Permanent and Transitory Technology Shocks and the Behavior of Hours: A Challenge for DSGE Models

Eric R. Sims*
University of Notre Dame and NBER

August 25, 2011

Abstract

This paper uses a quarterly version of the Basu, Fernald, and Kimball (2006) utilization-adjusted TFP series and extends the structural VAR analysis of Fisher (2006) to identify three different kinds of technology shocks – permanent neutral, transitory neutral, and investment specific. Positive transitory neutral shocks are associated with an expansion in hours worked; permanent neutral shocks lead to a reduction in hours. There is significant autocorrelation in growth rates conditional on a permanent neutral shock, so that much of the eventual rise in productivity is anticipated well in advance. Investment specific shocks lead to a significant expansion in hours worked. Overall, the three technology shocks combine to explain about 50 percent of the cyclical variation in output and hours. The paper asks how well standard medium scale DSGE models – with price and wage stickiness and a number of real frictions – can account for the conditional responses to these technology shocks. Overall, these models fit the responses better with fewer frictions than is typically found in the literature. In particular, the best-fitting parameter configuration features very low investment adjustment costs, no price or wage indexation, and comparatively little price and wage stickiness.

*Contact information: esims1@nd.edu, (574) 631-6309, 723 Flanner Hall, Notre Dame, IN 46556. I am grateful to Ruediger Bachmann for helpful comments.

1 Introduction

A central question in macroeconomics concerns the role of technology shocks in driving business cycle fluctuations. The conclusions reached in the literature range from technology shocks being the primary driver of cycles (the early RBC literature), to technology shocks being largely irrelevant (Gali, 1999, and Francis and Ramey, 2005), to somewhere in between (Fisher, 2006, and much of the estimated DSGE literature, e.g. Smets and Wouters, 2007). Regardless of how important technology shocks are as a source of variation in output over the business cycle, the conditional responses of endogenous variables to technology shocks can shed light on the underlying structural features of the economy and the appropriate modeling environments (e.g. sticky vs. flexible price models). The present paper contributes to this literature by identifying three different kinds of technology shocks in a structural VAR setting – transitory neutral, permanent neutral, and investment-specific – and then asks how well, and under what kind of parameter configuration, current state of the art DSGE models can match these responses.

The identification of technology shocks is made inherently difficult by the fact that "technology" is not directly observed. Rather, theoretical implications must be imposed on observed data in order to tease out measures of technology. The early RBC literature adopted growth accounting techniques based on neoclassical production functions and measured technology shocks via Solow residuals. This literature emphasizes persistent but transitory changes in the measured Solow residual as a source of fluctuations, and shows that they can account for a large share of output movements. These growth accounting techniques have been criticized as producing poor measures of true technical change, primarily due to unobserved input variation (e.g. Summers, 1986; Shapiro, 1993; Burnside and Eichenbaum, 1996; and Basu, 1996). Partly as a result, a second strand of the literature moved away from growth accounting techniques and towards the identification of technology shocks based on the behavior of measured labor productivity at frequency zero in a VAR setting. This literature exploits the shared prediction of the vast majority of models that only technology shocks can permanently impact productivity in the long run. With a subtle but potentially important difference from the earlier literature – permanent vs. persistent but transitory shocks – this literature typically finds that technology shocks are an unimportant source of fluctuations. Further, the conditional correlation between surprise technological improvement and hours worked is often found to be negative (e.g. Gali, 1999).[1] This conditional correlation is at odds with the prediction of relatively frictionless flexible price models, leading many to the conclusion that sticky price, Keynesian type models with an important role for "demand" are more promising than flexible price models.

Whereas the first two strands of the literature focus on neutral technological improvement, a third strand, popularized in Greenwood, Hurcowitz, and Krusell (1997), studies the role of investment specific technical change. Here the consensus has emerged that investment-specific technology shocks likely are an important source of fluctuations (Fisher, 2006). Finally, a fourth strand of the

[1] It should be pointed out that the negative impact effect on hours of positive technology shocks is far from universally accepted, and there is a debate over how hours should enter the VAR (first differences vs. levels). See Christiano, Eichenbaum, and Vigfusson (2006a) for a summary of the issues, as well as a discussion in Section 3.2.

literature, best exemplified in Basu, Fernald, and Kimball (2006), uses simple theoretical predictions to "purify" Solow residuals of movements owing to unobserved input variation. These authors reach conclusions similar to the VAR literature – in particular, they argue that technology shocks are permanent and contractionary, in the sense of improved technology leading to an immediate reduction in hours worked.

In the present paper I combine elements from all four of these strands of the literature to study the role of technology shocks. My analysis jointly considers the role of the persistent but transitory neutral shocks of the RBC literature, the permanent neutral shocks of the VAR literature, and investment-specific shocks. I measure neutral technological change with a quarterly version of the purified Solow residual of Basu, Fernald, and Kimball (2006), which I hereafter refer to as "adjusted total factor productivity (TFP)". Following Greenwood, Hercowitz, and Krussell (1997), I measure investment-specific technological change by the relative price of investment goods to consumption goods. I depart from Basu, Fernald, and Kimball (2006), who assume that adjusted TFP follows a univariate random walk, in allowing there to be both a permanent and a transitory component to neutral technology. This is motivated by the findings of Barsky and Sims (2011), who, in studying the role of "news shocks", find that surprise movements in the adjusted TFP series are largely temporary, whereas the permanent component of the series has an important predictable component. As a test of frictionless vs. sticky price models, the hours response to transitory technology shocks is more dispositive than is the response to permanent technology shocks. Whereas intertemporal substitution leads to an expansion in hours in response to a transitory technology shock in plausibly parameterized neoclassical models, the wealth effect can result in a decline in hours following a permanent technology shock provided the eventual rise in technology is sufficiently large relative to the impact effect.

I estimate medium-sized VARs featuring the adjusted TFP measure, the relative price of investment goods, and aggregate consumption, output, hours, inflation, and interest rates, imposing the cointegrating relationships implied by the neoclassical growth model. Following Fisher (2006), the investment specific technology shock is identified as the source of long run variation in the relative price of investment and the permanent neutral shock as the driver of adjusted TFP in the long run. The temporary neutral shock is identified as the innovation in adjusted TFP orthogonalized with respect to the permanent neutral shock.

The VAR results can be summarized as follows. First, there is an important transitory component to adjusted TFP. Both hours and output rise significantly in response to a positive realization of this shock, and it accounts for a non-trivial fraction of the variance of output and hours, particularly at higher frequencies. Second, the permanent neutral shock leads to a response of adjusted TFP that is highly autocorrelated in growth rates, so that much of the eventual technological improvement can be anticipated in advance. This shock is associated with an impact reduction in hours worked and a small positive response of output, which is in turn followed by significant growth. The investment specific shock leads to a significant expansion in hours worked. Both the permanent neutral and investment specific shocks are associated with significant disinflation. The three kinds of technology shocks combine to account for between 40 and 60 percent of the output variance at business cycle frequencies.

The paper then asks how well state of the art DSGE models can account for these conditional responses. To that end I construct a medium-scale DSGE model that builds off the seminal contributions of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). In addition to the three different technology shocks, the model features both price and wage rigidity as well as the usual real frictions – habit formation in consumption, investment adjustment costs, and variable capital utilization. I first consider a "standard" parameterization of the model, drawing parameter values from the existing literature. The model does a poor job at matching the empirical impulse responses, particularly with respect to the behavior of hours. Whereas in the data positive temporary neutral shocks raise hours but permanent neutral shocks lower hours, the exact opposite pattern obtains in the model. I then estimate the parameters of the model to minimize the distance between the model and data responses to all three kinds of technology shocks. The best-fitting parameter configuration differs in important ways from the "standard" one. In particular, investment adjustment costs disappear, there is no evidence of price and wage indexation, and there are comparatively low levels of nominal rigidity. The data also prefer a much higher Frisch labor supply elasticity than is commonly found in the literature.

The estimated level of investment adjust costs is so low because even moderate levels of these costs make it very unlikely that hours will rise following temporary neutral technological improvement. Francis and Ramey (2005) make this point forcefully, showing that investment adjustment costs and habit formation in consumption can make technological improvements "contractionary" without resorting to price-stickiness. The essential intuition is straightforward – households would like to smooth out a temporary productivity shock by increasing investment, but if it is very costly to undertake new investment, they will instead choose to consume more leisure. The model prefers relatively low levels of nominal rigidities and non-existent price and wage indexation because inflation responds quickly and without much inertia to technology shocks. Dupor, Han, and Tsai (2009) and Paciello (2011) make a similar point. The data prefer a high Frisch elasticity because hours respond rather strongly to technology shocks on impact. Most current DSGE models match hours volatility not with a high Frisch elasticity but rather with large and frequent preference shocks.

Several authors have made the point that there exists a tension between DSGE models designed to match responses to policy shocks (e.g. Christiano, Eichenbaum, and Evans, 2005) and those models designed to explain the responses to technology shocks (see especially Dupor, Han, and Tsai, 2009). The findings in the present paper generally support this conclusion. Whereas most of this literature focuses on the lack of an inertial response of inflation to technology shocks, the present paper shows that there is also a large tension resulting from the behavior of hours. Investment adjustment costs play a very important role in monetary DSGE models. Adjustment costs are necessary to break the connection between the real interest rate and the marginal product of capital – without adjustment costs, policy shocks which raise output also *raise* real interest rates, which is very much at odds with both evidence and intuition concerning the monetary transmission mechanism.[2] These adjustment costs also help to generate hump-shaped impulse responses and

[2] Models without capital, which form the basis of much of the New Keynesian literature as well as the literature

positive autocorrelation of output and its components in growth rates. But these adjustment costs exert such a strong effect on hours so as to make it essentially impossible to match the response of hours to technology shocks with these costs present in a significant amount. There does not appear to be a simple resolution of this tension. I speculate that models of informational rigidities seem promising in this regard.

The remainder of the paper is organized as follows. Section 2 describes the data used in the econometric applications. Section 3 presents the VAR evidence. Section 4 introduces a medium-scale DSGE model with a number of frictions. Section 5 estimates the model to match the VAR impulse responses and discusses some of the differences from a standard parameterization. The final section offers concluding thoughts.

2 Data

The variables used in the econometric analysis are a utilization-adjusted measure of total factor productivity (TFP), based on the adjustments in Basu, Fernald, and Kimball (2006); a measure of the relative price of investment to consumption goods; and hours worked, consumption, output, interest rates, and inflation. As the TFP measures is the least familiar, I begin with a discussion of it first.

The quarterly BFK (2006) adjusted TFP series presumes a constant returns to scale aggregate production function of the form:

$$y_t = a_t \left(u_t k_t\right)^{\alpha} \left(e_t n_t\right)^{1-\alpha} \tag{1}$$

k_t is capital input, u_t is utilization of capital, n_t is total labor hours, and e_t is labor effort. a_t is technology and α is capital's share. It is assumed that one can observe y_t, k_t, and n_t, but that a_t, u_t, and e_t are unobserved. A traditional Solow residual would be estimated as the residual of output less share-weighted observable inputs:

$$\Delta \ln TFP_t = \Delta \ln y_t - \alpha \Delta \ln k_t - (1-\alpha)\Delta \ln n_t \tag{2}$$

In many models u_t and e_t will vary substantially in response to non-technology shocks; as such the volatility and cyclicality of $\Delta \ln TFP_t$ may substantially overstate the true volatility and cyclicality of $\Delta \ln a_t$. BFK (2006) derive fairly general conditions under which variation in hours per worker, which is observed, can proxy for capital utilization and labor effort. Essentially the intuition is that a cost-minimizing firm would like to vary inputs along all margins simultaneously. Using industry level regressions, BFK aggregate up to a quarterly total utilization series:

$$\Delta \ln \widehat{U}_t = \alpha \Delta \ln u_t + (1-\alpha)\Delta \ln e_t \tag{3}$$

on optimal monetary policy, completely sidestep this issue since there is no condition relating the marginal product of capital to the real interest rate. Models without capital are often used with the justification that they are a good approximation to models with sufficient investment adjustment costs. The findings in this paper that investment adjustment costs are small suggest that this justification may be problematic.

Armed with this series, they can then construct an empirical counterpart of the model technology shifter as:

$$\Delta \ln \widehat{a}_t = \Delta \ln TFP_t - \Delta \ln \widehat{U}_t \tag{4}$$

I will refer to this series as "utilization-adjusted TFP" and the conventional Solow residual as "TFP". These corrections are important because the identification of transitory technology shocks in the VAR systems to be estimated below presumes that the empirical series measures true technology, and is not a conglomeration of true technology with unobserved input variation. For more detail on the construction of this series, please refer to Fernald (2009).

Figure 1 plots the cyclical components of both the adjusted and conventional TFP series, using an HP filter with smoothing parameter 1600. The data are quarterly from 1947q1 to 2009q4. The shaded gray regions in the figure are recessions as defined by the NBER Business Cycle Dating Committee. Two things from the figure visually stand out – first, the unadjusted TFP series appears more volatile, and, second, the two series are not that strongly correlated.

Table 1 shows some basic statistics for these alternative TFP series. It also uses data on output and hours. I define output as real output in the non-farm business sector divided by the population aged sixteen and over. Hours are total hours in the non-farm business sector divided by the population. Both series are available from the BLS. All series in this table are HP filtered with smoothing parameter 1600. So as to be consistent with the analysis that follows, these series are expressed in terms of consumption goods by multiplying by the appropriate price deflators. We see that the unadjusted TFP series is slightly more volatile than the adjusted series, though not remarkably so. Whereas the unadjusted TFP series is strongly procyclical (correlation with output of 0.8), the adjusted TFP series is actually slightly countercyclical. Further, the unadjusted series is positively correlated with hours (correlation coefficient 0.3), whereas the adjusted series is strongly negatively correlated with hours (correlation coefficient -0.4). Finally, the two TFP series are only weakly correlated with one another (correlation 0.3). These statistics indicate that these series are indeed very different at cyclical frequencies.

I measure the relative price of investment as the ratio of the chain-weighted price index of private fixed investment and durable goods to the chain-weighted price index of non-durable and services consumption, all from the NIPA accounts and available from the BEA. Figure 2 plots the log inverse of this series across time, with the shaded gray regions NBER dated recessions. Two things from the figure are apparent – first, the strong trend growth, particularly in the second half of the sample, and, two, the series visibly declines during most recessions, suggesting that investment specific technology falls during recessions. Table 2 shows some basic statistics, based on the HP filtered series. We see that the cyclical component of the (inverse) investment price is weakly procyclical and strongly positively correlated with hours worked. It is uncorrelated with unadjusted TFP and weakly negatively correlated with adjusted TFP.

There are additional ways in which researchers have measured the relative price of investment; for an extended discussion, see, for example, Fisher (2006). Most of these papers use price deflators for equipment with bias corrections due to Gordon (1989). A particular drawback of this series is that it is an annual and requires interpolation to convert to a quarterly frequency. As the results regarding the role of investment specific technical change reported below conform closely to those in Fisher (2006), I focus on the quarterly series so described above.

In addition to the series already introduced, the econometric analysis also makes use of data on consumption, interest rates, and inflation. I measure consumption as the sum of non-durable and services consumption, less imports of non-durables and services. These data are readily available from the NIPA accounts. The measure of interest rates is the Federal Funds Rate, converted from the underlying monthly frequency to a quarterly frequency by taking within-month averages. The measure of inflation is the percentage change of the GDP deflator.

3 VAR Evidence

This section presents the main VAR evidence. Section 3.1 estimates bivariate VARs with measures of TFP and hours and shows that unadjusted TFP innovations lead to increases in hours worked whereas adjusted TFP innovations are associated with decreases in hours worked. Section 3.2 estimates a larger model and uses restrictions implied by the neoclassical growth model to identify both permanent and temporary neutral technology shocks as well as investment specific technology shocks. It is shown that hours rise in response to positive temporary neutral and investment specific technology shocks, whereas hours decline in response to the permanent neutral shock.

