote an arbitrary 6×6 orthonormal matrix, $ww' = I_6$. Suppose \bar{A}_0 and $\bar{A}(L)$ is some set of structural form parameters that satisfies all our restrictions. Let the orthonormal matrix, W , be defined as follows:

$$W = \begin{bmatrix} 1 & 0 & 0 & 0 \\ & 1 \times 6 & & \\ 0 & w & 0 & 0 \\ & 6 \times 6 & 6 \times 1 & 6 \times 1 \\ 0 & 0 & 1 & 0 \\ & 1 \times 6 & & \\ 0 & 0 & 0 & 1 \\ & 1 \times 6 & & \end{bmatrix}. \quad (2.9)$$

It is easy to verify that the reduced form corresponding to the parameters, $W\bar{A}_0$, $W\bar{A}(L)$ also satisfies all our restrictions, and leads to the same reduced form:

$$Y_t = (W\bar{A}_0)^{-1} W\bar{A}(L)Y_{t-1} + (W\bar{A}_0)^{-1} W e_t.$$

To see this, note:

$$\begin{aligned} (W\bar{A}_0)^{-1} W\bar{A}(L) &= \bar{A}_0^{-1} W' W \bar{A}(L) = \bar{A}_0^{-1} \bar{A}(L) \\ E (W\bar{A}_0)^{-1} W u_t u_t' W' [(W\bar{A}_0)^{-1}]' &= E \bar{A}_0^{-1} W' W e_t e_t' W' [\bar{A}_0^{-1} W']' \\ &= E \bar{A}_0^{-1} W' W e_t e_t' W' W (\bar{A}_0^{-1})' \\ &= \bar{A}_0^{-1} W' W W' W (\bar{A}_0^{-1})' \\ &= \bar{A}_0^{-1} (\bar{A}_0^{-1})'. \end{aligned}$$

Recall that impulse response functions can be computed using the matrices in $B(L)$ and the columns of A_0^{-1} . It is easy to see that the impulse responses to ε_t^z , ε_t and e_{2t} are invariant to w . This is because:

$$(W\bar{A}_0)^{-1} = \bar{A}_0^{-1}W'.$$

It is easily verified that the first, 8th and last columns of $\bar{A}_0^{-1}W'$ coincide with those of \bar{A}_0^{-1} .

We conclude that there is a family of observational equivalent parameterizations of the structural form, which is consistent with our identifying assumptions on the monetary policy shock and the technology shock. We arbitrarily select an element in this family as follows. Let Q and R be orthonormal and lower triangular (with positive diagonal terms) matrices, respectively, in the QR decomposition of A_0^{22} . That is, $A_0^{22} = QR$. This decomposition is unique and guaranteed to exist given that A_0^{22} is non-singular, a property implied by our assumption that A_0 is invertible. The reasoning up to now indicates that we may, without loss of generality, select A_0 so that A_0^{22} is lower triangular with positive diagonal elements. This restriction does not restrict the reduced form in any way, nor does it restrict the set of possible impulse response functions associated with ε_t^z , ε_t or e_{2t} .

Thus, in (2.8) A_0^{22} is lower triangular. We seek consistent estimates of the parameters of (2.8), with this restriction imposed. Ordinary least squares will not work as an estimation procedure here because of simultaneity. To see this, consider the first equation in (2.8). Suppose the left hand variable is the first element in Y_{1t} . The only current period explanatory variable is Δy_t . But, note from the first equation in the structural form that Δy_t responds to Y_{1t} and, hence, to the innovations in Y_{1t} . That is, Δy_t is correlated with the first element in e_{1t} . We can instrument for Δy_t using ε_t^z , the (scaled) residual from the first structural equation. Clearly, this variable is correlated with Δy_t , and not with the first element in e_{1t} .

Now consider the second equation in (2.8). Think of the left hand variable as being the second variable in Y_{1t} . The current period explanatory variables in that equation are Δy_t and the first variable in Y_{1t} . Both these are correlated with the second element in e_{1t} . To see this, note that a disturbance in the second element of e_{1t} ends up in Δy_t via the first equation in the structural form, because Y_{1t} appears there. This explains why Δy_t is correlated with the second element of e_{1t} . But, the first element in Y_{1t} is also correlated with this variable because Δy_t is an ‘explanatory’ variable in the equation determining the first element in Y_{1t} , i.e., the first equation in (2.8). So, we need an instrument for Δy_t and the first element of Y_{1t} . For this, use ε_t^z and the residual from the first equation in (2.8). Thus, moving down the equations in (2.8), we use as instruments ε_t^z and the disturbances in the previous equations in (2.8).

With A_0 and $A(L)$ in hand, we are now in a position to compute the reduced form, using (2.4). In that reduced form, we find it convenient to refer to the shocks, e_{1t} , as Choleski shocks, because of the lower triangular normalization that underlies them. The dynamic

response of Y_t to technology and monetary policy shocks may be computed by simulating (2.2) with $i = 1, 8$, respectively.

2.2.2. Model-Based Identification of Other Shocks

We now consider identification of the other shocks in the model. The previous subsection discussed the computation of A_0 and $A(L)$ with the normalization that A_0^{22} is lower triangular and imposing our assumptions on monetary policy and technology shocks. From that discussion, we know that if W is an orthonormal matrix with structure (2.9), then WA_0 and $WA(L)$ is another parameterization of the structural form which satisfies the identification assumptions on monetary policy and technology shocks. That parameterization replaces the Choleski shocks, e_{1t} , with a linear combination, We_{1t} . The new vector of shocks, We_{1t} , has a different set of impulse response functions. The idea of model-based identification is to search over all possible such shocks, to identify the orthonormal rotation matrix, say W^* , such that the dynamic response in the VAR to W^*e_{1t} resembles the model's dynamic responses to shocks other than ε_t^z , ε_t and e_{2t} . Our metric for making precise 'resembles' is discussed further below.

It is useful, for later purposes, to develop some additional notation pertaining to model-based identification. Partition w in W defined in (2.9) as follows:

$$w' = [w_1, w_2, \dots, w_6].$$

By orthonormality of w , we require $w_i'w_i = 1$, $w_i'w_j = 0$, $i \neq j$, $i = 1, 2, \dots, 6$. Note that given a choice of w_i , the impulse response function to the i^{th} element in We_{1t} , e_{1t}^i , is determined, $i = 1, 2, \dots, 6$. Let W_i denote a 9×1 vector with w_i in locations 2 to 7 and zeros elsewhere, $i = 1, 2, 3$. Then, the dynamic response to the i^{th} element in e_{1t} , e_{1t}^i , is obtained by simulating

$$Y_t = B(L)Y_{t-1} + CW_i e_{1t}^i.$$

This procedure defines a family of responses to the i^{th} shock in some rotation of e_{2t} , for each i . Later, we show how we use the restrictions of the economic model to select an element in this family.

2.3. Impulse Response Functions: The Results for Monetary Policy and Technology

The procedure defined in the previous section allows us to determine the dynamic response to monetary policy and technology shocks independent of our dynamic equilibrium model. This is not so for the other shocks. Our model-based identification procedure is interactive

with our model. The remainder of this section discusses results for policy and technology shocks.

Figure 2 displays the response of our variables to a monetary policy shock. In each case, there is a solid line in the center of a gray area. The gray area represents a 95% confidence interval, and the solid line represents the point estimates.⁹ Note how all variables but the interest rate and money growth show zero response in the period of the shock. This reflects the identification assumption underlying our monetary policy shock. Note too, that the variables displayed in Figure 2 are transformations of the variables in Y_t , which are displayed in Figure 1. In all cases but inflation and the interest rate, the variables are in percent terms. Thus, the peak response of output is a little over 0.2 percent. The Federal Funds rate is in units of basis points, at an annual rate. So, the policy shock produces a 60 basis point drop in the federal funds rate. Inflation is expressed at a quarterly rate. In analyzing these results, we focus on the first 20 quarters' responses.

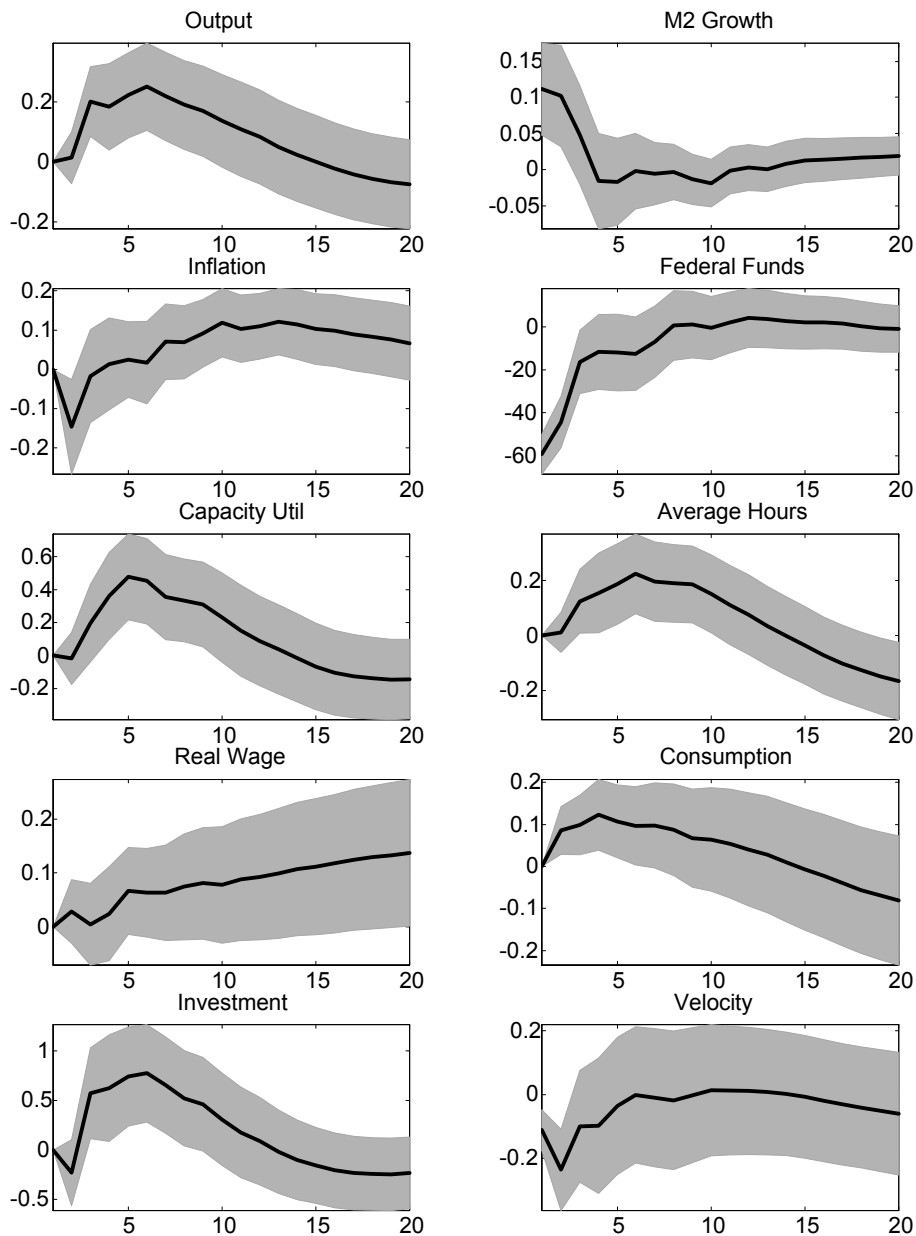
There are six features worth emphasizing here. First, however one measures the policy variable - whether by money growth or the interest rate - the policy variable has completed its movement within about one year. The other variables in the economy respond for a longer period of time. Clearly, any model that can explain these movements must exhibit a substantial amount of internal propagation. Second, inflation takes nearly 3 years to reach its peak response. This is a measure of the substantial inertia in this variable. Interestingly, the initial response of inflation to the monetary expansion is a marginally significant negative fall. In the literature, this has been referred to as the 'price puzzle', reflecting a presumption that no sensible model could reproduce it. The importance of working capital in the monetary transmission mechanism of the model, which causes the interest rate to enter marginal costs, ensures that our model can in principle account for this.¹⁰ Third, output, consumption, investment, hours worked and capacity utilization all display hump-shaped responses, that peak after roughly one year. Fourth, there is a significant liquidity effect. That is, the results indicate that to get the interest rate down, the policy authority must increase money growth. Fifth, velocity moves in the direction naive theory would predict, falling with the initial fall in the interest rate. Finally, the real wage exhibits responds positively, although

⁹The confidence intervals are constructed using standard error estimates of impulse responses obtained using bootstrap simulations.

