CFCM

CENTRE FOR FINANCE AND CREDIT MARKETS

Working Paper 11/06

Learning, Capital-Embodied Technology and Aggregate Fluctuations

Christoph Görtz and John Tsoukalas

Produced By:

Centre for Finance and Credit Markets
School of Economics
Sir Clive Granger Building
University Park
Nottingham
NG7 2RD

Tel: +44(0) 115 951 5619
Fax: +44(0) 115 951 4159
enquiries@cfcm.org.uk

CFCM

Learning, Capital-Embodied Technology and Aggregate Fluctuations*

Christoph Görtz[†] and John Tsoukalas[‡]

This version: May 2011.

Abstract

Recent evidence suggests that agents' expectations may have played a role in several cyclical episodes such as the U.S. "new economy" boom in the late 1990s, the real estate boom in Japan in the 1980s and the real estate boom in the U.S. which ended in 2008. One challenge in the expectations driven view of fluctuations has been to develop simple one sector models that can give rise to such fluctuations without a compromise on other dimensions. In this paper we propose a simple generalization of the Greenwood et al. (1988) one sector model and show it can generate fluctuations that arise as a result of agents difficulty to forecast productivity embodied in new capital. The two key assumptions in the model are: (1) the vintage view of capital productivity, whereby each successive vintage has (potentially) different productivity and (2) agents' imperfect information and learning about this productivity. The model is consistent with second and third moments from U.S. data. Simulations of the model suggest that, (a) noise amplifies fluctuations and (b) pure noise can trigger recessions that mimic in magnitude, duration and depth the typical post WW II U.S. recession.

Keywords: News shocks, expectations, growth asymmetry, Bayesian learning, business cycles.

JEL Classification: E2, E3, D83.

---

*We are grateful for comments from Thomas Lubik and from participants at the Money, Macro and Finance Research Group 41st Annual Conference in Bradford and the Economic Theory Workshop at the University of Nottingham. All remaining errors are our own.

[†]University of Nottingham, Sir Clive Granger Building – School of Economics, University Park, Nottingham NG7 2RD. Email: lexcgg@nottingham.ac.uk.

[‡]University of Nottingham, Sir Clive Granger Building – School of Economics, University Park, Nottingham NG7 2RD. Email: john.tsoukalas@nottingham.ac.uk.

1 Introduction

The idea that business cycles can be driven by shifts in expectations is a long-known concept and dates back to Beveridge (1909), Pigou (1926) and Clark (1935). According to this idea shifts in agents' expectations that are orthogonal to current fundamentals can generate a cycle whereby consumption, investment and hours worked co-move with economic activity. In early work, Barro and King (1994) (BK) pointed out that changes in beliefs about the future cannot generate empirically recognizable business cycles within the standard real business cycle model. Intuitively, news that future productivity will improve creates a wealth effect where agents finance the consumption of goods and leisure today from lower investment. Beaudry and Portier (2004) overcome the obstacles posed by the analysis of Barro and King (1994) with a carefully designed three sector model where beliefs about future total factor productivity (TFP) can induce responses of agents that resemble the typical pattern of co-movement observed in post-WW II U.S. business cycles.

One challenge in this area has been to develop simple one sector models that can give rise to expectations driven fluctuations without a compromise on other dimensions. In this paper we show that a simple one sector model with capital embodied productivity can generate fluctuations that arise as a result of agents difficulty to forecast this productivity. The two key assumptions in the model are: (1) the vintage view of capital productivity, whereby each successive vintage has (potentially) different productivity and (2) agents imperfect information and learning about this productivity. Essentially the model we propose can be interpreted as a generalization of Greenwood et al. (1988) with the incorporation of Bayesian learning. Our assumption of learning about vintage specific capital productivity is firmly rooted in empirical evidence based on learning in manufacturing plants (see for example Bahk and Gort (1993) and Sakellaris and Wilson (2004)). We consider capital embodied productivity as the sole driving force in the model given the evidence suggesting its importance as a major driving factor of U.S. macroeconomic fluctuations, especially during the investment boom and bust of the 1990s and early 2000s (see e.g. Fisher (2006) and Justiniano and Primiceri (2008)).[1]

In the model agents receive signals (or news) about the productivity of future capital vintages. However, signals are not necessarily accurate as they can be driven by pure noise and agents need to disentangle this noise from fundamentals. This is the signal extraction problem agents need to solve in order to decide on optimal investment, utilization, hours worked and output produced and consumed. Fluctuations in the model are thus driven by both noise and fundamentals. Thus learning productivity can potentially add another source of fluctua-

[1]Other recent theoretical work develops models with and without market frictions which overcome the Barro and King (1994) challenge. See, for example, Beaudry and Portier (2007), Christiano et al. (2008), Karnizova (2010), Gunn and Johri (2011) Keiichiro et al. (2007), Kobayashi and Nutahara (2010), Den Haan and Kaltenbrunner (2009) and Guo (2008).

tions in macroeconomic aggregates. Noise in the model—and the associated forecast errors it implies—can make agents over-invest, work and produce more when they are optimistic and cut investment, hours and output when they are pessimistic.

We ask whether this model is consistent with business cycle facts computed from U.S. data. We are not only interested in the usual set of moments (i.e. volatilities, serial correlations and correlations with output) but examine whether it is also in line with a measure of cyclical (more precise growth) asymmetry present in the data. This type of asymmetry, i.e. the fact that booms are more gradual than busts which are usually short and sharp, is a characteristic feature of recent U.S. cyclical episodes. For example, measuring the asymmetry of U.S. business cycles by computing the skewness of GDP we find that both the Information Technology (IT) (1991 Q1 - 2001 Q4) and the real estate ”bubble" (2001 Q4 - 2009 Q2) are consistent with highly asymmetric episodes exhibiting negative skewness, with a gradual expansion and a sharp recession phase. We find the model can replicate fairly well the usual business cycle moments and can generate business cycle (growth) asymmetries in line with the data.

The *impact* of noise—and the associated forecast errors it produces—in the model is substantial. According to a measure of large size forecast errors we adopt, which occurs in approximately 20% of time in the simulation, agents mistakes in forecasting productivity can give rise to substantial swings in investment. During periods of pessimism (when agents are underpredicting embodied productivity compared to the truth) we observe a decline in investment in the order of 4.0 percentage points **below** what one would observe in an economy without noise and forecast errors. During periods of optimism (when agents are overpredicting embodied productivity compared to the truth) we observe a boom in investment in the order of 8.9 percentage points **above** the level implied in the economy without noise and forecast errors. We find qualitatively similar differences in output, hours worked and utilization rates. Consequently, the learning mechanism in the model magnifies changes in fundamentals.

Noise does not only amplify changes in fundamentals but can also trigger recessions (when true productivity rises but agents forecast a decline) that would not occur in a perfect information economy. We find that noise triggered recessions can generate declines similar in magnitude to declines driven by un-favorable fundamentals. The share of recessions due to pure noise in the simulation equals 15%. Remarkably, noise can explain a large share of the (average) peak to trough decline in macroeconomic aggregates observed during U.S. downturns. It can account for the entire share in the decline of output and consumption and 57 percent of the decline in investment and hours worked.

Our model has similarities with several earlier studies that focus on expectations driven business cycles. We incorporate imperfect signals about productivity as in Beaudry and Portier (2004), although we go a step forward and allow agents to learn from these signals and importantly allow the signal's precision to vary in line with evidence from the Survey of Professional

Forecasters. Flodén (2007) uses a similar model—and the same vintage capital interpretation of investment specific technologies—as we do. He demonstrates how expectation driven cycles can arise naturally in this framework. However an important difference in our framework compared to his is the concept of learning we introduce; this feature can give rise to forecast errors that affect the economy's equilibrium and as we show amplify changes in fundamentals and trigger recessions independently of fundamentals. Jaimovich and Rebelo (2009) (henceforth JR) also develop a one sector model similar to ours. Their model differs from our setting in that we assume the vintage view of capital embodied productivity whereas they adopt the investment specific convention. In addition, relative to JR, we explore in greater detail the impact and importance of forecast errors in aggregate fluctuations and show how agents optimism and pessimism can amplify the economy's response to good and bad fundamentals respectively. Our model also shares similarities with the setup presented in Eusepi and Preston (2011) that focus on learning dynamics as a propagation mechanism. Similar to theirs, business cycles in our model are driven partly due to changes in fundamentals and partly due to agents forecast errors. One important difference is that the speed of learning varies in our model whereas it is constant in theirs. Lorenzoni (2009) develops a theory of demand (or noise) shocks based on a signal extraction problem and shows that they can account for a sizable fraction of demand side volatility. He uses a New Keynesian and focuses on total factor productivity shocks whereas we use a real model and analyze the impact of capital embodied shocks.

