

Can News About the Future Drive the Business Cycle?*

Nir Jaimovich[†] and Sergio Rebelo[‡]

September 2006

Abstract

We propose a model that generates an economic expansion in response to good news about future total factor productivity (TFP) or investment-specific technical change. The model has three key elements: variable capital utilization, adjustment costs to investment, and preferences that exhibit a weak short-run wealth effect on the labor supply. These preferences nest the two classes of utility functions most widely used in the business cycle literature as special cases. Our model can generate recessions that resemble those of the post-war U.S. economy without relying on negative productivity shocks. The recessions are caused not by contemporaneous negative shocks but rather by lackluster news about future TFP or investment-specific technical change.

J.E.L. Classification: E3.

*We thank Gadi Barlevy, Paul Beaudry, Larry Christiano, Fabrice Collard, Wouter Denhaan, Martin Eichenbaum, Zvi Hercowitz, Erin Hoge, Navin Kartik, Benjamin Malin, Franck Portier, and Laura Veldkamp for their comments.

[†]Stanford University.

[‡]Northwestern University, NBER and CEPR.

1. Introduction

There is an old literature, including work by Beveridge (1909), Pigou (1927), and Clark (1934), that proposes news about the future or changes in agents' expectations as important sources of business fluctuations. There is a revival of interest in this idea motivated in part by the investment boom of the late 1990s and the subsequent economic slowdown.

It is easy to tell a story in which high expectations about the prospects of new technologies, such as the internet, lead to high levels of investment and to an economic boom. When the new technologies fail to live up to what was expected, investment falls and a recession ensues. However, it is surprisingly difficult to make this story work in a standard business cycle model. Cochrane (1994), Danthine, Donaldson, and Johnsen (1998), and Beaudry and Portier (2004, 2005) find that many variants of the neoclassical growth model fail to generate a boom in response to expectations of higher future total factor productivity (TFP). When agents receive news that future TFP will be higher than previously expected, consumption rises, but output, investment, and hours worked fall. Good news about tomorrow generates a recession today! Barro and King (1984) anticipated this problem when they wrote “With a simple one-capital-good technology, no combination of income effects and shifts to the perceived profitability of investment will yield positive comovement of output, employment, investment, and consumption. Therefore, [...] changed beliefs about the future cannot be used to generate empirically recognizable business cycles.”

We propose a model that generates comovement not only in response to shocks to current fundamentals but also in response to news shocks. These shocks consist of information that is useful for predicting future fundamentals but does not affect current fundamentals.

There is ample evidence that economic agents receive and process news about the future. For example, agents receive advance information about future changes in TFP that are driven by new technologies, because it takes time for these technologies to diffuse throughout the economy (Rotemberg (2003) and Alexopoulos (2004)). Stock prices and consumer confidence, which naturally reflect agents' expectations about the future, lead the business cycle (Stock and Watson (1999)). Innovations to stock prices that are orthogonal to current TFP growth are correlated with future TFP growth (Beaudry and Portier (2006)). Future political events, such as the outcome of elections, affect investment flows (Bussie and Mulder (2000)).

Our model introduces three elements into the neoclassical growth model. The first element, variable capital utilization, increases output's response to news about the future. The second element, adjustment costs to investment or capital utilization, provides agents with an incentive to respond immediately to future technical progress.¹ The third element, preferences that exhibit a weak short-run wealth effect on the labor supply, allows hours worked to rise in response to positive news. We propose a class of preferences with this property that nests, as special cases, the two classes of utility functions most widely used in the business cycle literature; these are characterized in King, Plosser, and Rebelo (1988) and in Greenwood, Hercowitz, and Huffman (1988). We show that our model produces an expansion in response to positive news about future productivity for a wide range of parameter values, as long as the short-run wealth effects on the labor supply are small.

We use the model to illustrate how downward revisions to expectations about future technical progress can generate recessions. In these experiments, the econ-

¹The first two elements, variable capital utilization and adjustment costs to investment, are generally necessary to generate comovement in response to investment-specific technical change, see Greenwood, Hercowitz, and Krusell (2000).

omy appears to be too volatile because there are no contemporaneous fundamentals, other than news about the future, that can account for changes in output. We also consider a setting in which the impact of new technologies is uncertain. Agents form priors about the impact of new technologies and update these priors in a Bayesian manner. Optimistic priors generate an economic boom, but this boom carries with it the seeds of a future recession. As agents learn that the technology is not as promising as previously thought, investment falls, and the economy slips into a recession.

We simulate a version of our model driven by investment-specific technical change. We introduce news about the future by giving the representative agent output forecasts. He combines these forecasts, taking into account their precision, with current fundamentals in order to forecast future fundamentals. We choose the information content of the output forecasts received by the representative agent so as to match the predictive content of the Livingston survey six-months-ahead real-GDP forecasts.

We find that introducing news about the future preserves the model's ability to generate volatility, comovement, and persistence of macroeconomic aggregates that are empirically plausible. The model also can produce declines in the level of output, even though the rate of technical progress is always positive. In addition, the average recession and expansion generated by the model are similar to those in the post-war U.S. economy. Introducing news into the model yields a novel property. Recessions do not always coincide with bad current fundamentals. They can occur in response to lackluster news about future fundamentals.

There has been a decline in the volatility of business cycles and an increase in the persistence of output (see, for example, Stock and Watson (1999)). Our model is consistent with this secular change in business cycle characteristics under the assumption that there has been a secular increase in the quantity or quality

of news that is relevant for predicting the future. This increase in information about the future reduces the volatility and increases the persistence of output in our model.

We discuss two variants of the model that also generate comovement in response to news shocks. In the first variant, adjustment costs to investment are replaced by adjustment costs to utilization. In the second variant, we introduce labor adjustment costs.

Our work is related to several recent papers on the role of news and expectations as drivers of business cycles. Beaudry and Portier (2004) propose the first model that produces an expansion in response to news of high future TFP. Their model features two complementary consumption goods, one durable and one non-durable. Both of these goods are produced with labor and a fixed production factor. Christiano, Motto, and Rostagno (2005) show that habit persistence and investment adjustment costs produce comovement in consumption, employment, and investment in response to news about a future TFP shock. In their model, intertemporal substitution in the supply of labor is large enough to compensate for the negative wealth effect on labor of the news shock. One implication of their model is that hours worked fall when the shock materializes. This decline reflects the ongoing negative wealth effect on labor supply and the absence of a strong intertemporal substitution effect in the period when the shock materializes. Denhaan and Kaltenbrunner (2005) study the effects of news in a matching model. Since matching frictions are a form of adjustment costs, their model is related to the version of our model with adjustment costs to labor. Lorenzoni (2005) studies the case in which consumers have imperfect information about the level of aggregate productivity.

Our paper is organized as follows. In Section 2 we compare the response to news about future TFP or investment-specific technical change in both our model

and in variants of the one-sector neoclassical model. In Section 3 we explore the role that capital utilization, adjustment costs, and preferences play in our results. In Section 4 we discuss the robustness of our results by characterizing the range of parameters that produce an expansion today in response to news of higher future TFP or investment-specific technical change. We also explore versions of our model that incorporate adjustment costs to labor and capacity utilization. In Section 5 we study the model's response to news shocks under alternative information structures. We consider noisy news, news revisions, and Bayesian updating of beliefs about the future. In Section 6 we study simulations of a version of our model with investment-specific technological progress in which agents receive forecasts about future output growth. Section 7 concludes.

2. Our Model

Our model economy is populated by identical agents who maximize their lifetime utility (U) defined over sequences of consumption (C_t) and hours worked (N_t):

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \frac{(C_t - \psi N_t^\theta X_t)^{1-\sigma} - 1}{1-\sigma}, \quad (2.1)$$

where

$$X_t = C_t^\gamma X_{t-1}^{1-\gamma},$$

and E_0 denotes the expectation conditional on the information available at time zero. We assume that $0 < \beta < 1$, $\theta > 1$, $\psi > 0$, and $\sigma > 0$. The presence of the variable X_t implies that preferences are time non-separable in consumption and hours worked. These preferences nest as special cases the two classes of utility functions most widely used in the business cycle literature. When $\gamma = 1$ we obtain preferences in the class discussed in King, Plosser, and Rebelo (1988), which we refer to as KPR. When $\gamma = 0$ we obtain the preferences proposed by Greenwood,

Hercowitz, and Huffman (1988), which we refer to as GHH. We discuss below the properties of (2.1) that allow the model to produce comovement in response to news shocks.

Output (Y_t) is produced with a Cobb-Douglas production function using capital services and labor:

$$Y_t = A_t (u_t K_t)^{1-\alpha} N_t^\alpha. \quad (2.2)$$

Here A_t represents the level of TFP. Capital services are equal to the product of the stock of capital (K_t) and the rate of capital utilization (u_t). Output can be used for consumption or investment (I_t),

$$Y_t = C_t + I_t/z_t, \quad (2.3)$$

where z_t represents the current state of technology for producing capital goods. We interpret an increase in z_t as resulting from investment-specific technological progress, as in Greenwood, Hercowitz, and Krusell (2000). Combining (2.2) and (2.3) we obtain:

$$A_t (u_t K_t)^{1-\alpha} N_t^\alpha = C_t + I_t/z_t. \quad (2.4)$$

Capital accumulation is given by,

$$K_{t+1} = I_t \left[1 - \phi \left(\frac{I_t}{I_{t-1}} \right) \right] + [1 - \delta(u_t)] K_t. \quad (2.5)$$

The function $\phi(\cdot)$ represents adjustment costs that are incurred when the level of investment changes over time. We assume that $\phi(1) = 0$, $\phi'(1) = 0$, so there are no adjustment costs in the steady state, and that $\phi''(1) > 0$. This adjustment cost formulation is proposed in Christiano, Eichenbaum, and Evans (2004). They argue that this form of adjustment costs is better at mimicking the response of investment to a monetary shock than the specifications in Lucas and Prescott (1971), Abel and Blanchard (1983), and Hayashi (1982). Lucca (2006) shows

that, for an appropriate choice of the parameter values, the linearized investment first-order condition is identical when adjustment costs take the form (2.5) and when there is time-to-build in investment.