¹⁰The role of the working capital channel in providing a resolution to the price puzzle has been emphasized by Barth and Ramey.

the response is not significant over the 20 quarters indicated.

Figure 2: Impulse Responses to a Monetary Policy Shock



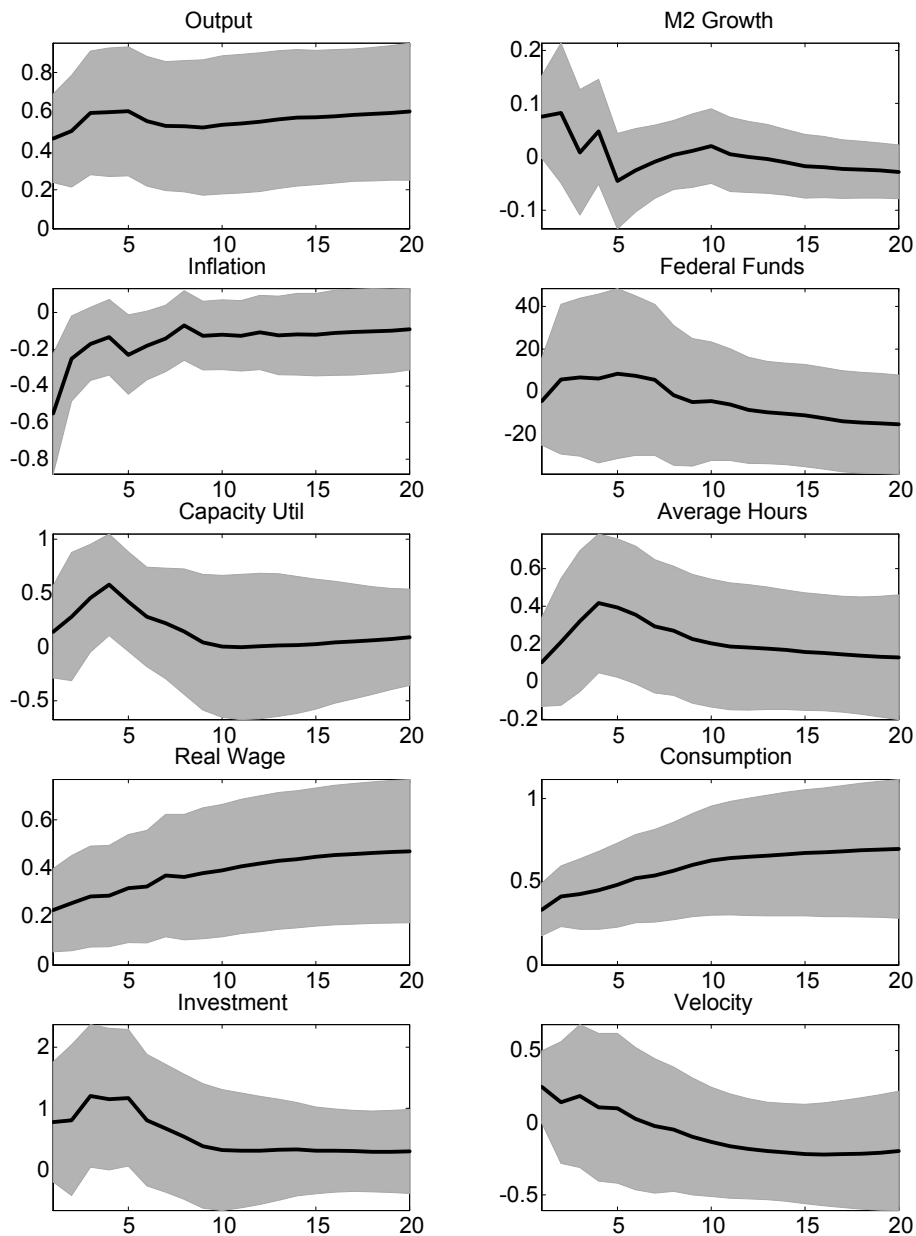
Next we discuss the response of the economy to a positive technology shock. These are

displayed in Figure 3. All responses are measured in the same units as in the previous figures. By construction, the impact of the technology shock on output, labor productivity, consumption, investment and the real wage can be permanent. Because the roots of our estimated VAR are stable, the impact of technology on the variables whose levels appear in Y_t must be temporary. These variables include capacity utilization, hours worked and inflation.

According to the results in Figure 3, the effect of a one-standard deviation positive technology shock is to increase output by about one-half of one percent. The initial reaction of capacity utilization and hours worked to a positive technology shock is (weakly) positive. Overall, our point estimates of the response of variables to a technology shock corresponds qualitatively to what a student of real business cycle models might expect. This contrasts with recent papers in the literature, which report point estimates which suggest that the labor input falls for a prolonged period of time in the wake of a positive technology shock.

The next subsection discusses the relationship between our work and this literature.

Figure 3: Impulse Responses to an Innovation in Technology



We now discuss the decomposition of variance of our two variables. The percent of

forecast error variance due to monetary policy shocks, at horizons 1, 4, 8, 12 and 30 quarters, is reported in Table 1. The key finding is that the monetary policy shock accounts for only a trivial component of the data. Of course, this is not to say that monetary policy is not important. The nature of the monetary policy rule may be very important in determining the forecast error variance. The results only pertain to the monetary policy shocks. It is interesting that even though monetary policy shocks appear to account for little of the variation in the data, the estimated dynamic responses to monetary policy shocks appear to contain a substantial amount of information about the parameters of our equilibrium model. We will see this later on.

Table 1: Variance in Forecast Errors at Different Horizons, in Indicated Variable Due to Monetary Policy Shocks

	Federal Funds				
	1	4	8	12	30
Output	0.0	3.1	5.4	4.4	2.5
M2 Growth	6.2	6.5	6.0	5.4	5.1
Inflation	0.0	1.7	1.8	3.4	4.4
Fed Funds	65.2	21.0	12.5	11.0	9.1
Capacity Util	0.0	2.6	7.2	5.6	4.4
Average Hours	0.0	1.9	5.0	5.1	6.0
Real Wage	0.0	0.2	0.9	1.3	3.2
Consumption	0.0	3.3	2.6	1.6	1.4
Investment	0.0	2.8	5.4	4.8	4.6
Velocity	2.3	2.1	1.1	0.9	0.8

The percent of forecast error variance due to technology shocks is displayed in Table 2. The results here indicate that technology shocks play a relatively important role in fluctuations. Almost half of the variance in the data at the business cycle horizon of 2-3

years is driven by technology shocks.

Table 2: Variance in Forecast Errors at Different Horizons, in Indicated Variable Due to Technology Shocks

	Technology				
	1	4	8	12	30
Output	48.4	48.8	47.1	45.7	62.7
M2 Growth	2.8	3.8	4.1	3.8	5.7
Inflation	41.1	32.1	28.5	25.8	17.4
Fed Funds	0.4	0.5	0.6	0.8	5.7
Capacity Util	1.4	9.7	8.2	5.3	4.5
Average Hours	4.1	15.6	19.3	17.2	11.2
Real Wage	27.7	32.1	35.6	36.9	44.8
Consumption	61.0	67.4	65.3	65.3	69.9
Investment	9.8	14.5	14.1	11.8	12.2
Velocity	11.4	3.0	1.7	2.2	3.7

Another way to assess the role of the identified monetary policy and technology shocks in driving the data, is presented in Figures 4-6. The thick line in Figure 4 displays a simulation of the ‘detrended’ historical data. The detrending is achieved like this. First, we simulated the estimated reduced form representation (2.1) using the fitted disturbances, \hat{u}_t , but setting the constant term, α , and the initial conditions of Y_t to zero. In effect, this gives us a version of the data, Y_t , in which any dynamic effects from unusual initial conditions (relative to the VAR’s stochastic steady state) have been removed, and in which the constant term has been removed. Second, the resulting ‘detrended’ historical observations on Y_t are then transformed appropriately to produce the variables reported in Figure 4. The high degree of persistence observed in output in Figure 4 reflects that our procedure for computing it makes it the realization of a random walk with no drift.

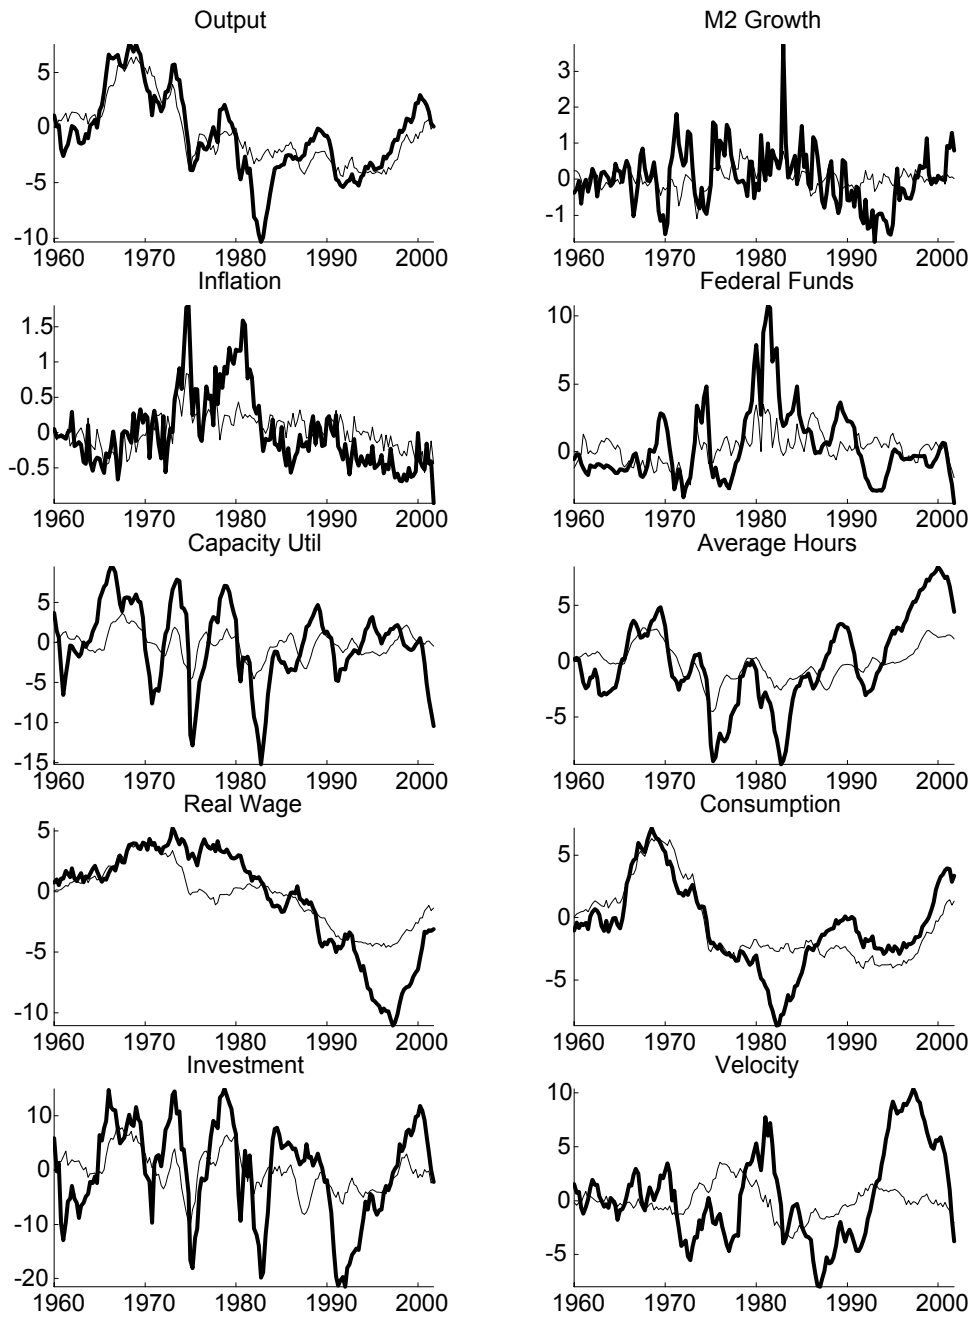
The procedure used to compute the thick line in Figure 4 was then repeated, with one change, to produce the thin line. Rather than using the historical reduced form shocks, \hat{u}_t , the simulations underlying the thin line use $C\hat{e}_t$, allowing only the 1st and 8th elements of \hat{e}_t to be non-zero. Here, \hat{e}_t is the estimated fundamental shocks, obtained from $\hat{e}_t = C^{-1}\hat{u}_t$. The results in Figure 4 give a visual representation of what is evident in Tables 1 and 2: our two shocks only account for about 50% of the fluctuations in the data.