Last, our model rationalizes the relationship between learning and output fluctuations similar to earlier work by Cagetti et al. (2001), Evans et al. (1998) and Kasa (2000). In this work, agents solve a signal extraction problem and output fluctuations, as in our model, can arise as a result of learning rather than actual changes in productivity.

The remainder of the paper is organized as follows: Section 2 describes the model. Section 3 describes calibration, computation and data. Section 4 presents results from simulations and section 5 concludes.

2 The Model

We develop a model close in spirit to Greenwood et al. (1988) (henceforth GHH) with two important differences. First, in contrast to GHH, we interpret capital embodied technology as vintage specific rather than enhancing the productivity of current investment expenditures. This difference implies the productivity of current investment is unknown until capital is installed and used in production. Second, agents receive imperfect signals about the productivity of future capital vintages and use Bayesian learning to form expectations about this productivity. This concept of learning we implement implies that agents make forecast errors that can in turn give rise to fluctuations that would not otherwise arise had agents possessed perfect information. The model can produce expectations driven cycles. The mechanism in the model works in much the same way as in Greenwood et al. (1988) and analyzed in detail in Flodén (2007). We only briefly describe it here. When agents expect the productivity of future capital to increase, they feel wealthier. They correctly infer current capital to have lower productivity compared to the newly installed capital. Hence, depreciation today is relatively cheaper than depreciation tomorrow. This encourages agents to use the existing capital stock more intensively and increase investment in order to benefit from the high productivity of capital's vintage to be installed tomorrow. Additionally, agents want to increase consumption due to the wealth effect. The simultaneous increase in consumption and investment is permitted by the higher output resulting from the intensified use of capital. More intensified capital usage also increases the marginal product of labor leading to an increase in hours worked if the substitution effect dominates the wealth effect. Hence, in response to higher expected productivity of future capital it is possible for output, investment, consumption, hours worked and capital utilization to all increase without any change in current fundamentals.

2.1 Firms

The economy comprises of a continuum of perfectly competitive identical firms with unit mass. Firms produce output, y_t, using a Cobb-Douglas production function with three inputs. The production function is given by,

$$y_t = (u_t k_t)^\alpha h_t^{1-\alpha}, \qquad 0 < \alpha < 1 \tag{1}$$

where k_t denotes the sum of all efficiency units of capital available for production in period t and is defined by:

$$\sum_{s=0}^{\infty} q_{t-s} k_{t,s} = k_t, \tag{2}$$

where $k_{t,s}$ is capital of vintage s that is available at time t. This formulation assumes that the aggregate capital stock contains distinct vintages of capital which are associated with different levels of productivity, q. In addition, capital can be utilized at different rates. The utilization rate is denoted by u_t and hours worked by h_t.

Denoting investment by i_t, the vintages of capital evolve according to:

$$k_{t+1,s} = \begin{cases} i_t & \text{for} \quad s = 0 \\ (1 - d(u_t))k_{t,s-1} & \text{for} \quad s \geq 1. \end{cases} \tag{3}$$

Using (3), the economy's capital accumulation constraint can be derived from equation (2) as

$$k_{t+1} = [1 - d(u_t)]k_t + q_{t+1}i_t, \qquad k_0 > 0 \text{ is given.} \tag{4}$$

Note that the capital accumulation constraint differs from the standard formulation such that capital in period $t + 1$ depends on the capital embodied shock q_{t+1}. Thus, the productivity of investment is unknown until the capital is actually installed.

One can interpret the expression q in the capital accumulation equation as the productivity of a new vintage of capital, whereas the productivity of installed capital remains constant, or as the efficiency of the production of investment goods. Both interpretations exist in the literature, but the timing differs. If q is interpreted as the efficiency of the production of investment goods, it makes sense to assume that there exists information about the production function of these investment goods at the time of the actual production. In the vintage specific case however, q is interpreted as the productivity of a new vintage of capital, the productivity of which is unknown in the period when investment occurs. The productivity may be known (or at least can be forecasted more accurately) in the period after the investment has been made, i.e. when the capital is actually installed and used in production.[2]

Finally, the depreciation rate of capital, $d(u_t)$, depends positively on the degree of capital utilization as follows,

$$d(u_t) = \delta + \mu(u_t^{\omega} - 1), \qquad \mu > 0, \quad \omega > 1, \quad 0 \leq \delta \leq 1.$$

Since $d(u_t)$ is strictly increasing and convex, more intensive use of capital accelerates depreciation exponentially. In this function, ω measures the costliness of varying the capital utilization in terms of capital depreciation and the elasticity of marginal capital utilization equals $\omega - 1$. The steady state depreciation rate is given by δ. The parameter μ allows to calibrate utilization and depreciation in the steady state independently from each other consistent with steady state utilization equal to unity.

Firms in this economy maximize profits period by period, that is $\max_{h_t, u_t, k_t} \Pi_t = y_t - r_t^k u_t k_t - w_t h_t$, by renting capital and labor services at the beginning of the period from households in perfectly competitive factor markets, subject to the production function (1). The rental rate of capital and the real wage rate are denoted by r_t^k and w_t, respectively.

[2]These interpretations of q and the associated timing assumptions are widespread in the literature. An exception is Greenwood et al. (1988). They interpret q_t as the productivity of the capital in period $t + 1$, which is already known at the beginning of period t.

2.2 Technology

We now describe the technology that determines capital productivity. Our goal in this section is not to develop a fully endogenized model that determines the productivity of future capital vintages (see for example Comin and Mulani (2009) for a growth theory that is based on the general innovation concept) but to use a parsimonious way to make learning about capital embodied productivity interesting in our one sector framework.[3] We assume that the state of productivity of each future vintage can take on two values, a high value, denoted by η^H and a low value denoted by η^L. We furthermore assume that the future state of productivity is influenced by the number of general innovations available for adoption. This assumption can be motivated by the fact that in the aggregate, sectors that produce capital equipment benefit from general innovations that are adopted widely across the economy. One such important innovation has been the advent and widespread use of Information Technology (IT).[4] Some empirical evidence for this channel is provided in Basu et al. (2003). Basu et al. (2003) report that both IT producing and IT using industries in the US have experienced significant acceleration of total factor productivity (TFP) growth in the post-1995 period, coinciding with the IT equipment investment boom of the 1990s.[5] We parameterize these considerations in the process below,

$$q_{t+1} = \eta_{t+1} v_t^{\kappa_t} + \epsilon_{t+1}, \qquad \text{with} \quad 0 < \kappa_t < 1. \tag{5}$$

where, v_t is the number of new general innovations available for adoption in period t and η_{t+1} is an ergodic two-state Markov process with $\eta_{t+1} \in \{\eta^L, \eta^H\}$. The term ϵ_{t+1} is $i.i.d.$ with mean zero and constant variance σ_ϵ^2. This latter term constitutes noise in our model.
The number of new general innovations available for adoption follows the process:

$$v_t = (1-\rho) + \rho v_{t-1} + \xi_t, \qquad \text{with} \quad v_0 = 1, \quad 0 < \rho < 1,$$

where ξ_t is i.i.d. with mean zero and variance σ_ξ^2. We assume only a fraction of the available innovations, $v_t^{\kappa_t}$, are adopted since there may be innovations that will not improve capital's

[3]This endeavor will require a fully fledged endogenous multi-sector technological change model which is beyond the scope of this paper.

[4]Some have argued that the advent of IT (the computer revolution) and its incorporation into production has slowly pushed the average rate of embodied technological change higher (see Greenwood and Yorukoglu (1997), Helpman and Trajtenberg (1994) among others), especially after 1973.

[5]Of course this acceleration of TFP assumes that the official price indices do not fully reflect quality embodiments. Examples of general innovations include, personal computers, internet search engines, the Ford assembly line, management practices, financial innovations and others.

productivity. We therefore require that $\kappa_t \in (0,1)$. We assume that κ_t is given by:

$$\kappa_t = \frac{1}{1 + exp\{-(\tau \frac{v_t - v_{t-1}}{v_{t-1}})\}}, \quad \tau > 0.$$

Comin (2009) suggests that the adoption behavior of general innovations is pro-cyclical over the business cycle. The formulation for κ_t above is consistent with this consideration. Productivity of future vintages of capital, q_{t+1} can thus change either as a result of a state change or an change in the number (or the rate of adoption) of new innovations.