The function $\delta(u_t)$ represents the rate of capital depreciation. We assume that depreciation is convex in the rate of utilization: $\delta'(u_t) > 0$, $\delta''(u_t) \geq 0$. The initial conditions of the model are $K_0 > 0$, I_{-1} , and $X_{-1} > 0$.

Parameter Values We solve the model by linearizing the equations that characterize the planner's problem around the steady state. We choose the following parameter values for our benchmark model. We set $\sigma = 1$, which corresponds to the case of logarithmic utility. We set θ set to 1.4, which corresponds to an elasticity of labor supply of 2.5 when preferences take the GHH form. We set the discount factor β to 0.985, implying a quarterly steady state real interest rate of 1.5 percent. The share of labor in the production function, α , is set to 0.64. We set the value of γ to 0.001, so preferences are close to a GHH specification. We choose the second derivative of the adjustment cost function evaluated at the steady state, $\phi''(1)$, to equal 1.3. Finally, we choose the second derivative of the adjustment-cost function ($\delta''(u)$) evaluated at the steady-state level of utilization, u , to equal 0.15. This value influences the degree of shock amplification present in the economy. When $\delta''(u)$ is high, the cost of utilization rises rapidly with the level of utilization. In this case, the rate of capital utilization is stable and the degree of shock amplification is small. When $\delta''(u)$ is zero, utilization costs are constant. In this case, the level of capital utilization is highly responsive to shocks, resulting in a powerful amplification mechanism. Since there is little guidance in the literature about appropriate values for $\phi''(1)$ and $\delta''(u)$, we discuss the robustness of our results to these parameters in Section 4.

Responding to News about the Future We illustrate the response of our model to news shocks with what we refer to as the baseline experiment. At time one, the economy is in a steady state with no technical progress. At time two, unanticipated news arrives. Agents learn that there will be a one-percent permanent increase in TFP beginning two periods later, in period four. Figure 1 depicts the response of the economy to this news. There is an expansion in periods two and three in response to positive news about future productivity. Consumption, investment, output, hours worked, average labor productivity, and capital utilization all rise in periods two and three even though the positive shock only occurs period four. Figure 2 shows the response to a version of the same experiment in which there is a future increase in z_t rather than in A_t .² Again consumption, investment, output, hours worked, average labor productivity, and capital utilization all rise before the technology shock materializes. With TFP shocks, the impact of news about future TFP is less important than the realization of those shocks. In contrast, with investment-specific technical change, most of the rise in output occurs in period two, when the news arrives, not in period four, when the investment-specific technical progress occurs.

One-sector Neoclassical Model Next we consider the response to news about future TFP in the standard one-sector neoclassical growth model. This model is a special case of our model with no adjustment costs ($\phi(x) = 0$ for all x), no variable capital utilization ($u_t = 1$), and $\gamma = 1$ (KPR preferences). The economy's

²Beaudry and Portier (2005) provide a useful characterization of the class of models that cannot generate an expansion today in response to future positive news. They emphasize that one-sector models with investment adjustment costs and variable capital utilization still fail to generate this type of expansion. Our model succeeds, despite its one-sector nature, because it embodies preferences and investment adjustment costs that are outside the class considered by Beaudry and Portier (2005).

technology is described by:

$$A_t K_t^{1-\alpha} N_t^\alpha = C_t + I_t, \quad (2.6)$$

$$K_{t+1} = I_t + (1 - \delta)K_t. \quad (2.7)$$

Figure 3 shows the response of a standard real-business-cycle model to the baseline experiment. Both hours worked and output fall at time two. This fall is driven by the properties of KPR preferences. These preferences imply that it is optimal to work a constant number of hours in a steady state in which the real wage rate grows at a constant rate. This property requires that the wealth and substitution effect of a permanent increase in the real wage rate be identical. Unfortunately, this property also implies that positive news about future TFP or investment-specific technical change reduce today's supply of labor. Positive news makes agents wealthier. Wealthier agents want to enjoy more leisure, so they reduce their labor supply. Since wages go up in the future but not in the present, there is no substitution effect today to counteract the wealth effect generated by positive news. As a result, today's labor supply falls, causing a drop in the level of output. At the same time, the positive wealth effect of the news shock drives consumption up. Agents feel wealthier, so they want to consume more at all future dates. Since consumption rises and output falls, investment must drop.³ The property that good news about the future fails to generate comovement holds for many versions of the RBC model, including versions with investment-specific shocks, capital utilization, and adjustment costs to investment.

Figure 4 shows the response of the same real-business-cycle model with GHH preferences ($\gamma = 0$) to our baseline experiment. With GHH preferences the optimal number of hours worked depends only on the contemporaneous real wage,

³A model with high intertemporal elasticity of substitution in consumption can generate a rise in investment, but this rise is associated with a fall in consumption.

which is equal to the marginal product of labor:

$$\theta\psi N_t^{\theta-1} = \alpha A_t K_t^{1-\alpha} N_t^{\alpha-1}.$$

News that wages are higher in the future does not depress the labor supply today through a wealth effect. This property makes it easier to obtain an expansion today in response to positive news about tomorrow. However, GHH preferences alone cannot generate an expansion in response to news about higher future values of A_t or z_t . Hours remain roughly constant, therefore output remains constant. The positive wealth effect dictates a decline in the marginal utility of consumption and a rise in the level of consumption. This rise in consumption implies a fall in investment.

3. The Elements of Our Model

Here we discuss the importance of the three elements that generate comovement between consumption, investment, output, and labor in response to news about the future TFP level, A_t , or the technology used for producing capital goods, z_t . In discussing the role of capital utilization and adjustment costs to investment it is useful to consider a version of the model with GHH preferences by setting γ to zero. In this case X_t is constant so, to simplify, we normalize the level of X to one. The first-order conditions for the planner's problem are:

$$(C_t - \psi N_t^\theta)^{-\sigma} = \lambda_t, \tag{3.1}$$

$$\theta\psi N_t^{\theta-1} = \alpha A_t (u_t K_t)^{1-\alpha} N_t^{\alpha-1}, \tag{3.2}$$

$$\lambda_t(1 - \alpha)A_t u_t^{-\alpha} K_t^{1-\alpha} N_t^\alpha = \eta_t \delta'(u_t) K_t, \tag{3.3}$$

$$\eta_t = \beta[\lambda_{t+1}(1 - \alpha)A_{t+1} u_{t+1}^{1-\alpha} K_{t+1}^{-\alpha} N_{t+1}^\alpha + \eta_{t+1}[1 - \delta(u_{t+1})]], \tag{3.4}$$

$$\lambda_t/z_t = \eta_t \left[1 - \phi \left(\frac{I_t}{I_{t-1}} \right) - \phi' \left(\frac{I_t}{I_{t-1}} \right) \frac{I_t}{I_{t-1}} \right] + E_t \left[\beta \eta_{t+1} \phi' \left(\frac{I_{t+1}}{I_t} \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \right], \quad (3.5)$$

where λ_t and η_t are the Lagrange multipliers associated with (2.4) and (2.5), respectively.

Investment Adjustment Costs The first-order condition for labor, (3.2), implies that, unless the rate of capital utilization changes, N_t does not respond to news about the future. The first-order condition for capital utilization, (3.3), implies that λ_t/η_t must increase in order for u_t to rise. A rise in λ_t/η_t requires the presence of adjustment costs to investment. Without adjustment costs, $\lambda_t z_t = \eta_t$ and the capital utilization equation reduces to:

$$(1 - \alpha) A_t u_t^{-\alpha} K_t^{1-\alpha} N_t^\alpha = z_t \delta'(u_t) K_t.$$

Since z_t and A_t both remain constant at time two, this equation along with (3.2) implies that both N_t and u_t remain constant. The dashed line in Figure 5 shows the response of a version of our model without investment adjustment costs to our baseline experiment. Investment initially falls and only rises in period four when TFP rises.

One potential objection to this model is that it requires a decline in Tobin's marginal q , λ_t/η_t . This decline can seem inconsistent with the stock market booms that often are associated with the advent of new technologies. We think this issue is not central for three reasons. First, we have no empirical evidence on marginal q . Marginal and average q are not proportional in our model, so we cannot use evidence on average q to infer the behavior of marginal q . Second, in Section 4 we discuss a version of the model that replaces adjustment costs in investment with adjustment costs in capital utilization. This formulation can generate an expansion in response to positive news about the future without a fall in marginal

q .⁴ Third, we show in Jaimovich and Rebelo (2006) that even in versions of the model in which marginal q falls, this fall does not necessarily imply a decline in the stock market value of firms.

Variable Capital Utilization To explain the role played by capital utilization, we consider a version of the model with constant utilization. To obtain the planner's first-order conditions for this model, we ignore the first-order condition for u_t , (3.3), and set $u_t = 1$ in equation (3.2):

$$\theta\psi N_t^{\theta-1} = \alpha A_t K_t^{1-\alpha} N_t^{\alpha-1}. \quad (3.6)$$

This equation implies that N_t does not respond to news about future changes in A_t or z_t . The positive wealth effect of future shocks reduces the marginal utility of consumption today, λ_t . Equation (3.1) implies that C_t rises. When $u_t = 1$, equation (2.4) implies that investment must fall. Therefore, labor and output do not respond to the news shock, consumption rises, and investment falls. These patterns are visible in the dotted line in Figure 5 which depicts the response of a version of our model with constant capital utilization to the baseline experiment.