Figure 4 also shows how well the shocks help to account for different frequencies of the data, as well as how well they work at accounting for the fluctuations in different subperiods. The shocks appear to do relatively well in the lower frequencies. In addition, technology and policy shocks together do well at accounting for the movements in output up to the late

1970s and in the late 1990s. These shocks do not account for much of the variation in the

data in the 1980s.

Figure 4: Detrended Historical Data (Thick Line) Versus Component Due to Monetary Policy and Technology Shocks Alone (Thin Line)



Figures 5 and 6 allow us to see how well our two shocks work individually at accounting for the fluctuations in the data. The thin line in Figure 5 is the VAR's estimates of what history would have looked like if there had been only monetary policy shocks. Consistent with the results in Table 1, the monetary policy shocks account for only a trivial amount of

the variation in the data.

Figure 5: Detrended Historical Data (Thick Line) Versus Component Due to Monetary Policy Shocks Alone (Thin Line)

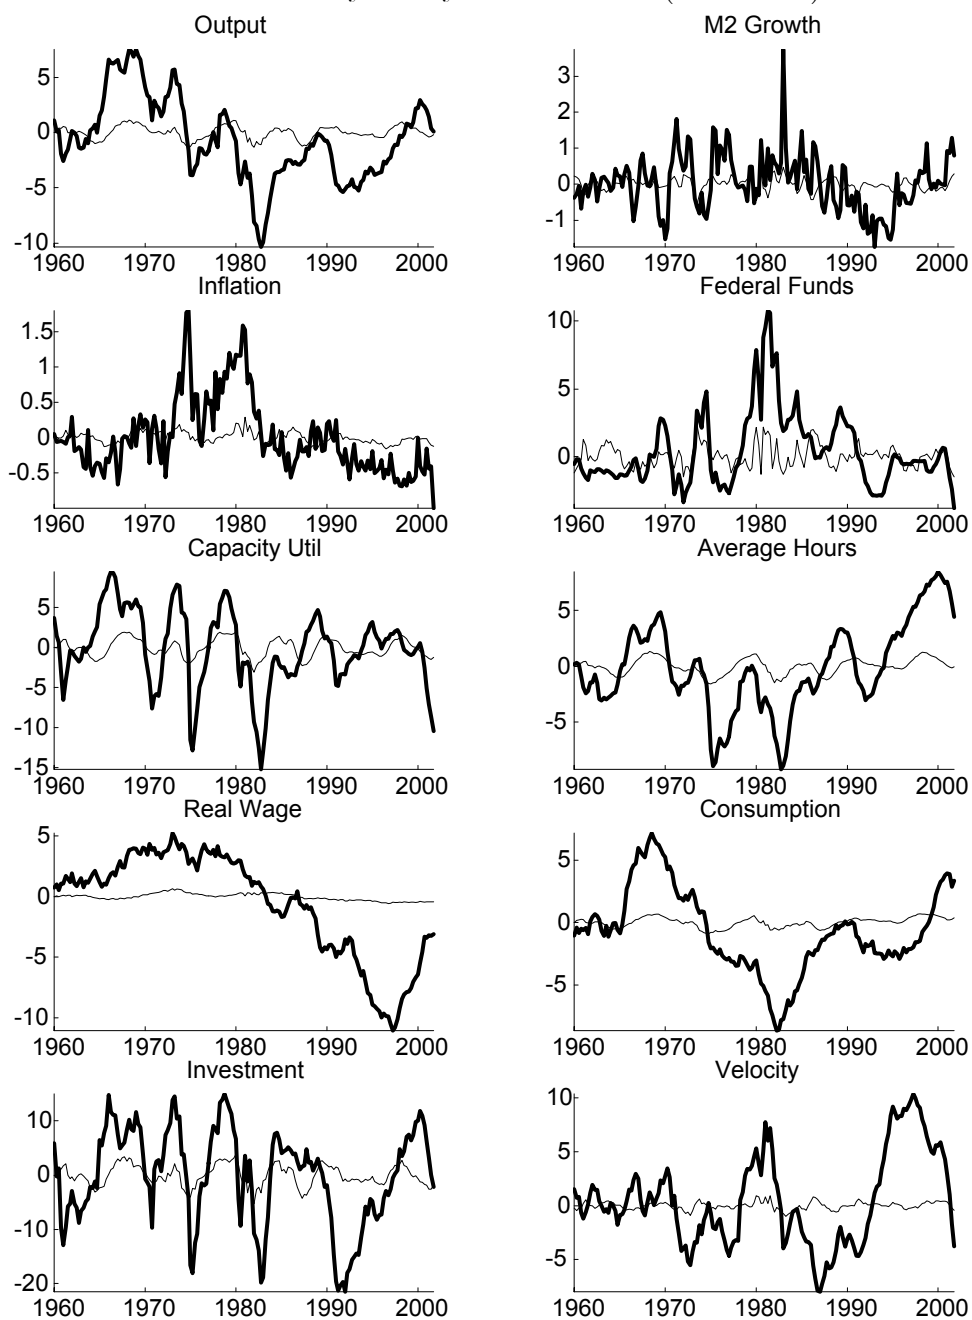
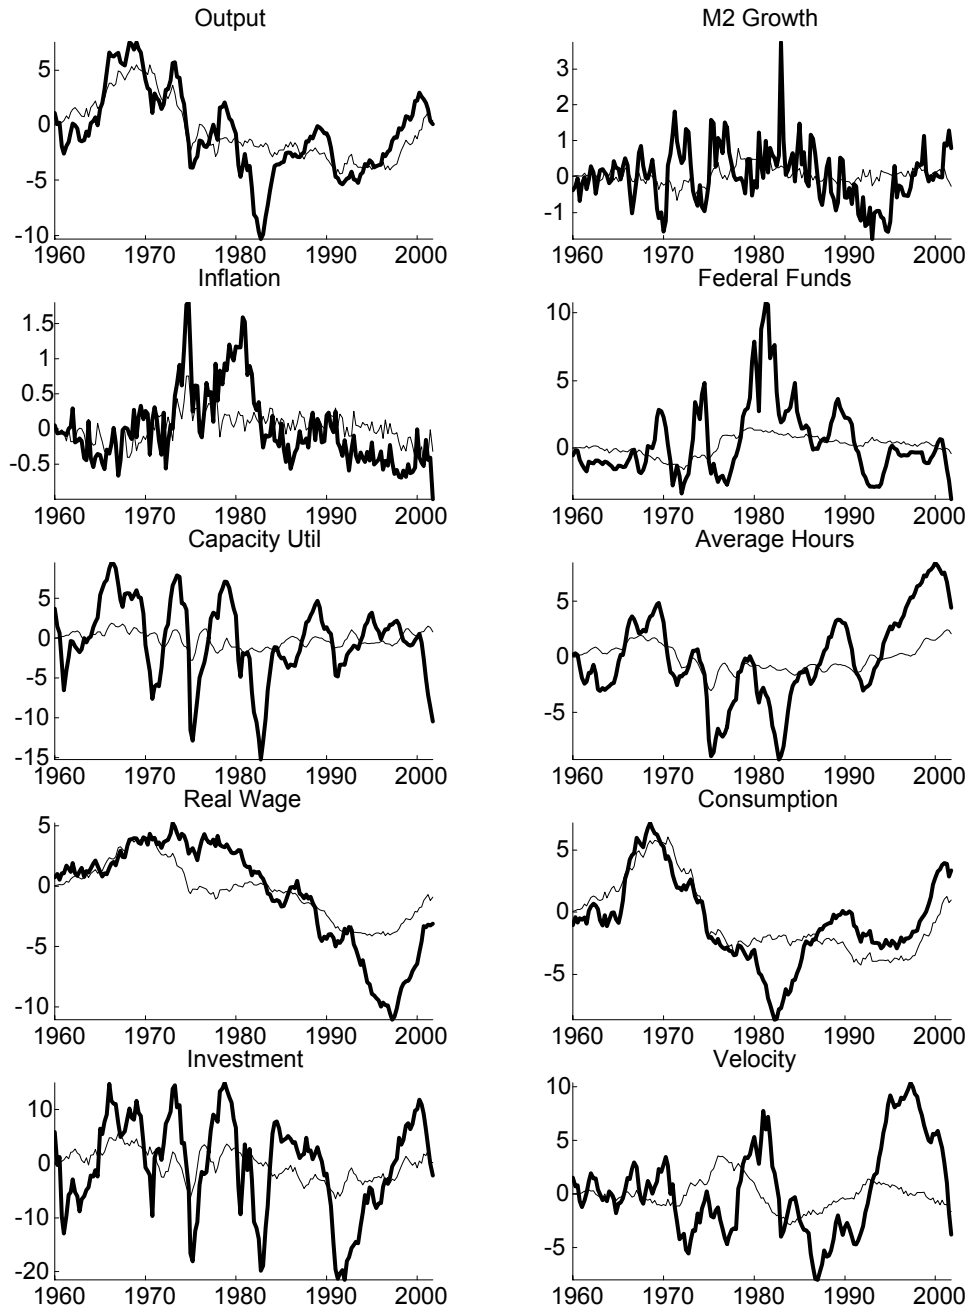


Figure 6 reports the analog of Figure 5, for the case of technology shocks. Consistent with the results in Table 2, the VAR analysis indicates that technology shocks account for a substantial portion of the fluctuations in the data. The portion of the data that it does best on, is the low frequency component. For example, technology shocks appear to play an important role in accounting for the rise and fall in output (relative to trend) before and after the 1970s, and the rise in the late 1990s. They also go a long way towards capturing the low frequency components of the consumption and investment data. Interestingly, it is not clear that technology has very much to do with the business cycle component in the data. Apart from the 1974 recession, technology shocks do not seem highly correlated with the major business cycle fluctuations. In particular, the simulations completely miss the 1970

recession, the recession in the early 1980s and the recession in the early 1990s.