2.3 Households

The economy is populated by a unit measure of identical, infinitely lived households. The representative household maximizes the discounted stream of expected utilities over its lifetime

$$\max_{c_t, k_{t+1}, h_t, u_t} E_0 \sum_{t=0}^{\infty} \beta^t U(c_t, h_t, x_t), \qquad 0 < \beta < 1. \tag{6}$$

subject to a flow budget constraint,

$$c_t + k_{t+1} = (1 - d(u_t))k_t + w_t h_t + r_t^k u_t k_t \tag{7}$$

and the capital accumulation equation, (4).

Households supply labor and capital in perfectly competitive markets and earn a wage rate w_t and a rental rate r_t^k.

The utility function is given by

$$U(c_t, h_t, x_t) = \frac{(c_t - \phi h_t^{1+\gamma} x_t))^{1-\sigma} - 1}{1-\sigma}, \qquad \text{with } \gamma \geq 0,\ \phi > 0,\ \sigma \geq 1,$$

where

$$x_t = c_t^{\chi} x_{t-1}^{1-\chi}, \qquad 0 \leq \chi \leq 1.$$

and c_t, denotes consumption and h_t denotes hours worked. The parameter γ is the inverse of the Frisch elasticity of labor supply and σ is the intertemporal elasticity of substitution parameter. The specification of the utility function follows Jaimovich and Rebelo (2009) and nests two preference classes. For $\chi = 0$ the utility function has the properties of the class proposed by Greenwood et al. (1988) and for $\chi = 1$ one obtains preferences as discussed in King et al. (1988). As long as $\chi > 0$ the utility is time-non-separable in consumption and hours worked. It further implies stationary hours worked. The household's optimality conditions will be presented collectively in the social planner's problem formulation in section 2.6.

2.4 Information and Forecasting

We now turn to describe the information assumptions and the expectation formation mechanism in this economy. Agents enter period t with information set $\Im_t \equiv \{k^t, q^t, v^t, \kappa^t, x^{t-1}, c^{t-1}, i^{t-1}, u^{t-1}, h^{t-1}, y^{t-1}, w^{t-1}, r^{t-1}\}$, where z^t denotes the infinite history of any variable z that belongs to the information set above. The agents in this economy face a simple signal extraction problem. They observe the whole history of q but do not observe the state, η or noise, ϵ separately. Agents know the distribution of the noise, ϵ, and are aware that the state (or signal), η, follows an ergodic two-state Markov process with states η^L and η^H and a transition matrix Π. For the agent's investment decision today it is essential to predict tomorrow's capital productivity. At the beginning of period t agents—conditional on $\Im_t$—form expectations about productivity in period $t+1$ using Bayesian updating.

Specifically, agents evaluate the posterior probability of η_t to be in a high state as follows:

$$P(\eta_t = \eta^H | \Im_t) = \frac{\Psi(q_t | \eta_t = \eta^H, \Im_t) P(\eta_t = \eta^H)}{\Psi(q_t | \eta_t = \eta^H, \Im_t) P(\eta_t = \eta^H) + \Psi(q_t | \eta_t = \eta^L, \Im_t)(1 - P(\eta_t = \eta^H))}. \quad (8)$$

Here, $\Psi(\cdot)$ denotes a normal probability density function. The probabilities of a state change are stated in the transition matrix

$$\Pi = \begin{bmatrix} p_{HH} & p_{LH} \\ p_{HL} & p_{LL} \end{bmatrix}, \quad (9)$$

where p_{ij} denotes the probability that the economy transits from state i to state j. From the ergodicity of the Markov chain it follows that $p_{ij} \in (0,1)$ and $p_{iH} + p_{iL} = 1$. We further assume the transition matrix Π to be symmetric in order to ensure that all asymmetry in the resulting dynamics is endogenous. This assumption and the previous equality implies that $p_{HL} = p_{LH}$ and $p_{HH} = p_{LL}$.

The product of the posterior probabilities that productivity was in state η^L, η^H in period t as computed in (8) above, with the transition matrix imply a prior belief about the probability of η to be in a certain state in period $t+1$:

$$[P(\eta_t = \eta^H | \Im_t), P(\eta_t = \eta^L | \Im_t)]\Pi = [P(\eta_{t+1} = \eta^H | \Im_t), P(\eta_{t+1} = \eta^L | \Im_t)]. \quad (10)$$

Finally, this prior belief allows agents to form an expectation for the productivity of capital in period $t + 1$. Since $E_t \epsilon_{t+1} = 0$, using (5) the expectation is given by

$$\tilde{q}_{t+1} = E_t \tilde{\eta}_{t+1} v_t^{\kappa_t}, \tag{11}$$

$$\text{with} \quad \tilde{\eta}_{t+1} = [P(\eta_{t+1} = \eta^H | \Im_t), P(\eta_{t+1} = \eta^L | \Im_t)] \begin{bmatrix} \eta^H \\ \eta^L \end{bmatrix},$$

where $\tilde{z}_{t+1}$ denotes the forecasted value in period t for the realization of any variable, z in $t+1$. The essence of the learning mechanism in the model works as follows. The process for productivity (5) implies that the productivity of a new vintage is determined by the amount of adopted innovations $v_t^{\kappa_t}$, a signal component η and a noise term, ϵ. The signal and the noise component cannot be observed by agents. However, by the Bayesian updating process described in (8) – (11) agents are able to make forecasts for η and therefore next period's productivity. This formulation implies that an increase in the amount of adopted innovations precipitates an increase in the precision of the posterior probability of the state of η. More precise information about the posterior probability implies an increase in the forecast precision for next period's productivity. At the beginning of a boom phase forecast precision is low (as the rate of adoption of general innovations is procyclical; evidence consistent with this is reported by Comin (2009)) but it rises with the amount of adopted innovations. The peak of the boom is characterized by very high productivity. This implies a very high signal precision and a relatively precise forecast for next period's productivity. At this stage, the response of capital's expected future productivity to a negative signal is much stronger than the reaction to a similar shock in a situation with low precision. At the peak of the boom, agents learn much faster and therefore a negative signal can potentially trigger a quick and sharp adjustment of the economy. Hence, the boom phase is, due to the learning mechanism, more gradual than the bust phase.

A key assumption in generating an empirically realistic (i.e. consistent with forecast moments from the Survey of Professional Forecasters) learning process is a pro-cyclical signal-to-noise ratio. This requires that the variance of the noise term (ϵ_{t+1}) rises at a slower rate than the variance of the signal, adjusted by the amount of adopted innovations ($\eta_{t+1} v_t^{\kappa_t}$) when productivity increases. In other words, during a boom, the impact of the noise on next period's productivity becomes relatively smaller compared to the impact of the signal and vice versa during a recession. The signal-to-noise ratio from (5) equals

$$\frac{var(\eta_{t+1} v_t^{\kappa_t})}{var(\epsilon_{t+1})} = var(v_t^{\kappa_t}) \frac{\sigma_\eta^2}{\sigma_\epsilon^2}.$$

Since both the variance of the Markov process σ_η^2 and the noise variance are constant, the signal-to-noise ratio is pro-cyclical if $var(v_t^{\kappa_t})$ increases in a boom and decreases in a recession.[6]

[6]Our calibration procedure ensures that $var(v_t^{\kappa_t})$ is procyclical. The noise variance is restricted to be constant

Discussion of learning. The key ingredient of the model is learning capital's productivity. Evidence of learning regarding vintages of capital embodied technological change is documented in Bahk and Gort (1993). Using data from 1973 to 1986 consisting of 2,000 firms from 41 industries, Bahk and Gort (1993) find that a plant's productivity increases by 15 percent over the first fourteen years of its life due to learning effects. The idea here is that the installation of new vintages of capital equipment is often associated with complementary investments in training workers as well as implementation of new organization structures or management practices and these take time to become fully productive. This process was coined by Arrow (1962), ”learning by doing". Uncertainty about capital's productivity before and after its installation is also documented elsewhere in the investment literature.[7] These considerations suggest learning about the productivity of future vintages as a natural assumption to incorporate in the model. In the model it also takes time for agents to learn the productivity of a new vintage, although there are no explicit ”learning by doing" effects. Agents learn over time and asymptotically know with certainty the true productivity of a specific vintage.