Preferences The second dashed line in Figure 5 shows the response of a version of our model with KPR preferences to the baseline experiment. Both hours worked and investment fall in response to news of higher TFP.

To isolate the role of preferences in generating the response to future news, we consider a version of the neoclassical growth model in which lifetime utility is given by (2.1) and technology by (2.6) and (2.7). We use this model to study the following simple experiment. At time zero, the economy is in the steady state.

⁴Christiano, Motto, and Rostagno (2005) propose an alternative way of circumventing the fall in marginal q . They consider a monetary model with sticky prices and wages in which the monetary authority follows a Taylor rule.

At time one, there is an unanticipated, permanent increase in TFP. The first panel of Figure 6 shows the response to this shock for three different values of γ . The strongest response of hours worked occurs with GHH preferences ($\gamma = 0$). However, in this case hours worked are not stationary; they rise permanently in response to the permanent increase in the real wage rate driven by the TFP shock.⁵ With KPR preferences, hours worked converge back to the steady state after the shock, but the short-run response of hours worked is weak. The third line in Figure 6 shows the response of hours worked when preferences are of the form in (2.1) and $\gamma = 0.25$. With these preferences, hours worked also converge to the steady state, but the short-run impact of the TFP shock falls in between that of GHH and KPR preferences. Lower (higher) values of γ produce short-run responses that are closer to those obtained with GHH (KPR) preferences. As long as $0 < \gamma \leq 1$, hours worked converge to the steady state.

To understand the response of hours worked for the three preference specifications it is useful to isolate the Hicksian wealth effect on the labor supply of the permanent TFP shock. We follow King (1991), who discusses a dynamic version of the Hicks decomposition into wealth and substitution effects. The permanent TFP shock raises lifetime utility from U to U^* . To calculate the wealth effect, we compute the path for labor supply of a household that receives an output transfer and faces wages and real interest rates that are constant at their steady-state levels. We compute the output transfer so that the agent's utility is U^* (without the transfer the agent's utility would be U). We compute U , U^* , and the output transfer for KPR, GHH, and our preference specification. The results are depicted

⁵A simple way to make hours stationary is to introduce a trend in the utility function such that the utility cost of supplying labor increases at the same rate as the real wage. This trend can be justified by appealing to home production. We found that in models with stochastic technical progress this formulation can generate large recessions through an implausible mechanism. In periods with low rates of technical progress, hours worked can fall significantly because the trend increase in the utility cost of supplying labor is not offset by increases in the real wage rate.

in panel 2 of Figure 6. The wealth effect on the labor supply is zero for GHH preferences and negative for KPR. In both cases the wealth effect is constant over time. With our preferences, the wealth effect is time varying. In the long run, the wealth effect is similar to that with KPR preferences. In the short-run, this effect is actually negative because consumption grows over time. This growth implies that the disutility of work is higher in the future than in the present (see (2.1)).

4. Robustness

In our model, news about future increases in z_t or A_t generate an expansion before the rise in z_t or A_t occurs. Table 1 illustrates the range of parameters for which our model generates an expansion in response to our baseline experiment. The first column corresponds to our benchmark calibration. Keeping all the other parameters the same, any value of $\sigma > 0.05$ produces an expansion. There is also a wide range of values for adjustment costs ($\phi''(1) > 0.51$), elasticity of labor supply ($1/(\theta - 1) > 0.3$), and elasticity of utilization ($\delta''(u)u/\delta'(u) < 2.85$), all consistent with an expansion in response to future positive news. The critical parameter is γ . We need $\gamma < 0.001$, that is preferences must be close to the GHH form. Otherwise the wealth effect of positive news about the future reduces hours worked today and generates a recession. The remaining columns of Table 1 report robustness results for three additional model parameterizations: infinite labor supply elasticity, high adjustment costs, and a high elasticity of utilization. When the elasticity of labor supply is high we need lower investment adjustment costs. In this case, the labor response generates enough additional output so that consumption can rise without causing investment to fall. With high adjustment costs, we need a low elasticity of labor supply ($1/(\theta - 1) > 0.003$) and a more moderate value of γ ($\gamma < 0.02$).

Table 1 shows that our model can generate an expansion in response to good

news about future productivity for many different parameter configurations. Figure 7 shows that this property also holds for different information leads. This figure shows the immediate response of output to unanticipated news that there will be a permanent increase in TFP (solid line) or in z (dashed line) in period $t+n$. News about events farther into the future (higher values of n) have a smaller impact on output today. However, the rate of decay of the strength of the immediate impact with respect to n is relatively small. News about events that will occur in ten periods still have an impact on today's output level.

The form of investment adjustment costs present in equation (2.5) is important in generating our results. We find that our model cannot generate an expansion in response to news about higher future technical progress when adjustment costs take the form proposed in Hayashi (1982):

$$K_{t+1} = \phi\left(\frac{I_t}{K_t}\right) K_t + [1 - \delta(u_t)]K_t.$$

Adjustment Costs to Capital Utilization An alternative to adjustment costs in investment are adjustment costs to capital utilization. These costs can be introduced by replacing equation (2.4) with the following equation:

$$Y_t = C_t + I_t/z_t + \psi(u_t/u_{t-1})u_t.$$

The function $\psi(\cdot)$, which represents adjustment costs to capital utilization, is increasing and convex with $\psi(1) = 0$. Figure 8 shows the response to our baseline experiment for a version of our model with adjustment costs in utilization. In order to produce a positive response to news about higher future productivity, we have to increase the elasticity of labor supply by reducing θ from 1.4 to 1.05. Adjustment costs to utilization reduce the extent to which utilization responds on impact. This weaker response of utilization reduces the incentive for hours worked to increase, and dampens the rise in output. The smaller output expansion can

be insufficient to allow a rise in both consumption and investment. Column 4 of Table 1 reports the range of parameters consistent with an expansion driven by positive news about future productivity for a version of the model with no adjustment costs to investment and with adjustment costs to utilization. The benchmark value of θ in this model is 1.05.

Adjustment Costs to Labor Next we discuss a version of our model that incorporates adjustment costs to labor, along the lines of Sargent (1978) and Cogley and Nason (1995). The only modification we introduce to the model is the replacement of equation (2.3) with the following equation:

$$Y_t = C_t + I_t/z_t + N_t\varphi(N_t/N_{t-1}),$$

where $\varphi(\cdot)$ is a function such that $\varphi(1) = \varphi'(1) = 0$ and $\varphi'(\cdot) > 0$, $\varphi''(\cdot) > 0$.

The introduction of labor adjustment costs allows the model to generate an expansion in response to our baseline experiment for a wider range of parameters, including much higher values of γ . Adjustment costs to labor provide an incentive to increase the labor supply immediately in anticipation of future increases in the labor supply that occur in response to the shock. When adjustment costs are significant it is not efficient to reduce the labor supply today and then increase it in the future once the shock occurs. As a result, the short-run wealth effect on the labor supply can be stronger than in the benchmark model.

Column 5 of Table 1 reports the range of parameters consistent with an expansion in our baseline experiment for a version of the model with adjustment costs to labor. This column corresponds to the case where adjustment costs are very moderate ($\varphi''(1) = 0.5$).⁶ The fact that an expansion occurs for any value of

⁶To measure the impact of labor adjustment costs, we study a version of the neoclassical growth model in which lifetime utility is given by (2.1) and technology by (2.6) and (2.7). In this model, calibrated with the parameters of our benchmark model, labor rises by 0.6 percent

$\gamma < 0.4$ is particularly striking. Recall that our benchmark model features very small short-run wealth effects on the labor supply ($\gamma = 0.001$). The presence of adjustment costs to labor allows us to have much higher short-run wealth effects on the labor supply. In fact, in a version of the model with high adjustment costs to investment ($\phi''(1) = 3$) and labor ($\varphi''(1) = 5$), we obtain an expansion in response to news for preferences that are close to KPR ($\gamma \leq 0.98$).

5. News-driven Fluctuations

We now discuss two types of news-related recessions. The first type occurs when the current rate of technical progress is lower than previously expected. The second type occurs when agents revise downward their expectations about the future in response to poor news about future technical progress. In both cases there are fluctuations in output that are not explained by conventional macro fundamentals. For this reason, the model economy appears to be too volatile relative to its fundamentals. Neither type of recession emerges in standard variants of the neoclassical growth model. In this model, when fundamentals are disappointing relative to expectations, and when expectations are downgraded, there is a negative wealth effect that contributes to a rise in the supply of labor and a fall in consumption.

Realized fundamentals are worse than expected An example of the first type of recession occurs when the economy receives noisy signals about future fundamentals. In this case a recession can occur because the actual realization of the fundamentals is worse than what was expected, given the signal that had been received by agents. Figure 9 illustrates this possibility with the following

in response to a 1 percent i.i.d. TFP shock. Introducing labor adjustments costs to this model with $\varphi''(1) = 0.5$ reduces the labor response to the same shock to 0.5 percent.

experiment. At time one the economy is in a steady state with no technical progress. At time two the economy receives unanticipated news that in two periods the level of z_t will either stay the same, increase by one percent, or increase by two percent. These events occur with equal probability, so the expected change in z_t is one percent. The solid line in Figure 9 shows the time path for the economy when the realized change in z_t is equal to the expected change. In this case, the economy undergoes a smooth expansion. The dashed line shows the case where the change in z_t is two percent. In this case there is an acceleration in the rate of expansion of the economy. The dotted line shows the case where the realized change in z_t is zero percent. In that case the economy goes into a recession even though there is no realized fall in z_t . Fundamentals remain as strong as in the past, but they are lower than previously expected. The same forces that cause the economy to expand in periods two and three in anticipation of an increase in z_t in period four are set in reverse once realized fundamentals fail to live up to what was expected.