Figure 6: Detrended Historical Data (Thick Line) Versus Component Due to Technology Shocks Alone (Thin Line)



To summarize, the evidence suggests that monetary policy shocks account for only a trivial part of the variation in the data, while technology shocks account for nearly 50 percent of the variation. Although the role of technology shocks appears to be quantitatively large, this role seems to be confined to explaining the relatively low frequency component of the data. In particular, technology shocks appear to play only a small role in triggering business cycle fluctuations. It is unclear at this point how to interpret this. A possibility is that technology shocks have two components. One component has a long run impact on productivity, and the other one has only a transitory impact. Our identification strategy, as we have implemented it so far, is designed to pick up the first one and not the second. So, it is still possible that technology shocks play an important role in driving the business cycle, if the driving force is the stationary component of technology. As noted above, the analysis of this paper is being extended to other shocks, including a stationary shock to technology. When that analysis is completed we will hopefully have a more complete assessment of the role of technology shocks in business cycles.

2.4. Related Literature

There is a growing literature, started by Gali (1999), which attempts to identify the dynamic effects of technology shocks using reduced form methods. In particular, Gali makes the assumption - which is a feature in many models - that innovations to technology are the only disturbances that have an effect on the level of labor productivity in the long run. He showed - and we have followed his lead in our own analysis above - that this assumption is sufficient to permit identifying the dynamic effects of technology shocks in the data. When he did this, he obtained results very different from the ones we reported above. He found that hours worked fall after a positive technology shock. The fall is so long and protracted that, according to his estimates, technology shocks are a source of negative correlation between output and hours worked. Reasoning from the observation that hours worked in fact are procyclical, Gali concluded that some other shock or shocks must be playing the predominant role in business cycles. Thus, he concludes that technology shocks at best play a small role in fluctuations. Moreover, he argues that standard real business cycle models shed little light on whatever small role they do play, because they do not generally imply a protracted fall in employment after a positive technology shock. In effect, real business cycle models are doubly dammed: they address things that are unimportant, and they do it badly at that. Other recent papers reach conclusions that complement Gali's in various ways (see, e.g., Shea (1998), Basu, Kimball and Fernald (1999), and Francis and Ramey (2002).) Francis and Ramey perhaps do not overstate too much when they say (p.2) that Gali's argument is a '...potential paradigm shifter'.

Our results differ from those in the literature in that our point estimates imply a rise

in hours after a policy shock. Confidence intervals are wide, so the disagreement is not as sharp as the point estimates themselves suggest. Still, there is disagreement. Regarding the importance of technology shocks in the cycle, our results so far are qualitatively consistent with Gali's view that they are not important. However, for the reasons noted above, Gali's conclusion may not survive our analysis when we extend it to include other types of shocks to technology.

The remainder of this section reports on our efforts to understand why we find that hours rises after a technology shock, while others (primarily, Gali and Francis-Ramey) find that it falls. The difference in results is surprising, since their fundamental identification assumption - that shocks to technology are the only shocks that have a long-run impact on labor productivity - is also adopted in our analysis. Still, there are a variety of differences between our VARs and those used by Gali and Francis-Ramey. One difference is that the number of variables used in the analysis differs. They tend to work in much smaller, typically bivariate, systems. Various experiments that we have done suggest that this is not the basic source of the difference. We have worked to minimize any possible role played by differences in data definition. For example, we use the same measure of hours, population and output that Francis and Ramey use.¹¹

Our preliminary results suggest that at least part of the reason for the difference in results may lie in the way per capita hours are modeled. In our analysis, it is the level of hours worked that appears in Y_t . This specification is consistent with the property of our model, that hours worked are a stationary variable.¹² By contrast, the work of Gali and Francis-Ramey incorporates the assumption that hours worked require first differencing to induce stationarity.¹³ When we replace log hours in Y_t with its first difference, then we obtain results like those of Gali and Francis-Ramey. That is, we find that hours worked decline after a positive monetary policy shock. The results are presented in Figure 7. The lines indicated by x 's in each panel indicate our point estimates. Note how the response in hours worked is

¹¹As Francis and Ramey point out, the measures of hours and output used to compute productivity in our analysis are not without flaws. The output measure covers the entire economy, while the hours measure excludes government and farm work. Francis and Ramey show that results are not sensitive to a change in variables in which output and hours that match each other more closely.

¹²The eigenvalues of the relevant characteristic equation of our estimated VAR has eigenvalues that are less than unity in absolute value. This implies that all variables in Y_t , including hours worked, are covariance stationary.

¹³No doubt, a test of the null hypothesis that hours worked is an integrated variable would not be rejected. Of course, this is not a basis for modeling hours as having a unit root. First, this would imply rejecting our point estimates, which seems hard to justify. Second, even if there were a unit root in the hours data, then specifying hours in levels in the VAR does not prevent the VAR from finding a unit root.

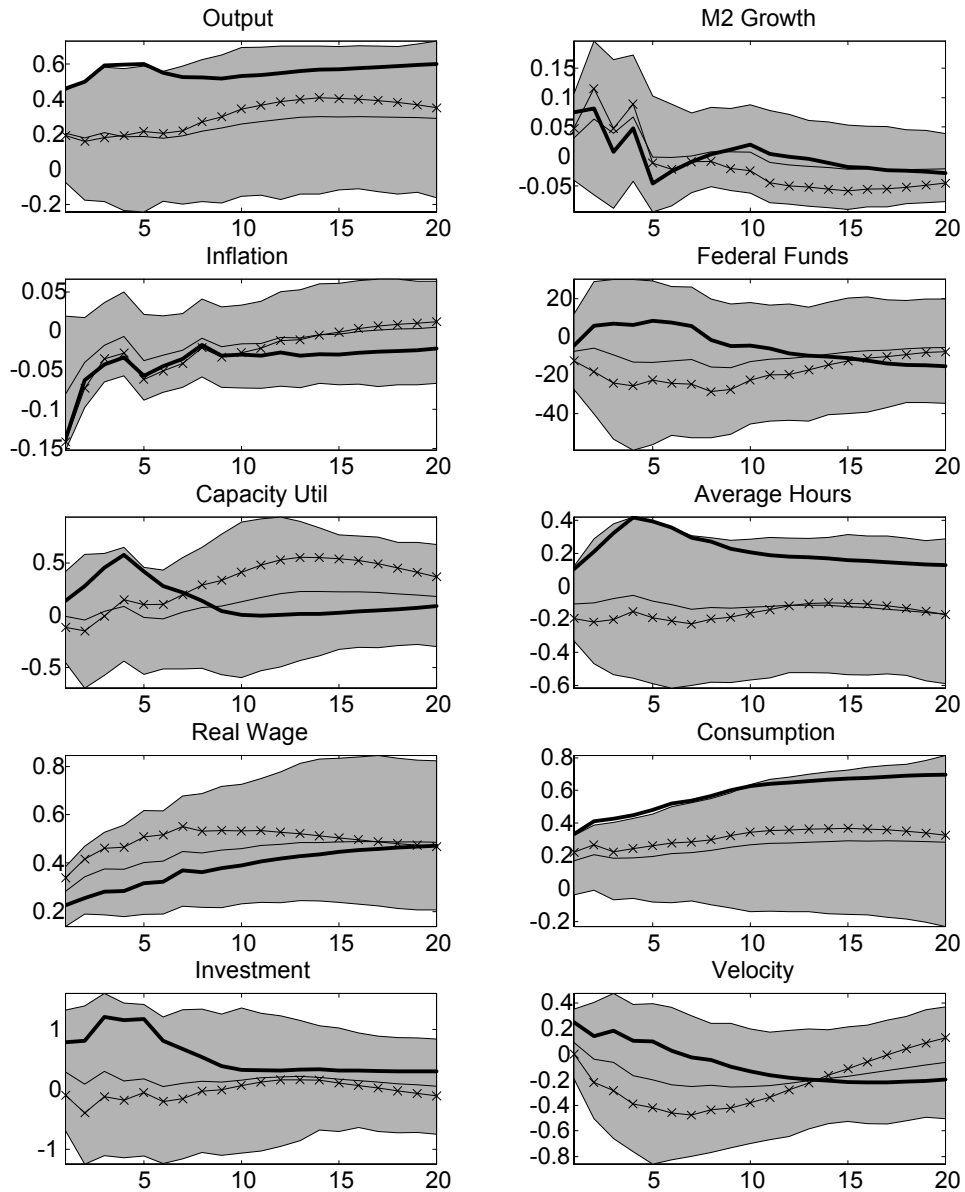
negative for each of the 20 quarters of responses shown. For comparison, the thick dark line in Figure 7 reproduce our baseline point estimates displayed in Figure 3 (we discuss the other lines momentarily). Under our assumption about hours worked, the VAR estimated by Gali and Francis-Ramey is misspecified because hours worked are over-differenced. When there is over-differencing of a variable, a VAR is (at least in population) is a poor approximation to the true data generating mechanism because over-differencing produces a unit moving average root, which is not invertible.

We investigated whether the Gali and Francis-Ramey finding that hours fall after a positive technology shock could simply reflect distortions due to over-differencing. We find that it can. We determined this by generating numerous samples of artificial data from our estimated VAR in which hours appears in level form in Y_t . In each artificial data sample, we fit the version of the VAR in which hours worked appears in growth rate form. We then computed the impulse responses to a technology shock. The mean impulse responses appear as the thin line in Figure 7. The gray area represents the 95 confidence interval of the simulated impulse response functions.¹⁴ The key thing to note is how close the thin line and the line indicated by x 's are. This means that our VAR with hours specified in levels is a very good predictor (its prediction is indicated by the thin line) of what one gets when one estimates impulse response functions based on a VAR with hours specified in growth rate form. Put differently, our levels specification passes an encompassing test: it successfully

¹⁴That is, for each lag we ordered the impulse responses from smallest to largest. The confidence interval is defined by the interval from the 25th element in this ranking to the 975th element (the number of data sets that were simulated is 1000.)

predicts the results obtained by Gali and Francis-Ramey.