3.1 Two Variable VARs with TFP and Hours

This section replicates the main empirical results of BFK (2006). Consider a simple bivariate model with either adjusted or unadjusted TFP and the level of hours worked per capita. The system can be written (abstracting from the constant terms):

$$\Delta \ln x_t = \sum_{j=1}^{p} \gamma_{1,j} \Delta \ln x_{t-j} + \sum_{j=1}^{p} \gamma_{2,j} \ln n_{t-j} + \varepsilon_{1,t} \tag{5}$$

$$\ln n_t = \sum_{j=1}^{p} \gamma_{3,j} \Delta \ln x_{t-j} + \sum_{j=1}^{p} \gamma_{4,j} \ln n_{t-j} + a_0 \Delta \ln x_t + \varepsilon_{2,t} \tag{6}$$

Here x_t is either TFP_t or adjusted TFP, a_t. p is the lag length. $\varepsilon_{1,t}$ can be interpreted as a technology shock, while $\varepsilon_{2,t}$ is a non-technology shock. The fact that the contemporaneous value of $\Delta \ln x_t$ shows up in (6) reflects the underlying assumption that technology is exogenous; given this a_0 can be estimated consistently via least squares.

Figure 3 shows impulse responses of the technology series and hours to a technology shock. The VARs use $p = 4$ lags. The left panel shows responses using unadjusted TFP as the measure of technology, while the right panel shows the responses using the BFK adjusted TFP series. The shaded gray regions are +/- one standard error confidence bands, using the bias-corrected bootstrap after bootstrap of Kilian (1998). When technology is measured using the crude TFP measure, we see that hours significantly increase following a positive technology shock. There is also some reversion of measured TFP to the technology shock before leveling off, perhaps suggesting that part of what $\varepsilon_{1,t}$ is picking up are "demand" shocks. The responses in the right panel using the BFK adjusted TFP series are very different. Here we see that hours worked actually decline following a positive technology shock; there is also no evidence of reversion of adjusted TFP to its own innovation – its response is consistent with an exact random walk. This pattern of responses leads BFK (2006) to conclude that technology improvements are "contractionary", in the sense of exogenous improvements in productivity leading to reductions in hours worked. They argue that this finding lends support to sticky price models of the business cycle over flexible price models.

3.2 Larger Dimensional Systems and Multiple Technology Shocks

This section estimates a larger dimensional multivariate system and identifies three different kinds of technology shocks – permanent and transitory neutral shocks and investment specific shocks. A larger system is necessary in order to identify multiple kinds of technology shocks. The restrictions used to empirically identify the three technology shocks of interest are implied by a simple neoclassical growth model. Most modern DSGE models, such as the one presented in Section 4, deviate substantially from the neoclassical benchmark in the short run, but almost all behave according to the predictions of the basic growth model in the long run. Since the VAR restrictions are based on the long run properties of the data, it suffices to discuss the growth model here.

The model can be written as a planner's problem. I abstract from population growth both here and for the remainder of the paper; one can think of all variables as being per capita. The objective is to pick consumption, investment, future capital, and hours of work to maximize the present discounted value of flow utility, subject to an accounting identity, the law of motion for capital, and the exogenous stochastic processes:

$$\max_{c_t, I_t, k_{t+1}, n_t} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln c_t - \theta \frac{n_t^{1+\eta}}{1+\eta} \right)$$

s.t.

$$a_t^p a_t^s k_t^\alpha n_t^{1-\alpha} = c_t + I_t \tag{7}$$

$$k_{t+1} = \chi_t I_t + (1-\delta) k_t \tag{8}$$

$$\Delta \ln a_t^p = (1-\rho_{a^p}) g_a + \rho_{a^p} \Delta \ln a_{t-1}^p + \varepsilon_{a^p,t} \tag{9}$$

$$\Delta \ln \chi_t^p = (1-\rho_\chi) g_\chi + \rho_\chi \Delta \ln \chi_{t-1} + \varepsilon_{\chi,t} \tag{10}$$

$$\ln a_t^s = \rho_{a^s} \ln a_{t-1}^s + \varepsilon_{a^s,t} \tag{11}$$

(7) is an accounting identity and (8) is the law of motion for capital. a_t^p is a permanent neutral technology shifter while a_t^s is a stationary neutral technology shifter. The product of these two components, $a_t^p a_t^s$, corresponds to the a_t that is in principle measured by the adjusted TFP series. The permanent component of technology follows an AR(1) in growth rates as given by (9), with g_a

the trend growth rate. $\rho_{a^s} = 0$ would correspond to the familiar random walk case. The stationary component of neutral technology obeys a mean zero AR(1) in the log. χ_t is the investment specific technology shifter; the bigger is χ the more efficient the economy is at transforming investment into capital goods. It corresponds to the inverse relative price of investment to consumption goods. It is assumed to follow an AR(1) in the growth rate, given by (9).[3] g_χ is the trend growth rate of investment specific technology.

The model as written is consistent with balanced growth. It is straightforward to verify that, along the balanced growth path, hours will be stationary but all the other variables will inherit trend growth from the trends in both a_t^p and χ_t. In particular, the following transformed variables will be stationary:

$$\widehat{y}_t = \frac{y_t}{a_t^{p\frac{1}{1-\alpha}} \chi_t^{\frac{\alpha}{1-\alpha}}}, \quad \widehat{c}_t = \frac{c_t}{a_t^{p\frac{1}{1-\alpha}} \chi_t^{\frac{\alpha}{1-\alpha}}}, \quad \widehat{I}_t = \frac{I_t}{a_t^{p\frac{1}{1-\alpha}} \chi_t^{\frac{\alpha}{1-\alpha}}}, \quad \widehat{k}_{t+1} = \frac{k_{t+1}}{a_t^{p\frac{1}{1-\alpha}} \chi_t^{\frac{1}{1-\alpha}}}$$

These transformations imply cointegrating restrictions on the data. In particular, $\ln y_t - \frac{1}{1-\alpha} \ln a_t^p - \frac{\alpha}{1-\alpha} \ln \chi_t$ should be stationary, as should $\ln c_t - \ln y_t$ and $\ln I_t - \ln y_t$.

One can measure the product $a_t^p a_t^s$ in the data with a_t, the BFK adjusted TFP measure. One can measure χ_t as the inverse relative price of investment. Given observations on these data, the specification written above implies natural restrictions which can be used to identify these shocks in a VAR setting. The investment specific shock (i) should not affect adjusted TFP and (ii) should have a permanent effect on the relative price of investment. The permanent neutral shock should permanently affect adjusted TFP, while the transitory neutral shock should not. In principle, the model as written implies another restriction – that neutral TFP shocks not affect the relative price of investment. This would not hold (in the short run) in a more complicated model in which there is curvature in the transformation of investment to consumption goods, but would continue to hold in the long run – see Fisher (2009) for a simple example. It would also not hold, even in the long run, if the two kinds of shocks happen to be correlated, which appears to be a feature of the data.

I estimate a VAR(p) with the following variables: the growth rate of the BFK adjusted TFP measure, the growth rate of the (inverse) relative price of investment to consumption goods, hours per capita, the cointegrating term relating output to the levels of neutral and investment specific technology, the log ratio between consumption and output, the Federal Funds rate, and inflation as measured by the GDP deflator.[4] Formally:

$$Y_t = A(L)Y_{t-1} + \nu_t, \qquad \nu_t = B\varepsilon_t \tag{12}$$

[3] In principle one could also entertain stationary investment specific shocks. This would be difficult to identify empirically due to the potential short run endogeneity of the relative price of investment in a more general setting. I follow most of the rest of the literature in assuming that the investment specific technology shifter has a stochastic trend.

[4] Technically this cointegrating relationship should apply to $\ln a_t^p$, but since $\ln a_t^s$ is normalized to be mean 1, $\ln y_t - \frac{1}{1-\alpha} \ln a_t - \frac{\alpha}{1-\alpha} \ln \chi_t$ will also be stationary, where $\ln a_t = \ln a_t^s + \ln a_t^p$.

$$Y_t = \begin{pmatrix} \Delta \ln a_t \\ \Delta \ln \chi_t \\ \ln n_t \\ \ln y_t - \frac{1}{1-\alpha} \ln a_t - \frac{\alpha}{1-\alpha} \ln \chi_t \\ \ln c_t - \ln y_t \\ i_t \\ \pi_t \end{pmatrix}$$

$A(L)$ is a lag polynomial of order p and ν_t is a vector of reduced form innovations. I assume that there is a linear mapping between structural shocks, ε_t, which are defined as being uncorrelated with one another, and the reduced form innovations, given by the square matrix B. The variables of the empirical model are all expressed in terms of consumption goods using the chain-weighted deflator for non-durable and services consumption. Given the forward-looking nature of the consumption to output ratio, it is important to include this variable to help ensure invertibility. The inclusion of the nominal interest rate and inflation help to control for monetary factors which may be important in identification.

The data included in the analysis run from 1955q1 to 2009q4.[5] Construction of the model-based cointegrating relationship requires a value of α. I use $\alpha = 0.31$, which is the average capital share provided by Fernald (2009). Figure 4 plots three of the series used in the VAR: (i) the model-based cointegrating relationship between output and the two kinds of technology, (ii) the consumption-output ratio, and (iii) hours per capita. The shaded gray areas are NBER defined recessions. All three series appear roughly stationary, consistent with the implications of balanced growth, though there is some evidence of a slight upward trend in the consumption-output ratio beginning in the early 1980s. In the data the average growth rate of adjusted TFP is 0.18 percent per quarter and the average growth rate of the inverse relative price of investment is 0.28 percent. With a value of $\alpha = 0.31$, this would imply that the average growth rate of output ought to be about 0.4 percent. The actual average growth rate of output per capita in the sample is 0.39 percent, so this is very close. The model-based cointegrating relationship is quite procyclical; this means that output typically falls by more than a_t and χ_t in a recession. The consumption-output ratio, in contrast, is countercyclical. This follows from the fact that consumption falls by less than output in a typical recession; see also Cochrane (1994). Hours per capita are naturally procyclical.

The exact restrictions used to identify the structural shocks follow from the implications of the growth model discussed above and are as follows. First, the adjusted TFP measure reacts within period only to (i) the temporary neutral technology shock and (ii) the permanent neutral technology shock. This imposes zero restrictions on B and suffices to identify these two shocks from the remainder of the shocks in the system. The neutral technology shocks are differentiated from one another with the long run restriction that the temporary technology shock have no effect on the level of adjusted TFP in the long run, and imposes another restriction on B which can be

[5] Most of these data go back to 1947q1. The sample for the empirical model begins in 1955 to omit (i) the immediate aftermath of World War II and (ii) the Korean War. This is a common sample in the business cycle literature. The results are not sensitive to beginning the sample in 1947.

implemented via the methods proposed in Blanchard and Quah (1989) and Shapiro and Watson (1988). The investment specific shock is identified with the long run restriction that it is the only of the remaining shocks in the system that can affect the level of the relative price of investment in the long run. These restrictions together uniquely identify the three columns of B corresponding with these shocks. The remaining shocks in the system are left unidentified.[6]

Figure 5 shows impulse responses to the temporary neutral technology shock. Adjusted TFP jumps up by about 0.6 percent and then reverts back to zero. Crucially, hours worked rise significantly on impact, before rising even further. The hump-shaped response has hours returning back to the starting point after roughly 30 quarters. Output rises significantly on impact and also follows a hump shape, with its dynamic response similar in shape to the response of hours. Consumption jumps mildly on impact and eventually reverts back. Consistent with the intuition from the simple permanent income hypothesis, the consumption response to the temporary neutral technology shock is considerably smaller than the output response, suggesting an important response of investment. Neither the nominal interest rate nor the inflation rate react significantly on impact to the temporary technology shock, with responses that are mildly positive a number of quarters after the shock. The relative price of investment does not significantly react at any horizon.

The impulse responses to the permanent neutral technology shock are shown in Figure 6. Adjusted TFP jumps up on impact but then is expected to continue to grow for a number of quarters. In fact, the long horizon response of adjusted TFP to the permanent shock is about three times as large as the impact effect (1.5 percent versus 0.5 percent). This suggests that a substantial fraction of the low frequency component of adjusted TFP is anticipated, a finding which comports with many of the results from the "news" literature (see, e.g., Beaudry and Portier, 2006, or Barsky and Sims, 2011, for a discussion).[7] One observes that hours worked decline significantly on impact in response to the favorable permanent neutral technology shock. Hours continue to fall for a few quarters before rising and going back to the pre-shock level. Output essentially does not react on impact, which is consistent with technology improving but hours declining. After impact output grows for a number of quarters before reaching a new, permanently higher level. Consumption jumps up on impact, though it substantially undershoots its long run level. Both the nominal interest rate and inflation fall significantly in response to the permanent neutral technology shock.

The last panel of Figure 6 shows the response of the inverse relative price of investment to the permanent neutral shock. We see that the shock that permanently raises adjusted TFP is associated with a significant (and permanent) reduction in the inverse relative price of investment (equivalently a reduction in the efficiency of turning investment goods into capital goods). Nothing in the VAR identification prevents this from happening – in the benchmark identification the permanent neutral

[6]The neoclassical growth model provides further overidentifying restrictions that are not explicitly imposed in the VAR identification. These are that the measure of adjusted TFP not react to any other shock in the system at any horizon (as opposed to just on impact) and that the relative price of investment not react to the permanent neutral shock in the long run, provided $\varepsilon_{a^p,t}$ and $\varepsilon_{\chi,t}$ are assumed to be uncorrelated. More will be made of this point below. A further restriction that the relative price of investment not react to the neutral technology shocks in the short run only holds in the special case in which there is no curvature in the investment-consumption goods frontier.

[7]I do not separately consider news shocks in the identification. To the extent to which news shocks are present and permanently impact adjusted TFP, they will be reflected in the permanent neutral shock. The responses of output, hours, and TFP are consistent with an important news component.

and investment specific shocks are not identified via a long run restriction but rather a short run restriction that the investment specific shock not affect adjusted TFP contemporaneously. Nothing in the benchmark theoretical model rules this out, either – it simply means that that neutral and investment specific technology shocks (i.e. $\varepsilon_{a^p,t}$ and ε_{χ_t}) are correlated, evidently negatively so. The VAR identification attributes all of the common component of these two shocks to the permanent neutral shock. This means that the responses in Figure 6 are driven both by the increase in neutral technology and a reduction in investment specific technology.

A natural next step is therefore to isolate just the role of the neutral technological improvement. This can be done by dropping the orthogonality restriction between investment specific and neutral shocks and replacing it with another long run restriction that the neutral shock have no long run impact on the relative price of investment. This is consistent with a world in which neutral and investment specific shocks are indeed correlated; it simply attributes the common component to the investment specific shock, thereby isolating the role of neutral technological progress. The responses under this alternative orthogonalization are shown in Figure 7. These are fairly similar to those shown in Figure 6. In particular, adjusted TFP significantly undershoots its permanently higher value, again implying that a substantial fraction of the ultimate impact on the level of TFP is anticipated. Hours again decline. The main difference relative to Figure 6 is that the decline in hours is much less persistent; here we see that the hours response turns mildly positive after a number of quarters, whereas in Figure 6 the hours response is much more persistently negative.

Figure 8 shows the impulse responses to an investment-specific technology shock. As with the permanent neutral shock, we see that there is significant growth in the (inverse) relative price of investment after a relatively small initial jump. Hours essentially do not react on impact but then grow robustly for a number of quarters, with a peak response greater than 0.6 percent after about six quarters. Output increases slightly on impact and then grows for a number of quarters, over-shooting its long run value. Both the substantial increase in hours and the over-shooting of output are consistent with the responses estimated in Fisher (2006). The interest rate essentially does not react to the investment specific shock, while inflation falls. Adjusted TFP does not significantly react to the shock at any horizon.