News-based downward revisions in expectations about the future One example of this type of recession is the model discussed in Section 4 with perfect signals about the future. The first type of news-related fluctuation is the one we discuss in the previous section and illustrate in Figures 1 and 2. In this case, the economy receives a perfect signal about the future. Bad (good) news about future TFP or investment-specific technical change causes a recession (expansion) before the changes in fundamentals occur.

Another example which includes news updates is illustrated in Figure 10. At time one, the economy is in a steady state with no technical progress. In period two, the economy receives the same noisy signal about z_t described above. In period three, the economy receives an update about the value of z_t in period four.

The solid line corresponds to the case where the economy learns in period three that the change in z_t at time four will coincide with the expected change. In this case the economy continues on a smooth expansion. The dashed line shows the case where the economy learns at time three that the change in z_t in period four will be higher than previously expected. This good news generates a stronger expansion even though current fundamentals have not changed. The dotted line corresponds to the case where the economy learns in period three that the change in z_t in period four will be lower than expected. This bad news plunges the economy into a recession.

Learning about the future: rational “optimism” and “pessimism” Several authors, including Cochrane (1994), Danthine, Donaldson, and Johnsen (1998), and Blanchard (2006), emphasize the potential role of changes in expectations about the future in driving economic fluctuations. Cochrane (1994) and Danthine, Donaldson, and Johnsen (1998) stress that this potential is not fulfilled in variants of the neoclassical growth model. Here we now explore the impact of rational changes in expectations. Agents make the best possible use of the available information to resolve fundamental uncertainty about the economy. Recessions can occur because, as agents observe the realization of fundamentals, they learn that the future is not as bright as expected.

For concreteness, we consider the following example. The economy is in steady state at time one. At time two, agents learn that a new wave of technology such as the internet will be available from time four onwards. Agents form priors about the effects of the new technology on the rate of change of z_t . From period four on, agents observe realizations of z_t and update their expectations in a Bayesian fashion. Suppose that the change in $\log(z_t)$ is generated by a normal i.i.d process with true mean θ and variance σ^2 . To simplify, we assume that agents know the

variance but do not know the mean. Agents form a prior about the distribution of μ . This prior is normally distributed with mean μ and variance V . The posterior distribution is a normal with mean μ_t^* ,

$$\mu_t^* = \frac{\bar{\varepsilon}_t(n/\sigma^2) + \mu/V}{n/\sigma^2 + 1/V},$$

and variance, V_t^* ,

$$V_t^* = \frac{1}{n/\sigma^2 + 1/V}.$$

The variable $\bar{\varepsilon}_t$ represents the average change in z_t in the sample up to time t , while n denotes the number of observations in this sample. We assume that the initial prior is relatively informative ($V = 10^{-5}$), i.e. agents have confidence in their beliefs. We consider three different priors. In the first case $\mu = \theta$, so expectations are ex-post “realistic.” In the second case $\mu < \theta$, so expectations are ex-post “pessimistic.” In the third case $\mu > \theta$, so agents are ex-post “optimistic.” We simulate the model 100 times for each of the three different priors that we consider. Figure 11 shows the average simulation for each of the three priors. The solid line corresponds to the case where $\theta = \mu$. While there is still some updating and resolution of uncertainty that goes on, output fluctuations are small in this case. The dashed line corresponds to the case where expectations are ex-post “pessimistic”, i. e. $\mu < \theta$. In this case, as agents update their expectations, the economy goes into an expansion. In the third case, agents are ex-post optimistic, $\mu > \theta$. Optimism about θ generates an initial expansion, but this expansion carries with it the seeds of a future recession. Agents gradually realize that they have been gearing up for an increase in the level of investment-specific technical progress that will not occur. As they lower their expectations about the future, the economy falls into a recession. This recession takes place with no changes in observed fundamentals. It does not take much imagination to see in this scenario some of the elements that may have played a role in the large boom of the late

1990s and the subsequent slowdown.

6. Model Simulations

We have shown that our model can generate expansions and contractions in response to news about the future. One natural question is whether the model, calibrated with the parameters used in the experiments discussed so far, can generate empirically recognizable business cycles. To answer this question we simulate a version of our model driven by stochastic, investment-specific technical progress to compute the standard set of business-cycle statistics.⁷

We assume that $\log(z_t)$ follows a random walk:

$$\log(z_{t+1}) = \log(z_t) + \varepsilon_{t+1}.$$

We use the method proposed by Tauchen and Hussey (1991) to estimate a two-point Markov chain for ε_t . We measure z_t using quarterly data on the U.S. real price of investment for the period 1947.I to 2004.IV. These data were constructed by Fisher (2006) using National Income and Product Accounts series for the consumption deflator and Cummins and Violante's (2003) updated series for Gordon's (1989) quality-adjusted producer durable-equipment deflator.⁸ The support of the estimated Markov chain is: $\{0.00, 0.0115\}$. The transition matrix is:

$$\pi = \begin{bmatrix} 0.7378 & 0.2622 \\ 0.2622 & 0.7378 \end{bmatrix}. \quad (6.1)$$

We generate 1000 model simulations with 230 periods each. For each simulation, we detrend the logarithm of the relevant time series with the Hodrick-Prescott filter using a smoothing parameter of 1600.⁹

⁷Fisher (2006) and Justiniano and Primiceri (2005) argue that investment-specific technical progress is the most important determinant of output variability.

⁸We thank Ricardo Di Ceccio for providing us with an updated version of this time series.

⁹We verify that hours worked are stationary in the simulated data.

We consider versions of the model in which agents receive news about the future with different degrees of precision. In our main calibration we consider as our measure of news the Livingston survey of output forecasts. The Livingston survey pools professional forecasters to obtain forecasts of different economic variables. Two-quarter ahead GDP forecasts are available for the period 1971:IV – 2003:IV.¹⁰

It is difficult to choose the information lead, n , with which agents receive news about the future. We set $n = 2$ motivated by the observation that output leads investment-specific technical progress by two quarters (see Fisher (2006)) and by the fact that the Livingston survey output forecasts, which we use below to calibrate the information content of the signal, are available for a six-month horizon. Our results are generally robust to other values of n .

Perfect Signal Column 7 of Table 2 summarizes the business cycle properties of our model when at time t agents receive perfect signals about ε_{t+2} , the growth rate of z_t in two periods. This model generates business cycle moments that are similar to those in postwar U.S. data reported in column 1. Consumption, investment, and hours worked are procyclical. Investment is more volatile than output, consumption is less volatile than output, and the volatility of hours is similar to that of output. Output volatility in the model is 64 percent of that in the data.

No Signal Column 4 in Table 2 summarizes the business cycle properties of a version of our model in which the economy receives no news. Forecasts of future values of ε_t are solely based on the Markov chain (6.1). This version of the model generates patterns of volatility and comovement that are similar to those of the model with perfect signals. The main difference between the two models is in

¹⁰See Croushore (1993) for a description of the Livingston survey.

the level of volatility generated and in the persistence of output movements. The economy without news shocks is more volatile than the one with news shocks. News shocks make it easier to forecast the future, which reduces cyclical volatility and make output more persistent. Columns 4 and 6 show that our model is robust to changes in the information structure. Providing the economy with news about the future does not alter the baseline patterns of comovement or relative volatility of the major macroeconomic aggregates.

Noisy News Next we consider two settings in which agents receive noisy news about the future. In our first setting, agents receive a signal about the value of ε_{t+2} at time t . The signal can be high (H) or low (L). The signal's precision, p_i , is the probability that ε_{t+2} will be high (low) given that the signal is high (low):

$$p_i = \Pr(S = i | \varepsilon_{t+2} = i), \quad i = H, L.$$

The precision of the signal can be different in the two states of nature. Column 5 of Table 2 reports statistics for a version of the model in which agents receive a signal that has precision 0.8 in both states. Here the main result is that the volatility of output falls between the case of the perfect signal and the case in which there is no signal or the signal is uninformative.

Forecasts of future rates of investment-specific technical change are not available for our sample, so it is difficult to choose the precision of signals about ε_{t+2} . For this reason, we consider a second setting in which we provide agents with a signal, S^y , for whether the growth rate of output two periods later is going to be above or below the average. The signal has two values, high (H) or low (L). We choose the signal to have the same precision as the Livingston survey of output

forecasts. The precision of these forecasts is as follows:¹¹

$$\begin{aligned}\Pr(g_{t+2}^y \geq \text{Average}(g^y)|S^y = H) &= 0.70, \\ \Pr(g_{t+2}^y < \text{Average}(g^y)|S^l = L) &= 0.58,\end{aligned}\tag{6.2}$$

where g_{t+2}^y represents the growth rate of output at time $t+2$. The forecast precision is higher in expansions than in recessions.¹²

To provide agents in the model with a signal on output with the same precision as the Livingston survey forecast, we implemented the following algorithm. First, we assumed values q_1 and q_2 for the following conditional probabilities,

$$\begin{aligned}\Pr(S^y = H|\varepsilon_{t+2} = H) &= q_1, \\ \Pr(S^y = L|\varepsilon_{t+2} = L) &= q_2.\end{aligned}$$

We simulate time series for ε_t and generate S^y according to q_1 and q_2 . Agents receive these signals and forecast ε_{t+2} using both the signal and the current realization of ε_t :

$$\Pr(\varepsilon_{t+2} = H|S^y = i, \varepsilon_t) = \frac{\Pr(S^y = i|\varepsilon_{t+2} = H) \Pr(\varepsilon_{t+2} = H|\varepsilon_t)}{\sum_{j=H,L} \Pr(S^y = H|\varepsilon_{t+2} = j) \Pr(\varepsilon_{t+2} = j|\varepsilon_t)}.$$

We simulate the model and compute:

$$\begin{aligned}\Pr(g_{t+2}^y \geq \text{Average}(g^y)|S^y = H), \\ \Pr(g_{t+2}^y < \text{Average}(g^y)|S^l = L).\end{aligned}$$

We then revise the values of q_1 and q_2 until the precision of S^y in the model coincides with the precision (6.2) estimated in the data. We obtain $q_1 = 0.99$ and

¹¹To obtain a discrete signal with two possible values we use the Tauchen and Hussey (1991) method to estimate a two-point Markov chain for the Livingston survey forecasts.