Figure 7 Response of Variables to Technology Shock
 Data Generating Mechanism: VAR in Level of Hours
 CI and Mean Generated Using H DGP



36 Notes:

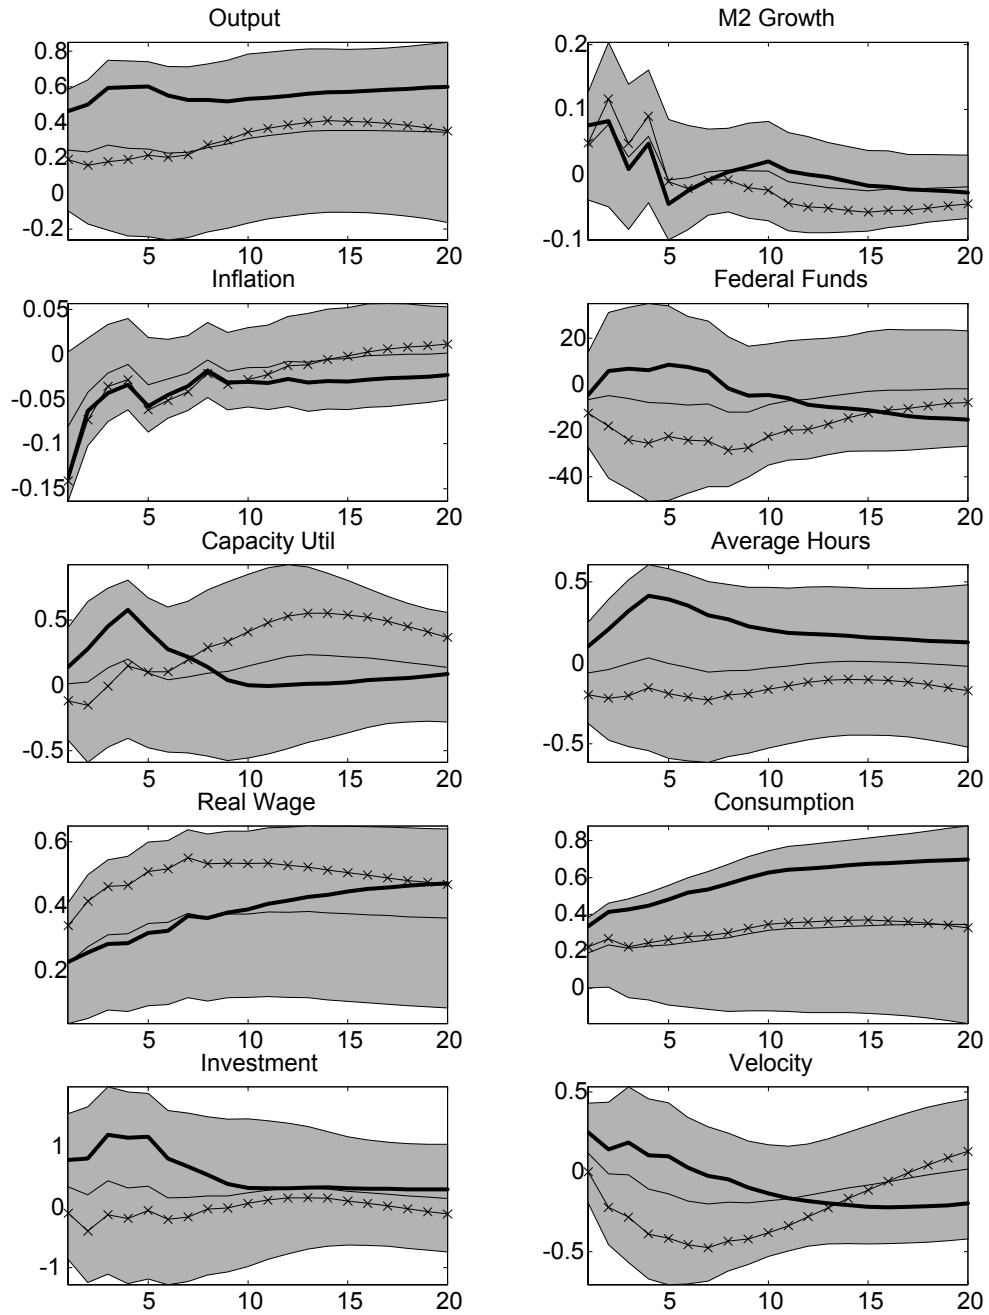
- Thick Line - Baseline Point Estimates Based on VAR with Level Hours.
- X's - Point Estimates Based on VAR with Log First Difference Hours.
- Thin line - Average Responses, Impulse Responses Computed From VAR with Log First Difference Hours, in Artificial Data Generated from Baseline VAR.

We also did the reverse test. We investigated whether the type of VAR specification adopted by Gali and Francis-Ramey can predict the results we obtained when we analyze the differenced data. The results are reported in Figure 8. The thick, solid line and line composed of x 's reproduce the analogous lines from Figure 7 for convenience. The thin line in Figure 8 is the prediction of the VAR with hours in first differences for the impulse responses one obtains with a VAR with hours in levels. The gray area is the associated confidence interval. What is notable about the results, is that the first differenced hours-based VAR has trouble explaining our results that hours rises in response to a positive technology shock. To see this, note that the thin solid line is negative at all lags displayed, and so is qualitatively different from our baseline results, indicated by the thick, solid line. In this sense, the first difference VAR fails the econompassing test. These observations fail to take into account the wide confidence interval. According to the confidence interval, the results obtained based on the level specification are not extremely out of line from what one expects from a VAR with first differenced hours.

We conclude that the data are consistent with the results that the Gali and Francis-Ramey findings are the result of overdifferencing of hours. However, we have not completely ruled out the possibility that Gali and Francis-Ramey are right, while our results are an artifact of not first differencing the hours data when first differencing is the correct thing to

do.

Figure 8: Response of Variables to Technology Shock
 Data Generating Mechanism: VAR in First Difference of Hours



38 Notes

Thick Line - Baseline Point Estimates Based on VAR with Level Hours.

X's - Point Estimates Based on VAR with Log First Difference Hours.

Thin line - Average Responses, Impulse Responses Computed From VAR with Log Level Hours, in Artificial Data Generated from VAR with Level Hours.

3. Results

This section reports our parameter estimates and diagnoses model fit by evaluating how well the model’s impulse responses match those estimated in the data.

We divide the parameters into those whose values are estimated here and those whose values are taken from elsewhere. The latter are reported in Table 1. For the most part, the values used are standard. The parameter governing market power of household labor suppliers, λ_w , is arbitrarily set to 1.05. In future drafts, we plan to include this parameter in the list of parameters to be estimated.

Table 1: Parameters that Do Not Enter Formal Estimation Criterion		
discount factor	β	1.03 ^{-.25}
capital’s share	α	0.36
capital depreciation rate	δ	0.025
markup, labor suppliers	λ_w	1.05
mean, money growth	μ	1.017
labor utility parameter	ψ_0	set to imply $L = 1$
real balance utility parameter	ψ_q	set to imply $Q/M = 0.44$
fixed cost of production	ϕ	set to imply steady state profits = 0

The 13 model parameters that we estimate here are:

$$\gamma \equiv (\lambda_f, \xi_w, \xi_p, \sigma_q, S'', b, \sigma_a, \\ \text{6 parameters governing exogenous shocks}).$$

As a reminder, $\lambda_f \geq 1$ is the markup set by monopolist intermediate good suppliers, ξ_w is the probability that a household cannot reoptimize the wage for its differentiated labor service, ξ_p is the probability that the monopoly supplier of a differentiated intermediate good cannot reoptimize its price, σ_q is a curvature parameter related to money demand, S'' is a curvature parameter related to adjustment costs on investment, b is the habit parameter, and σ_a is the parameter controlling the curvature on costs of capital utilization. The list of ‘parameters governing monetary policy and technology’ are simply the parameters in () and (1.3).

Corresponding to each γ , we compute a set of model impulse response functions, $\psi(\gamma)$. Denote the impulse response functions for the data by $\hat{\psi}$. This is the list of numbers reported

in Figures 2 and 3.¹⁵ We have not yet implemented our procedure for also estimating the other shocks.¹⁶ So, the vector, $\hat{\psi}$, summarizes the first 20 lags in the response function of our 9 variables to technology and monetary policy. Our estimator of γ minimizes the distance between $\psi(\gamma)$ and $\hat{\psi}$:

$$\hat{\gamma} = \arg \min_{\gamma} (\hat{\psi} - \psi(\gamma))' V^{-1} (\hat{\psi} - \psi(\gamma)),$$

where V is the diagonal matrix composed of our estimate of the sample standard deviation in $\hat{\psi}$. Essentially, our estimation procedure tries to get the model's impulse responses as close to the thick line in Figure 2 and 3. It pays most attention to impulses where the gray area is the thinnest. We computed standard errors for the estimated values of γ using the usual delta function method.

The results are reported in the following two tables. Table 2 reports the values of the

¹⁵There are 360-6 elements in $\hat{\psi}$: nine variables, 20 lags and 2 shocks. We subtract 6 from the total to take into account the 6 variables whose contemporaneous responses to a monetary policy shock are assumed to be zero under our identifying assumptions.

¹⁶We do have some intriguing, preliminary results. With model-based estimation, a subset of the elements of $\hat{\psi}$ is a function of unknown parameters: the elements of the 6 dimensional orthonormal matrix, w , discussed in section 2.2.2. We began model-based identification, by working with one shock in e_{1t} alone. We interpreted the first element of e_{1t} as the shock to the preference for leisure (or, to labor market power). The only part of w that is relevant for this is w_1 , the first column (so, the only restriction to implement is that the length of w_1 be unity). Let the subset of $\hat{\psi}$ that corresponds to the dynamic response to a leisure shock be denoted $\hat{\psi}'(w_1)$. We attempted to estimate $\hat{\psi}'(w_1)$ by minimizing its distance from the corresponding part of $\psi(\gamma)$, which we denote by $\psi'(\gamma)$. Regardless of starting values, the estimation procedure always chose w_1 and the variance of the preference shock in the model to make $\hat{\psi}'(w_1)$ and $\psi'(\gamma)$ close to zero.

We conjecture that this finding reflects a problem with estimating just one element in a list of several shocks, by our model-based approach. To see the problem, recall the finding in the existing literature on dynamic factor analysis, which suggests that a small number of shocks account for a large amount of the variation in the data. A corollary of this is that a large number of shocks explain very little. Apparently, our model-based procedure, when applied to only one shock out of potentially several, undertakes a 'race to the bottom', by choosing w_1 to produce the shock with *least* variance. In effect, the distance between $\hat{\psi}'(w_1)$ and $\psi'(\gamma)$ is minimized by setting each close to zero (setting $\hat{\psi}'(w_1)$ exactly to zero is impossible, since that would be inconsistent with Y_t having full rank). We suspect that the right way to go is estimate all the shocks subject to model-based identification at the same time. This is because the overall variation in e_{1t} is fixed. The estimation procedure could still set the variance of the leisure shock to zero, but only at the cost of setting other variances higher. If one did not actually want to apply a label to all the elements in e_{1t} , one could avoid the trivial outcome described above by associating some of the w_i 's in w with the shocks that explain the least variance in output, and attributing them to measurement error, or such.

economic parameters, while results for the parameters of the exogenous shock processes are reported in Table 3.

We divide our discussion into three parts. We begin with the benchmark estimation results, in which γ is chosen make the model match all the impulses simultaneously. To gain an understanding for the role played by impulse responses to technology and to policy in the results, we perform two other analyses. First, we re-estimate γ by including only responses to policy shocks in the estimation. Then, we re-estimate γ by including only responses to technology shocks. In each of these two cases, we must delete from γ the components pertaining to the impulses not included.

3.1. Benchmark Results

We now turn to the benchmark estimation. The first row of Table 2 exhibits the resulting model parameter values. Our impression is that these are all reasonable. The estimated value of λ_f implies a steady state markup of 14 percent. The estimated value of ξ_w implies that wage contracts last on average a little over 4 quarters, while the estimated value of ξ_p implies that price contracts last a little under 2 quarters. By comparison with existing survey evidence on the degree of sticky wages and prices, our estimated amount of stickiness is quite modest. The habit parameter, b , is very similar to the value used in Boldrin, Christiano and Fisher (2001), using a non-monetary version of the model here, to match basic asset pricing facts such as the equity premium.

Table 2: Estimated Economic Parameter Values (Standard Errors)							
Estimation Based Estimated Responses to:	λ_f	ξ_w	ξ_p	σ_q	S''	b	σ_a
Policy and Technology Shocks Simultaneously	1.14 (.016)	.78 (0.04)	.42 (0.12)	14.13 (1.74)	7.69 (1.33)	0.73 (0.07)	0.05 (0.01)
Policy Shocks Only	1.15 (0.27)	0.73 (0.03)	0.45 (0.06)	12.33 (0.61)	9.97 (3.36)	0.77 (0.04)	0.03 (0.01)
Technology Shocks Only	1.65	0.99	0.08	18.67	20.00	0.60	0.02

The estimated parameters of the exogenous shocks for the benchmark run are reported in the first column of Table 3. The first order autocorrelation of the growth rate of technology is estimated to be 0.80. The standard deviation of the innovation is 0.12 percent. This corresponds to an overall unconditional standard deviation of 0.2 percent for the growth rate of technology. These results differ somewhat from Prescott (1986), who estimates the properties of the technology shock process using the Solow residual. He finds the shock is

roughly a random walk and its growth rate has a standard deviation of roughly 1 percent.¹⁷ Our results are potentially consistent with Prescott’s findings, for three reasons. First, we model technology shocks as having two components, a temporary one and a permanent one. The analysis up to now has only included the permanent one. Second, from the perspective of our model, Prescott’s estimate of technology confounds technology with variable capital utilization. Both these factors may explain why our technology shock standard deviation is one-fifth the size of Prescott’s. They may also explain why we find so much more persistence. A more conclusive finding on this dimension awaits our analysis of the model with the additional shocks.

