

Can News About the Future Drive the Business Cycle?*

Nir Jaimovich[†] and Sergio Rebelo[‡]

November 2005
(preliminary version)

Abstract

In this paper we propose a model that generates an 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 income effect on the labor supply. These preferences nest, as special cases, the two classes of utility functions most widely used in the business cycle literature. Even though our model abstracts from negative productivity shocks, it generates recessions that resemble those in the post-war U.S. economy. Recessions are caused not by contemporaneous negative shocks but by lackluster news about the future TFP or investment-specific technical change.

J.E.L. Classification: E3.

*We thank Gadi Barlevy, Larry Christiano, and Erin Hoge for their comments.

[†]University of California at San Diego.

[‡]Northwestern University, NBER and CEPR.

1. Introduction

There is an old literature, including work by Beveridge (1909), Clark (1917), and Pigou (1927), that proposes news about the future or changes in agent's expectations as important sources of business fluctuations. There is 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 an economic boom. When the new technologies fail to live up to what was expected, investment falls and a recession ensues. But it is surprisingly hard to make this story work in a standard business cycle model. Cochrane (1994), Danthine, Donaldson, and Johnsen (1998), and Beaudry and Portier (2004b, 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! These difficulties were anticipated by Barro and King (1984) 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 in response to news about future technical progress. For most of the paper we assume that expectations are rational but agents receive "news shocks" (we briefly investigate the implications of expectation biases in section 7). News shocks consist of information that is useful to predict future fundamentals but does not affect current fundamentals.

There is evidence that the economic agents receive and process news about the future. Agents receive advance information about future changes in TFP driven by new technologies, since it takes time for these technologies to diffuse throughout the economy (Rotemberg (2003) and Alexopoulos (2004)). Stock prices and consumer confidence, which naturally reflect agent's expectations about the future, are known to 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 (2004a)). Future political events, such as the outcome of elections, often affect investment flows (see e.g. Bussie and Mulder (2000)).

Our model incorporates three elements into the neoclassical growth model. The first element, variable capital utilization, increases the extent to which output can respond to news about the future. The second element, adjustment costs to investment or capital utilization, gives agents an incentive to respond immediately to future technical progress. The third element, preferences that exhibit a weak short-run income effect on the labor supply, is necessary so that hours worked 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, those 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 the short-run wealth effects on the labor supply are small.

We use the model to illustrate how downward revisions in expectations about future technical progress can generate recessions. In these experiments the economy appears to be too volatile since there are no contemporaneous fundamentals, other than news about the future, that can account for changes in output. We

also study a setting in which the impact of new technologies is uncertain. Agents form priors about the impact of each new technology 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 progress, in which agents receive news about future technical progress. We choose the information content of this news to match the predictive content of the Livingston's survey six-months-ahead real-GDP forecasts. We find that the same calibration that produces an expansion today in response to news of higher technical progress tomorrow can generate business cycles with statistical properties that resemble those of U.S. data. We discuss the model's implications for the standard set of statistics on volatility, comovement, and persistence of macroeconomic aggregates. In addition, we show that the average recession and expansion are similar in the model and in post-war U.S. economy. One interesting feature of the model is that it can produce declines in the level of output even though the rate of technical progress is always positive.

It is well known that there has been a secular 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 are relevant to predict the future. This increase in precision reduces the volatility and increases the persistence of output in our model.

We explore the business cycle implications of two well-known psychological biases, optimism and overconfidence. Optimistic agents over-estimate the precision

of the signals that they receive. Both biases results in more frequent mistakes, generating higher business-cycle volatility.

Our work is related to several recent papers on the role of news and expectations as drivers of business cycles. Beaudry and Portier (2004b) 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) emphasize the role of monetary policy in generating booms in response to news about higher future productivity. Denhaan and Kaltenbrunner (2005) focus on the role played by search in the labor market in the response to news shocks. Lorenzoni (2005) studies the case where 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. We also 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 TPF or investment-specific technical change. We also discuss alternative specifications for investment adjustment costs and capital 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 information about the rate of technical progress two periods in advance. In Section 7 we study the business cycle implications of two psychological biases, optimism and overconfidence. In Section 8 we review our main results.

2. Our Model

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

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

where

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

We assume that $0 < \beta < 1$, $\theta > 1$, $\psi > 0$ and $\sigma > 0$. Note that 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.

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 exogenous 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),

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

where z_t represents the current state of technology to produce capital goods. We interpret declines 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:

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

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 to investment. We assume that $\phi(1) = 0$, $\phi'(1) = 0$, and $\phi''(1) > 0$. These conditions imply that there are no adjustment costs in the steady state and that adjustment costs are incurred when the level of investment changes over time. This adjustment cost formulation is proposed in Christiano, Eichenbaum and Evans (2004). These authors 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).