An implication of learning in the model is that forecast precision is positively correlated with productivity. A possible explanation for this correlation might be that an increase in productivity enables firms to devote more resources to market research which leads to improvements in forecast accuracy. Some corroborative evidence for this implication of the model is available using data from the Survey of Professional Forecasters (SPF). We regress the cyclically (using the HP filter) adjusted labor productivity on the forecast errors for nominal GDP. Consistent with the prediction of the model the size of the *absolute* forecast errors in output are negatively correlated with labor productivity.[8]

In sum, the speed of learning—measured as the speed in which beliefs converge to the truth—increases during booms and decreases during recessions due to varying levels of forecast precision over the business cycle. The specific mechanism used in this paper is similar in spirit to the one used in Van Nieuwerburgh and Veldkamp (2006); both rely on changes in signal quality over the business cycle. Signals at the peak of the business cycle contain more precise information and cause stronger reactions by agents compared to signals in the trough, hence a bad signal at the peak will have a more pronounced effect compared to a good signal at the trough. However, while in the set up of Van Nieuwerburgh and Veldkamp (2006) signal quality increases during recovery due to higher factor inputs, in our model signal quality increases due

only for simplicity. This restriction can be relaxed: The noise variance can vary over time as long as it is guaranteed that the signal-to-noise ratio remains pro-cyclical.

[7] Dixit and Pindyck (1994) for example argue that the combination of irreversibility and uncertainty about capital's productivity creates an option value of waiting to invest. Jovanovic and Lach (1989) describe that due to learning-by-doing the uncertainty about capital's productivity decreases after its installation.

[8] We regress the log absolute forecast error of one to five quarters ahead forecasts for nominal GDP on HP filtered labor productivity. Forecast errors are constructed using data from the Survey of Professional Forecasters (1968 Q4-2009 Q2).

to an increase in the amount of adopted innovations.

2.5 Equilibrium

Equilibrium in the decentralized economy described above is a sequence of quantities and prices that solve: (1) firms' problem, (2) households' problem and (3) satisfy market clearing. Market clearing implies the aggregate resource constraint,

$$y_t = c_t + i_t. \tag{12}$$

2.6 The Social Planner Problem

The decentralized economy has a social planner analog. We work with this formulation. A benevolent social planner maximizes the utility of the representative agent (6), subject to the capital accumulation constraint (4) and the resource constraint (12). The planner's problem can be formulated in a recursive way. At the beginning of period t, η_t and ϵ_t are realized but cannot be observed. However, the social planner observes the productivity of capital installed in period t, q_t. The planner uses the forecasting mechanism described in (8) – (11) to form an expectation about the productivity of the vintage in period $t+1$, $\tilde{q}_{t+1}$. Hence, the social planner enters the period with state variables $s_t = (k_t, x_{t-1}, \tilde{q}_{t+1})$. The state variables determine the choice of h_t, k_{t+1} and u_t. Since the choice of investment depends on $\tilde{q}_{t+1}$, the value of the state variable k_{t+1} can differ from the realized capital stock in period $t+1$. This depends on the difference between $\tilde{q}_{t+1}$ and q_{t+1} and hence on forecast precision. Consumption in turn is determined from the recourse constraint. Formally, the planner solves:

$$\begin{aligned}
V(k_t, x_{t-1}, \tilde{q}_{t+1}) = \max_{h_t, u_t, k_{t+1}} & \; U(c_t, x_{t-1}, h_t) + \beta E_{t|k_t, x_{t-1}, \tilde{q}_{t+1}}[V(k_{t+1}, x_t, \tilde{q}_{t+2})] \\
\text{s.t. } k_{t+1} &= (1 - d(u_t))k_t + i_t q_{t+1} \\
c_t &= (u_t k_t)^\alpha h_t^{1-\alpha} - i_t \\
x_t &= c_t^\chi x_{t-1}^{1-\chi}
\end{aligned}$$

with x_{-1}, q_1 and k_0 given. V denotes the value function.

This yields the first-order conditions:

$$(c_t - \phi h_t^{1+\gamma} x_t)^{-\sigma} - \chi \psi_t c_t^{\chi-1} x_t^{1-\chi} = \lambda_t, \tag{13}$$

$$\psi_t - (c_t - \phi h_t^{1+\gamma} x_t)^{-\sigma} \phi h_t^{1+\gamma} = \beta E_t \psi_{t+1} (1-\chi) E_t c_{t+1}^{\chi} x_t^{-\chi}. \tag{14}$$

$$(c_t - \phi h_t^{1+\gamma} x_t)^{-\sigma} \phi (1+\gamma) h_t^{\gamma} x_t = \lambda_t (1-\alpha)(u_t k_t)^{\alpha} h_t^{-\alpha}, \tag{15}$$

$$\frac{\pi_t}{\lambda_t} = \frac{1}{E_t q_{t+1}}, \tag{16}$$

$$\alpha u_t^{\alpha-1} k_t^{\alpha} h_t^{1-\alpha} = \frac{\pi_t}{\lambda_t} \mu \omega u_t^{\omega-1} k_t, \tag{17}$$

$$\pi_t = \beta E_t \Big\{ \lambda_{t+1} \alpha u_{t+1}^{\alpha} k_{t+1}^{\alpha-1} h_{t+1}^{1-\alpha} + \pi_{t+1}(1-\delta-\mu(u_{t+1}^{\omega}-1)) \Big\}, \tag{18}$$

where π_t is the multiplier on the capital accumulation equation, λ_t is the multiplier on the resource constraint, and ψ_t the multiplier on the equation that defines the auxiliary variable x_t. Equations (13) and (14) determine optimal consumption. Equation (15) sets the household's marginal rate of substitution between consumption and hours worked equal to the real wage and determines labor supply. Note that for $\chi > 0$, the intertemporal decision for optimal hours worked depends on the real wage rate as well as on consumption. Equation (16) determines the real price of investment and is given by the ratio of the two multipliers. Equation (17) determines the optimal rate of capital utilization by setting the marginal user cost equal to the marginal benefit of capital services. The marginal user costs of capital on the right hand side of the equation consists of the partial derivative of $d(u_t)$ with respect to u_t, which represents the marginal cost in terms of increased depreciation of using capital at a higher rate. This cost is scaled by $1/q_{t+1}$, which determines current replacement costs of old capital in terms of new capital. Finally, equation (18) determines optimal investment. Note that we interpret q as the productivity of a new vintage of capital rather than as the efficiency of the production of investment goods. This preserves the model's ability to generate business cycles that are driven by expectations shifts. We will show below that the model is able to generate the comovement pattern of the macroeconomic aggregates which is typical for expectations driven business cycles.

It is important to note that the planner's problem described above is based on the assumption that the planner does not take into account the effect of optimal choices on the evolution of beliefs. Thus there is no feedback between actions and beliefs in this economy and learning is passive. This is similar to Van Nieuwerburgh and Veldkamp (2006) but different from Eusepi and Preston (2011) who allow actions to affect beliefs. The possibility of active learning would invalidate the Welfare theorems in the social planning economy and hence there will be no decentralized counterpart to the planner's equilibrium.[9] The passive learning is reflected in the iteration process of the social planner described above: Expectations of capital's productivity are formed at the beginning of the period. Based on these expectations and the endogenous state variables (k_t, x_{t-1}), optimal actions are chosen. Given these, expectations are updated at the beginning of the next period. This process is repeated until the expectations coincide with the actual policies.

3 Calibration, computation and data

3.1 Calibration

Table 1 reports the parameter values used for calibrating the model. The model is calibrated on a quarterly basis. We assume the depreciation rate of capital, δ = 0.025, quarterly discount factor, β = 0.99 and the capital share of production, α = 0.36. These are all standard values in the literature.

Our calibration of the inverse of the Frisch labor supply elasticity and the parameter which determines the costliness of varying the capital utilization are based on the values used in Jaimovich and Rebelo (2009). Setting the inverse of the Frisch labor supply elasticity γ = 0.4 is a value widely used in the literature implying an intertemporal elasticity of substitution for labor supply of approximately 2.5. In general there are no widely accepted guidelines in the empirical literature about the magnitude for the parameter which determines the costliness of capital utilization. Setting ω = 1.15 implies an elasticity of marginal capital utilization of 0.15. We set σ = 1.0 corresponding to logarithmic utility. The parameter that determines the wealth effect on labor supply, χ is set equal to 0.001, (almost) corresponding to GHH preferences. Finally, ϕ and μ are free parameters and we calibrate these to guarantee that capital utilization is equal to unity and hours worked are equal to one third of the total time endowment in the steady state.[10]

The following parameters are specific to the learning and the productivity process. In order to compute the probability of a state change in productivity we first re-write the ergodic two-state Markov chain as an AR(1) process. Since the transition matrix is symmetric the autoregressive parameter is given by $(2p_{HH} - 1)$. The relative price of investment (i.e. the price of investment relative to consumption goods) should provide a good empirical measure of the quality improvements embodied in new capital. Hence we use this relative price in order to

[9] In an economy with active learning the provision of information is a public good and information externalities emerge in this case. This implies that the provision of information will collapse as no agent will have an incentive to confer benefits on other agents.