Table 3 shows the forecast error variance decomposition for the benchmark orthogonalization. The first set of rows shows the fraction of the forecast error variance of the VAR variables attributable to the temporary neutral shock at various horizons. The temporary neutral shock explains the majority of the innovation in adjusted TFP and about one quarter of the business cycle variances of both output and hours. It appears inconsequential for movements in consumption, interest rates, and inflation. The next set of rows shows the variance decomposition due to the permanent neutral shock. This shock explains the bulk of the variances of both adjusted TFP and the relative price of investment at lower frequencies. It accounts for about 25 percent of the variance of hours at business cycle frequencies and is also an important driver of inflation and interest rates. The permanent neutral shock explains a large share of the output and consumption variances at longer horizons. The investment specific shock accounts for around 20 percent of the business cycle variance in output, hours, and consumption and also has important implications for inflation. The final set of rows shows the total fraction of the forecast error variance in the variables due to the three technology shocks combined. These shocks essentially explain all of the variance in adjusted TFP at all horizons, which is consistent with the idea that this series represent a measure of actual technology. Though the shocks explain virtually all of the lower frequency movements in the relative price of investment, there are some higher frequencies movements in this series for which the technology shocks fail to account. The technology shocks combine to account for between 40 and 60 percent of the variance of output, consumption, and hours at business cycle frequencies. Though this does leave open an important channel for demand shocks, one must conclude, as in the title of the working paper version of Fisher (2006), that "technology shocks matter".

Based on the balanced growth path implication of the stationarity of hours, the systems estimated so far feature hours worked per capita in log levels. There is a large debate in the literature over how hours should enter the VAR system. In bivariate productivity-hours systems, the effect on hours of technology shocks identified using long run restrictions depends crucially on whether hours enter in levels, differences, or deviations from a trend. The typical finding in the literature is that hours rise in response to a positive technology shock when they enter such systems in levels, whereas hours decline when they enter the system in first differences or as deviations from trend.

It is therefore natural to investigate the robustness of the above results to how hours enter the system. Figure 9 plots out the impulse responses of hours to the three kinds of technology shocks for three different cases: hours enter in levels (solid line), hours enter in first differences (dashed line), or hours enter as deviations from a low frequency HP filter (dotted line).[8] The responses are broadly similar across the different specifications. In particular, hours worked rise in response to the temporary neutral shock and fall in response to the permanent neutral shock regardless of how they enter the system. The most substantial difference is in the behavior of hours following an investment specific shock. While the responses under the levels and first difference specifications are virtually the same, when HP filtered hours essentially do not react at all to the investment specific shock.

I conducted a number of additional robustness checks. These are omitted in the interest of space conservation. The inclusion of interest rates, inflation, and the cointegrating relationships are not necessary to identify the shocks of interest. The basic results are virtually the same without all or some of these additional variables. The lag length in the VAR system appears largely irrelevant for the qualitative conclusions. Estimating the system on sub-samples of the data (e.g. pre and post "Great Moderation") leaves unaffected the main conclusions regarding the general pattern of movements in response to the technology shocks, though there are some important differences, particularly with respect to the behavior of inflation (see, e.g., Paciello, 2011).

[8] Francis and Ramey (2009) emphasize that there are important low frequency movements in hours per worker, and suggest correcting for this by detrending with an HP filter with a larger smoothing parameter than is typically used for quarterly data (16,000 vs. 1600). I follow them here.

4 The Model

The model considered here is a by now relatively standard medium scale model. It features wage and price stickiness and a number of real frictions. It builds off of the canonical models of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007), with a few modifications.

The model features five types of actors of interest: households, final goods firms, labor-packing firms, intermediate goods firms, and the government. The text presents the optimization problems of each of these types of agents in the model, gives the exogenous stochastic processes, and describes the equilibrium. The Appendix gives the first order conditions and discusses the solution methodology.

4.1 Final Goods Firm

There is a representative final goods firm. It is competitive and bundles intermediate goods into a final good using a CES technology. There are a continuum of intermediate goods producers of measure 1, indexed by $j \in (0,1)$. The technology mapping intermediate inputs into the final good is:

$$y_t = \left(\int_0^1 y_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}} \tag{13}$$

It is assumed that $\varepsilon > 1$. Profit maximization by the representative final goods firm yields a downward sloping demand curve for each intermediate good and an aggregate price index, where $p_{j,t}$ is the price of variety j:

$$y_{j,t} = \left(\frac{p_{j,t}}{p_t} \right)^{-\varepsilon} y_t \tag{14}$$

$$p_t = \left(\int_0^1 p_{j,t}^{1-\varepsilon} dj \right)^{\frac{1}{1-\varepsilon}} \tag{15}$$

4.2 Labor-Packing Firm

In the model households are monopoly suppliers of labor. There exists a representative labor-packing firm that is competitive and bundles household labor supply into a labor input which is then rented to intermediate goods firms. There are a continuum of households of measure 1, indexed by $l \in (0,1)$. The technology mapping household labor supply into the packed labor input is given by:

$$n_t = \left(\int_0^1 n_{l,t}^{\frac{\eta-1}{\eta}} dl \right)^{\frac{\eta}{\eta-1}} \tag{16}$$

Profit maximization by the labor-packing firm gives rise to a downward sloping demand curve for each type of labor and an aggregate real wage index, where $w_{l,t}$ is the real wage of labor of household l:

$$n_{l,t} = \left(\frac{w_{l,t}}{w_t}\right)^{-\eta} n_t \tag{17}$$

$$w_t = \left(\int_0^1 w_{l,t}^{1-\eta} dl\right)^{\frac{1}{1-\eta}} \tag{18}$$

4.3 Intermediate Goods Firms

Intermediate goods firms produce output using capital services, labor, and aggregate technology. Technology is common to all firms, and is composed of both a permanent and stationary component.

$$y_{j,t} = a_t \widetilde{k}_{j,t}^{\alpha} n_{j,t}^{1-\alpha} \tag{19}$$

$$a_t = a_t^p a_t^s \tag{20}$$

$\widetilde{k}_{j,t}$ is the amount of capital services (the product of utilization and the physical capital stock). Intermediate goods firms rent capital services from households and labor from the representative labor-packing firm each period.[9] a_t is total technology, the product of both a permanent, a_t^p, and stationary, a_t^s, component.

Given its monopoly power, intermediate producers can choose their prices. They are subject to pricing frictions a la Calvo (1983), facing a fixed probability, $1 - \phi_p$, of being able to adjust their price in any period. This probability is independent of when the firm last updated its price. With probability ϕ_p a firm must charge the price it had in the previous period plus some adjustment for indexation to aggregate inflation. Regardless of whether the firm can adjust price so as to maximize profits, it will always find it optimal to choose inputs to minimize cost, given a price. As such, we can break the problem down into two parts. Let w_t and R_t be real factor prices for labor and capital services, respectively. These are common across intermediate goods firms. The firm's objective is to pick labor and capital services to minimize nominal costs, subject to the restriction of producing enough to meet demand:

$$\min_{\widetilde{k}_{j,t}, n_{j,t}} \quad w_t p_t n_t + R_t p_t \widetilde{k}_{j,t}$$

$$\text{s.t.}$$

$$a_t \widetilde{k}_{j,t}^{\alpha} n_{j,t}^{1-\alpha} \geq \left(\frac{p_{j,t}}{p_t}\right)^{-\varepsilon} y_t$$

As part of the first order conditions of the cost minimization problem one can construct a variable real marginal cost, mc_t, which is equal to the multiplier on the constraint divided by the aggregate price level. Real marginal cost depends only on factor prices, and so is common across

[9] Therefore, households, not firms, choose capital utilization. This is an unimportant detail; the problem could be modified so that firms choose utilization and the solution would be the same.

firms (hence no j subscript). It is straightforward to show that within period real profits can then be written:

$$\Pi_{j,t} = \frac{p_{j,t}}{p_t} y_{j,t} - mc_t y_{j,t} \tag{21}$$

Now consider the problem of a firm given the opportunity to update its price in period t. When setting its price, it must take into account that the probability that it will not have been able to re-optimize its price by period $s \geq 1$ is ϕ_p^s. As such, the pricing problem is dynamic. Non re-optimizing firms are able to partially index their period $t+s$ price to lagged inflation. Hence, with probability ϕ_p^s a firm that re-optimizes at time t will have price at $t+s$: $p_{j,t+s} = \prod_{m=1}^{s}(1+\pi_{t+m-1})^{\zeta_p} p_{j,t}$. π_t is aggregate inflation and $\zeta_p \in (0,1)$ is an indexation parameter, with the boundaries corresponding to no indexation and full indexation, respectively. Let β be the subjective discount factor of a representative household and λ_{t+s} be the expected marginal utility of an extra dollar of income at time $t+s$. The price optimization problem for an updating firm is:

$$\max_{p_{j,t}} \quad E_t \sum_{s=0}^{\infty} (\phi_p \beta)^s \lambda_{t+s} \left(\left(\frac{\prod_{m=1}^{s}(1+\pi_{t+m-1})^{\zeta_p} p_{j,t}}{p_{t+s}} \right)^{1-\varepsilon} y_{t+s} - mc_{t+s} \left(\frac{\prod_{m=1}^{s}(1+\pi_{t+m-1})^{\zeta_p} p_{j,t}}{p_{t+s}} \right)^{-\varepsilon} y_{t+s} \right)$$

The solution is an optimal reset price, $p_t^{\#}$, that will be common across all updating firms. This follows from the fact that marginal cost is common across firms, due to common factor markets.

4.4 Households

There are a continuum of households indexed by $l \in (0,1)$. Households choose consumption, how much capital to accumulate, how much to save in riskless government bonds, how intensively to utilize their existing capital, and how much to work. Given the downward sloping demand for labor from above, they can also choose their wage. Households are not freely able to adjust their nominal wage each period, however, with staggered contracts due to Calvo (1983). As is standard in the literature, following the arguments set forth in Erceg, Henderson, and Levin (2000), I assume that there exist state contingent securities so as to eliminate idiosyncratic wage risk. This implies that households will be heterogenous with respect to wages and labor supply, but homogeneous along all other dimensions. So as to economize on notation, I will impose these features in writing down the household problem and will omit the state contingent claims from the budget constraint. I abstract from the money holding decision, but could include real balances as an argument in the utility function without complication.

Given these notational assumptions, preferences are given by the following separable utility function:

$$E_0 \sum_{t=0}^{\infty} \beta^t \psi_t \left(\ln(c_t - \gamma c_{t-1}) - \theta_t \frac{n_{l,t}^{1+\xi}}{1+\xi} \right)$$

ψ_t is stochastic intertemporal preference shock, while θ_t is a stochastic intratemporal preference shock. γ measures the degree of habit persistence. ξ is the inverse of the Frisch labor supply elasticity. Given the separability between consumption and labor, it is useful to break the problem into two parts – the first concerning the choices of consumption, capital, bonds, utilization, and investment and the other wages and hours. The first part of the problem is given by:

$$\max_{c_t, I_t, k_{t+1}, u_t, B_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \beta^t \psi_t \left(\ln(c_t - \gamma c_{t-1}) \right)$$

$$\text{s.t.}$$

$$c_t + I_t + \frac{B_{t+1} - B_t}{p_t} \leq w_{l,t} n_{l,t} + R_t u_t k_t - \left(\Psi_0 (u_t - 1) + \frac{\Psi_1}{2} (u_t - 1)^2 \right) \frac{k_t}{\chi_t} + i_{t-1} \frac{B_t}{p_t} + \frac{\text{Profit}_t}{p_t} - \frac{T_t}{p_t}$$

$$k_{t+1} = \chi_t \left(1 - \frac{\tau}{2} \left(\frac{I_t}{I_{t-1}} - \Lambda_I \right)^2 \right) I_t + (1 - \delta) k_t$$

Here χ_t is an investment-specific technology shock, and Λ_I is the balanced growth path gross growth rate of investment. It is straightforward to verify that χ_t is equal to the inverse relative price of investment to consumption. τ is a parameter governing adjustment costs to investment. $\left(\Psi_0 (u_t - 1) + \frac{\Psi_1}{2} (u_t - 1)^2 \right)$ is the resource cost of capital utilization. It is assumed to be proportional to the capital stock and is measured in terms of consumption goods by dividing by χ_t. B_t is nominal holdings of one period government bonds, and i_t is the safe interest rate on these assets. T_t denotes nominal lump sum taxes/transfers from the government, and Profit_t denotes distributed nominal profits from intermediate goods firms. u_t is utilization, with $u_t k_t$ denoting capital services. Only the wage and employment have l subscripts, as per the discussion above.

Next consider the wage-setting problem of the household. Households get to change their nominal wage in each period with probability $1 - \phi_w$. Between price changes, I allow for partial indexation to aggregate inflation, given by the parameter $\zeta_w \in (0, 1)$. With probability ϕ_w^s, then, a household that updates its nominal wage at time t will have real wage at time $t + s$ equal to: $w_{l,t+s} = \prod_{m=1}^{s} (1 + \pi_{t+m-1})^{\zeta_w} (1 + \pi_{t+m})^{-1} w_{l,t}$. The problem of an updating household at time t can then be expressed in the following dynamic form:

$$\max_{n_{l,t+s}, w_{l,t}} \quad E_t \sum_{s=0}^{\infty} (\beta \phi_w)^s \left(-\psi_{t+s} \theta_{t+s} \frac{n_{l,t+s}^{1+\xi}}{1 + \xi} + \lambda_{t+s} n_{l,t+s} \prod_{m=1}^{s} \frac{(1 + \pi_{t+m-1})^{\zeta_w}}{1 + \pi_{t+s}} w_{l,t} \right)$$

$$\text{s.t.}$$

$$n_{l,t+s} = \left(\frac{\prod_{m=1}^{s} \frac{(1+\pi_{t+m-1})^{\zeta_w}}{1+\pi_{t+s}} w_{l,t}}{w_{t+s}} \right)^{-\eta} n_{t+s}$$

λ_{t+s} is the multiplier on the household's budget constraint and is equal to the marginal utility of an additional dollar of income. Because households are homogeneous with respect to consumption and utility is separable this does not vary across l. The problem as written reflects the probability that a household will be stuck with its wage chosen at time t several periods out into the future, plus the adjustment for indexation. The problem is subject to the constraint that the household supplies as much labor each period as is demanded from the labor-packing firm. The solution is an optimal reset wage, $w_t^{\#}$ that will be common across updating households. This reset wage, along with (17) and (18), suffices to characterize the behavior of the aggregate wage and employment. Because of the Calvo assumption, it is not necessary to keep track of individual labor supply and wages.

4.5 The Government

The government is composed of both a fiscal and a monetary authority. The monetary authority sets interest rates according to a modified Taylor (1993) type interest rate rule which allows for partial adjustment:

$$i_t = (1-\rho_i)i^* + \rho_i i_{t-1} + (1-\rho_i)\varphi_\pi (\pi_t - \pi^*) + (1-\rho_i)\varphi_y \left(\frac{y_t}{y_{t-1}} - \Delta_y \right) + \varepsilon_{i,t} \tag{22}$$

π^* is the inflation target of the Fed and is an exogenous parameter of the model. i^* is the steady state nominal interest rate; this is not a free parameter, as it is determined by π^*, household preferences, and trend growth rates. $\varphi_\pi > 1$ is the response to deviations of inflation from target and $\varphi_y \geq 0$ is the response to deviations of output growth from trend, where Δ_y is the balanced growth path gross output growth rate. $\varepsilon_{i,t}$ is a monetary policy shock. As noted early, I do not explicitly model money. If I did, the central bank would print the amount of money required to meet household demand at the interest rate given above, and would remit all seignorage revenue to the fiscal authority.

The fiscal authority consumes a stochastic share, ω_t, of total output each period:

$$g_t = \omega_t y_t \tag{23}$$

ω_t is the stochastic share of private output that the government consumes, and follows a stochastic process to be discussed below. Modeling the spending process in this way (in terms of a share of output, as opposed to just a level) facilitates working with a model that grows over time.

The fiscal authority's budget constraint in nominal terms is:

$$p_t g_t - D_t \leq T_t + (1 + i_{t-1})D_{t-1} \tag{24}$$

Given a time path for g_t, nominal debt, D_t, and nominal lump sum taxes, T_t, will adjust so that this constraint holds with equality, though the mix between debt and taxes is indeterminate.