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 solved the model by linearizing the equations that characterize the planner's problem around the steady state. We chose the following parameter values for our benchmark model. We set $\sigma = 1$, which corresponds to the case of logarithmic utility. The parameter θ is set to 1.4, which corresponds to an elasticity of labor supply of 2.5. The discount factor is set to $\beta = 0.985$ 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, $\phi''(1)$, equal to 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 be equal to 0.15. This value influences the degree of shock amplification present in the economy. When the value of $\delta''(u)$ 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 in Section 4 the robustness of our results to these parameters.

Responding to News about the Future We illustrate the response of our model to news shocks with what we refer to as our basic experiment. At time zero the economy is in a steady state with no technical progress. At time one unanticipated news arrives. Agents learn that there will be a once-and-for-all one-percent permanent increase in TFP starting two periods later, in period three. Figure 1 depicts the response of the economy to this news. There is an expansion in periods one and two in response to positive news about future productivity. Consumption, investment, output, hours worked, and capital utilization all rise in periods one and two even though the positive shock only occurs in the future. Since in Section 3 we consider a version of the model with investment-specific technical progress, Figure 2 shows the response to a version of the same experiment in which there is a future increase in z_t instead of in A_t . Once again we see an expansion, with consumption, investment, output, hours worked, and capital utilization rising before the shock materializes. Note that with TFP shocks the impact of news about future TFP is less important than the realization of TFP shocks. In contrast, with investment-specific technical change most of the rise in output occurs in period one, when the news arrives, not in period 3, when the

investment-specific technical progress occurs.

One-sector Neoclassical Model We now consider the response to news about future TFP in the standard one-sector neoclassical growth model that has been the workhorse of the real business cycle literature. 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 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 our basic experiment. Note that both hours worked and output fall at time one. 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 where the real wage rate grows at a constant rate. This property requires that the income and substitution effect of a permanent increase in the real wage rate be exactly offsetting. Unfortunately, this property implies that positive news about future TFP or investment-specific technical change reduce today's supply of labor. Positive news make 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 income 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 has to drop. 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 basic experiment. With GHH preferences the optimal number of hours worked depends only on the contemporaneous real wage, 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 do not depress the labor supply today through an income effect. This property makes it easier to obtain an expansion today in response to positive news about tomorrow. But 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, so output also remains constant. The positive wealth effect dictates a decline in 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

We now discuss the importance of the three elements that generate comovement between consumption, investment, output, and labor in response to news about future A_t or z_t . To discuss 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 to the planner's problem are:

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

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

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

$$\eta_t \left[1 - \phi \left(\frac{I_t}{I_{t-1}} \right) - \phi' \left(\frac{I_t}{I_{t-1}} \right) \frac{I_t}{I_{t-1}} + \phi' \left(\frac{I_{t+1}}{I_t} \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \right] = \lambda_t z_t, \quad (3.4)$$

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

Investment Adjustment Costs We first consider the role played by 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 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 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 remain constant at time one, this equation together with (3.2) imply that both N_t and u_t remain constant. The dashed line in Figure 5 depicts the response of a version of our model with no investment adjustment costs to our basic experiment. Note that investment falls initially and rises only in period 3 when the rise in TFP occurs.

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 problem 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):

$$\theta \psi N_t^{\theta-1} = \alpha A_t K_t^{1-\alpha} N_t^{\alpha-1}. \quad (3.5)$$

This equation implies that N_t does not respond to news about future changes in A_t or z_t . The positive income 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 has to fall. So, on impact, 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 utilization to our basic experiment.

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

To isolate the role played by 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 experiment. At time zero the economy is in the steady state. At time one there is an unanticipated, permanent increase in TFP. Figure 6 shows the response to this shock for three different values of γ . The strongest response of hours worked occurs with GHH preferences ($\gamma = 0$). But 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.¹ 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 take the form (2.1) and $\gamma = 0.25$. Note that with these preferences hours worked also converge to the steady state but the short-run impact of the TFP shock is in between that of GHH and KPR preferences. Lower

¹A simple way to make hour 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.

(higher) values of γ produce short-run responses that are closer to those obtained with GHH (KPR) preferences. But, as long as $0 < \gamma \leq 1$, hours worked converge to the steady state.

Beaudry 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).