¹²Using the Survey of Professional Forecasters, Van Nieuwerburgh and Veldkamp (2006) also find that forecast precision is higher in expansions than in recessions.

$q_2 = 0.62$. Column 6 of Table 2 shows the results for this version of the model. The main result is that the volatility of output falls between the two extremes of uninformative signals and perfect signals.¹³

News and Volatility It is well-known that output volatility has declined over the past 60 years in virtually all developed countries. At the same time the persistence of output has increased. These facts are documented for the U.S. in Table 2. This table reports moments for the main macroeconomic time series detrended with the HP filter with a smoothing parameter of 1600. Columns 2 and 3 provide statistics for the U.S. for the period 1947-1982 and 1983-2003. The volatility of output declines from 1.88 in the first sample to 0.97 in the second sample. The persistence of output, as measured by the sum of the four estimated coefficients in an AR(4) process for output, rises from 0.65 to 0.86.

Stock and Watson (2003) document both the reduction in output volatility and the increase in persistence for the G7 countries and discuss several possible explanations, including better monetary policy, changes in sectoral composition toward sectors with lower volatility, and declines in the volatility of the shocks that buffet the economy. Our model provides a complementary explanation for the decline in business cycle volatility. Advances in telecommunications and computer technology have led to dramatic increases in the volume of information available and in the ability to process this information. Let us assume that the increase in information volume has made it easier to forecast the future. Under this assumption, we can think of the increased volume of information as moving the economy

¹³The Survey of Professional Forecasters (SPF) is an alternative source of output growth forecasts for the U.S. economy. The SPF produces forecasts for each of the following four quarters. So in every period there is a four-period ahead forecast and a revision of the one-, two- and three-period ahead forecasts. This information allows us to introduce news revisions into the model. We computed the precision of the SPF forecasts and used them to calibrate our model. The results were similar to those obtained with the Livingston forecasts.

from Column 4 of Table 2 (no news) toward Column 7. The availability of news makes it easier to forecast the future, thus reducing economic volatility. Evidence from the Livingston survey is consistent with the idea that business cycles have become easier to forecast. The survey contains unemployment forecasts at a six-month horizon from the fourth quarter of 1961 to the fourth quarter of 2003. The average absolute percentage forecast error is 3.3 percent in the first part of the sample (1961:IV-1982:IV) but only 1.5 percent in the second part of the sample (1983:I-2003:IV).

Recessions According to our estimated Markov chain, (6.1), the rate of technical progress is always positive. This is a good approximation to the behavior of investment-specific technical progress in the data. Declines in z_t are rare (they occur in only 6 percent of the quarters in our sample) and small in magnitude. The average percentage decline in z_t in quarters in which z_t falls is 0.8 percent.

The absence of technical regress in our calibration raises the question of whether the model can generate recessions.¹⁴ To study this question we first describe the average recession in U.S. data. Our strategy is similar to that used by the Business Cycle Dating Committee of the National Bureau of Economic Research (NBER) for comparing different recessions (see Hall, Feldstein, Frankel, Gordon, Romer, Romer, and Zarnowitz (2003)). It is also reminiscent of the methods used by Burns and Mitchell (1946) in their study of the properties of U.S. business cycles.

To date the beginning of U.S. recessions, we compute trend output using the HP filter with a smoothing parameter of 1600. We identify periods in which output is below trend for at least two consecutive quarters, say, t and $t+1$. Recessions are dated as starting at time $t-1$. This timing method produces recession dates that

¹⁴King and Rebelo (1999) propose a real business cycle model that generates recessions in the absence of negative technology shocks. Their model shares one key feature with our model, which is variable capital utilization, but it relies on a much higher elasticity of labor supply.

are similar to those chosen by the NBER dating committee. The HP procedure produces six recessions whose starting dates coincide with those chosen by the NBER: 1948.IV, 1957.III, 1960.II, 1980.I, 1981.III, 1990.III. There are four other recessions in which the HP procedure produces recession dates that are within two quarters of the NBER dates (indicated in parentheses): 1953.III (1953.II), 1969.III (1969.IV), 1974.II (1974.III), and 2001.II (2001.I). The HP procedure identifies four additional recessions starting in 1962-II, 1967-II, 1986-III, and 1994.III. None of the latter episodes involved a fall in output, which suggests that our procedure corresponds to a broader definition of recession than that of the NBER.

Once we identify the 14 recessions in post-war U.S. data, we compute the average time series for different macroeconomic variables during recession periods. The solid line in Figure 12 shows the average behavior during recessions of the HP-detrended logarithm of real GDP, real consumption of nondurables and services, real private investment, and hours worked. Time zero is the quarter in which the recession begins. The dashed lines represent the 95 percent confidence interval around the average. The fall from peak to trough in consumption, output, investment and hours is 0.7 percent, 1.8 percent, 4.3 percent, and 1.7 percent, respectively.

The dashed line in Figure 12 shows the average recession in our model. The model captures the salient features of recessions in the data. Figure 13, which displays the behavior of investment-specific technical change in the average recession, shows an interesting feature of the recessions generated by the model. On average, recessions occur when there is a high contemporaneous rate of change in investment-specific technical progress but the economy learns that two periods later technical change will slow down. It is impossible to identify what causes recessions in our model by lining up the usual suspects—contemporaneous shocks to the economy. Recessions are driven not by bad shocks today but by lackluster

news about the future.

The model only generates nine recessions, as opposed to 14 in the data. In addition, recessions are more shallow in the model than in the data. We view our model as suggesting an additional channel through which recessions can occur, not as providing an explanation for all the recessions in the data. While we emphasize news about future investment-specific technical change, the same mechanism is likely to produce recessions in response to bad news about the future values of other fundamentals, such as tax rates and oil prices.

Figure 14 compares the average expansion in the model and in the U.S. data. It shows that the model comes close to reproducing the average expansion in U.S. data.

7. Conclusion

In this paper we propose a model that generates an expansion (recession) in response to positive (negative) news about future TFP or investment-specific technical change. The model has three key elements: variable capital utilization, adjustment costs to investment, and a new form of preferences. These preferences combine the desirable features of the specifications proposed by Greenwood, Hercowitz, and Huffman (1988) and by King, Plosser, and Rebelo (1988). Our preferences share with the Greenwood, Hercowitz, and Huffman (1988) specification the ability to generate a strong short-run response of hours worked to movements in the wage rate. They share with the King, Plosser, and Rebelo (1988) specification the ability to generate a constant supply of labor in the steady state of a model with labor-augmenting technical progress or investment-specific technical change. The version of the model with investment-specific technical change accounts for roughly 60 percent of cyclical output fluctuations in the U.S. economy. The model can generate recessions that resemble those of U.S. data despite featuring no tech-

nical regress. Recessions are caused not by contemporaneous negative shocks, but by lackluster news about the future rate of technical progress.

The introduction of news about the future reduces the volatility of output relative to a model with no news. This suggests that improvements in the quantity and quality of information that is useful for forecasting the future may have contributed to the observed secular decline in business cycle volatility.

References

- [1] Abel, Andrew and Olivier Blanchard, “An Intertemporal Equilibrium Model of Savings and Investment,” *Econometrica* 51: 675-692, 1983.
- [2] Alexopoulos, Michelle “Read All about it!! What Happens Following a Technology Shock?,” mimeo, University of Toronto, 2004.
- [3] Barro, Robert J. and Robert G. King “Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles,” *The Quarterly Journal of Economics*, 99: 817-839, 1984.
- [4] Beaudry, Paul, and Franck Portier “An Exploration into Pigou’s Theory of Cycles,” *Journal of Monetary Economics*, 51: 1183–1216, 2004.
- [5] Beaudry, Paul and Franck Portier “When Can Changes in Expectations Cause Business Cycle Fluctuations in Neo-Classical Settings?,” mimeo, University of British Columbia, 2005.
- [6] Beaudry, Paul, and Franck Portier “News, Stock Prices and Economic Fluctuations,” Working Paper 10548, National Bureau of Economic Research, forthcoming *American Economic Review*, 2006.

- [7] Beveridge, William H. “Unemployment: A Problem of Industry,” Longmans Green, London, 1909.
- [8] Blanchard, Olivier “Adjustment Within the Euro. The Difficult Case of Portugal,” mimeo, MIT, 2006.
- [9] Burns, Arthur and Wesley Mitchell *Measuring Business Cycles*, National Bureau of Economic Research, New York, 1946.
- [10] Bussie, Matthieu and Christian Mulder, “Political Instability and Economic Vulnerability,” *International Journal of Finance and Economics* 5: 309–330, 2000.
- [11] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113: 1–45, 2005
- [12] Christiano, Lawrence, Roberto Motto, and Massimo Rostagno “Monetary Policy and a Stock Market Boom-Bust Cycle,” mimeo, Northwestern University, 2005.
- [13] Clark, John Maurice “Strategic Factors in Business Cycles,” National Bureau of Economic Research, 1934.
- [14] Cochrane, John H. “Shocks,” *Carnegie-Rochester Conference Series on Public Policy*, 41: 295-364, 1994.
- [15] Cogley, Timothy and James M. Nason “Output Dynamics in Real-Business-Cycle Models,” *American Economic Review*, 85: 492-511, 1995.
- [16] Croushore, Dean “Introducing: The Survey of Professional Forecasters,” Federal Reserve Bank of Philadelphia Business Review, 3–15, 1993.