4.6 Exogenous Processes

The exogenous state variables of the model are the two neutral technology shifters, a_t^p and a_t^s; the investment specific shock, χ_t; the two preference shocks, θ_t and ψ_t; and the government spending share, ω_t. These series obey the following autoregressive processes:

$$\ln a_t^s = \rho_{a^s} \ln a_{t-1}^s + \varepsilon_{a^s,t} \tag{25}$$

$$\Delta \ln a_t^p = (1-\rho_{a^p})g_a + \rho_{a^p}\Delta \ln a_{t-1}^p + \varepsilon_{a^p,t} \tag{26}$$

$$\Delta \ln \chi_t = (1-\rho_\chi)g_\chi + \rho_\chi \Delta \ln \chi_{t-1} + \varepsilon_{\chi,t} \tag{27}$$

$$\ln \theta_t = (1-\rho_\theta)\theta^* + \rho_\theta \ln \theta_{t-1} + \varepsilon_{\theta,t} \tag{28}$$

$$\ln \psi_t = \rho_\psi \ln \psi_{t-1} + \varepsilon_{\psi,t} \tag{29}$$

$$\omega_t = (1-\rho_g)\omega^* + \rho_g \omega_{t-1} + \varepsilon_{g,t} \tag{30}$$

The stationary component of neutral technology, a_t^s, and the intertemporal preference shifter, ψ_t, are both normalized to be mean 1 (0 in logs). g_a is the trend growth rate of neutral technology and g_χ is trend growth in investment specific technology. θ^* is the steady state value of the intratemporal preference shock on labor and ω^* is the steady state government consumption share of private output. The model has seven structural shocks, given by the vector: $\varepsilon_t = (\varepsilon_{a^s,t} \;\; \varepsilon_{a^p,t} \;\; \varepsilon_{\chi,t} \;\; \varepsilon_{\theta,t} \;\; \varepsilon_{\psi,t} \;\; \varepsilon_{g,t} \;\; \varepsilon_{i,t} \;\;)'$. These are drawn from mean zero normal distributions with constant variances.

4.7 Aggregation, Equilibrium, and Balanced Growth

In equilibrium total demand for labor from intermediate goods firms must equal the amount supplied by the labor-packing firm and total demand for capital services must equal total supply from households:

$$n_t = \int_0^1 n_{j,t} dj$$

$$u_t k_t = \int_0^1 \widetilde{k}_{j,t} dj$$

Combining these market-clearing conditions with the demand curve for intermediate goods, and using the fact that all firms will hire capital and labor in the same ratio (since they face the same factor prices and marginal cost), gives rise to the aggregate production function:

$$y_t = \frac{a_t \left(u_t k_t\right)^\alpha n_t^{1-\alpha}}{v_t} \tag{31}$$

$$v_t = \int_0^1 \left(\frac{p_{j,t}}{p_t}\right)^{-\varepsilon} \tag{32}$$

Integrating over households' budget constraints, integrating over firm profits, and combining with the government budgets constraint gives rise to the aggregate accounting identity:

$$y_t = c_t + I_t + g_t + \left(\Psi_0 (u_t - 1) + \frac{\Psi_1}{2} (u_t - 1)^2 \right) \frac{k_t}{\chi_t} \tag{33}$$

Since steady state utilization is normalized to 1, on average this accounting identity is the standard one that output equals the sum of consumption, investment, and government consumption. Outside of steady state there is an adjustment for the resource cost of utilization.

The characterization of equilibrium in the model economy is straightforward. Given the exogenous shocks, an equilibrium is a set of prices (wages, interest rates, etc.) and quantities (consumption, hours, etc.) such that all optimality conditions hold for all agents and all markets simultaneously clear. To solve for the equilibrium the model must be written in stationary form. As in the simple growth model of Section 2, define $z_t = (a_t^p)^{\frac{1}{1-\alpha}} \chi_t^{\frac{\alpha}{1-\alpha}}$. Then, for most of the trending variables the following is a stationary transformation: $\widehat{x}_t = \frac{x_t}{z_t}$. There are some exceptions to this, such as the capital stock, where the stationary transformation is: $\widehat{k}_t = \frac{k_t}{z_t \chi_t}$. The model is solved by log-linearizing the equations characterizing the equilibrium about the balanced growth path using standard techniques.

5 How Well Can the Model Fit the Data?

In this section I ask how well the benchmark model of the previous section, variants of which are now in wide use by policymakers and academics alike, can match the conditional responses to technology shocks estimated in the data in Section 3. In the first subsection I simply take a "standard" paramaterization of the model and show that it performs poorly. In the second subsection I apply a minimum distance estimator, choosing the parameters to optimize the model's fit to the empirical impulse responses. The best-fitting parameter configuration turns out to be quite different than the "standard" one, and in the third subsection I discuss some of the reasons why and speculate on possible improvements to the model.

5.1 A "Standard" Paramaterization

I set $g_a = 0.0018$ and $g_\chi = 0.0028$, which are equal to the sample averages of the growth rates of adjusted TFP and the relative price of investment in the data. I set $\pi^* = 0.0086$, which is equal to the average value of inflation in the data. $\alpha = 0.31$, which is the capital share number used in the construction of the TFP series. The average value of the Fed Funds rate in the data is 1.4 percent at a quarterly frequency (non-annualized). This together with the other parameters implies a value of $\beta = 0.999$. I choose a conventional value for the depreciation rate of $\delta = 0.025$. The steady state share of government purchases is set to $\omega^* = 0.2$, consistent with the data. ε and η are both set to 10, implying steady state price and wage markups of about 10 percent. The steady state value of θ^* is set to ensure that steady state hours worked are one-third of the available time endowment. Ψ_0 is set equal to the steady state marginal product of capital, which normalizes steady state utilization to be unity.

The choices for the remaining non-stochastic parameters are loosely based on those estimated in the literature (see, for example, Christiano, Eichenbaum, and Evans, 2005; Smets and Wouters, 2007; or Fernandez-Villaverde, 2010). The habit formation parameter, γ, is set to 0.7. The Calvo parameters for both wage and price-stickiness are set to 0.6, which is broadly consistent with most of the existing estimates in the literature. Likewise, the indexation parameters for both wage and price stickiness are set to 0.5. $\xi = 1$, implying a Frisch labor supply elasticity of 1. $\tau = 2.5$, which is near the central estimate of Christiano, Eichenbaum, and Evans (2005). $\Psi_1 = 0.01$. This implies that the costs of capital utilization are very nearly linear. A low value for this parameter provides significant amplification. The parameters of the Taylor rule are $\rho_i = 0.8$, $\varphi_\pi = 1.5$, and $\varphi_y = 0.25$.

The stochastic parameters of the model are set as follows. The persistence parameters for both preference shocks, the stationary neutral technology shock, and the government spending shock are set to 0.9. The persistence parameters governing the growth rates of the permanent neutral and investment-specific technology are set to 0.5. The standard deviation of the investment-specific shock is set to 0.25 percent; the remaining shock volatilities are all set to 0.5 percent. These parameter values are shown in Table 4.

I conduct the following experiment. I generate 1000 different data sets with 220 observations each using the parameterized model as described above; this is the same number of observations used to estimate the empirical VARs. On each generated data set I estimate exactly the same VAR as described in Section 3 and use the same restrictions to identify the impulse responses to the three technology shocks of interest. Figures 10-12 show the true response to each kind of technology shock from the model (dashed line) as well as the average VAR estimated response from the 1000 simulations (dotted line), along with the responses estimated in the data (solid line) and associated confidence bands (gray area). The parameterization of the shock processes for the technology processes are loosely chosen to match the estimated empirical responses, though no attempt is made (as of yet, see below) to match the apparent correlation between neutral and investment specific shocks.

In Figure 10 we see that the model provides a very poor fit to the actual response of hours to the temporary neutral shock. Whereas hours rise in the estimated responses in the data, hours decline on impact rather significantly in the model. The ensuing hump-shaped response of hours in the model is far too large relative to the data, which leads to the hump-shape in the output response being too large. The model also does a poor job with respect to inflation, as inflation essentially does not react at any horizon in the data, whereas it falls sharply on impact in the model.

The fit between model and data is arguably worse in Figure 11, though care needs to be taken because in the data the permanent neutral shock leads to a significant reduction in investment specific technology, though this is not a feature of the model. The shape of the hours response is qualitatively correct, but the sign is flipped – in the data hours decline in response to the permanent neutral shock, whereas in the model they rise. The output and consumption responses are too large on impact in the model relative to the data. Output significantly overshoots its long run response in the model; this is not a feature of the impulse responses estimated in the data. The impact drop in inflation is roughly the same in the model as in the data; but whereas the inflation response in the data is persistently negative, in the model it quickly turns positive and follows a hump-shape.

As noted, in the data the inverse relative price of investment declines significantly in response to the permanent neutral shock; this is not a feature of the model as written, but will be for the purposes of estimation in the next subsection.

As can be seen in Figure 12, the model generally does a better job of matching the responses to the investment specific shock than it does for the other two shocks. Nevertheless, it struggles to generate as large an hours response as there is in the data and it fails to deliver as much disinflation as there is in the data. Also, in the data consumption jumps up in response to improved investment specific technology, whereas in the model it declines very slightly on impact before rising to its new long run level.

In short, the parameterized model does rather poorly in generating the observed impulse responses to the three technology shocks in the data. A point worth emphasizing is that, in Figures 10 through 12, the dashed and dotted lines lie very close to one another at most horizons for almost all the variables included in the VAR. Recall that the dashed lines are the theoretical impulse responses in the model, while the dotted lines are the average estimated responses on data simulated from the model, using the VAR specification and identification of Section 3. Put differently, the VAR evidently performs very well. This is important, as there is a literature which questions the reliability of structural VARs, particularly those identified with long run restrictions (e.g. Christiano, Eichenbaum, and Vigfusson, 2006b; Chari, Kehoe, and McGrattan, 2008). An important difference in the current paper, and evidently the source of the very small biases relative to what some others have found in slightly different exercises, is that I employ a measure of TFP as a measure of technology, as opposed to most of the extant literature, which uses labor productivity. Another important difference is that I condition on significantly more information in estimating a medium-sized VAR model, as opposed to most of the Monte Carlo exercises which are conducted on bivariate systems. I have experimented with many different parameterizations of the model and the good performance – both quantitative and qualitative – of the VAR identification of Section 3 obtains in almost all of them. This good observed performance of the VAR forms the basis of the estimation procedure employed in the next subsection.

5.2 Estimating the Model

In this section I choose a subset of the parameters of the model to minimize the distance between the impulse responses to the technology shocks from those responses estimated in the data. This exercise is informative as it will reveal (i) the extent to which the model as written can account for the observed responses and (ii) how the best-fitting parameters differ from those commonly used in the literature.

I estimate a subset of the model's parameters using the approach employed in Christiano, Eichenbaum, and Evans (2005). Let Θ be a $q \times 1$ vector of parameters to be estimated. Let $\mathbf{M}$ be an $f \times 1$ vector of impulse response point estimates to the three kinds of technology shocks as estimated in the data. Let $\mathbf{M}(\Theta)$ be a vector of the same size of impulse responses to the three kinds of technology shocks in the model, given a parameter vector Θ. The estimator, $\widehat{\Theta}$, is the solution to the following minimization problem:

$$\widehat{\Theta} = \underset{\Theta}{\text{argmin}} \quad (\mathbf{M} - \mathbf{M}(\Theta))' \mathbf{W} (\mathbf{M} - \mathbf{M}(\Theta))$$

Here $\mathbf{W}$ is an $f \times f$ weighting matrix. In practice, I include in $\mathbf{M}$ the impulse responses of seven variables – utilization adjusted TFP, the inverse relative price of investment, hours worked per capita, output per capita, consumption per capita, the nominal interest rate, and inflation – to all three technology shocks from impact up to a horizon of 20 quarters. This means that there are $f = 7 \times 3 \times 20 = 420$ elements in both $\mathbf{M}$ and $\mathbf{M}(\Theta)$. I set $\mathbf{W}$ equal to a diagonal matrix, with diagonal elements equal to the inverse of the variance of the estimated impulse response functions from the data. This matrix is thus computed using the same bootstrap procedure used to generate the confidence bands for the impulse responses. This weighting scheme puts the most weight on those responses which are estimated with the most precision.

Before discussing which parameters are to be estimated, it is first necessary to slightly alter the model. In particular, as can be seen in Figure 6 and Table 3, the relative price of investment responds significantly to the permanent neutral shock. As noted in Section 3, this is not a violation of the assumptions underlying identification; rather it just means that the two shocks are likely correlated. As such, I modify the exogenous process for the level of investment specific technology as follows:

$$\Delta \ln \chi_t = (1 - \rho_\chi) g_\chi + \rho_\chi \Delta \ln \chi_{t-1} + \varsigma \varepsilon_{a^p,t} + \varepsilon_{\chi,t} \tag{34}$$

The parameter ς governs the extent to which the two kinds of permanent technology shocks are correlated. Given the estimated responses, one would expect $\varsigma < 0$.

A number of parameter values are fixed at their values discussed above in Section 5.1 and given in Table 4. These are labor's share in the Cobb-Douglas production function, α; the growth rates of neutral and investment specific technology, g_a and g_χ; the depreciation rate on capital, δ; the parameters governing steady state markups in both prices and wages, ε and η; steady state inflation, π^*; and the government purchase share of output, ω^*. Because I am matching the model's theoretical impulse responses to those estimated in the data, the parameters governing the stochastic processes for other the exogenous state variables need not be specified or estimated.[10] The parameters left to be estimated are then: $\Theta = \begin{pmatrix} \tau & \gamma & \Psi_1 & \xi & \phi_p & \phi_w & \zeta_p & \zeta_w & \varsigma & \varphi_\pi & \varphi_y & \rho_i & \rho_{a^s} & \rho_{a^p} & \rho_\chi & \sigma_{\varepsilon_{a^s}} & \sigma_{\varepsilon_{a^p}} & \sigma_{\varepsilon_\chi} \end{pmatrix}'$.

Table 5 shows the estimated parameter values and associated standard errors in parentheses.[11] The first three parameters in the first row of the table govern the extent of real frictions in the

[10] Another approach would be to simulate data from the model, estimate VARs on simulated data, and match the average estimated responses from the simulation to those in the data. In addition to being more computationally burdensome, this approach requires estimating the parameters of the stochastic processes which do not affect the theoretical responses to the technology shock, and therefore requires looking at additional moments beyond the impulse responses to the technology shocks. In practice these approaches are likely to yield similar results, given the good performance of the VARs in the Monte Carlo experiment of the previous section.

[11] The standard errors are computed as follows. Under regularity conditions the distribution of the estimator is approximately: $\sqrt{T}\left(\widehat{\Theta} - \Theta_0\right) \rightarrow N(0, V)$, where $V = \left(\widehat{D}\mathbf{W}\widehat{D}'\right)^{-1}$ and T is the sample size. $\widehat{D}$ is the numerical Jacobian of $\mathbf{M} - \mathbf{M}(\widehat{\Theta})$ evaluated at the estimated parameter values. See, e.g., Dejong and Dave (2007). The standard errors are then the square roots of the diagonal elements of V, divided by the square root of T.

model – investment adjustment costs, habit formation in consumption, and variable capital utilization, respectively. The parameter governing investment adjustments costs, τ, is very small and statistically and economically indistinguishable from 0. This differs substantially from most of the DSGE literature, which finds values of this parameter ranging anywhere from 2.5 (Christiano, Eichenbaum, and Evans, 2005) to close to 10 (Fernandez-Villaverde, 2010). The habit formation parameter, estimated at $\gamma = 0.89$, is large, but within the range of most empirical estimates. As is common in the literature, the parameter governing the curvature of the utilization cost function, Ψ_1, is estimated to be very nearly equal to zero. This implies that the costs of capital utilization are essentially linear, and provides an important amplification mechanism in the model. The next parameter, ξ, is the inverse Frisch labor supply elasticity. Estimated at nearly zero, this means that the Frisch elasticity is nearly infinite. Put differently, household preferences are estimated to be very nearly linear in labor, so that the model is effectively observationally equivalent to the indivisible labor models of Hansen (1985) and Rogerson (1988). This estimate differs a good deal from other estimates in the DSGE literature, which often find estimates of the Frisch elasticity in the neighborhood of unity (Fernandez-Villaverde, 2010).