4. Robustness

We have seen that in our model positive news about future rises 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 the model generates an expansion in response our basic 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 ($\delta''(u)u/\delta'(u) < 2.85$) that are consistent with an expansion in response to future positive news. One critical parameter is γ . We need $\gamma < 0.001$, i.e. preferences must be close to a GHH form. When this is not the case the positive 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 high elasticity of utilization. We see that there is a trade-off between the different features of the model. 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 a significant impact on today's level of output.

The form of investment adjustment costs present in equation (2.5) is important in generating our results. We found 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):

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

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

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

The function $\phi(\cdot)$, which represents adjustment costs to capital utilization is convex with $\phi(1) = 1$. Figure 8 shows the response to our basic experiment of a version of our model with adjustment costs in utilization. In order to produce a

positive response to news about higher future productivity we had to increase the elasticity of labor supply from 2.5 ($\theta = 1.4$) to 20 ($\theta = 1.05$). 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 the magnitude of the rise in output. The lower rise in output can be insufficient to be consistent with a rise in both consumption and investment.

5. News-driven Fluctuations

We now discuss four scenarios under which our model can generate news-related fluctuations: perfect news about the future, noisy news about the future, news revisions, and Bayesian updating of beliefs about the future. We note that in all these scenarios, variants of the standard real business cycle model generate a negative correlation between consumption and output.

Perfect Signals about the Future The first type of news-related fluctuations is the one we discussed in the previous section and illustrated in Figures 1 and 2. In this case the economy receives a perfect signal about the future. Good (bad) news about future TFP or investment specific technical change causes an expansion (recession) before the changes in fundamentals occur.

Noisy Signals about the Future The second type of news-related fluctuations occurs when the economy receives noisy signals about changes in future fundamentals. In this case an expansion (recession) can occur because the actual realization of the fundamentals turns out to be better (worse) than what agents expected conditional on the signal they received. Figure 9 shows the outcome of the following experiment. At time zero the economy is in a steady state with no technical progress. At time one the economy receives unanticipated news that

within 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 larger than expected. 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 lower than expected. In this case the economy goes into a recession even though there is no realized fall in z_t . Fundamentals remain as good as in the past, but they are worse relative to expectations. The same forces that cause the economy to expand in periods one and two in anticipation of an increase in z_t are set in reverse once realized fundamentals fail to live up to what was expected.

News Updates The third type of news-related fluctuations is illustrated in Figure 10. The setup of this experiment is the same as the previous one. In period two the economy receives the same noisy signal about z_t described above. However, there is one modification. In period three the economy receives an update about the value of z_{t+4} . The solid line corresponds to the case where in period one the economy learns that the change in z_t at time three will coincide with the expected change. In this case the economy continues on a smooth expansion. The dashed line shows the case where at time one the economy learns that the change in z_t 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 one that the change in z_t will be lower than expected. This bad news plunges the economy into a recession.

Rational “Optimism” and “Pessimism” Authors such as Cochrane (1994) and Danthine, Donaldson and Johnsen (1998) emphasize the potential role of changes in expectations about the future in driving economy fluctuations, while stressing that this potential is not fulfilled in variants of the standard real business cycle model. Our fourth type of news-related fluctuations relies on changes in expectations. We consider changes in expectations that are rational. Agents make the best possible use of the available information to resolve fundamental uncertainty about the economy. For concreteness, suppose that when a new wave of technology such as the internet is introduced, agents form priors about the effects of the new technology on the rate of change of z_t . They then observe realizations of z_t and update these 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 have a prior normal distribution about μ and this prior has 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 assumed that the initial prior is relatively informative ($V = 10^{-5}$), i.e. agents have confidence in their beliefs. 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 expectations are ex-post "realistic", i.e. $\mu = \theta$. While there is still some updating and resolution of uncertainty that goes on, output fluctuations are small. 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 is capable of generating 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.²

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 for the U.S. real price of investment for the period 1947.I to 2004.IV. These data were constructed by Fisher (2004), using National Income and Product Accounts series for the

²Fisher (2004) and Justiniano and Primiceri (2005) argue that the most important shock in explaining the variability of output is investment-specific technical progress.

consumption deflator and Cummins and Violante’s (2003) updated series for Gordon’s (1989) quality-adjusted producer durable-equipment deflator.³ 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.

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 (2004)) and by the fact that the Livingston survey output forecasts, which we use below in one of our calibrations, are available for a six-month horizon. The results we emphasize are generally robust to other values of n .

Perfect Signal Column 6 of Table 2 summarizes the business cycle properties of our model when agents receive at time t 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 of 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.

Uninformative Signal Column 4 in Table 2 summarizes the business cycle properties of a version of our model in which the economy receives no news, or

³We thank Ricardo Di Ceccio for providing us with an updated version of this time series.

the signal is informative. 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 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 basic patterns of comovement or relative volatility of the major macroeconomic aggregates.