- [17] Cummins, Jason and Giovanni Violante “Investment-Specific Technical Change in the United States (1947-2000): Measurement and Macroeconomic Consequences,” *Review of Economic Dynamics*, 5: 243-284, 2002.
- [18] Danthine, Jean Pierre, John B. Donaldson and Thore Johnsen “Productivity Growth, Consumer Confidence and the Business Cycle,” *European Economic Review*, 42, 1113-1140, 1998.
- [19] Denhaan, Wouter J. and Georg Kaltenbrunner “Growth Expectations and Business Cycles,” mimeo, London Business School, 2005.
- [20] Fisher, Jonas “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 114: 413-451, 2006.
- [21] Gordon, Robert, *The Measurement of Durable Goods Prices*, Chicago: University of Chicago Press, 1989.
- [22] Greenwood, Jeremy, Zvi Hercowitz and Gregory Huffman, “Investment, Capacity Utilization and the Real Business Cycle,” *American Economic Review* 78, 402-417, 1988.
- [23] Greenwood, Jeremy Zvi Hercowitz and Per Krusell “The Role of Investment-specific Technological Change in the Business Cycle,” *European Economic Review*, 44: 91-115, 2000.
- [24] Hall, Robert, Martin Feldstein, Jeffrey Frankel, Robert Gordon, Christina Romer, David Romer, and Victor Zarnowitz “The NBER’s Recession Dating Procedure,” mimeo, NBER, 2003.
- [25] Hayashi, Fumio, “Tobin’s Marginal q and Average q: A Neoclassical Interpretation,” *Econometrica* 50: 213–24, 1982.

- [26] Jaimovich, Nir and Sergio Rebelo “News about the Future and the Value of the Firm,” mimeo, Northwestern University, 2006.
- [27] Justiniano, Alejandro and Giorgio Primiceri “The Time Varying Volatility of Macroeconomic Fluctuations,” mimeo, Northwestern University, 2005.
- [28] King, Robert G. “Value and Capital in the Equilibrium Business Cycle Program,” in Lionel McKenzie and Stefano Zamagni, *Value and Capital Fifty Years Later*, MacMillan, London, 1991.
- [29] King, Robert G., Charles Plosser, and Sergio Rebelo “Production, Growth and Business Cycles: I. the Basic Neoclassical Model,” *Journal of Monetary Economics*, 21: 195-232, 1988.
- [30] King, Robert G. and Sergio Rebelo “Resuscitating Real Business Cycles,” in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, volume 1B, 928-1002, 1999.
- [31] Lorenzoni, Guido “Imperfect Information, Consumers’ Expectations and Business Cycles,” mimeo. MIT, 2005.
- [32] Lucas, Robert and Edward Prescott, “Investment Under Uncertainty,” *Econometrica*, 39, 659-681, 1971.
- [33] Lucca, David “Resuscitating Time-to-Build,” mimeo, Northwestern University, 2006.
- [34] Veldkamp, Laura and Stijn Van Nieuwerburgh “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53: 753-772, 2006.
- [35] Pigou, Arthur *Industrial Fluctuations*, MacMillan, London, 1927.

- [36] Rotemberg, Julio “Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles,” *American Economic Review* 93: 1543-1559, 2003.
- [37] Sargent, Thomas “Estimation of Dynamic Labor Demand Schedules under Rational Expectations,” *Journal of Political Economy* 86, 1009-1044, 1978.
- [38] Stock, James H. and Watson, Mark W. “Business Cycle Fluctuations in US Macroeconomic Time Series," in Woodford and Taylor (eds.) *Handbook of Macroeconomics*, Amsterdam, Holland, 1999.
- [39] Stock, James H. and Watson, Mark W. “Has the Business Cycle Changed? Evidence and Explanations,” FRB Kansas City Symposium, Jackson Hole, Wyoming, 28-30, 2003.
- [40] Tauchen, George and Robert Hussey “Quadrature Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica*, 592: 371–396, 1991.

TABLE 1: ROBUSTNESS ANALYSIS

| | Benchmark Model | High Adjustment Costs, $\phi''(1) = 4.5$ | Infinite Labor Supply Elasticity ($\theta=1$) | Low Elasticity of Utilization | Adjustment costs to labor ($\phi''(1) = 0.5$) |
|---|-----------------|--|---|-------------------------------|---|
| Curvature in utility, σ | >0.05 | >0 | >0.5 | >0.25 | >0.05 |
| Utility parameter γ | <0.0001 | <0.02 | <0.0001 | <0.13 | < 0.4 |
| Adjustment costs, $\phi''(1)$ | >0.51 | | >0.05 | >0.525 | < 1.5 |
| Elasticity of labor supply ($1/(\theta-1)$) | >0.3 | >0.003 | | >1 | < 21 |
| Elasticity of utilization | <2.85 | <3.3 | <3.9 | | < 1.5 |

TABLE 2: BUSINESS CYCLE MOMENTS

| | DATA | | | MODEL | | | |
|------------------------------------|-----------|-----------|-----------|----------------------|----------------------------------|--|----------------|
| | 1947-2004 | 1947-1983 | 1983-2004 | Uninformative Signal | Imperfect Signal (precision 0.8) | Imperfect Signal (Livingston survey precision) | Perfect Signal |
| Std. Dev. Output | 1.56 | 1.88 | 0.97 | 1.10 | 1.06 | 1.00 | 0.94 |
| Std. Dev. Hours | 1.51 | 1.88 | 1.00 | 0.78 | 0.75 | 0.71 | 0.67 |
| Std. Dev. Investment | 4.84 | 5.41 | 3.69 | 3.45 | 3.37 | 3.33 | 3.30 |
| Std. Dev. Consumption | 1.11 | 1.22 | 0.75 | 0.81 | 0.78 | 0.75 | 0.73 |
| Correlation Output and Hours | 0.86 | 0.88 | 0.87 | 1.00 | 1.00 | 1.00 | 1.00 |
| Correlation Output and Investment | 0.89 | 0.75 | 0.91 | 0.97 | 0.96 | 0.93 | 0.85 |
| Correlation Output and Consumption | 0.77 | 0.68 | 0.75 | 0.94 | 0.94 | 0.92 | 0.89 |
| Sum of 4 coefficients in AR(4) | 0.77 | 0.65 | 0.86 | 0.71 | 0.74 | 0.78 | 0.8 |
| Number of Recessions | 14 | | | 9 | 10 | 9 | 9 |

Figure 1: Response to TFP News Shock, Our Model



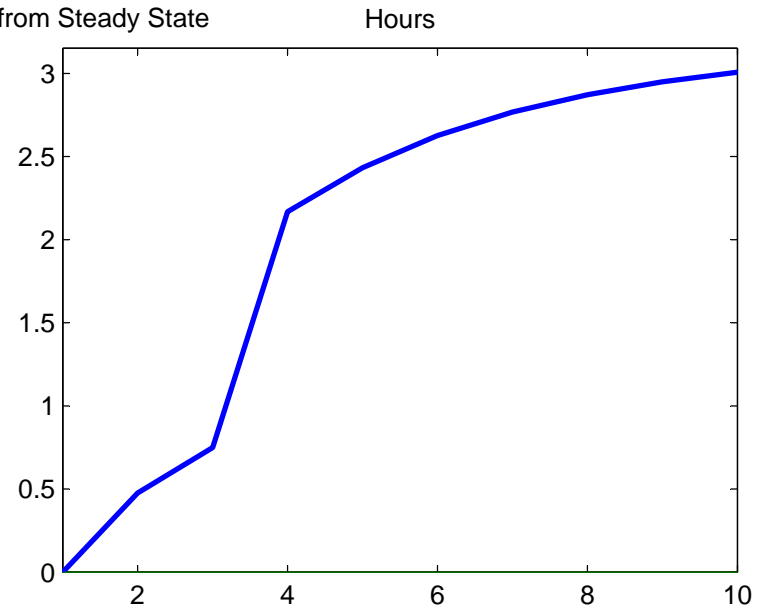
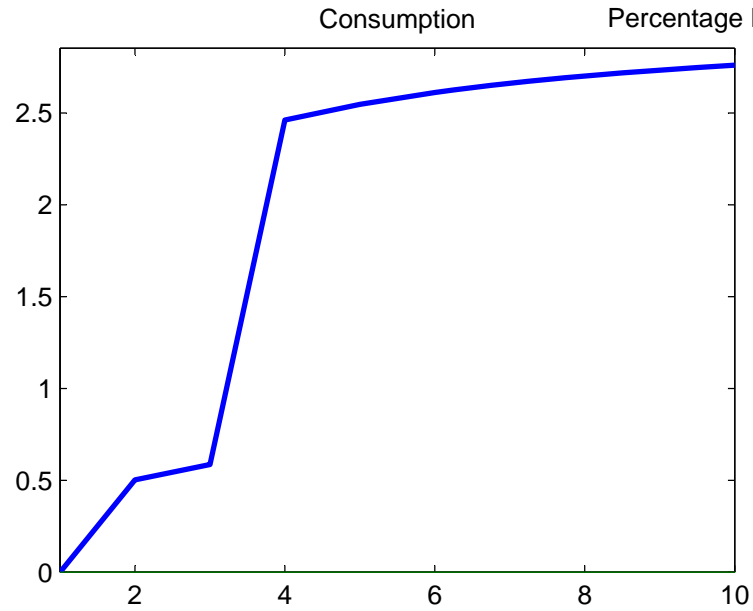