[10] The derivation of the expressions for ϕ and μ can be found in Appendix 1.

calibrate the parameters of the productivity process. Specifically, we use the (detrended) measure of relative price of investment constructed by Fisher (2006) which has an autocorrelation of 0.99. There are other estimates (e.g. Greenwood et al. (2000) using a different relative price series) indicating a first-order serial correlation equal to 0.64. We give more weight to Fisher's measure and set $p_{HH} = (0.9+1)/2 = 0.95$. It then follows from the structure of the transition matrix that the probability of a state change equals 0.05.

The parameter τ in the equation that describes κ_t governs the impact of the growth rate of general innovations on their adoption rate. While the empirical literature provides indications about the qualitative changes (see Comin (2009)) in the adoption behavior of general innovations over the business cycle, it is silent about the quantitative changes. We choose τ = 45 which ensures that changes in the growth rate of general innovations cause a substantial change in the adopted share. While this value is somewhat arbitrary it nevertheless guarantees that κ_t visits all values in its domain during the simulations equally likely.

The next objective is to calibrate the standard deviations of the three processes and the autocorrelation parameter of the adoption process, ρ. Ideally we want to strike a balance between the size of the noise variance and the variance of the signal such that learning about capital productivity is difficult. The relation between these variances implies a certain signal precision since it determines the difficulty to learn: the noise variance must be high enough to make a boom look like a recession. If it is very low, learning is trivial. However, if the noise variance is very high, estimates about the current state of η will be quite inaccurate and this makes learning almost impossible. Estimates for the signal precision of investment-specific technological change are not available due to a lack of forecast data for this variable. We calibrate σ_η, σ_ϵ, σ_ε and ρ in order to match as close as possible three moments from the Survey of Professional Forecasters: signal precision (mean absolute forecast error), standard deviation and serial correlation of forecast errors for GDP. This choice guarantees the average ”difficulty" of learning in the model is similar to that observed in the data.[11]

The calibration above implies a standard deviation for the noise, σ_ϵ = 0.01. As we compare percentage deviations in the model the absolute values of η^H and η^L are not relevant. However, the distance is important since it has an impact on the volatility of the Markov chain. Assigning the values [0.93, 1.07] to η^L and η^H implies a standard deviation σ_η = 0.07. Finally, this calibration procedure implies σ_ξ = 0.035 and ρ = 0.8.[12] The calibration of these parameters guarantees that the model generates procyclical rate of adoption consistent with Comin (2009).[13]

[11]The targeted/model moments are: signal precision (0.39/0.32), standard deviation (0.85/1.05), serial correlation (-0.023/-0.035). For this calculation we use the one quarter ahead forecasts for nominal GDP from 1968:4 to 2009:2. This is the longest forecast series available from this survey.

[12]We also run simulations of the model with ρ = 0.9 or $\rho = 0.7$, values that without any material change in our results.

[13]In the simulations of the model with the learning mechanism, the correlation between output and the number of adopted general innovations is 0.61.

Table 1: Parameter values—baseline calibration

β	=	0.99	ω	=	1.15	τ	=	45
α	=	0.36	χ	=	0.001	ρ	=	0.8
δ	=	0.025	ϕ	=	0.5822	σ_ξ	=	0.035
γ	=	0.4	μ	=	0.0305	σ_η	=	0.07
σ	=	1.00	p_{HH}	=	0.95	σ_ε	=	0.01

3.2 Computational details and data

The model is solved using value function iteration. We use the policy functions to simulate the model 500 times over 255 periods. The first 50 periods of each simulation are discarded to avoid influences due to the choice of the starting values. Statistics are calculated over the remaining 205 periods corresponding to the sample size (1958 Q2 to 2009 Q2). Second moments are calculated from HP filtered series. Since the model is calibrated on a quarterly basis the smoothing parameter is 1600. Skewness is calculated from first-differenced series.

The U.S. data for output is real GDP (GDPC96). Investment is defined as gross private domestic investment (GPDIC96) and consumption is real personal consumption expenditures (PCECC96). These series are quarterly, seasonally adjusted and in billions of chained 2005 dollars from the US Department of Commerce, Bureau of Economic Analysis (BEA). The data for hours of all persons in the non-farm business sector (HOANBS) and the series of civilian non-institutional population (CNP16OV), used to derive per-capita time-series, are from the US Department of Labor, Bureau of Labor Statistics. The forecast error statistics we use for nominal GDP (NGDP) are from the Survey of Professional Forecasters. This survey pools professional forecasters to obtain one to five quarter ahead predictions for different variables. Data availability issues restrict us from using the same window as for the macroeconomic aggregates. We use forecasts from 1968 Q4 to 2009 Q2 to calculate the forecast errors. The forecast error for a given quarter is the log absolute difference between the median of all forecasters predictions for nominal GDP and the final revised value of nominal GDP as it appears today.

4 Results

Our first goal is to evaluate the model ability to match a set of business cycle statistics. We then evaluate the effects of learning and focus on (a) the ability of the model to generate asymmetry of cycles, (b) the effects of forecast errors and (c) characteristics of recessions. To do this we compare the outcomes of a model which allows agents to learn over the business cycle with the outcomes of a model without learning (i.e.perfect information case). Business cycle asymmetry is measured by the skewness of macroeconomic aggregates. The more gradual the boom and the sharper the recession, the more negative is the skewness measure. In Appendix 3 we briefly discuss the ability of the model to generate co-movement in response to a shift in expectations of capital productivity. Flodén (2007) has extensively analyzed the theoretical restrictions that need to be satisfied for this vintage capital model to generate this type of co-movement.

4.1 Business cycle statistics

In this section we evaluate the ability of the model to match business cycle statistics computed from U.S. data. We focus on relative volatilities, serial correlations, co-movement and growth asymmetry. We compute second moments from HP filtered series. We evaluate asymmetry by computing a variable's skewness from its log first difference. If negative changes are larger than positive changes as in the data then variables will exhibit negative skewness.[14] We simulate two versions of the model: a no-learning, perfect information version and the full version which incorporates learning about productivity. In the no-learning version, agents observe the state of η_t at the beginning of period t and hence have a perfect signal about productivity (except for the i.i.d. noise), whereas in the learning version the state of η_t is not revealed.[15]

Table 2 reports various moments from the data (panel A) and the two versions of the model (panel B and C). Both versions of the model match reasonably well the relative volatilities and correlations with output. Specifically, both correctly rank investment to be more volatile than output and consumption to be less volatile than output. However, they under-predict the volatility of hours worked which is more volatile than output in the data.[16] They also match reasonably close the serial correlations, although the full model generates slightly lower serial correlations compared with the data. This is a direct consequence of the difference in the serial correlation between the actual, q and forecasted productivity, $\tilde{q}$. In the learning version, the latter's serial correlation is markedly lower compared to the true process; for learning to be realistic (i.e. neither impossible nor trivial) the noise shock has to be big enough to make a boom look like a recession. This however implies that agents' may wrongly infer a state change in capital's productivity when none has occurred. Thus conceptually, agents' forecasted productivity is ”changing state" more often than true productivity and this imparts a lower

[14] Since the HP filter is a two-sided filter, information from the past as well as the future are used. De-trending with this filter implies that agents have information about the future which can have an impact on their decision today. Using a two-sided filter diminishes the filtered values prior to a downturn. This reduces the magnitude of the bust and influences our evaluation of business cycle asymmetry. To avoid this distortionary effect of two-sided filters – such as the commonly used HP or bandpass filter – we calculate the variable's skewness from the log first-differences.

[15] The no learning version of the model differs from Greenwood et al. (1988) only by the fact that productivity of the newly installed capital is subject to the additive $i.i.d.$ shock ϵ which cannot be observed by the agents.

[16] The low relative volatility of hours is a well known problem of RBC models. It can be addressed by introducing for example the Hansen (1985) indivisible labor approach into the utility function.

autocorrelation in forecasts for $\tilde{q}$.