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

$$p_i = \Pr(\varepsilon_{t+2} = i | S = i), \quad i = \text{high, low}.$$

Note that the signal precision 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. The main result here is that the volatility of output is in between the case of the perfect signal and the case in which there is no signal or the signal is uninformative.

In our second setting we provide agents with a signal, S^y , on whether the growth rate of output in two periods is going to be above or below the average. The signal takes two values, high (H) or low (L). We choose the signal to have the same precision as the Livingston survey 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.⁴ The precision of these forecasts is as follows:⁵

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

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

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:

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

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 :

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

We simulate the model and compute,

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

⁴See Croushore (1993) for a description of the Livingston survey.

⁵To 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.

We then revise the values of q_1 and q_2 until the precision of S^y in the model coincides with the precisions (6.2) estimated in the data. We obtained $q_1 = 0.99$ and $q_2 = 0.7874$. Column 6 of Table 2 shows the results for the version of the model. The main result is that the volatility of output is in between the two extremes of uninformative signals and perfect signals.

News and Volatility It is well-known that output volatility has declined over the past thirty years in virtually all the 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 decline 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 lead 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 from Column 4 of Table 2 (no news) toward Column 7. The availability of news, makes it easier to forecast the future reducing economic volatility. There

is evidence from the Livingston survey that 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 until 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) as compared to 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. Falls in z_t are rare (they occur in only 6 percent of the quarters in our sample) and small in magnitude. The average change 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.⁶ 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) to compare different recessions (see Hall et al. (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 are similar to those chosen by the NBER dating committee. The HP procedure produces six recessions whose starting date coincides with that

⁶King and Rebelo (2000) propose a real business cycle model that generates recessions in the absence of negative technology shocks. Their model shares two key features with our model, variable capital utilization and a highly responsive labor supply.

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. This suggests that our procedure corresponds to a broader definition of recession than that of the NBER.

Once we identified 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 starts. 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 0.71 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. Note that, on average, recessions occur when there is a high contemporaneous rate of change in investment-specific technical progress but the economy learns that in two periods 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 9 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 emphasized news about future investment-specific technical change, the same mechanisms are likely to produce recessions in response to bad news about the future value 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. Expectation Biases

So far we have discussed environments in which agents have rational expectations about the future. Here we investigate the impact on business cycles of two well-known psychological biases, optimism and over-confidence. Optimistic agents overestimate the probability of good outcomes. Over-confident agents overestimate the precision of the signals that they receive.⁷

To investigate the impact of optimism we consider the following modified version of our benchmark model with noisy signals. The values of ε_t are generated by the Markov chain matrix (6.1). However, agents forecast future values of ε_t according to the following distorted version of the matrix (6.1):

$$\pi^* = \begin{bmatrix} 0.90 & 0.10 \\ 0.41 & 0.59 \end{bmatrix}.$$

The diagonal elements of π^* have been changed relative to those in π by 20 percent so as to increase the probability of the good shock.

⁷Brunnermeier and Parker (2005) discuss the literature on these biases and provide a model of the optimal level of optimism.

To incorporate over-confidence into the model we consider the noisy signal version of our benchmark model in which the precision of the signal is 0.8 in both states but assume that agents view the precision of the signal as perfect. Table 3 shows that the main impact of both behavioral biases is to increase the volatility of the economy. Output volatility increases from 1.06 in the fully rational case to 1.15 (1.12) in the overconfidence (optimism) case. In both cases agents make more errors about the future. Optimistic agents consistently over-invest and the realization of ε_t are on average disappointing. In the over-confidence case there is both over-investment and under-investment.

8. 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, Huffman and Hercowitz (2000) and by King, Plosser, and Rebelo (1988). Our preferences share with the Greenwood, Huffman and Hercowitz (2000) 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 technical 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 reduce the volatility of the output relative to a model with no news. This suggests that improvements in the quantity and quality of information that is useful to forecast the future may have contributed to the observed secular decline in business cycle volatility.

We show that behavioral biases such as optimism (over-estimation of the probability of good outcomes) and over-confidence (over-estimation of information precision) increase the frequency of mistakes and make the economy more volatile.