Figure 2: Response to Investment-specific Technical Progress News Shock, Our Model







Figure 3: Response to TFP News Shock, Benchmark RBC with KPR Preferences

Percentage Deviations from Steady State







Figure 4: Response to TFP News Shock, Benchmark RBC with GHH Preferences

Percentage Deviations from Steady State







Figure 5: Response to TFP News Shock, Variants of Our Model

Percentage Deviations from Steady State




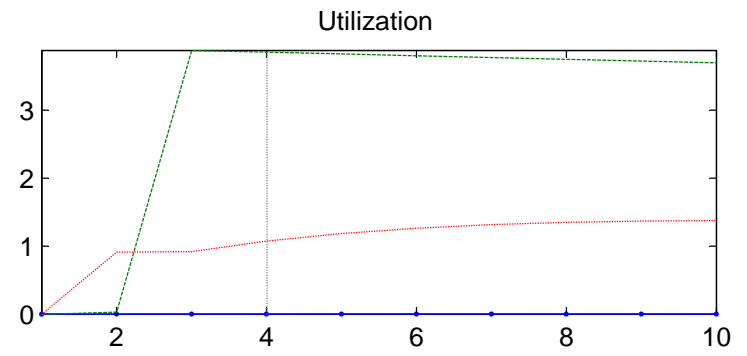

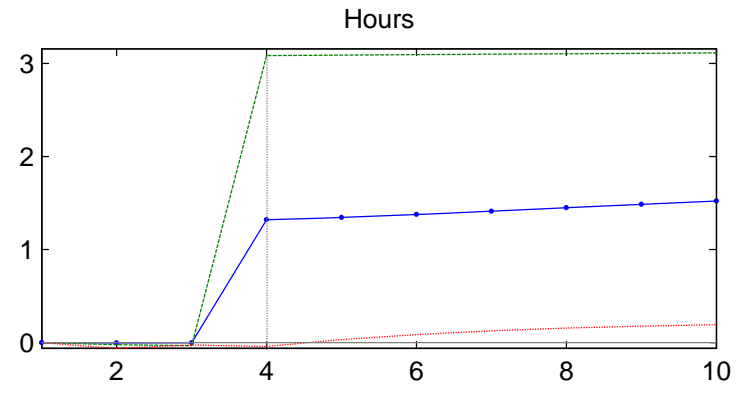
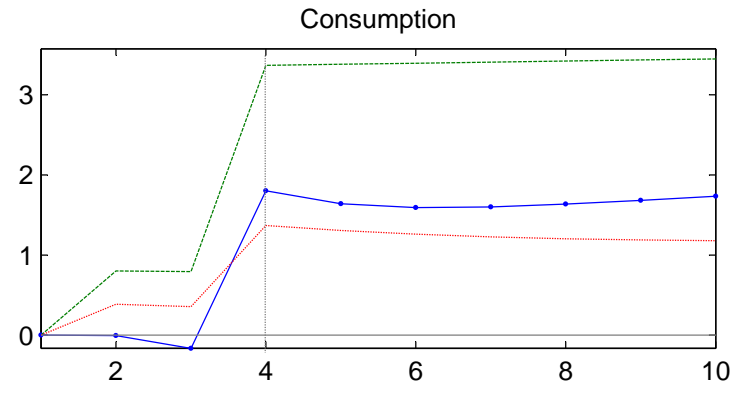


Figure 6: Response of Hours to Permanent TFP Shock at Time One, Standard RBC Model

Percent Deviations from Steady State

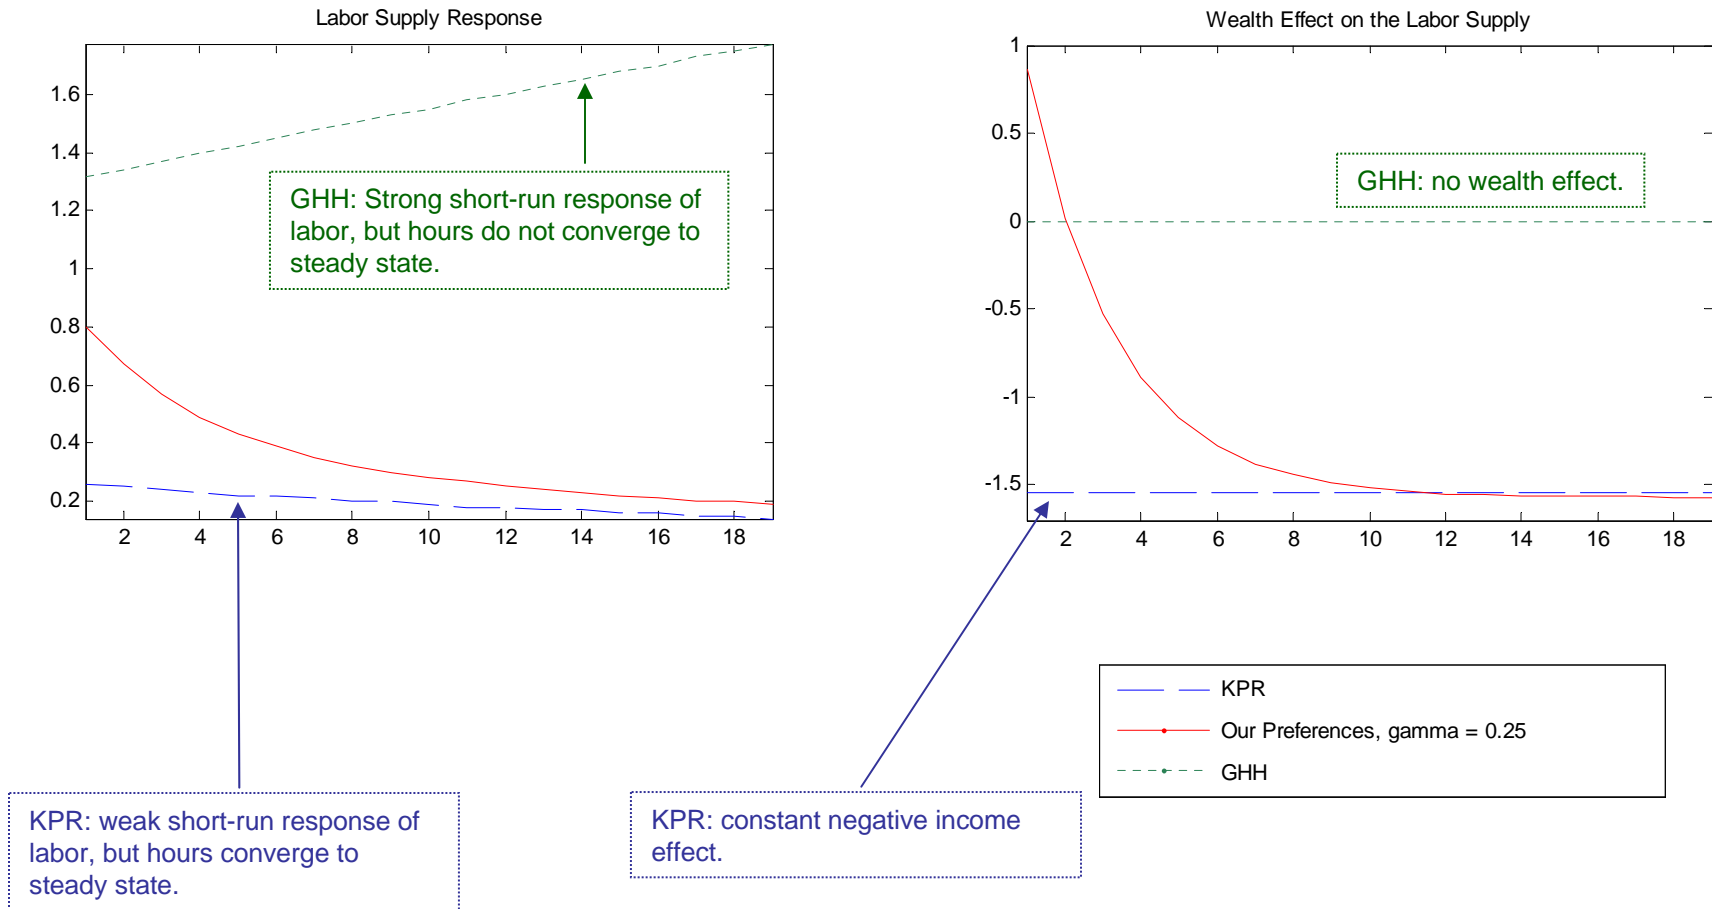


Figure 7: Response of Time t Output to News of Permanent Increase in z or TFP at time t+n