The next set of parameters in the table concern the degree of nominal rigidities. The Calvo parameter for price-setting by intermediate goods firms is $\phi_p = 0.41$. This means that prices last on average less than two quarters, and is substantially lower than most other estimates, which typically range from 0.5 to 0.8. The indexation parameter for price-setting, ζ_p, is almost identically zero. The Calvo parameter for wage-stickiness is estimated to be $\phi_w = 0.513$, implying wage contracts with an average duration of one half a year. This is also somewhat on the low side of existing estimates. Like the price indexation parameter, the indexation parameter for wage contracts is estimated to be almost exactly zero. The parameters of the monetary policy rule are fairly standard, with the coefficient on inflation equal to 1.21 and the coefficient on output growth equal to 0.3. Perhaps surprisingly, there is almost no evidence of an explicit interest smoothing desire, with ρ_i estimated to be close to zero.

There is strong persistence in both the permanent and stationary components of neutral technology, with $\rho_{a^s} = 0.97$ and $\rho_{a^p} = 0.66$. The persistence of the stationary component of technology is consistent with many RBC calibrations (e.g. King and Rebelo, 1999), while the persistence of the permanent component is similar to the estimates in Altig, Christiano, Eichenbaum, and Linde (2011). The estimated persistence of the investment specific shock, $\rho_{\chi} = 0.05$, means that investment specific technology is estimated to follow a process close to a random walk. We observe that positive permanent neutral shocks are negatively correlated with the state of investment specific technology, with $\varsigma = -2.89$. The standard deviations of the two neutral technology shocks are about a third of a percent, while the standard deviation of the investment specific shock is about 0.6 percent. The J statistic for the estimated model is 90.02; with the large number of degrees of freedom, it is not possible to reject the over-identifying restrictions.

Figures 13 through 15 show the impulse responses in the model at the estimated parameter values (dashed line) along with the responses estimated in the data (solid line) and associated confidence bands (shaded gray regions). Figure 13 shows the responses to the transitory neutral shock. The model does a very good job of matching the responses of quantities, with the model responses of hours, output, and consumption all lying very close to their counterparts in the data and within the confidence regions at most horizons. Of particular interest is the fact that hours rise on impact in the model. The estimated impact effect on adjusted TFP is too small in the model relative to the data, but the model matches the dynamic response well. The estimated model does less well at matching the responses of the interest rate and inflation. In particular, it is very difficult for the model to not generate some disinflation on impact to a positive temporary neutral shock, whereas there is little response of inflation in the data.

Figure 14 shows the model and data responses to a permanent neutral shock. Here the fit of the model is very good. The model responses of adjusted TFP, consumption, output, hours, and the relative price of investment nearly lie on top of the data responses at all horizons. The model does substantially better at matching the responses of interest rates and inflation. In particular, both inflation and the interest rate are estimated to fall significantly on impact, and stay persistently below their pre-shock value for a number of quarters.

Figure 15 shows model and data responses to the investment specific shock. Here the fit is quite good as well. In particular, the model responses of output, hours, investment specific technology, and the interest rate lie very close to their counterparts from the estimated VAR. The model has some difficulty in generating the persistence of the disinflation observed in the data, and does a poor job at matching the consumption response. In the data consumption rises on impact in response to the investment specific shock; in the model it essentially does not react, and only very slowly approaches its new long run value.

5.3 Discussion

While the fit of the estimated model to the empirical impulse responses to the three technology shocks is far from perfect, it nevertheless appears to represent a substantial improvement over the "standard" parameterization considered in Section 5.1. A visual comparison of Figures 13 through 15 to Figures 10 through 12 reveals that the estimated model does a substantially better job at matching the responses from a qualitative perspective, particularly so with respect to the behavior of hours.

Formally, one can conduct a likelihood ratio type test of the restricted and estimated models, with $LR = T\left(J\left(\Theta_0\right) - J\left(\widehat{\Theta}\right)\right)$, where Θ_0 is the vector of "standard" parameter values given in Table 4 and $\widehat{\Theta}$ is the vector of estimated parameters given in Table 5. This test statistic follows a chi-squared distribution with degrees of freedom equal to the number of restrictions, which in this case equals the number of parameters, 18. The value of the test statistic is 517, leading to an overwhelming rejection of the "standard" parameterization.

The three areas along which the estimated parameters differ the greatest from the "standard" parameterization are (i) the level of investment adjustment costs, (ii) the very high Frisch labor supply elasticity, and (iii) the lack of price and wage indexation. The estimated parameters also show less levels of nominal rigidity than is commonly found, though this difference is less substantial.

Investment and/or capital adjustment costs play an important role in many modern DSGE models, especially where the main point of interest is in understanding the effects of monetary policy shocks on the real economy. These adjustment costs play three important roles: (i) they break the connection between the real interest rate and the marginal product of capital; (ii) they help to generate autocorrelation in output and investment growth rates; and (iii) they help generate hump-shaped impulse responses to shocks, particularly monetary policy shocks.

As to point (i), when there are no adjustment costs, the real interest rate and the rental rate on capital must always be approximately equal. The rental rate is in turn related to the price markup and the marginal product of capital. Most observers think of contractionary monetary policy as associated with rising nominal and real interest rates, and this is indeed what the VAR evidence shows. But without adjustment costs, rising real rates necessitate an increase in the marginal product of capital in the model, which can only come about through an increase in hours and utilization, which in turn causes output to expand. Hence, absent some kind of adjustment costs, policies to raise interest rates will typically lead to an expansion in output, which is deeply at odds with the data and most of our intuition.

Roles (ii) and (iii) of investment adjustment costs are essentially the flip side of the same coin. Cogley and Nason (1995) forcefully make the point that output and its components display significant positive autocorrelation in growth rates. Standard real business cycle models, with a weak internal propagation mechanism, cannot generate this degree of autocorrelation. Adjustment costs – be they to investment, labor, or consumption – can help do this. The intuition is straightforward. When there are convex costs associated with adjusting some activity (e.g. investment), we will tend to observe under, and then over, shooting in response to shocks. This leads to hump-shaped impulse responses and higher unconditional autocorrelations in growth rates.

Investment adjustment costs play another role in the model, which is to severely weaken the hours response to a technology shock. Habit formation in consumption plays a complementary role here. Consider a transitory increase in neutral technology. In a simple RBC model intertemporal substitution would lead households to increase their consumption a little and work more; permanent income motives would lead them to substantially increase investment in an attempt to smooth out the shock. With significant investment adjustment costs, however, smoothing via increased investment is not sensible. With habit formation in consumption, it does not make sense to increase consumption by much either. Households are forced to "spend" the gains from higher technology on leisure, and hence hours are very likely to fall when technology improves.

The point that investment adjustment costs and consumption habit formation can lead to "contractionary" technology shocks is made in Francis and Ramey (2005). Figure 16 plots out the impact effect on hours of a positive temporary neutral technology shock for different values of τ and γ, holding all other parameters fixed at the "standard" values given in Table 4. The values of τ range from 0 to 10, while the values of γ range from 0 to 1. One clearly observes that the impact jump in hours when neutral technology improves is strictly decreasing in both of these parameters, for the reasons cited in the paragraph above. Hence, other things being equal, the model would better fit the hours response to a temporary neutral shock with very little investment adjustment costs and very little habit formation. However, in order to match the small impact responses of consumption to all of the shocks estimated in the data, the model needs habit formation in consumption. Hence, the only way for the model to match the impact rise in hours in response to a temporary neutral technology shock is to have investment adjustment costs essentially vanish.

The second area in which the estimated parameters differ from their "standard" values is in the high Frisch labor supply elasticity, which is estimated to be nearly infinite. The reason for this high estimated elasticity is because hours react significantly on impact in response to both permanent and transitory neutral shocks in the VAR identifications. For the model to match this feature, it needs a high elasticity. Similar DSGE models estimated using full information approaches typically find much lower elasticities, but they compensate for this with extremely volatile preference shocks (Fernandez-Villaverde, 2010), which are difficult to interpret. Models estimated using limited information approaches such as that used here, but focusing on monetary policy shocks (e.g. Christiano, Eichenbaum, and Evans, 2005), also find lower elasticities, but this is because the response of hours to policy shocks is quite inertial. In response to technology shocks, there does not seem to be much inertia in the hours response.

A third important difference in the estimated model here is that there is no evidence of wage or price indexation to lagged wages and prices. High levels of indexation are important in generating inertial behavior of inflation in response to a monetary policy shock (Christiano, Eichenbaum, and Evans, 2005).[12] Here, however, particularly in response to the two permanent technology shocks, the behavior of inflation is not very inertial at all. In fact, as can be seen clearly in Figures 6 and 8, the maximal response of inflation to these shocks is on impact. This lack of inertia drives the estimates of the indexation parameters to zero.

The model also features lower levels of nominal rigidity than is typically found, with price and wage contracts having average durations of two quarters or less. There are competing forces at work in the estimated levels of these rigidities. The large impact response of inflation to both the permanent neutral shock and the investment specific shock argues in favor of very low levels of price rigidity. The increase of hours in response to the temporary neutral shock is also consistent with highly flexible prices. Figure 17 plots in the left panel the impact response of inflation to the permanent neutral shock as a function of ϕ_p while the right panel plots the impact response of hours to the temporary neutral shock, also as a function of ϕ_p. The impact decline in inflation is largest when ϕ_p is low, while the hours response is larger the lower is ϕ_p. Both of these these facts tend to push the estimate of ϕ_p towards zero (price flexibility). What counterbalances this is the lack of an inflation response to the temporary neutral shock; this is more consistent with extreme price rigidity. There is also a tension with respect to the level of wage rigidity – the impact increase in hours in response to the temporary neutral shock is consistent with fairly sticky wages, while the impact decline following a permanent neutral shock fits better with flexible wages.

In summary, it seems fair to conclude that the model matches the empirical responses to technology shocks better with fewer frictions than is commonly assumed. In particular, the model fit is better with relatively little nominal rigidity and no investment adjustment costs. There exists a

[12] This kind of inertia can take different forms other than strict inflation indexation. Backward-looking firms (Gali and Getler, 1999) and sticky information (Mankiw and Reis, 2006) will have similar effects. See Dupor, et al (2009) for a discussion.

tension between these findings and the parameterizations that are needed in order to understand the responses to monetary policy shocks. Previous authors have proposed informational frictions as a potential resolution of this apparent tension (Paciello, 2010). The basic intuition is that, if gathering information is costly, it may be optimal to not pay much attention to monetary shocks relative to technology shocks, because the benefits of full optimization are relatively minor following monetary shocks. If agents optimally choose to not pay much attention to monetary shocks, the economy may behave with a great deal of inertia, and therefore look like an economy with significant adjustment costs and other rigidities. Recent empirical work by Coibion and Gorodnichenko (2011a and 2011b) seems promising and fruitful. More work is needed, however, in incorporating these kinds of informational frictions in medium-scale models.

6 Conclusion

This paper contributes to the study of technology shocks in the business cycle. Its main novelty is that it simultaneously considers three different kinds of technology shocks – transitory neutral, permanent neutral, and investment specific. Empirically, the three technology shocks combine to account for about half of the business cycle variance of output. Positive transitory neutral shocks raise hours worked while permanent neutral shocks lower hours; investment specific shocks raise hours worked significantly with some delay. Standard parameterizations of popular medium scale DSGE models do a poor job of accounting for the pattern of responses to the three kinds of technology shocks, particularly with regard to the behavior of hours. Parameterizations with relatively fewer frictions tend to match the data better. This raises an important puzzle because frictions are needed to understand the dynamic responses to monetary policy shocks. While models of informational frictions seem promising, more work is needed. This task is left to future research.

References

[1] Altig, David, Lawrence Christiano, Martin Eichenbuam, and Jesper Linde. 2011. "Firm-Specific Capital, Nominal Rigidities, and the Business Cycle." *Review of Economic Dynamics*, 14(2), 225-247.

[2] Barsky, Robert and Eric Sims. 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics*, forthcoming.

[3] Basu, Susanto. 1996. "Procyclical Productivity: Increasing Returns or Cyclcial Utilization?" *Quarterly Journal of Economics* 111(3), 719-751.

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

[5] Beaudry, Paul and Franck Portier. 2006. "Stock Prices, News, and Economic Fluctuations." *American Economic Review* 96(4), 1293-1307.

[6] Blanchard, Olivier and Danny Quah. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79(4), 655-673.

[7] Burnside, Craig and Martin Eichenbaum. 1996. "Factor Hoarding and the Propagation of Business Cycle Shocks." *American Economic Review* 86(5), 1154-1174.

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

[9] Chari, VV, Patrick Kehoe, and Ellen McGrattan. 2008. "Are Structural VARs with Long Run Restrictions Useful in Developing Business Cycle Theory?" *Journal of Monetary Economics* 55(8), 1337-1352.

[10] Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson. 2006a. "Alternative Procedures for Estimation Vector Autoregressions Identified with Long Run Restrictions." *Journal of the European Economic Association* 4(2), 475-483.

[11] Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson. 2006b. "Assessing Structural VARs." *NBER Macroeconomics Annual 2006*, Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds.

[12] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1), 1-45.

[13] Cochrane, John H. 1994. "Permanent and Transitory Components of GNP and Stock Prices." *Quarterly Journal of Economics*, 109(1), 241-266.

[14] Cogley, Timothy and James Nason. 1995. "Output Dynamics in Real Business Cycle Models." *American Economic Review* 85(3), 492-511.

[15] Coibion, Olivier and Yuri Gorodnichenko. 2011a. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." Working paper, University of

California at Berkeley.

[16] Coibion, Olivier and Yuri Gorodnichenko. 2011b. “What Can Survey Forecasts Tell Us about Informational Rigidities?” Working paper, University of California at Berkeley.

[17] Dejong, David and Chetan Dave. 2007. *Structural Macroeconometrics*. Princeton, NJ: Princeton University Press.

[18] Dupor, William, Jing Han, and Yi-Chan Tsai. 2009. “What do Technology Shocks Tell Us About the New Keynesian Paradigm?” *Journal of Monetary Economics*, 56(4), 560-569.

[19] Edge, Rochelle, Thomas Laubach, and John Williams. 2003. “The Responses of Wages and Prices to Technology Shocks.” Finance and Economics Discussion Series 2003-65, Federal Reserve Board of Governers.

[20] Erceg, Christopher, Dale Henderson, and Andrew Levin. 2000. “Optimal Monetary Policy with Staggered Wage and Price Contracts.” *Journal of Monetary Economics* 46(2), 281-313.

[21] Fernald, John. 2009. “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity.” Federal Reserve Bank of San Francisco working paper.

[22] Fernandez-Villaverde, Jesus. 2010. “The Econometrics of DSGE Models.” *SERIES: The Journal of the Spanish Economic Association* 1(1), 3-49.

[23] Fisher, Jonas. 2009. “Comment on: ‘Letting Different Views about Business Cycles Compete.” *NBER Macroeconomics Annual 2009*, Daron Acemoglu, Ken Rogoff, and Michael Woodford, eds., 457-474.

[24] Francis, Neville and Valerie Ramey. 2005. “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited.” *Journal of Monetary Economics*, 52(8), 1379-1399.

[25] Francis, Neville and Valerie Ramey. 2009. “Measures of Hours per Capita and Their Implications for the Technology-Hours Debate.” *Journal of Monetary Economics*, 41(6), 1071-1097.

[26] Gali, Jordi. 1999. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review* 89(1), 249-271.

[27] Hansen, Gary. 1985. “Indivisible Labor and the Business Cycle.” *Journal of Monetary Economics* 16(3), 309-327.

[28] Gali, Jordi and Mark Gertler. 1999. “Inflation Dynamics: A Structural Econometric Analysis.” *Journal of Monetary Economics* 44(2), 195-222.

[29] Gordon, Robert. 1989. *The Measurement of Durable Goods Prices.* Chicago: University of Chicago Press.

[30] Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 1997. “Long-Run Implications of Investment Specific Technological Change.” *American Economic Review* 87(3), 342-362.

[31] Kilian, Lutz. 1998. “Small Sample Confidence Intervals for Impulse Response Functions.” *Review of Economics and Statistics*, 80(2), 218-230.