The main difference between the two versions of the model in Table 2 is with respect to the generated asymmetry. In particular, only the learning version can generate asymmetry in all variables in line with the data (panel C)—the no-learning version fails in this dimension. More precisely the point estimate in the data skewness measure is within two standard deviations of the model's skewness. This can be seen in the last column of Table 2. In panel B the skewness of output is close to zero, indicating that boom and recession phases are symmetric. Moreover the future productivity of capital, q, has skewness close to zero. Since agents perfectly observe the signal— as the state of η_t is revealed at the beginning of period t—their forecast for productivity, $\tilde{q}$, differs from q only by the additive noise shock, ϵ_{t+1}. This noise shock on its own is not a source of asymmetry which explains why the skewness of $\tilde{q}$ is close to zero and very similar to the one for q. Since there is no other mechanism in the model to make booms longer and more gradual than recessions, all other macroeconomic aggregates exhibit skewness which is close to zero as well. The main reason for the generated growth asymmetry is that in the learning version the skewness of agent's forecast for productivity, $\tilde{q}$, is negative in contrast to the no-learning version. The introduction of agent's learning over the business cycle is the crucial mechanism to generate growth asymmetries in line with the data. Booms tend to be more gradual than recessions because agent's speed of learning varies procyclically over the business cycle. The asymmetry in agent's forecast for productivity imparts negative skewness in the remaining macro-aggregates in the learning version. This effect is very strong for output, investment, hours worked and capital utilization while it is less so for consumption. These results demonstrate the importance of the learning mechanism to generate the growth asymmetries present in the data.

4.2 The role of forecast errors: optimism and pessimism

This section provides a more detailed analysis about the functioning of the learning mechanism by evaluating the role of forecast errors. Specifically we wish to examine the effects of optimism and pessimism on the cyclical fluctuations of the model. We define an agent as pessimistic (optimistic) when we observe a ”large" (to be defined below) negative (positive) forecast error in the simulation. A negative (positive) forecast error implies that agents underpredict (overpredict) capital's productivity.

We use the simulation set-up described in section 4.1 in order to study how forecast errors can affect the equilibrium allocations in the model. We simulate the learning and no-learning economies using an identical sequence of the shocks that determine the productivity of next period's capital (signal, noise and adoption process shocks).

Table 3 reports the behavior of the model in periods during which agents make forecast er-

Table 2: Key moments of macroeconomic aggregates

	Relative std deviation	First-order autocorrelation	Correlation with y	Skewness
Panel A: Data				
y	1.00	0.85	1	-0.27
i	4.62	0.79	0.90	-0.76
h	1.16	0.91	0.87	-0.74
c	0.80	0.87	0.87	-0.69
Panel B: Model without learning				
y	1	0.824	1	0.004
	(0.000)	(0.042)	(0.000)	(0.332)
i	3.537	0.782	0.948	-0.036
	(0.276)	(0.047)	(0.021)	(0.394)
h	0.746	0.811	0.992	0.022
	(0.009)	(0.044)	(0.002)	(0.331)
c	0.387	0.813	0.772	0.026
	(0.052)	(0.066)	(0.047)	(0.319)
u	1.488	0.818	0.991	0.019
	(0.016)	(0.044)	(0.002)	(0.333)
q	0.707	0.817	0.944	0.018
	(0.018)	(0.041)	(0.015)	(0.318)
$\tilde{q}$	0.692	0.832	0.970	0.012
	(0.015)	(0.039)	(0.008)	(0.332)
Panel C: Model with learning				
y	1	0.708	1	-0.171
	(0.000)	(0.078)	(0.000)	(0.312)
i	3.499	0.670	0.954	-0.170
	(0.260)	(0.079)	(0.017)	(0.335)
h	0.747	0.695	0.993	-0.157
	(0.008)	(0.078)	(0.002)	(0.297)
c	0.379	0.742	0.771	-0.043
	(0.052)	(0.090)	(0.043)	(0.341)
u	1.485	0.701	0.991	-0.194
	(0.016)	(0.078)	(0.002)	(0.322)
q	0.686	0.812	0.885	0.011
	(0.038)	(0.045)	(0.031)	(0.334)
$\tilde{q}$	0.689	0.710	0.972	-0.189
	(0.013)	(0.078)	(0.006)	(0.337)

Notes. Sample is 1958 Q2 to 2009 Q2. Values reported in parentheses are standard deviations. The model is simulated 500 times over 255 periods. The first 50 periods are discarded. Second moments are calculated from HP filtered series. Skewness is calculated from (log) first-differenced series. Variables included: Output (y), investment (i), hours worked (h), consumption (c), capital utilisation (u), productivity (q) and the forecast for productivity ($\tilde{q}$).

rors of all sizes (small and large). These periods account on average for 85% of the simulation. Table 3 reports the growth rate of variables under both positive and negative forecast errors. The columns ”Absolute distance” reports the absolute difference between the two economies. The main conclusion from inspecting these numbers is that forecast errors amplify the fluctuations observed in the perfect information economy, although this amplification is relatively modest (with the exception of investment). In periods with negative forecast errors agents underpredict productivity compared to the truth; variables decline more in the learning compared to the no learning economy. Similarly in periods with positive forecast errors agents overpredict productivity and variables rise more in the learning compared to the no learning economy.

Table 3 reports the behavior of the model when all forecast errors are considered (small and large). However, there are periods in the simulation when agents make big forecast errors. We also want to examine this special case where forecast errors can potentially have a large impact on the allocations of the model.

We examine the distribution of forecast errors obtained from the simulation and choose to examine forecast errors that exceed one standard deviation above or below the average forecast error. We label those errors as ”large”. This threshold generates forecast errors that occur in approximately 19% of the simulation. We observe large negative forecast errors in 10% of the simulation and large positive forecast errors in 9% of the simulation.[17] We calculate the mean growth in variables from the two economies. These results are summarized in Table 4.

We draw attention to the following facts from Table 4. First, agents are pessimistic when the true growth rate of productivity is negative and optimistic when the true growth rate of productivity is positive. The presence of noise makes it difficult for agents to accurately predict true productivity when the latter is changing and agents make substantial forecast errors when trying to predict the true process. Second, agents in the no learning economy always forecast capital productivity perfectly, thus no forecast errors occur in this economy. Changing fundamentals cause fluctuations in macroeconomic aggregates in both economies but errors in forecasting productivity amplify those fluctuations. The magnitude of amplification is quite substantial. In order to demonstrate this we look at the absolute distance, for each variable, between the learning and no-learning economies. This distance quantifies by how much equilibrium allocations differ due to forecast errors.

The distance in investment growth rates is larger among all variables followed by utilization, output and hours. The distance in investment growth is equal to 4.0% for negative forecast errors and 8.9% for positive forecast errors. In the learning version, pessimistic agents cut investment on average by 6.2% relative to a modest 2.2% when they possess perfect information. When agents are optimistic they raise investment by 25.4% compared to 16.5% in the perfect

[17] Using the forecast errors for GDP (one through four quarters ahead) from the SPF we compute that forecast errors exceeding the average by one standard deviation occur in approximately the same range as in the simulation, from 20% to 25% of the sample period.

information economy. In this case agents over-invest. This is an interesting finding because it has a parallel with the boom in investment rates observed during the IT boom-bust cycle in the 1990s. When agents are pessimistic, output in the learning economy declines by 3.5% compared to 1.6% in the no learning economy, while in periods of optimism output in the learning economy rises by 6.0% compared to 4.0% in the no learning economy. Similar differences occur in utilization rates and hours worked, while the difference in consumption allocations is relatively small.

4.3 Characteristics of recessions

We also want to examine the nature of recessions in the model economy. We define recessions in the model as periods with at least two quarters of negative output growth. Table 5 reports characteristics of recessions from the model and compares them with recessions from the data. Several findings are worth highlighting. First, the average length of the recession in the model is four quarters, very similar to that in the data (4.25 quarters). Second, recessions in the model cannot only be driven by **un-favorable fundamentals** but also by **noise** (with no change in fundamentals). The share of recessions that occur purely due to noise equals 15%. The remaining 85% of recessions are caused by unfavorable fundamentals. The noise triggered episodes coincide with agents mistakenly forecast productivity to be declining when true productivity is actually rising at the onset of the recession.

Table 5 reports two measures: the average growth of variables and the peak to trough changes for both types of recessions. There are two interesting findings. First, both noise and fundamental triggered recessions generate declines that are very similar in magnitude. For example, output declines on average by 1.7% in the noise triggered compared to 1.8% in the fundamentals triggered recession. Investment declines on average by 4.9% in the noise triggered compared to 5.2% in the fundamentals triggered recession. Second, the model's average growth declines match reasonably well the average growth declines in macro-aggregates observed during U.S. recessionary episodes.[18] For example, the model generates very similar average growth declines in hours worked, consumption and investment although overpredicts to some extent the output growth decline. We view these findings as a success of the model given it is driven by a single disturbance.