Figure 8: Response to TFP News Shock, Model with Adjustment Costs in Utilization



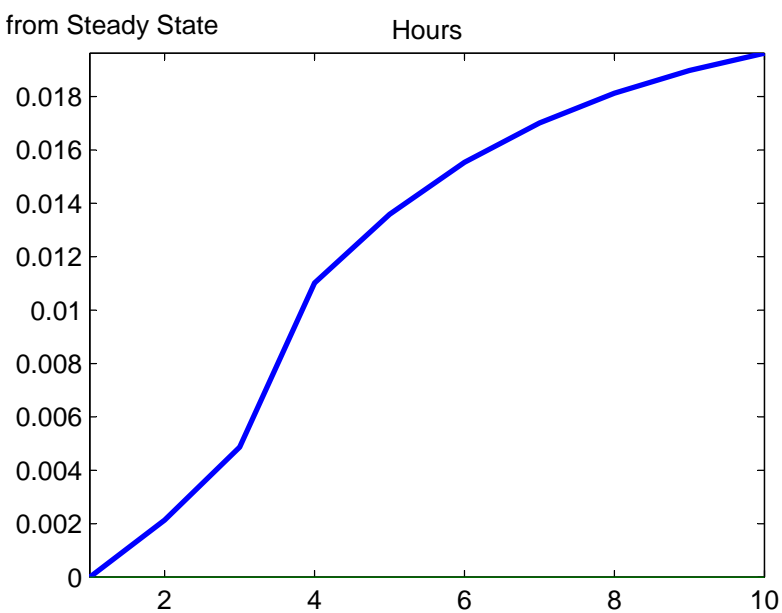



Figure 9: The Effects of Noisy Signals

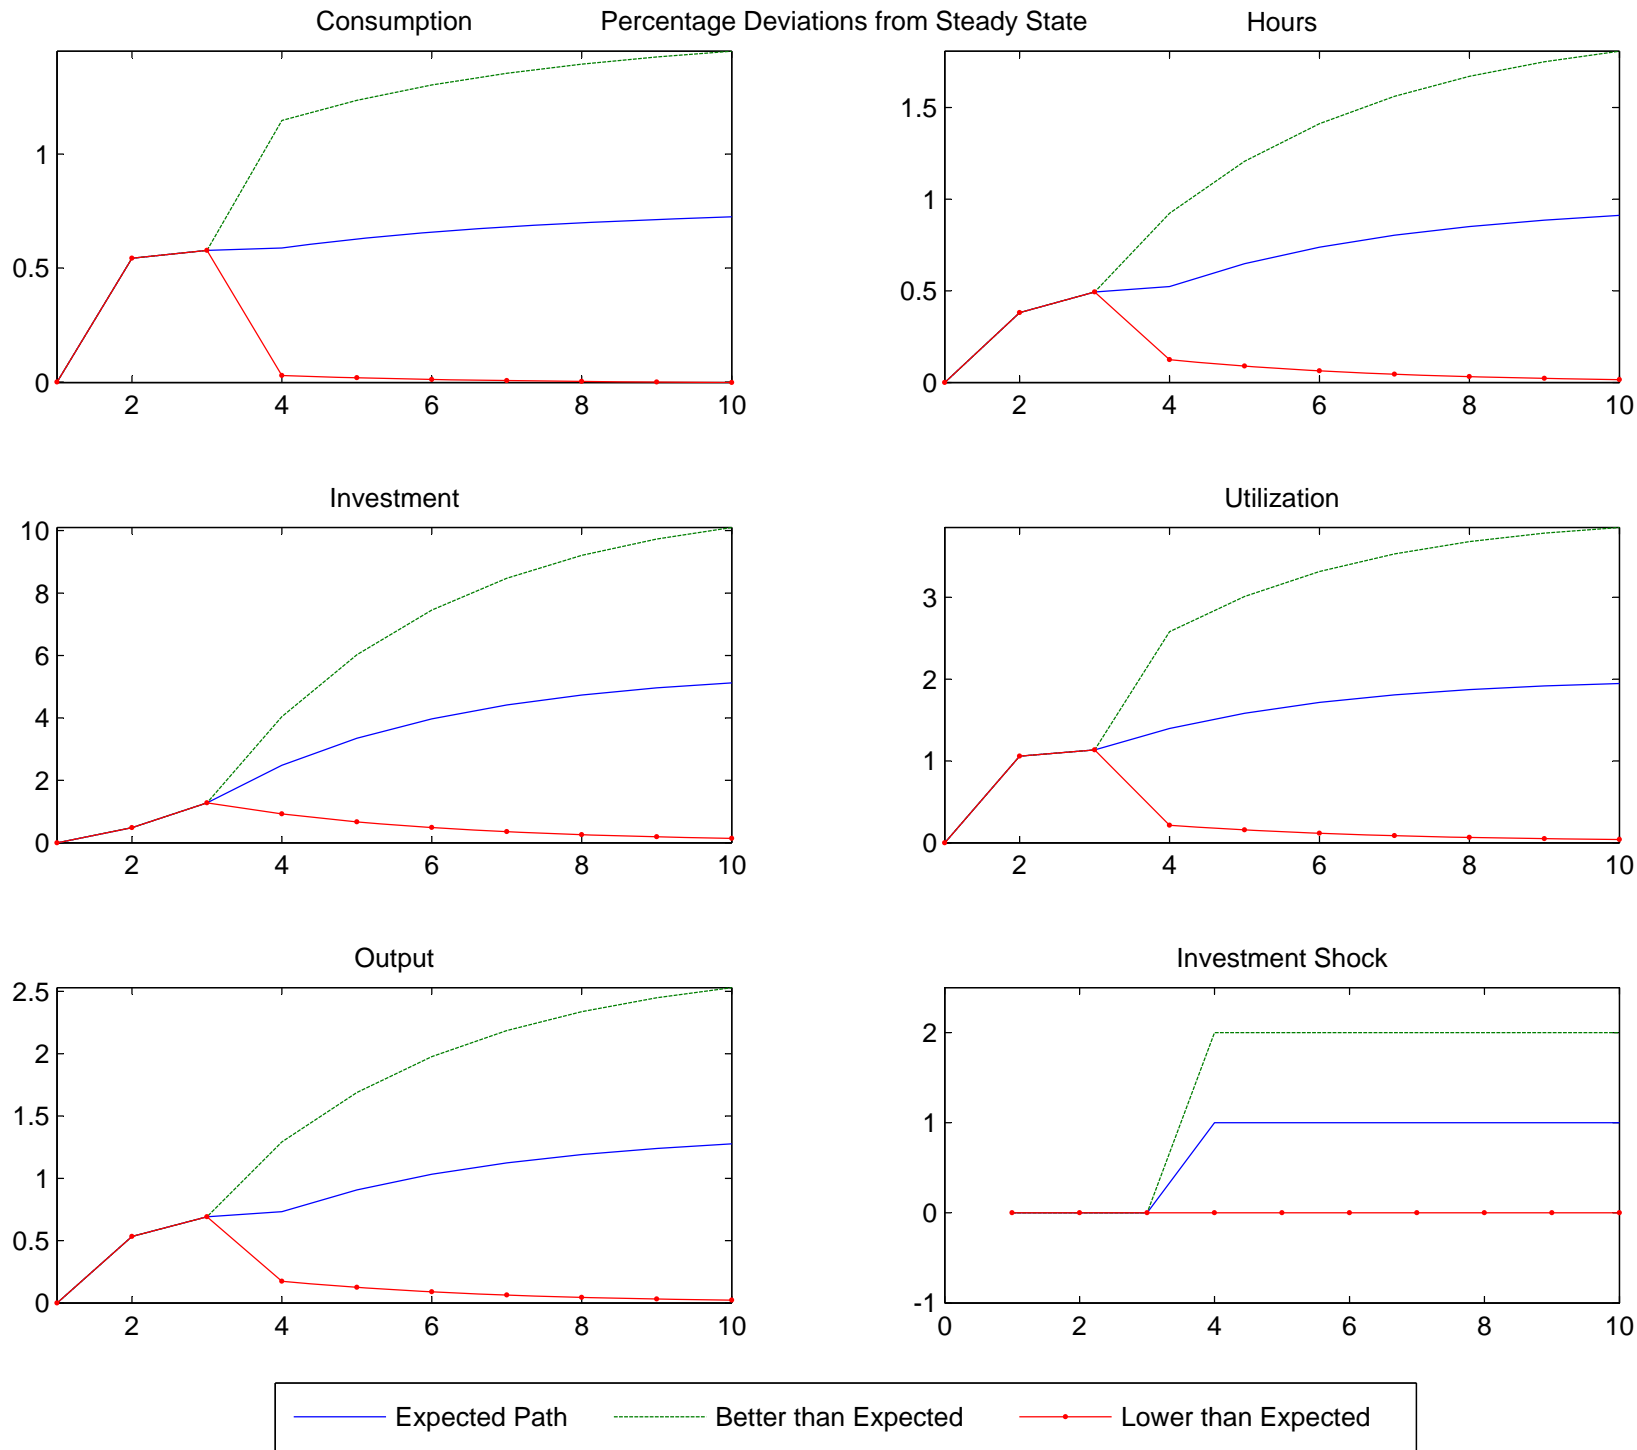


Figure 10: The Effects of News Updating

Percentage Deviations from Steady State


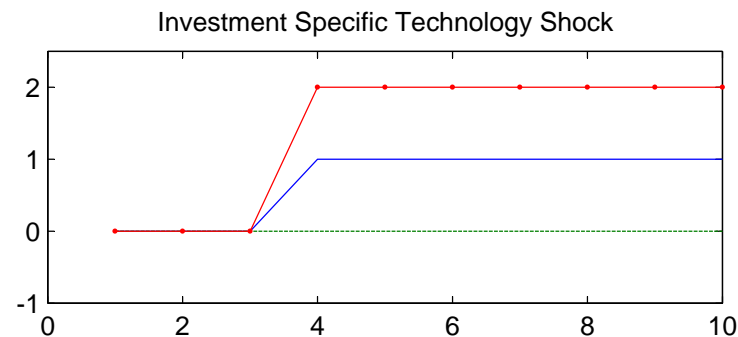







Figure 11: Bayesian Updating

Percentage Deviations from Steady State










Figure 12: Average Recession in the Model and U.S. Data




Figure 13: Behavior of z in Average Recession in Model and U.S. Data




Figure 14: Average Expansion in Model and U.S. Data