[32] King, Robert and Sergio Rebelo. 1999. "Resuscitating Real Business Cycles." *Handbook of Macroeconomics*, John Taylor and Michael Woodford, eds., 1(1), 927-1007.

[33] Mankiw, N. Gregory and Ricardo Reis. 2003. "Sticky Information vs. Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 107(4), 1295-1328.

[34] Paciello, Luigi. 2010. "Monetary Policy and Price Responsiveness to Aggregate Shocks Under Rational Inattention." Working paper, Einaudi Institute for Economics and Finance.

[35] Paciello, Luigi. 2011. "Does Inflation Adjust Faster to Aggregate Technology Shocks than to Monetary Policy Shocks?" Forthcoming, *Journal of Money, Credit, and Banking.*

[36] Rogerson, Richard. 1988. "Indivisible Labor, Lotteries, and Equilibrium." *Journal of Monetary Economics* 21(1), 3-16.

[37] Shapiro, Matthew. 1993. "Cyclical Productivity and the Workweek of Capital." *American Economic Review* 83(2), 229-233.

[38] Shapiro, Matthew and Mark Watson. 1988. "Sources of Business Cycle Fluctuations." *NBER Macroeconomics Annual 1988*, Stanley Fischer, ed., 111-148.

[39] Smets, Frank and Raf Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3), 586-606.

[40] Summers, Lawrence. 1986. "Some Skeptical Observations on Real Business Cycle Theory." *Federal Reserve Bank of Minneapolis Quarterly Review* 10(4), 23-27.

Table 1: Cyclicality and Volatility of TFP Series

	Std. Deviation	Corr w/ GDP	Corr w/ Hrs	Corr w/ Unadjusted TFP
Unadjusted TFP	0.014	0.811	0.337	1
Adjusted TFP	0.012	-0.008	-0.378	0.321

Notes: Data are 1947q1 to 2009q4 and are HP filtered with smoothing parameter 1600.

Table 2: Cyclicality and Volatility of Relative Investment Price

	Std. Deviation	Corr w/ GDP	Corr w/ Hrs	Corr w/ TFP	Corr w/ Adj. TFP
Rel. Invest Price	0.020	0.054	0.433	0.011	-0.154

Notes: Data are 1947q1 to 2009q4 and are HP filtered with smoothing parameter 1600.

Table 3: Forecast Error Variance Decomposition

	$h=1$	$h=4$	$h=8$	$h=16$	$h=24$	$h=40$
Temp. Neutral						
Adj. TFP	0.58	0.43	0.31	0.18	0.12	0.06
Rel. Inv. Price	0.00	0.01	0.01	0.01	0.01	0.00
Hours	0.30	0.18	0.16	0.22	0.22	0.21
Output	0.65	0.28	0.23	0.28	0.22	0.14
Consumption	0.12	0.05	0.05	0.06	0.05	0.03
Interest Rate	0.05	0.04	0.03	0.10	0.15	0.17
Inflation	0.02	0.01	0.02	0.09	0.14	0.17
Permanent Neutral						
Adj. TFP	0.42	0.53	0.64	0.77	0.84	0.89
Rel. Inv. Price	0.28	0.70	0.77	0.78	0.76	0.73
Hours	0.44	0.47	0.36	0.26	0.22	0.20
Output	0.04	0.01	0.02	0.14	0.29	0.49
Consumption	0.07	0.08	0.10	0.15	0.19	0.26
Interest Rate	0.10	0.45	0.63	0.65	0.60	0.50
Inflation	0.21	0.33	0.46	0.44	0.37	0.33
Investment Specific						
Adj. TFP	0.00	0.00	0.03	0.03	0.04	0.05
Rel. Inv. Price	0.05	0.04	0.11	0.16	0.19	0.25
Hours	0.01	0.06	0.18	0.20	0.27	0.29
Output	0.01	0.14	0.22	0.20	0.18	0.14
Consumption	0.52	0.28	0.29	0.32	0.33	0.32
Interest Rate	0.00	0.02	0.05	0.05	0.05	0.08
Inflation	0.28	0.25	0.18	0.18	0.18	0.19
Total Technology						
Adj. TFP	1.00	0.96	0.98	0.98	1.00	1.00
Rel. Inv. Price	0.33	0.74	0.89	0.95	0.96	0.98
Hours	0.75	0.71	0.55	0.68	0.71	0.70
Output	0.70	0.43	0.47	0.62	0.69	0.77
Consumption	0.71	0.41	0.44	0.53	0.57	0.61
Interest Rate	0.15	0.51	0.71	0.80	0.80	0.75
Inflation	0.51	0.59	0.66	0.71	0.69	0.69

Notes: These numbers are the fraction of the forecast error variance of each shock accounted for by the different structural shocks. The columns, labeled with h, refer to the forecast horizon. The final set of rows given the fraction of the total forecast error variance of the three technology shocks combined (just the sum of the numbers in the table).

Table 4: Standard Parameterization

α	β	g_a	g_χ	δ	π^*
0.31	0.999	0.0018	0.0029	0.025	0.008
τ	γ	Ψ_1	ξ	ϕ_p	ϕ_w
2.5	0.7	0.010	0.01	0.6	0.6
ζ_p	ζ_w	φ_π	φ_y	ρ_i	ρ_{a^s}
0.5	0.5	1.5	0.25	0.80	0.90
ρ_{a^p}	ρ_χ	ρ_g	ρ_θ	ρ_ψ	ω^*
0.5	0.5	0.9	0.9	0.9	0.2
$\sigma_{\varepsilon_{a^s}}$	$\sigma_{\varepsilon_{a^p}}$	$\sigma_{\varepsilon_\chi}$	σ_{ε_g}	$\sigma_{\varepsilon_\theta}$	$\sigma_{\varepsilon_\psi}$
0.005	0.005	0.005	0.005	0.005	0.005
σ_{ε_i}					
0.005					

Notes: These are the parameter values, largely taken from the literature, used in the simulation exercises in Section 5.1.

Table 5: Estimated Parameters

τ	γ	Ψ_1	ξ	ϕ_p	ϕ_w
0.01 (0.17)	0.897 (0.03)	0.009 (0.01)	0.009 (0.16)	0.419 (0.17)	0.513 (0.26)
ζ_p	ζ_w	ς	φ_π	φ_y	ρ_i
0.010 (0.81)	0.010 (2.09)	-2.899 (0.75)	1.210 (0.12)	0.309 (0.20)	0.01 (0.78)
ρ_{a^s}	ρ_{a^p}	ρ_χ	$\sigma_{\varepsilon_{a^s}}$	$\sigma_{\varepsilon_{a^p}}$	$\sigma_{\varepsilon_\chi}$
0.969 (0.02)	0.659 (0.08)	0.052 (0.15)	0.0032 (0.0006)	0.0031 (0.0012)	0.0059 (0.0008)

Notes: These are the point estimates for the benchmark estimation as described in Section 5.2. The numbers are parentheses are standard errors. Non-estimated parameter values are set at the values given in Table 4. The J statistic for the test of overidentifying restrictions is 90.02. There are 402 degrees of freedom – 420 impulse response point estimates and 18 parameters.

Figure 1: Adjusted and Unadjusted TFP

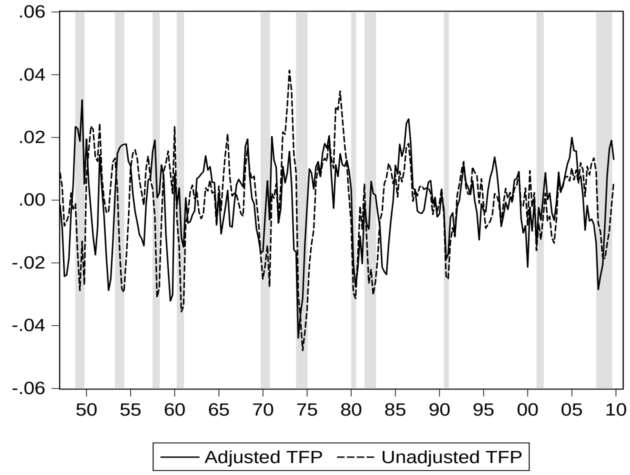

Notes: The solid line is the HP filtered adjusted TFP series and the dashed line is the unadjusted TFP series, with smoothing parameter 1600. The shaded gray areas are NBER defined recessions.

Figure 2: Inverse Relative Price of Investment

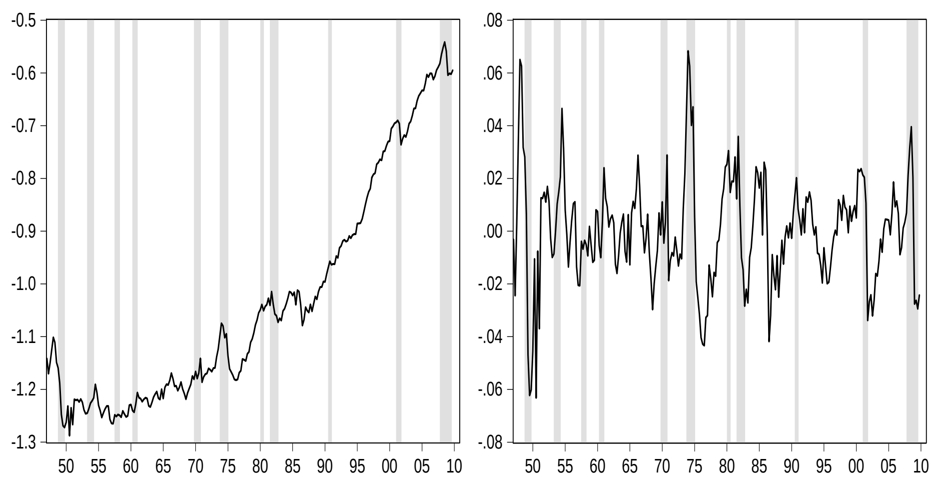

Notes: The left panel plots the inverse level of the price of investment relative to consumption goods. The right panel plots the cyclical component from an HP filter with smoothing parameter 1600. The shaded gray areas are NBER defined recessions.

Figure 3: Technology and Hours: Bivariate VARs

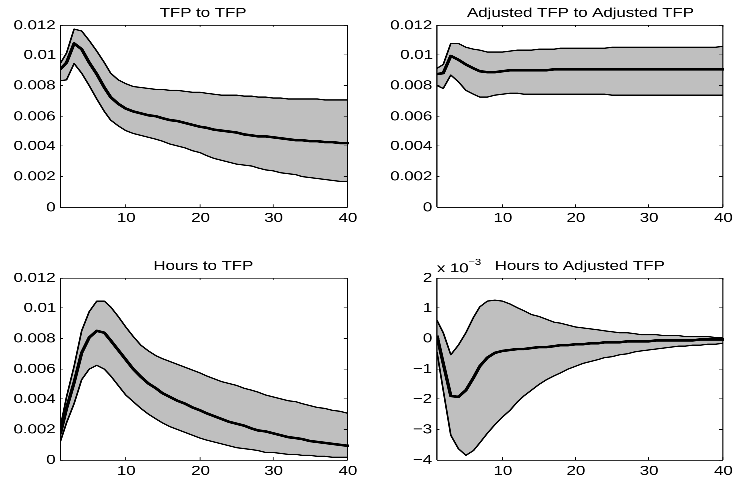

Notes: The left panel show impulse responses to an innovation to unadjusted TFP from a bivariate system with TFP and hours; the right panel show the same responses from a system using the BFK adjusted TFP series. Shaded gray regions are +/- one standard error confidence bands.

Figure 4: VAR Variables

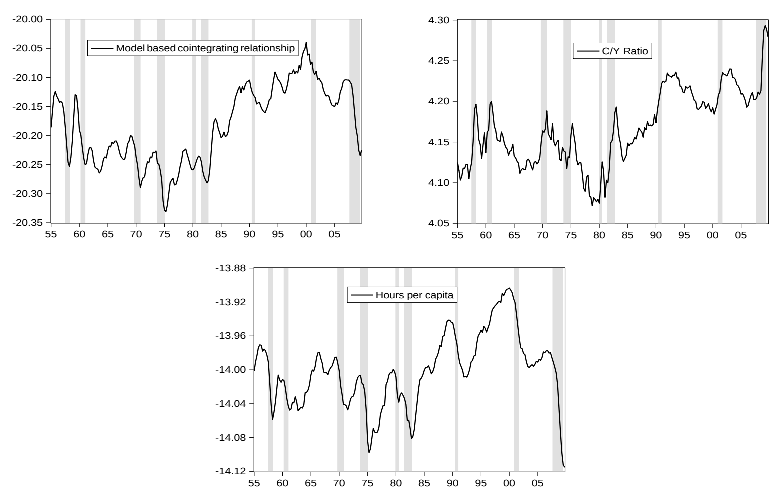

Notes: These figures plot some of the variables included in the larger VAR systems. Shaded gray regions are NBER defined recessions.

Figure 5: Impulse Responses to Temporary Neutral Technology Shock

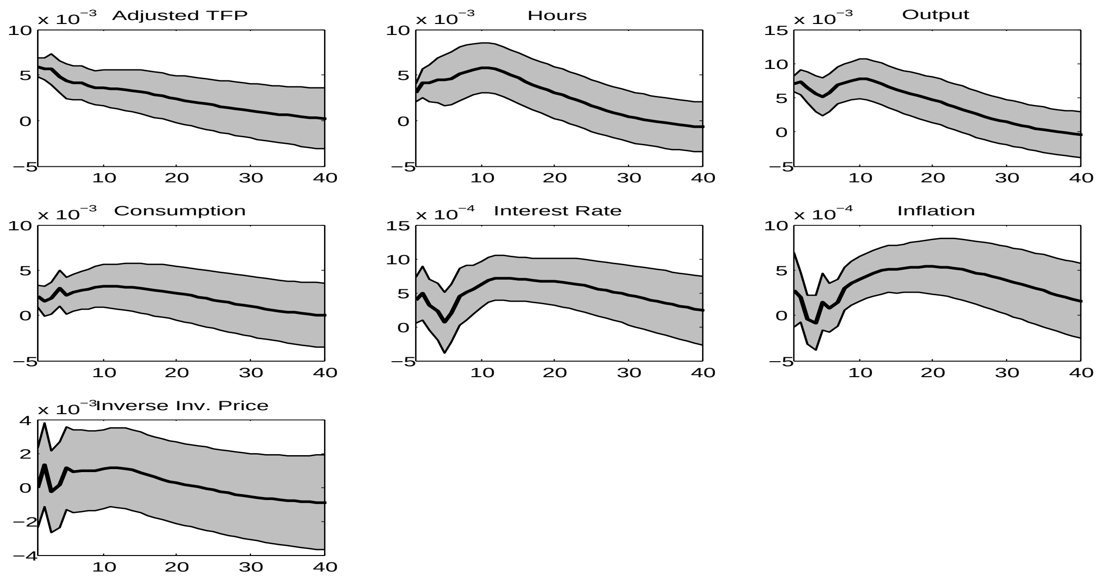

Notes: These are impulse responses to a temporary neutral technology shock. Shaded gray regions are +/- one standard error confidence bands.

Figure 6: Impulse Responses to Permanent Neutral Technology Shock

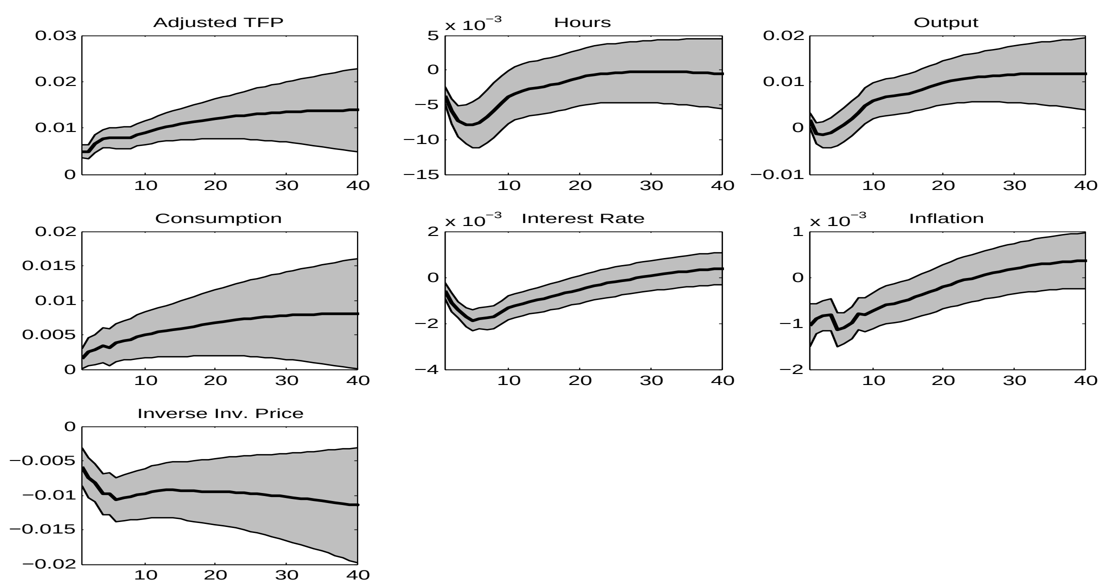

Notes: These are impulse responses to a permanent neutral technology shock. Shaded gray regions are +/- one standard error confidence bands.