The declines from peak to trough are also similar for the two types of recessions in the model. For example, the decline in output is 3.7% in the noise triggered compared to 3.8% in the fundamentals triggered episode, whereas the peak to trough decline in investment is 10% and 10.3% in the noise and fundamental triggered episodes respectively. Interestingly,

[18] We have identified 8 recessions from the U.S. data, based on the NBER procedure (sample 1958 Q2 to 2009 Q2): 1960 Q2 to 1961 Q1, 1969 Q4 to 1970 Q4, 1973 Q4 to 1975 Q1, 1980 Q1 to 1980 Q3, 1981 Q3 to 1982 Q4, 1990 Q3 to 1991 Q1, 2001 Q1 to 2001 Q4, 2007 Q4 to 2009 Q2.

the noise driven recession can explain a large fraction of the peak to trough change in macroeconomic aggregates computed from the data. The last column of Table 5 reports the share of peak to trough decline in the data that can be accounted for by noise in the model. A noise triggered recession can account for all of the decline in output and consumption (arithmetically it accounts for over 100% of the decline in those aggregates) and 57% of the decline in investment and hours worked. This is a remarkable finding given that the model's exogenous processes were calibrated to match the time series behavior of forecast errors for GDP and not calibrated to match any statistic from Table 5 or statistics from aggregate macroeconomic variables (e.g. volatility and persistence of GDP) that could potentially overweight the model's ability to match recessions observed the data.

Table 3: The impact of all forecast errors

	Negative FE			Positive FE		
	No-Learning	Learning	Absolute distance	No-Learning	Learning	Absolute distance
$\Delta y/y$	-0.006	-0.010	0.004	0.011	0.017	0.007
$\Delta i/i$	0.004	0.002	0.002	0.060	0.083	0.023
$\Delta h/h$	-0.005	-0.008	0.003	0.008	0.012	0.004
$\Delta c/c$	-0.002	-0.004	0.002	0.003	0.005	0.002
$\Delta u/u$	-0.008	-0.013	0.005	0.019	0.028	0.009
$\Delta q/q$	-0.005	-0.005	0.000	0.008	0.008	0.000
Forecast error	0.000	-0.038	0.038	0.000	0.037	0.037

Notes. Variables included: Output (y), investment (i), hours worked (h), consumption (c), capital utilisation (u), productivity (q) and forecast error for productivity, computed as $E_{t-1}q_t - q_t$. The model is simulated 500 times over 255 periods each. The first 50 periods are discarded and the mean growth rate in variables of the economy with and without learning is calculated over the remaining periods.

5 Conclusions

Barro and King (1994) (BK) pointed out that changes in beliefs about the future cannot generate empirically recognizable business cycles within the standard real business cycle model. The BK challenge has been difficult to address. Specifically, it has been difficult to develop simple one sector models that can give rise to expectations driven fluctuations without a compromise on other dimensions. In this paper we propose a simple generalization of the Greenwood et al. (1988) framework and show it can generate fluctuations that arise as a result of agents difficulty

Table 4: The impact of large forecast errors

	Large negative FE Pessimism			Large positive FE Optimism		
	No-Learning	Learning	Absolute distance	No-Learning	Learning	Absolute distance
$\Delta y/y$	-0.016	-0.035	0.019	0.040	0.060	0.020
$\Delta i/i$	-0.022	-0.062	0.040	0.165	0.254	0.089
$\Delta h/h$	-0.012	-0.027	0.015	0.029	0.043	0.014
$\Delta c/c$	-0.006	-0.012	0.006	0.010	0.016	0.006
$\Delta u/u$	-0.023	-0.049	0.026	0.060	0.093	0.033
$\Delta q/q$	-0.013	-0.013	0.000	0.028	0.028	0.000
Forecast error	0.000	-0.091	0.091	0.000	0.092	0.092

Notes. Variables included: Output (y), investment (i), hours worked (h), consumption (c), capital utilisation (u), productivity (q) and the forecast error for productivity, computed as $E_{t-1}q_t - q_t$. The model is simulated 500 times over 255 periods each. The first 50 periods are discarded and the mean growth in variables of the economy with and without learning is calculated over the remaining periods. Forecast errors are defined to be large when their absolute value exceeds one standard deviation of the average forecast error.

Table 5: Recession statistics

	Average share of recessions						
	Model recessions due to noise (0.15)		Model recessions due to fundamentals (0.85)		U.S. data average recession†		Share explained by noise
	average growth	peak to trough change	average growth	peak to trough change	average growth	peak to trough change	
y	-0.017	-0.037	-0.018	-0.038	-0.007	-0.034	1*
i	-0.049	-0.100	-0.052	-0.103	-0.048	-0.176	0.57
h	-0.012	-0.028	-0.013	-0.029	-0.011	-0.049	0.57
c	-0.007	-0.016	-0.007	-0.016	-0.003	-0.011	1*
u	-0.020	-0.041	-0.021	-0.042	n.a.	n.a.	
q	-0.003	0.003	-0.003	0.003	n.a.	n.a.	
$\tilde{q}$	-0.007	-0.014	-0.007	-0.014	n.a.	n.a.	

Notes. Variables included: Output (y), investment (i), hours worked (h), consumption (c), capital utilisation (u), productivity (q) and the forecast for productivity ($\tilde{q}$). Share explained by noise is defined as peak to trough change (due to noise) over peak to trough change in U.S. data. †: growth and peak to trough declines computed as the average from all U.S. recessions. * Shares that exceed one in the Table above are set equal to one.

to forecast capital embodied productivity. The two key assumptions in the model are: (1) the vintage view of capital productivity, whereby each successive vintage has (potentially) different productivity and (2) agents' imperfect information and learning about this productivity. The main findings from simulations of the model are as follows. First, noise amplifies fundamentals in both directions, upward and downward. Fluctuations in the model are larger when agents make forecast errors compared to the perfect information case. Second, recessions can arise purely due to noise, i.e. even without a change in fundamentals. Third, pure noise can trigger recessions that mimic in magnitude, duration and depth the typical post WW II U.S. recession.

6 Appendix

Appendix 1: Determination of the Parameters ϕ and μ

The steady state value for hours worked is unity. This choice is based on the fact that the total amount of time per period is normalised to 3, and that agents use about 1/3 of their amount of time to work. Under the assumption that h = 1, the steady state capital stock can be derived from the steady state Euler equation (18):

$$\begin{aligned} 1 =& [\beta(\alpha u^{\alpha} k^{\alpha-1} q + (1 - d(u))] \\ \Leftrightarrow k =& \left\{ \left[\frac{1}{\beta} - (1 - (\delta + \mu(u^{\omega} - 1)) \right] \frac{1}{\alpha u^{\alpha} q} \right\}^{\frac{1}{\alpha-1}}. \end{aligned}$$

We want to calibrate the model in a way that u = 1 when h = 1. Considering the steady state expression for capital utilisation — which can be derived from the first order condition (17) — and using the equation above, one can derive a formulation of steady state utilisation solely depending on parameters:

$$\begin{aligned} u =& \left\{ \frac{\alpha}{\mu\omega} k^{\alpha-1} q \right\}^{\frac{1}{\omega-\alpha}} \\ \Leftrightarrow u =& \left\{ \frac{\alpha q}{\mu\omega} \left[\left\{ \left(\frac{1}{\beta} - (1 - (\delta + \mu(u^{\omega} - 1))) \right) \frac{1}{\alpha u^{\alpha} q} \right\}^{\frac{1}{\alpha-1}} \right]^{\alpha-1} \right\}^{\frac{1}{\omega-\alpha}} \\ \Leftrightarrow u =& \left\{ \left(\frac{1}{\beta} - (1 - \delta) - \mu \right) \frac{1}{\mu(\omega - 1)} \right\}^{\frac{1}{\omega}}. \end{aligned} \tag{.1}$$

From this equation one can derive an expression for μ as the steady state utilisation is chosen to equal unity:

$$\begin{aligned} 1 =& \left[\frac{1}{\mu(\omega-1)}\left(\frac{1}{\beta}-(1-\delta)-\mu\right)\right]^{\frac{1}{\omega}} \\ \Leftrightarrow \mu =& \frac{1}{\omega}\left(\frac{1}{\beta}-(1-\delta)\right). \end{aligned}$$

Note that the choice of u to equal unity allows to calibrate capital utilisation and depreciation in the steady state independently form each other, as $d(u) = \delta + \mu(u^{\omega} - 1) = \delta$.