According to the estimates in Table 3, monetary policy responds immediately to a positive realization of the technology shock. For every one percent innovation in technology, the money stock jumps by 2 percent, according to the point estimates. At the same time, the standard error on this parameter is estimated particularly imprecisely, with a standard error of 1.17. The autoregressive parameter on the response of money to technology indicates that money growth increases not just in the period of a technology shock, but jumps again in the period afterward.

We now consider the results in Table 3 pertaining to monetary policy shocks. These indicate that a monetary policy shock drives up the money stock by 0.11 percent, with an extremely tight standard error. The increase in money growth is autocorrelated over time.

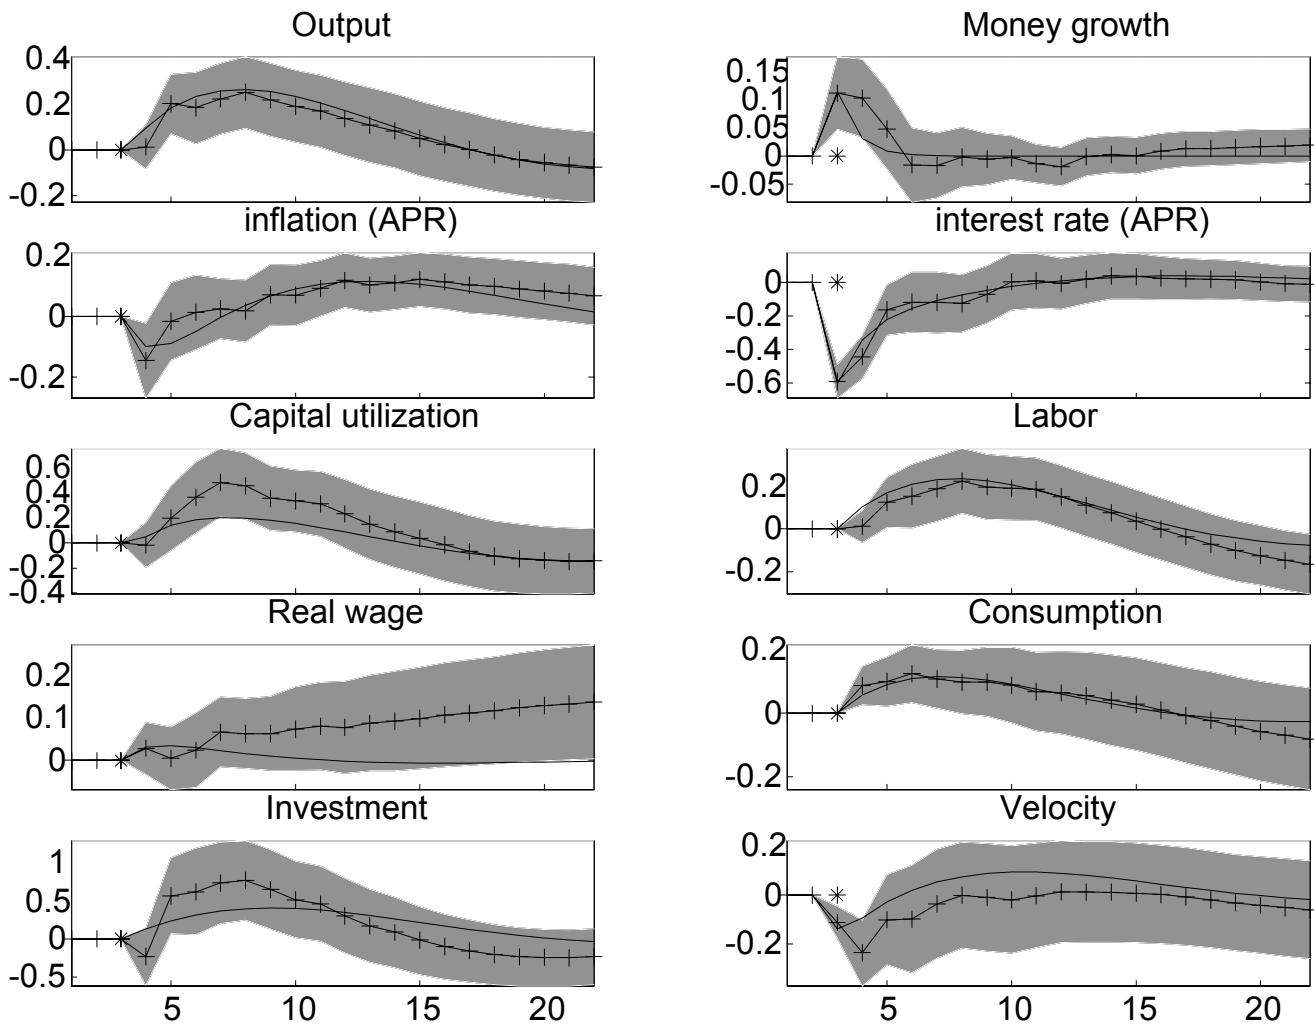
	Estimation Based on Estimated Responses to:		
	Policy and Technology	Policy	Technology
Parameter	Shocks Simultaneously	Shocks Only	Shocks Only
ρ_x	0.80 (0.11)	na	0.92
σ_{ε_x}	0.12 (0.06)	na	0.05
ρ_{μ_x}	0.47 (0.10)	na	0.29
c_{μ_x}	2.07 (1.17)	na	3.59
ρ_{μ_p}	0.27 (0.07)	0.27 (0.10)	na
$\sigma_{\varepsilon_{\mu_p}}$	0.11 (0.005)	0.13 (0.01)	na

Figures 9 and 10 display the dynamic response of the model variables (see the continuous

¹⁷Prescott (1986) actually reports a standard deviation of 0.763 percent. However, he adopts a different normalization for the technology shock than we do, by placing it in front of the production function. Instead, our technology shock multiplies labor directly in the production and is taken to a power of labor’s share. The value of labor’s share that Prescott uses is 0.70. When we translate Prescott’s estimate into the one relevant for our normalization, we obtain $0.763/.7 \approx 1$.

lines) at the estimated parameter values. The period of the shock is indicated by a ‘*’. For convenience, we have included the empirical impulse responses (see the lines marked by ‘+’) and 95% confidence intervals (see the grey areas) estimated in the data and reported in Figures 2 and 3. In our view, the fit is very good. The response of capital utilization is slightly weak, though still inside the confidence intervals everywhere. Velocity misses the confidence interval very slightly in the period after the shock.

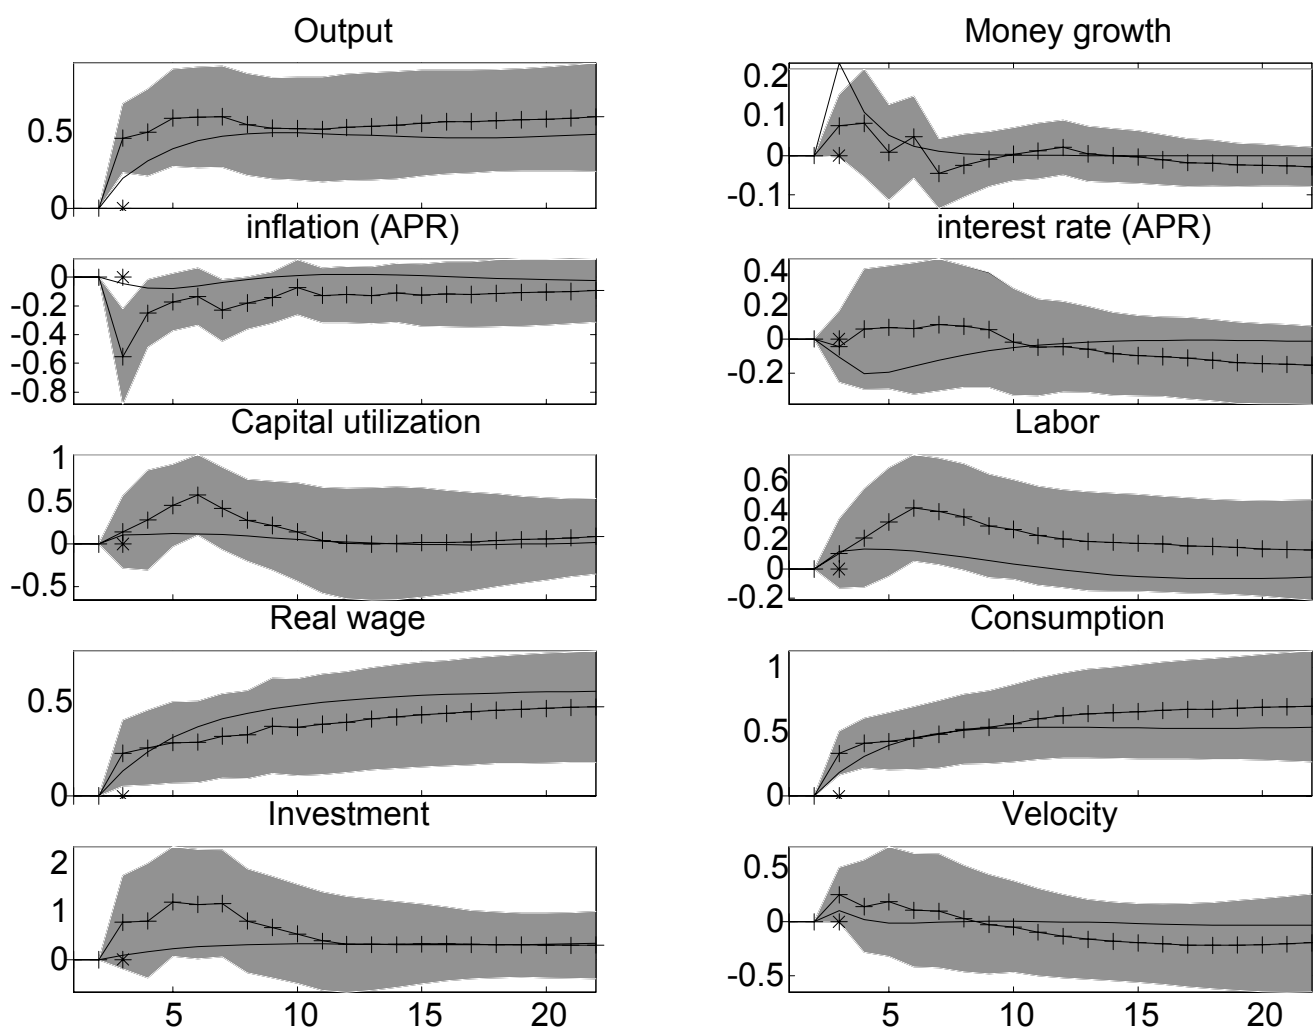
Figure 9: Properties of Benchmark Estimated Model - Dynamic Response to Monetary Policy Shock



Now consider the responses to a technology shock, reported in Figure 10. Here too the model mimics the impulse responses in the data reasonably well. However, it is easier to find fault with the model in Figure 10 than it is to do so in Figure 9. Inflation in the model does not quite fall enough, and the response in capital utilization, labor and output is somewhat on the weak side. Finally, the response of money growth is too strong.