References

- [1] Abel, Andrew and Olivier Blanchard, “An Intertemporal Equilibrium Model of Savings and Investment, *Econometrica* 51, 675-692, 1983.
- [2] Barro, Robert J. and Robert G. King “Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles,” *The Quarterly Journal of Economics*, Vol. 99, No. 4. (Nov., 1984), pp. 817-839.
- [3] Beaudry, Paul, and Franck Portier “An Exploration into Pigou’s Theory of Cycles,” *Journal of Monetary Economics*, 51(6), 1183–1216, 2004a.
- [4] Beaudry, Paul, and Franck Portier “News, Stock Prices and Economic Fluctuations,” Working Paper 10548, National Bureau of Economic Research, 2004b.
- [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] Brunnermeier, Markus and Jonathan Parker “Optimal Expectations,” *American Economic Review*, Vol 95 No 4, (September 2005) 1092-1118.
- [7] Bussie, Matthieu and Christian Mulder, “Political Instability and Economic Vulnerability,” *International Journal of Finance and Economics* 5: 309–330, 2000.
- [8] Christiano, Lawrence, Roberto Motto, and Massimo Rostagno “Monetary Policy and a Stock Market Boom-Bust Cycle,” mimeo, Northwestern University, 2005.

- [9] Cochrane, John H. "Shocks," *Carnegie-Rochester Conference Series on Public Policy*, 41: 295-364, 1994.
- [10] Croushore, Dean (1993): "Introducing: The Survey of Professional Forecasters," *Federal Reserve Bank of Philadelphia Business Review*, pp. 3–15.
- [11] Cummins, Jason and Giovanni Violante, "Investment-Specific Technical Change in the United States (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, 5(2), April 2002, 243-284.
- [12] Danthine, Jean Pierre, John B. Donaldson and Thore Johnsen "Productivity Growth, Consumer Confidence and the Business Cycle," *European Economic Review*, 42, 1113-1140, 1998.
- [13] Denhaan, Wouter J. and Georg Kaltenbrunner "Growth Expectations and Business Cycles," mimeo, London Business School, 2005.
- [14] Fisher, Jonas "Technology Shocks Matter," mimeo, Federal Reserve Bank of Chicago, 2004.
- [15] Gordon, Robert, *The Measurement of Durable Goods Prices*, Chicago: University of Chicago Press, 1989.
- [16] Greenwood, Jeremy, Zvi Hercowitz and Gregory Huffman, "Investment, Capacity Utilization and the Real Business Cycle," *American Economic Review* 78, 402-417, 1988.
- [17] Greenwood, Jeremy Zvi Hercowitz and Per Krusell "The Role of Investment-specific Technological Change in the Business Cycle," *European Economic Review*, Volume 44, Issue 1, January 2000, Pages 91-115.

- [18] Hayashi, Fumio, 1982, "Tobin's Marginal q and Average q : A Neoclassical Interpretation," *Econometrica* 50.
- [19] Justiniano, Alejandro and Giorgio Primiceri "The Time Varying Volatility of Macroeconomic Fluctuations," mimeo, Northwestern University, 2005.
- [20] 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.
- [21] 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.
- [22] Lorenzoni, Guido "Imperfect Information, Consumers' Expectations and Business Cycles," mimeo. MIT, May 2005.
- [23] Lucas, Robert and Edward Prescott, "Investment Under Uncertainty," *Econometrica*, 39, 659-681, 1971.
- [24] Rotemberg, Julio, "Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles," *American Economic Review* 93 (December 2003), 1543-1559.
- [25] 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.
- [26] Tauchen, George and Robert Hussey "Quadrature Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models" *Econometrica*, 592:371-396, 1991.

Table 1

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

Table 2

	Data			Our Model			
	1947-2004	1947-1983	1985-2004	Uninformative Signal	Imperfect Signal (precision 0.8)	Imperfect Signal (precision calibrated to Livingston survey)	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

Table 3

Imperfect Signal, Precision 0.8

	Data		Our Model	
	1947-2004	Rational	Optimism	Overconfidence
Std. Dev. Output	1.56	1.06	1.13	1.15
Std. Dev. Hours	1.51	0.75	0.80	0.82
Std. Dev. Investment	4.84	3.37	3.54	3.37
Std. Dev. Consumption	1.11	0.78	0.98	0.92
Correlation Output and Hours	0.86	1.00	1.00	1.00
Correlation Output and Investment	0.89	0.96	0.87	0.90
Correlation Output and Consumption	0.77	0.94	0.89	0.93
Sum of 4 coefficients in AR(4)	0.77	0.74	0.70	0.70
Number of Recessions	14	10	10	10