Figure 7: Impulse Responses to Permanent Neutral Technology Shock: Alt. Orthogonalization

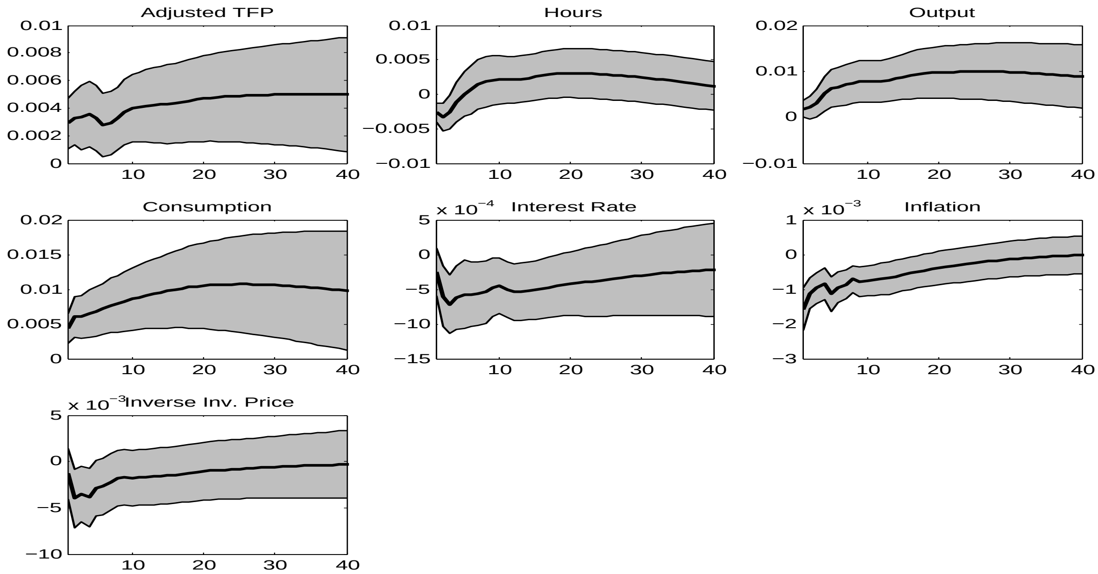

Notes: These are impulse responses to a permanent neutral technology shock under the alternative orthogonalization described in the text. Shaded gray regions are +/- one standard error confidence bands.

Figure 8: Impulse Responses to Investment Specific Technology Shock

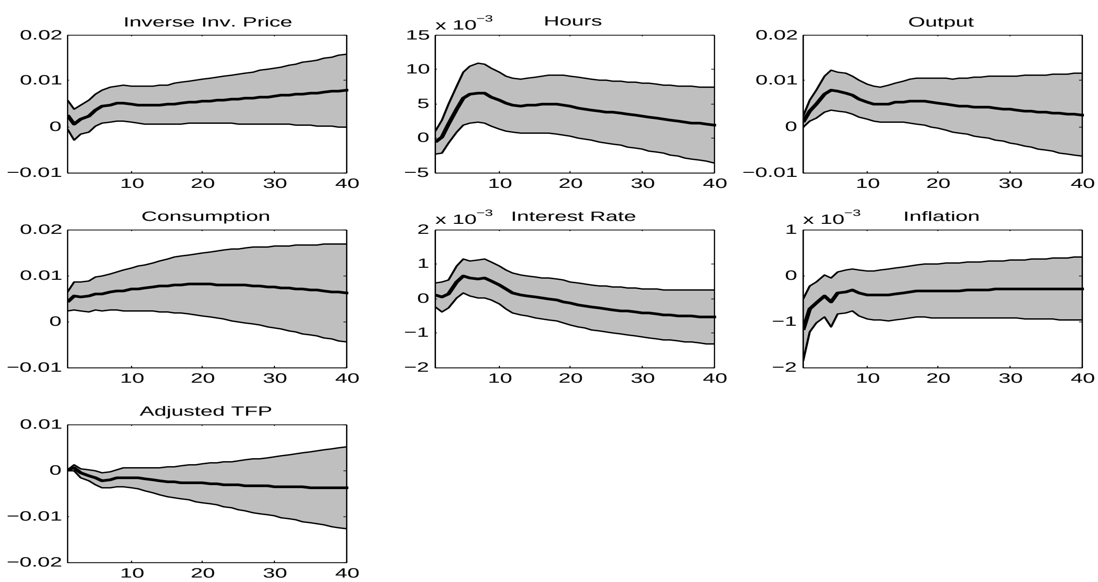

Notes: These are impulse responses to an investment specific technology shock. Shaded gray regions are +/- one standard error confidence bands.

Figure 9: Impulse Responses of Hours Under Different Stationarity Assumptions

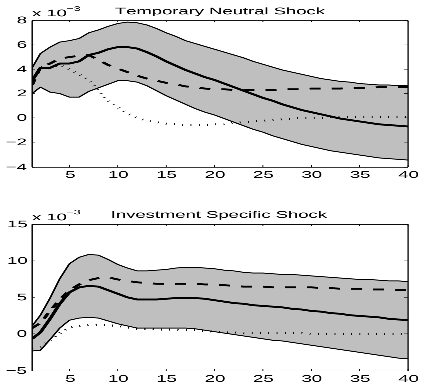

Notes: These are impulse responses of hours to the three different technology shocks under different assumptions about how hours enter the VAR. The solid lines are responses when hours enter in levels, the dashed lines when hours enter in first differences (these responses are cumulated), and the dotted lines when hours are HP filtered with smoothing parameter 16,000.

Figure 10: Benchmark Model vs. Data Responses: Temporary Neutral Shock

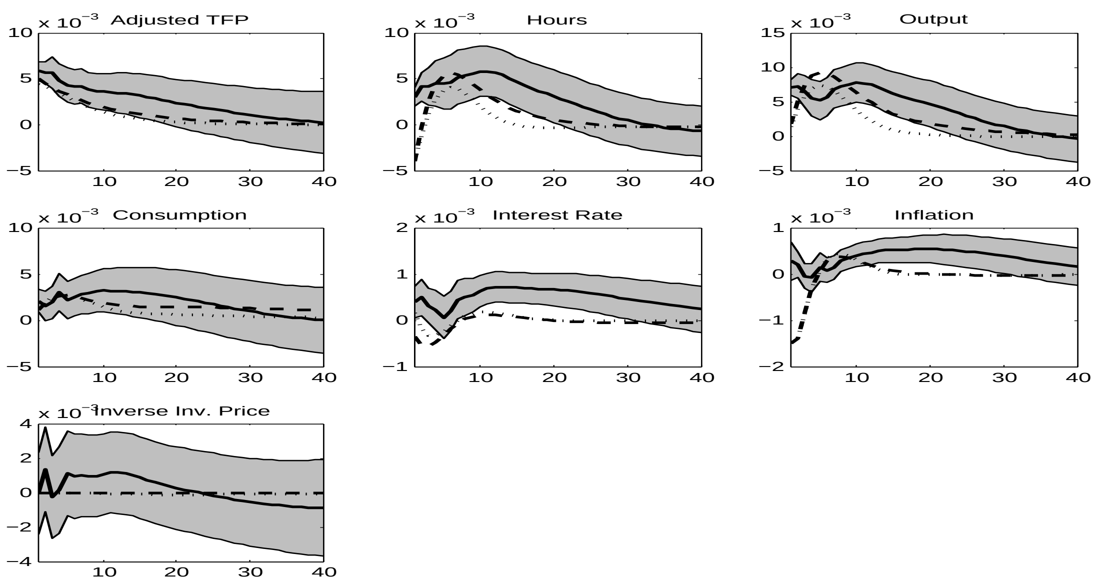

Notes: The dark lines are the impulse responses to a temporary neutral shock as estimated in the data, with the shaded gray area the +/- one standard error confidence bands. The dashed lines are theoretical impulse responses to the shock from the model of Section 4. The dotted lines are the average VAR estimated responses across 1000 different simulations of the data from the model.

Figure 11: Benchmark Model vs. Data Responses: Permanent Neutral Shock

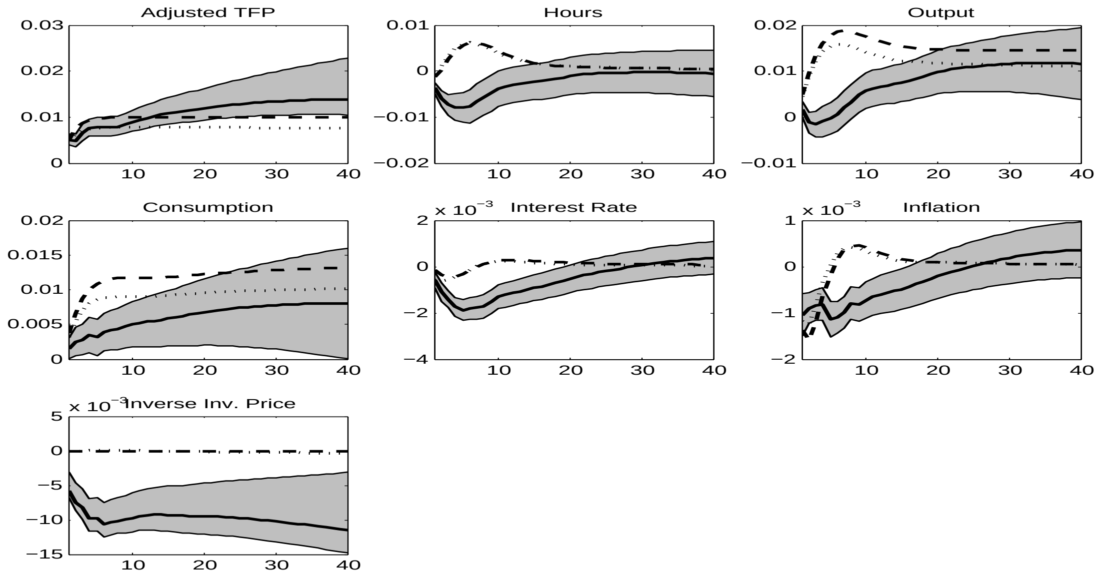

Notes: See the notes to Figure 10.

Figure 12: Benchmark Model vs. Data Responses: Investment Specific Shock

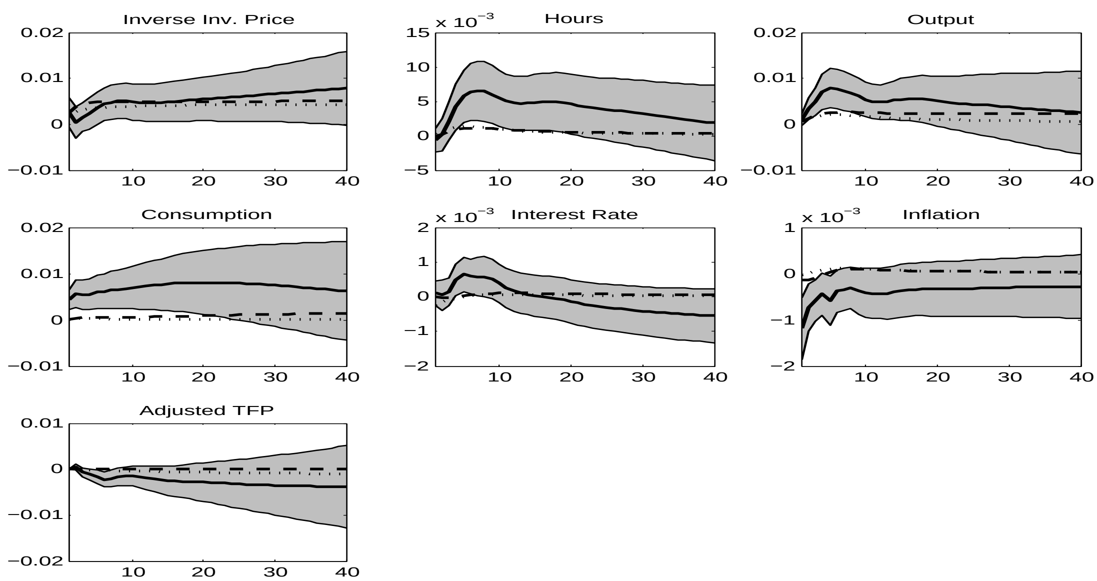

Notes: See the notes to Figure 10.

Figure 13: Estimated Model vs. Data Responses: Temporary Neutral Shock

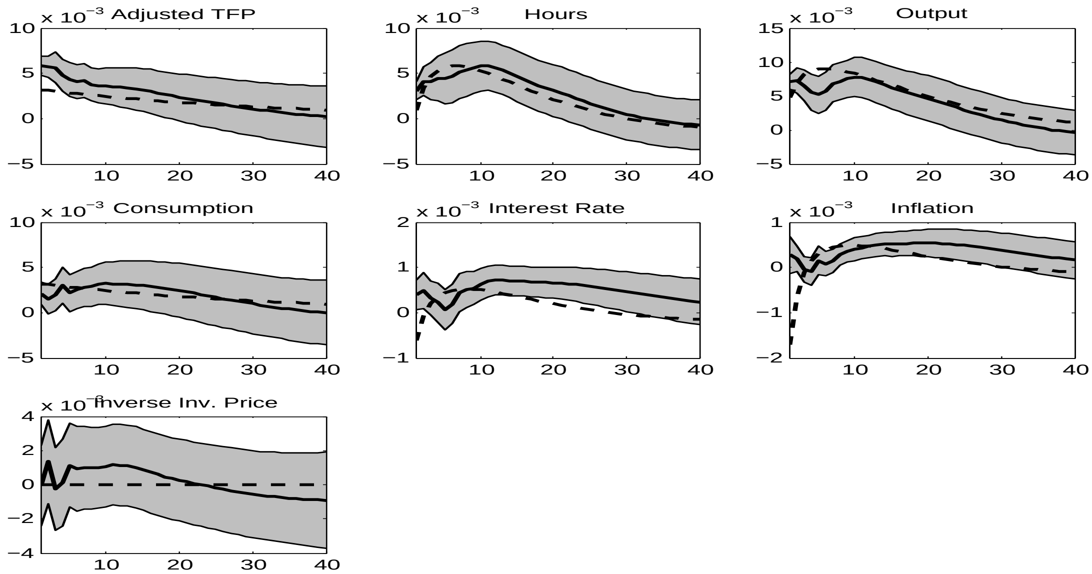

Notes: The dark lines are the impulse responses to a temporary neutral shock as estimated in the data, with the shaded gray area the +/- one standard error confidence bands. The dashed lines are theoretical impulse responses to the shock at the estimated parameter values.