Figure 10: Properties of Benchmark Estimated Model - Dynamic Response to Technology Shock

Figure 2: Model and Data Impulse Response Functions to a Non-stationary Technology Shock



To better understand the reasons for these estimation results, we turn to estimation based

on only policy and technology shocks, in the next two subsections.

3.2. Estimation Based on Policy Shocks Alone

We can obtain insight into what is driving the results by considering what happens when model parameters are estimated using only the impulse responses to a monetary policy shock. For this experiment, the parameters governing the univariate representation of the technology shock and the parameters governing the response of monetary policy to technology were held fixed at the benchmark estimates. A notable feature of the results, is that they are not very different. For example, the estimated parameter values in Tables 2 and 3 are very similar for the benchmark run, and the run pursued here. In terms of the responses to a policy shock, the improvements are nearly imperceptible. Similarly, in terms of the response to technology shocks, the deterioration in the performance of the model is quite small. This can be seen by comparing the results in Figure 12 with those in Figure 10. It appears that the benchmark estimation results have been driven by the empirical estimates to a monetary policy shock, and that those estimates work reasonably well for the response of technology shocks too.

Figure 11: Properties of Model Fit to Policy Impulse Responses Only

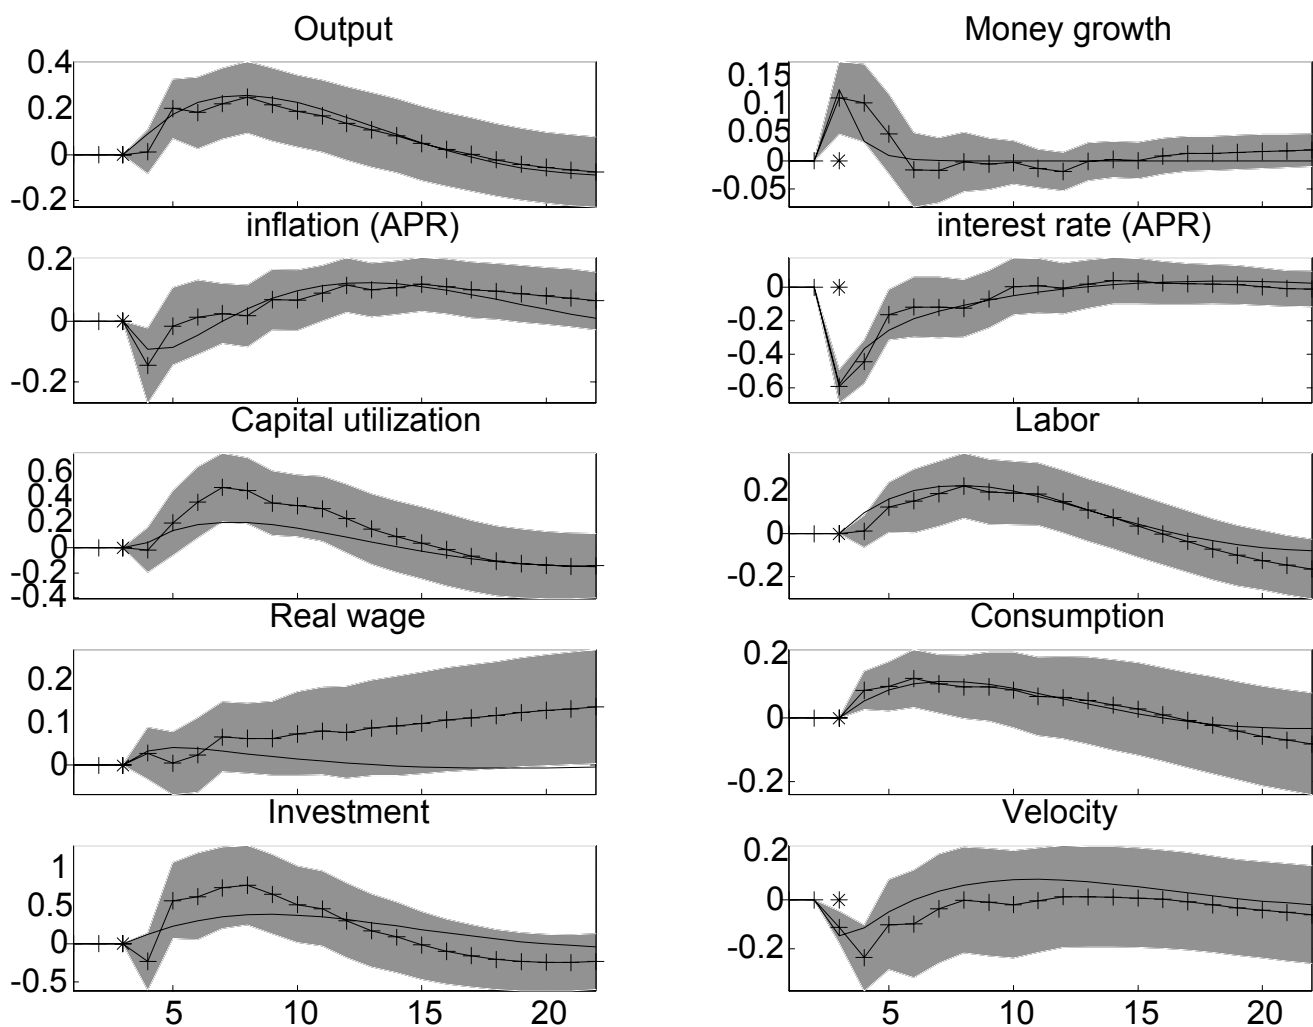
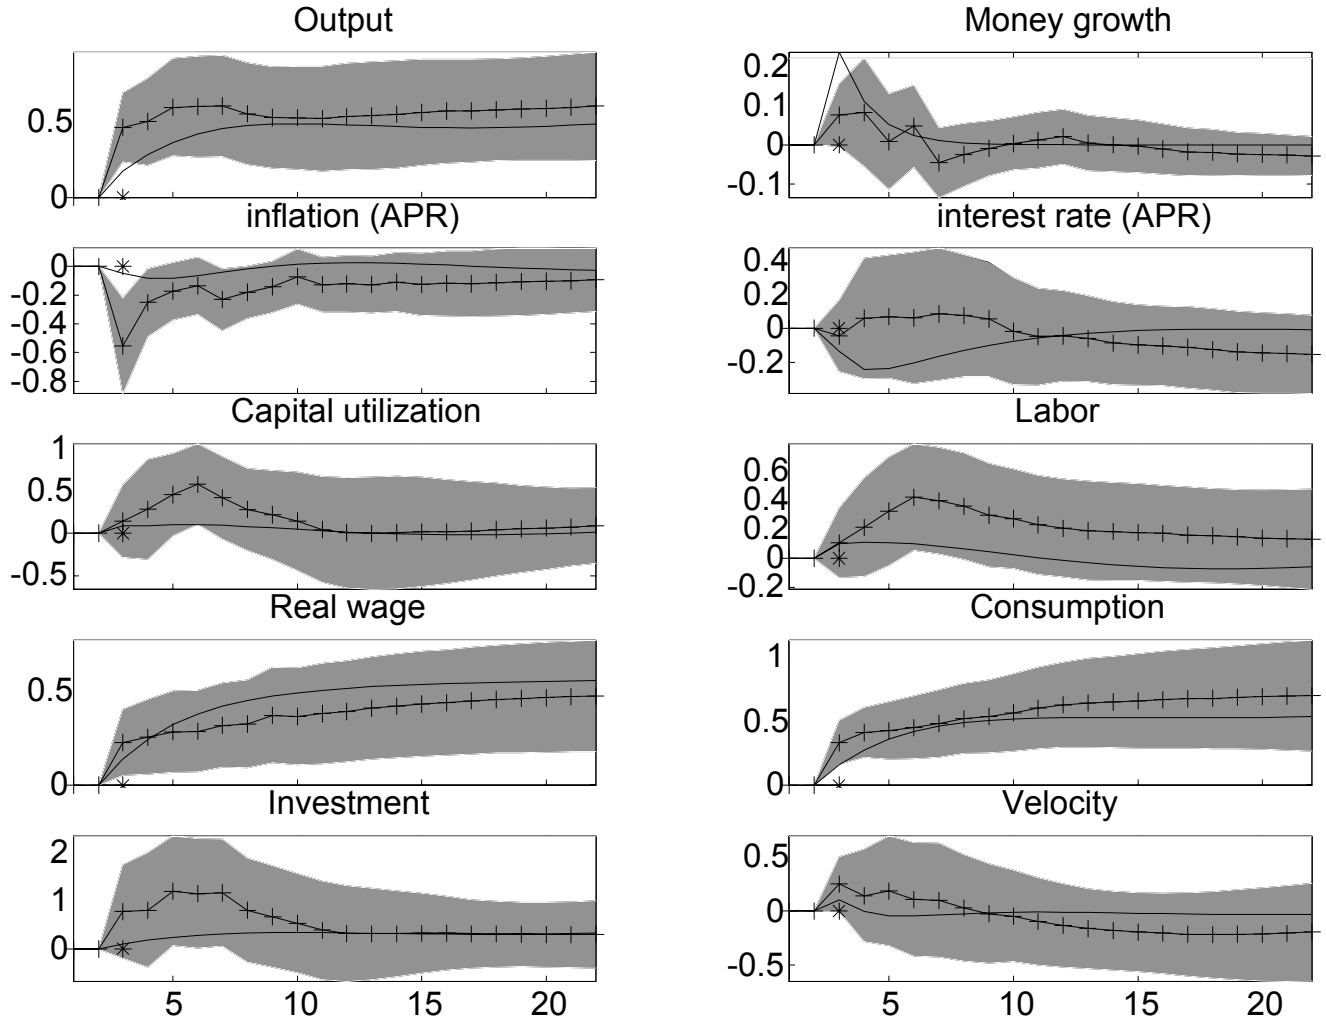


Figure 12: Properties of Model Fit to Policy Impulse Responses Only

Figure 2: Model and Data Impulse Response Functions to a Non-stationary Technology Shock



3.3. Estimation Based on Technology Shocks Alone

We now turn to the results based on estimating the model on technology shocks alone. These results are quite different from our benchmark findings. Table 2 reports the new parameter values. Stickiness in prices has been almost completely eliminated, while the degree of stickiness in wages has moved to its upper bound of 0.99. Adjustment costs in

investment and the degree of market power of intermediate good producers have increased substantially.¹⁸ According to Table 3, the standard deviation of the technology shock was cut in half, and the response of money to technology was increased from about 2 to about 3.

Figures 13 and 14 indicate what the consequences of these new parameter values are. Figure 14 shows that the new parameters correct the main failures of the benchmark model in reproducing the dynamic responses to technology. However, these improvements come at great cost in terms of being able to fit the dynamic response to a monetary policy shock. The effect of the shock on inflation in the first 20 quarters is now completely dominated by cut in the interest rate. With the fall in prices and the rise in nominal demand, labor, capital utilization, consumption, investment and output surge. In the case of output and labor, the increase is far too great. The enormous stickiness in the nominal wage rate relative to intermediate good prices implies that the real wage stays low.

In a later draft we will more fully diagnose the implications of these model results. For now, we note that these results confirm the conclusion of the previous subsection: the benchmark results are principally driven by the empirical responses to monetary policy. It is interesting that the monetary policy shock, which has so little impact on the dynamics of the data, plays such an important role in pinning down model parameters.

¹⁸Standard errors have not yet been computed for these parameter values.