Figure 1: Response to TFP News Shock, Our Model

Percentage Deviations from Steady State







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

Percentage Deviations from Steady State







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




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

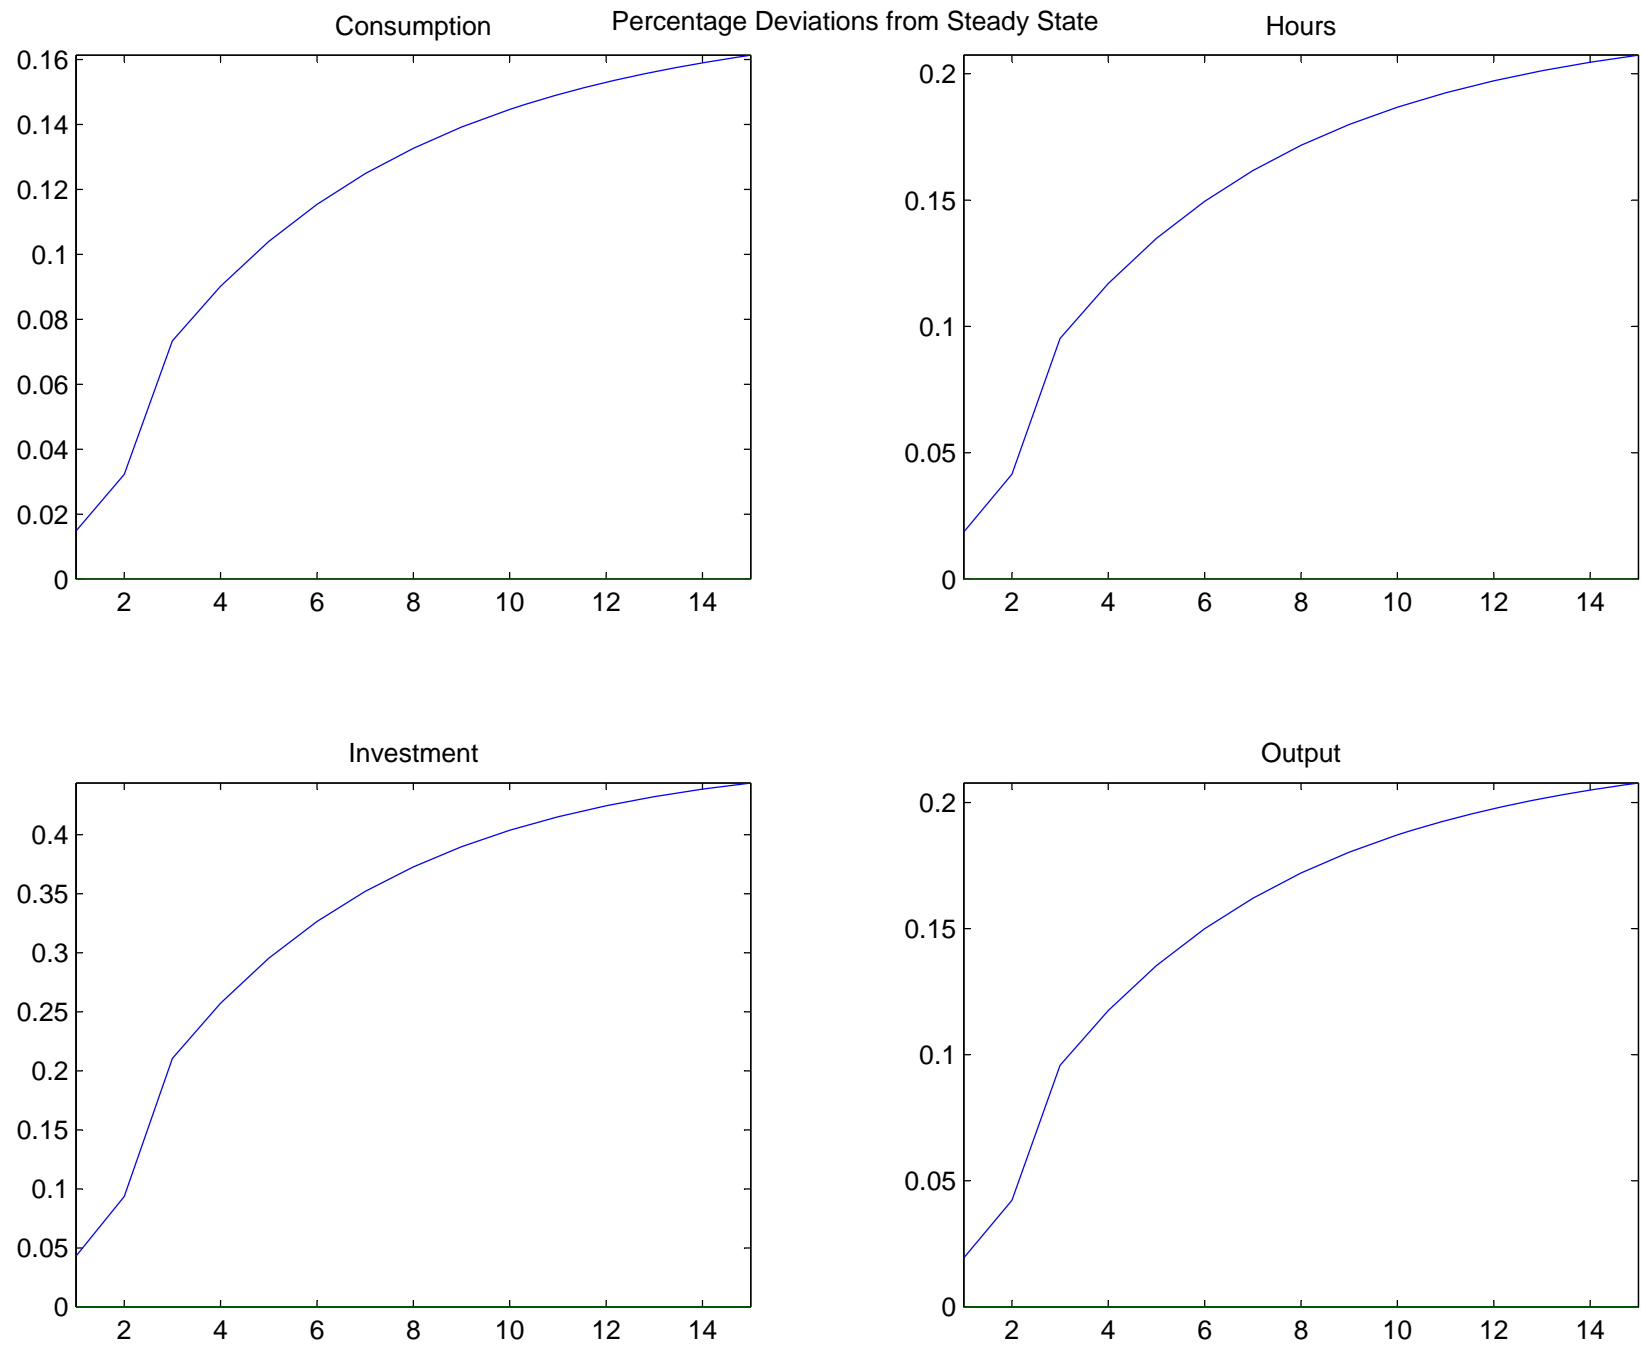


Figure 9: The Effects of Noisy Signals