Figure 14: Estimated Model vs. Data Responses: Permanent Neutral Shock

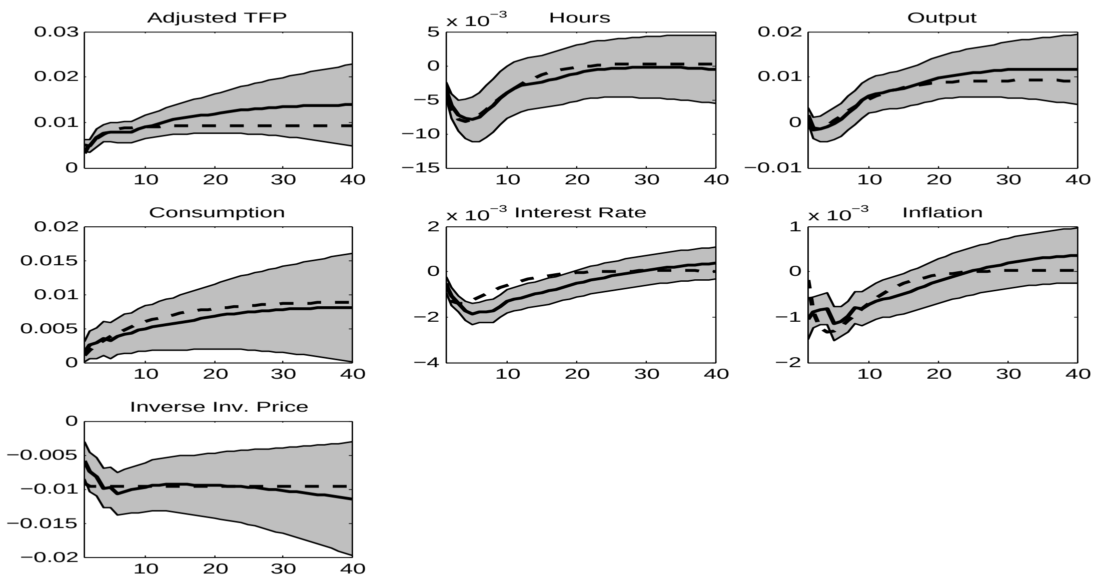

Notes: The dark lines are the impulse responses to a permanent neutral shock as estimated in the data, with the shaded gray area the +/- one standard error confidence bands. The dashed lines are theoretical impulse responses to the shock at the estimated parameter values.

Figure 15: Estimated Model vs. Data Responses: Investment Specific Shock

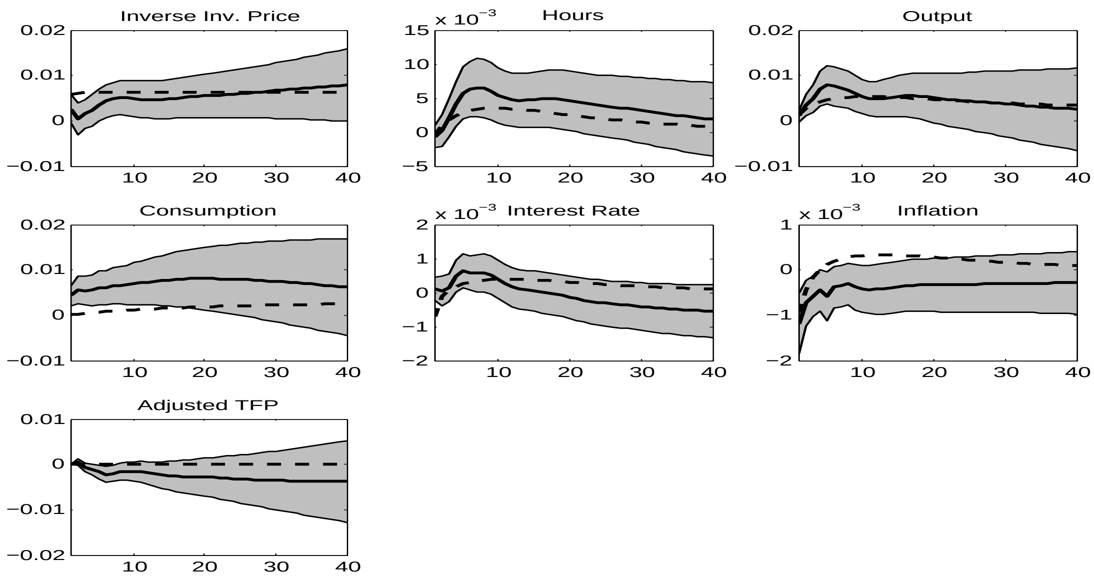

Notes: The dark lines are the impulse responses to an investment specific shock as estimated in the data, with the shaded gray area the +/- one standard error confidence bands. The dashed lines are theoretical impulse responses to the shock at the estimated parameter values.

Figure 16: Impact Response of Hours to Temporary Neutral Shock

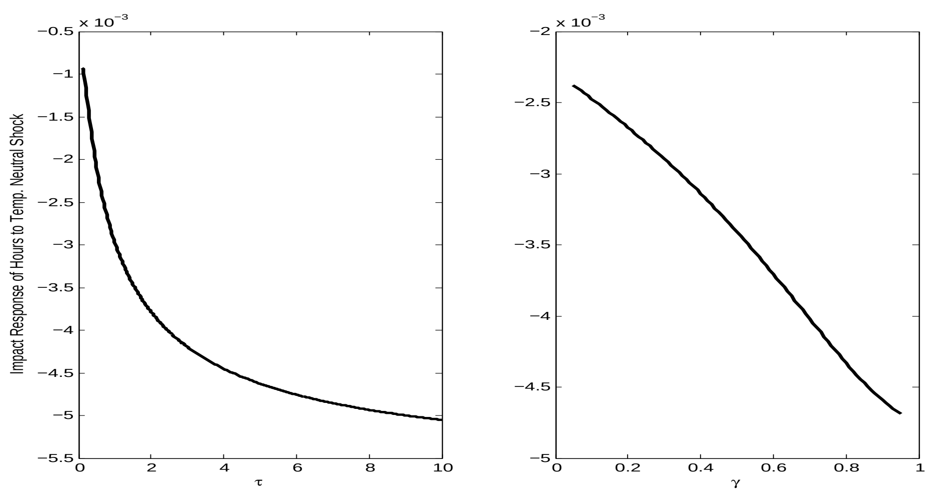

Notes: These figures plot the impact response of hours in the model to a positive neutral technology shock for different values of τ, the parameter governing investment adjustment costs, and γ, the habit formation parameter. The remaining parameters of the model are held fixed at their "standard" values as given in Table 4.

Figure 17: Impact Responses of Hours and Inflation

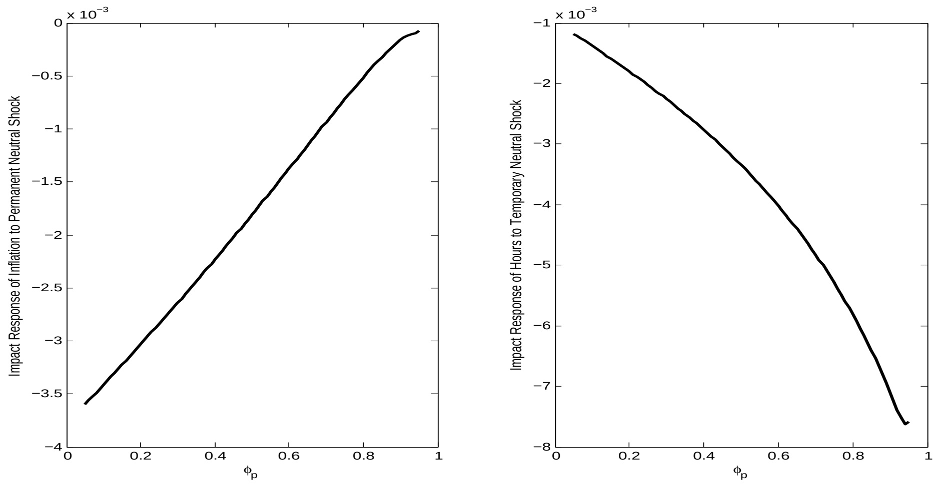

Notes: These figures plot the impact responses of inflation to a permanent neutral shock (left panel) and hours to a temporary neutral shock (right panel). The values of ϕ_p range from 0.05 to 0.95. The remaining parameters of the model are held fixed at their "standard" values as given in Table 4.

A Appendix

This Appendix lists the conditions characterizing the equilibrium of the model of Section 4.

$$\frac{\psi_t}{c_t - \gamma c_{t-1}} - E_t \frac{\beta \gamma \psi_{t+1}}{c_{t+1} - \gamma c_t} = \lambda_t \tag{35}$$

$$R_t = \Psi_0 + \Psi_1 \left(u_t - 1 \right) \tag{36}$$

$$\lambda_t = \beta E_t \lambda_{t+1} (1 + i_t)(1 + \pi_{t+1})^{-1} \tag{37}$$

$$\mu_t = \beta E_t \left(\lambda_{t+1} R_{t+1} u_{t+1} - \lambda_{t+1} \left(\Psi_0 (u_{t+1} - 1) + \frac{\Psi_1}{2} (u_{t+1} - 1)^2 \right) \chi_{t+1}^{-1} + \mu_{t+1} (1 - \delta) \right) \tag{38}$$

$$\lambda_t = \mu_t \chi_t \left(1 - \frac{\tau}{2} \left(\frac{I_t}{I_{t-1}} - \Lambda_I \right)^2 - \tau \left(\frac{I_t}{I_{t-1}} - \Lambda_I \right) \frac{I_t}{I_{t-1}} \right) + \ldots$$

$$\beta E_t \mu_{t+1} \chi_{t+1} \tau \left(\frac{I_{t+1}}{I_t} - \Lambda_I \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \tag{39}$$

$$k_{t+1} = \chi_t \left(1 - \frac{\tau}{2} \left(\frac{I_t}{I_{t-1}} - \Lambda_I \right)^2 \right) I_t + (1 - \delta) \, k_t \tag{40}$$

$$q_t = \frac{\mu_t}{\lambda_t} \tag{41}$$

$$w_t^{\#(1+\eta\xi)} = \frac{\eta}{\eta - 1} \frac{X_t}{V_t} \tag{42}$$

$$X_t = \psi_t \theta_t w_t^{\eta(1+\xi)} n_t^{1+\xi} + \beta \phi_w E_t \left(\frac{(1 + \pi_t)^{\zeta_w}}{1 + \pi_{t+1}} \right)^{-\eta(1+\xi)} X_{t+1} \tag{43}$$

$$V_t = \lambda_t w_t^{\eta} n_t + \beta \phi_w E_t \left(\frac{(1 + \pi_t)^{\zeta_w}}{1 + \pi_{t+1}} \right)^{1-\eta} V_{t+1} \tag{44}$$

$$w_t = mc_t (1 - \alpha) a_t^s a_t \left(\frac{u_t k_t}{n_t} \right)^{\alpha} \tag{45}$$

$$R_t = mc_t \alpha a_t^s a_t \left(\frac{u_t k_t}{n_t} \right)^{\alpha - 1} \tag{46}$$

$$1 + \pi_t^{\#} = (1 + \pi_t) \frac{\varepsilon}{1 - \varepsilon} \frac{A_t}{D_t} \tag{47}$$

$$A_t = \lambda_t y_t mc_t + \phi_p \beta \left((1 + \pi_t)^{\zeta_p} \right)^{1-\varepsilon} E_t \left(1 + \pi_{t+1} \right)^{\varepsilon} A_{t+1} \tag{48}$$

$$D_t = \lambda_t y_t + \phi_p \beta \left((1 + \pi_t)^{\zeta_p} \right)^{-\varepsilon} E_t \left(1 + \pi_{t+1} \right)^{\varepsilon - 1} D_{t+1} \tag{49}$$

$$1 + \pi_t = \left((1 - \phi_p) \left(1 + \pi_t^{\#} \right)^{1-\varepsilon} + \phi_p (1 + \pi_{t-1})^{\zeta_p (1-\varepsilon)} \right)^{\frac{1}{1-\varepsilon}} \tag{50}$$

$$y_t = \frac{a_t^s a_t (u_t k_t)^{\alpha} n_t^{1-\alpha}}{\nu_t} \tag{51}$$

$$\nu_t = (1 - \phi) \left(\frac{1 + \pi_t^{\#}}{1 + \pi_t} \right)^{-\varepsilon} + \phi \left(\frac{(1 + \pi_{t-1})^{\zeta_p}}{1 + \pi_t} \right)^{-\varepsilon} \nu_{t-1} \tag{52}$$

$$i_t = (1-\rho_i)i^* + \rho_i i_{t-1} + (1-\rho_i)\varphi_\pi(\pi_t - \pi^*) + (1-\rho_i)\varphi_y\left(\frac{y_t}{y_{t-1}} - \Lambda_y\right) + \varepsilon_{i,t} \tag{53}$$

$$y_t = c_t + I_t + g_t + \left(\Psi_0(u_t - 1) + \frac{\Psi_1}{2}(u_t-1)^2\right)\frac{k_t}{\chi_t} \tag{54}$$

$$g_t = \omega_t y_t \tag{55}$$

$$\ln a_t^s = \rho_{a^s} \ln a_{t-1}^s + \varepsilon_{a^s,t} \tag{56}$$

$$\Delta \ln a_t^p = (1-\rho_{a^p})g_a + \rho_{a^p}\Delta \ln a_{t-1}^p + \varepsilon_{a^p,t} \tag{57}$$

$$\Delta \ln \chi_t = (1-\rho_\chi)g_\chi + \rho_\chi \Delta \ln \chi_{t-1} + \varepsilon_{\chi,t} \tag{58}$$

$$\ln \theta_t = (1-\rho_\theta)\theta^* + \rho_\theta \ln \theta_{t-1} + \varepsilon_{\theta,t} \tag{59}$$

$$\ln \psi_t = \rho_\psi \ln \psi_{t-1} + \varepsilon_{\psi,t} \tag{60}$$

$$\omega_t = (1-\rho_g)\omega^* + \rho_g \omega_{t-1} + \varepsilon_{g,t} \tag{61}$$

Equation (35) defines the marginal utility of income, where λ_t is the multiplier on the household budget constraint. (36) is the first order condition for the choice of capital utilization. (37) is the first order condition for bonds and implicity defines the Fisher relationship. (38) is the first order condition for the choice of the future capital stock, while (39) is the optimality condition for current investment. (40) is the capital accumulation equation. (41) defines marginal q, which is the price of capital in terms of consumption goods. (42)-(44) characterize the labor market, with (42) defining the optimal reset wage for updating households. (45)-(46) are the first order conditions for cost-minimization by intermediate goods firms, and implicitly define real marginal cost in terms of real factor prices. (47)-(49) characterize optimal price-setting by intermediate goods firms, and take a form similar to (42)-(44), with $\pi_t^{\#}$ denoting reset price inflation for updating firms. (51) is the aggregate production function, with ν_t a distortion term related to price dispersion. It can be written in recursive form as in (52). (53) is the nominal interest rate rule. (54) is the aggregate accounting identity, and (55) is the government spending process. (56)-(61) are the exogenous processes.

Both a_t^p and χ_t^p are trending. Many of the endogenous variables of the model will inherit that trend. As in the simple growth model of Section 2, it is straightforward to verify that dividing many of the variables by $a_t^{p\frac{1}{1-\alpha}}\chi_t^{\frac{\alpha}{1-\alpha}}$ will render them stationary. The endogenous variables for which this transformation renders them stationary are: y_t, c_t, I_t, g_t, w_t, and $w_t^{\#}$. k_t divided by $a_t^{p\frac{1}{1-\alpha}}\chi_t^{\frac{1}{1-\alpha}}$ will be stationary. R_t and q_t are stationarized by multiplying them by χ_t, while λ_t is stationarized by multiplying, instead of dividing, by $a_t^{p\frac{1}{1-\alpha}}\chi_t^{\frac{\alpha}{1-\alpha}}$. μ_t is stationarized by multiplying by $a_t^{p\frac{1}{1-\alpha}}\chi_t^{\frac{1}{1-\alpha}}$. Hours worked, the interest rate, inflation, utilization, and the price dispersion parameter are stationary. Setting Ψ_0 equal to the steady state of the transformed marginal product of capital imposes the normalization that steady state utilization is equal to 1. Finally, the trend growth rates of output and investment are equal: $\Delta_I = \Delta_y = \exp(g_a)^{\frac{1}{1-\alpha}}\exp(g_\chi)^{\frac{\alpha}{1-\alpha}}$.