Figure 13: Properties of Model Fit to Technology Impulse Responses Only

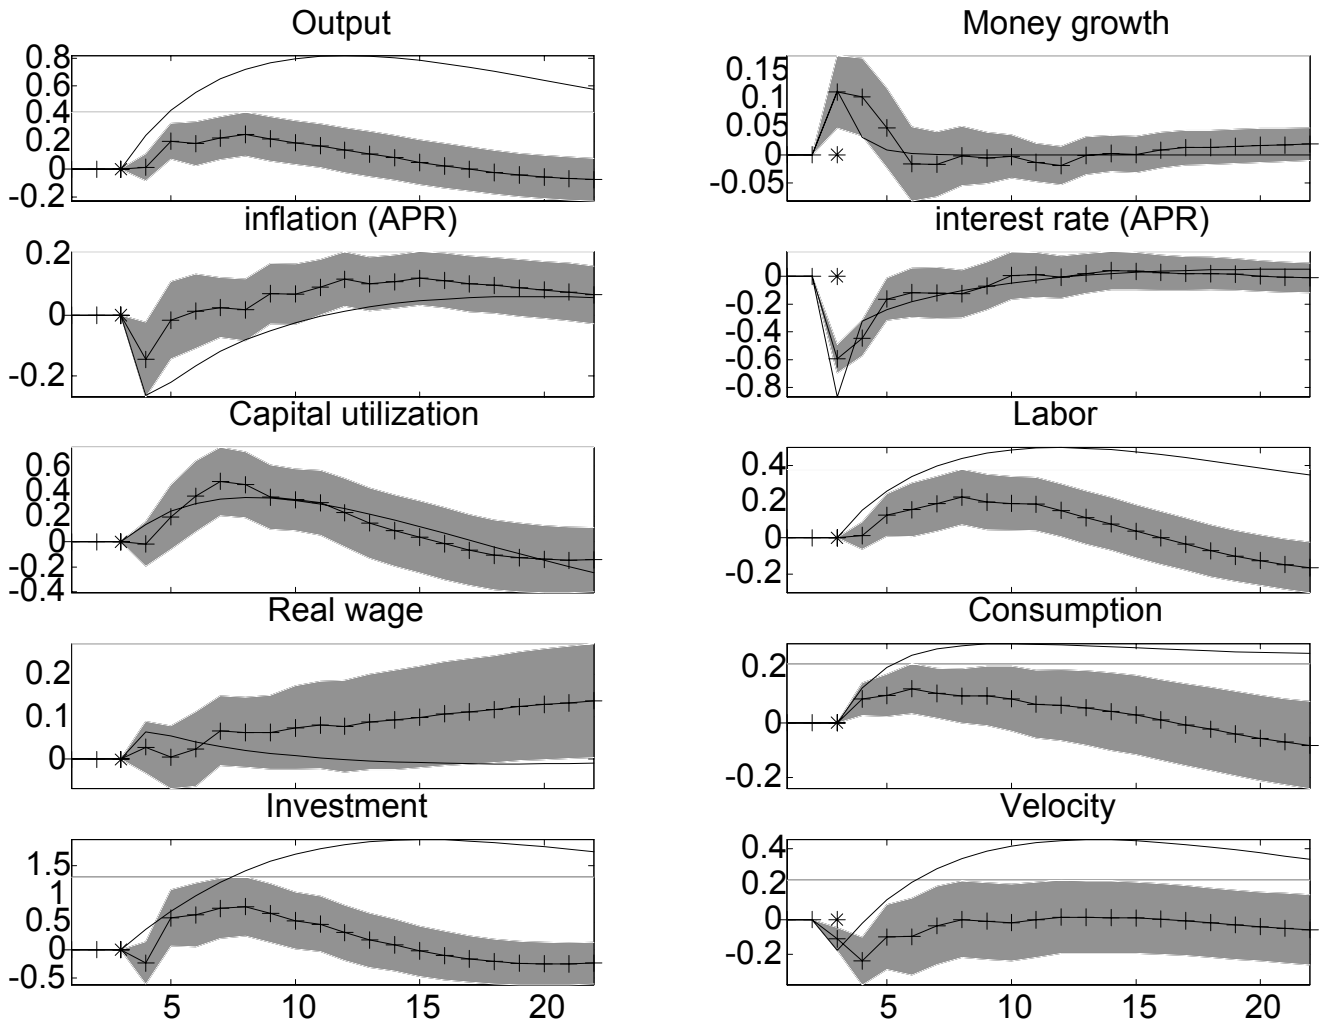
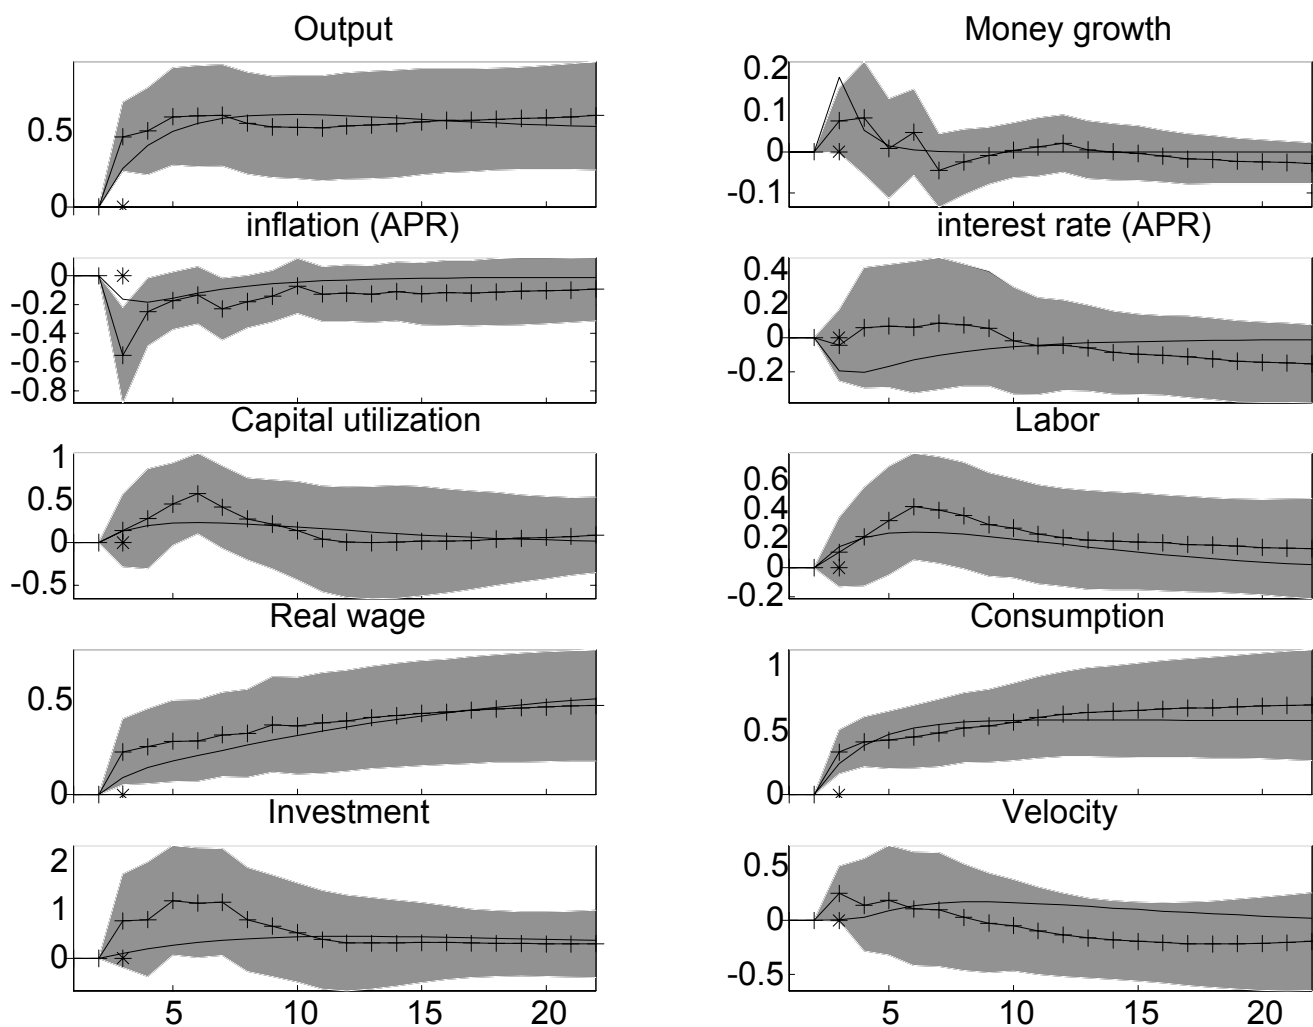


Figure 14: Properties of Model Fit to Technology Impulse Responses Only

Figure 2: Model and Data Impulse Response Functions to a Non-stationary Technology Shock



4. Conclusion

[to be added later]

References

- [1] Barth, Marvin and Valerie Ramey,
- [2] Boldrin, Michele, Lawrence J. Christiano and Jonas Fisher, 2001, ‘Asset Pricing Lessons for Modeling Business Cycles,’ *American Economic Review*.
- [3] Calvo, G.A., 1983, ‘Staggered Prices in a Utility-Maximizing Framework’, *Journal of Monetary Economics* 12:383-398.
- [4] Christiano, Lawrence J., 2002 ‘Solving Dynamic Equilibrium Models by a Method of Undetermined Coefficients’, forthcoming, *Computational Economics*.
- [5] Christiano, Lawrence J., Martin Eichenbaum and Charles Evans, 1999, Monetary Policy Shocks: What Have We Learned, and to What End?, in Taylor and Woodford, *Handbook of Macroeconomics*.
- [6] Christiano, Lawrence J., Martin Eichenbaum and Charles Evans, 2001, ‘Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy’, manuscript.
- [7] Erceg, Christopher, J., Henderson, Dale, W. and Andrew T. Levin, 2000, ‘Optimal Monetary Policy with Staggered Wage and Price Contracts’, *Journal of Monetary Economics*, 46(2), October, pages 281-313.
- [8] Francis, Neville, and Valerie A. Ramey, 2001, ‘Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,’ manuscript, UCSD.
- [9] Gali, Jordi, 1999, ‘Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?’ *American Economic Review*, vol. 89, no. 1, 249-271.
- [10] Gali, Jordi, Mark Gertler and J. David Lopez-Salido, 2001, ‘Markups, Gaps and the Welfare Costs of Business Fluctuations,’ May.
- [11] Gali, Jordi, J. David Lopez-Salido, and Javier Valles, 2002, ‘Technology Shocks and Monetary Policy: Assessing the Fed’s Performance’, National Bureau of Economic Research Working Paper 8768.

- [12] Greenwood, Jeremy, Zvi Hercowitz and Per Krusell, 1998, 'Long-Run Implications of Investment-Specific Technological Change,' *American Economic Review*, 87:3, 342-362.
- [13] Greenwood, Jeremy, Zvi Hercowitz and Per Krusell, 1998a, 'The Role of Investment-Specific Technical Change in the Business Cycle,' *European Economic Review*.
- [14] Hall, Robert, 1991, 'Macroeconomic Fluctuations and the Allocation of Time,' *Journal of Labor Economics*, 15, S223-S250.
- [15] Shapiro, Matthew, and Mark Watson, 1988, 'Sources of Business Cycle Fluctuations,' NBER *Macroeconomics Annual*, pp. 111-148.
- [16] Uhlig, Harald, 2001, 'What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure,' February 27, Humboldt University, Berlin.
- [17] Uhlig, Harald, 2002, 'What Moves Real GNP?', April 8, Humboldt University, Berlin.