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

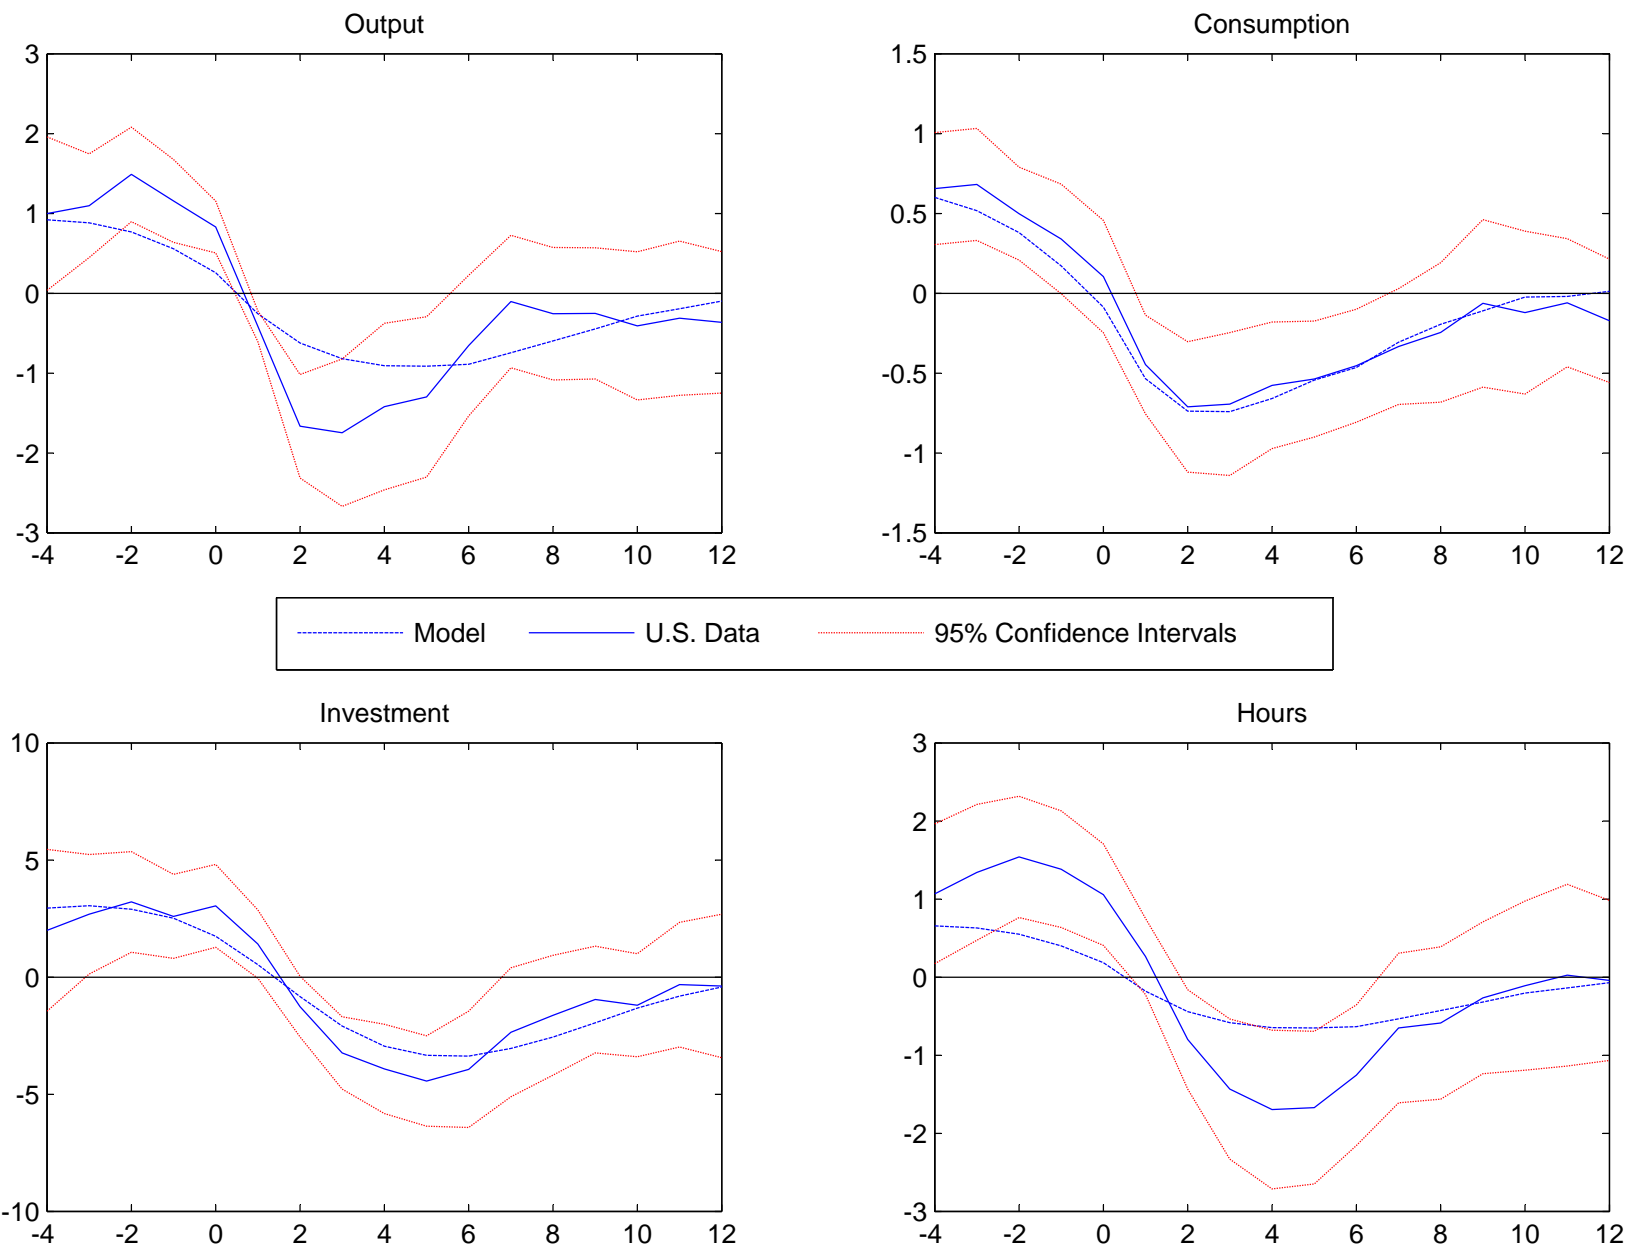


Figure 13: Behavior of z in Average Recession in Model and U.S. Data




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