A steady state capital utilisation of unity and a steady state productivity of q = 1 implies that the steady state capital stock can be expressed as:

$$k = \left\{\frac{1}{\beta\alpha} - \frac{(1-\delta)}{\alpha}\right\}^{\frac{1}{\alpha-1}}. \tag{.2}$$

Given u = 1, q = 1 and the steady state capital stock, one can numerically solve the steady state versions of equations (13) – (15) for ϕ, λ and ψ so that $h = 1$. Note that steady state capital utilisation and steady state employment are independent of the value of steady state productivity q. For capital utilisation this is shown in the derivation of equation (.1). The fact that steady state employment is unaffected by the value of q follows from the properties of the chosen utility function which guarantee stationary employment.

Appendix 2: Derivation of the Elements of the Bayesian Updating Formula

The Unconditional Probability of η to be in a Certain State:
The stochastic process for η is designed to be an ergodic two-state Markov process. By definition, an ergodic Markov chain has exactly one eigenvalue which equals unity. All other eigenvalues lie inside the unit circle. The eigenvector which is associated with the unit eigenvalue is therefore unique and can be interpreted as the vector of unconditional probabilities. For the above described ergodic two-state Markov chain the eigenvector associated with the unit eigenvalue turns out to be

$$P\left\{\begin{array}{c} \eta_t = \eta^H \\ \eta_t = \eta^L \end{array}\right\} = \left(\begin{array}{c} \frac{1-p_{LL}}{2-p_{HH}-p_{LL}} \\ \frac{1-p_{HH}}{2-p_{HH}-p_{LL}} \end{array}\right). \tag{.3}$$

The assumptions of a symmetric transition matrix and an ergodic two-state Markov chain imply that $p_{HH} = p_{LL}$. It follows from expression (eq:uncProb) that the unconditional probability of η to be in a high (low) state is 0.5. Thus, the formula for Bayesian updating (8) depends solely on the probability of q_t conditional on $\eta_t = \eta^H$ and $\eta_t = \eta^L$, respectively.

The Normal Probability Density Function:
There is always some uncertainty about the state of η. Therefore, all inference about η takes the form of statements of probability. Uncertainty is described in terms of the normal probability density function $\Psi(\cdot)$. Given the mean $E(q_t|\eta_t = \eta^H)$ and the variance $V(q_t|\eta_t = \eta^H)$ the normal probability density for every outcome q_t, given η_t being in a high (low) state, can be calculated according to

$$\Psi(q_t|\eta_t = \eta^H) = \frac{1}{V(q_t|\eta_t = \eta^H)\sqrt{2\Pi}} \exp\left(- \frac{(q_t - E(q_t|\eta_t = \eta^H))^2}{2(V(q_t|\eta_t = \eta^H))^2} \right). \tag{.4}$$

As $E_t\epsilon_t = 0$, the mean and the variance can be derived by using the formulation for productivity (5):

$$\begin{aligned} E(q_t|\eta_t = \eta^H) &= E_t[\eta^H v_{t-1}^{\kappa_{t-1}} + \epsilon_t] \\ &= \eta^H v_{t-1}^{\kappa_{t-1}}. \end{aligned}$$

$$\begin{aligned} V(q_t|\eta_t = \eta^H) &= E_t[(\eta^H v_{t-1}^{\kappa_{t-1}} + \epsilon_t)^2] - (E_t[\eta^H v_{t-1}^{\kappa_{t-1}} + \epsilon_t])^2 \\ &= E_t[(\eta^H v_{t-1}^{\kappa_{t-1}})^2 + 2\eta^H v_{t-1}^{\kappa_{t-1}} \epsilon_t + \epsilon_t^2] - (\eta^H v_{t-1}^{\kappa_{t-1}})^2 \\ &= \sigma_\epsilon^2. \end{aligned}$$

Knowing the unconditional probability of η to be in a certain state as well as the normal probability density (.4), allows — by using the Bayesian updating formula (8) — to derive the posteriori probability for η_t to be in a high (low) state.

Appendix 3: Expectations Driven Business Cycles

In this section we evaluate the model's ability to generate cycles that exhibit co-movement. We consider the model's equilibrium as given in equations (13) – (18) abstracting learning. For this exercise we solve the model in DYNARE using a local perturbation method. In this case we assume productivity q_t to follow an AR(1) process with persistence 0.64 and error term standard deviation of 0.035 as estimated by Greenwood et al. (2000). Figure 2 shows the case in which agents receive a signal in period 2 that the newly installed vintage of capital in period 3 will be more productive. There is no change in the productivity of capital in period 2. In response to this shift in expectations, output, consumption, investment, hours work and capital utilization all increase in period 2.

The mechanism that enables the co-movement of these variables works as follows. The expected increase in the productivity of tomorrow's capital, agents want to increase investment to benefit from the high productive vintage of capital. Simultaneously they want to increase consumption due to the wealth effect. The first order condition for capital utilization (17) in combination with (16) indicate that a rise in $E_t q_{t+1}$ triggers an increase in capital utilization. Agents intensify the use of today's capital stock to increase investment as depreciation today is relatively cheaper than depreciation tomorrow. Higher capital utilization leads to an increase in output which allows for a simultaneous increase in investment and consumption. It also raises the marginal product of labor leading to an increase in hours worked.

Figure 2. The economy's response to news in period 2 that the new vintage of capital in period 3 will be more productive

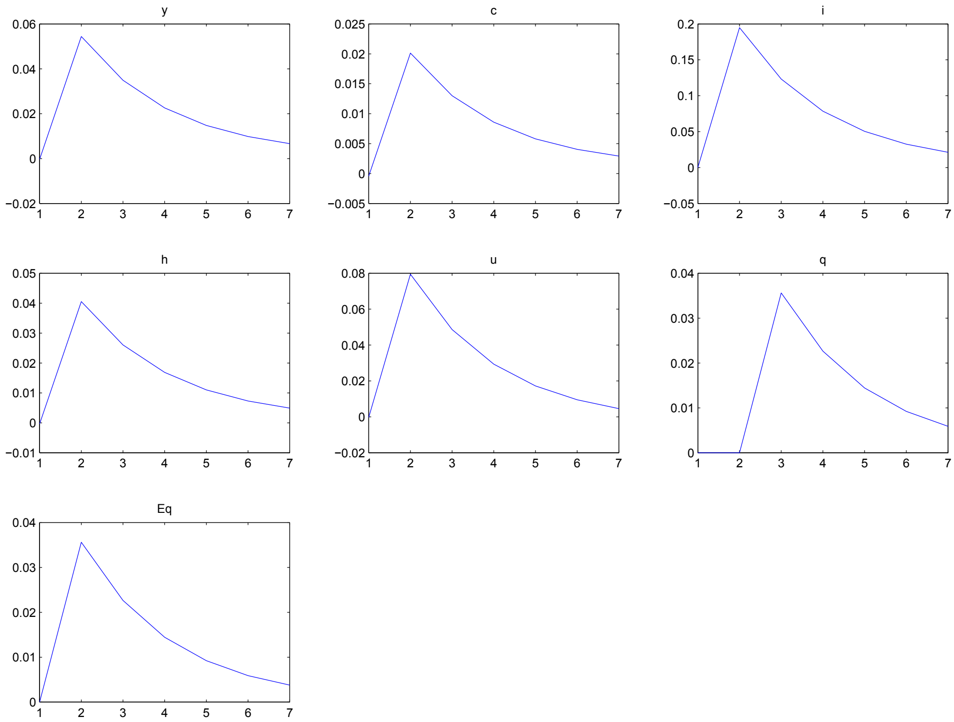

Notes. Notes. Variables included: Output (y), investment (i), hours worked (h), consumption (c), capital utilisation (u), productivity (q) and the forecast for productivity (Eq)).

References

Arrow, K. (1962). The economic imlications of learning by doing. *Review of Economic Studies*, 29(3):155–173.

Bahk, B.-H. and Gort, M. (1993). Decomposing learning by doing in new plants. *Journal of Political Economy*, 101(4):561–83.

Barro, R. and King, R. (1994). Time-separable preferences and intertemporal-substitution models of business cycles. *Quarterly Journal of Economics*, 99(4):817–839.

Basu, S., Fernald, J., Oulton, N., and Srinivasan, S. (2003). The case of the missing productivity growth: Or, does information technology explain why productivity accelarated in the United States but not the United Kingdom? *NBER Macroeconomics Annual*.

Beaudry, P. and Portier, F. (2004). An exploration into Pigou's theory of cycles. *Journal of Monetary Economics*, 51:1183–1216.

Beaudry, P. and Portier, F. (2007). When can changes in expectations cause business cycle fluctuations in neo-classical settings? *Journal of Economic Theory*, 135:458–477.

Beveridge, W. H. (1909). *Unemployment: A Problem of Industry*. Longmans Green, London.

Cagetti, M., Hansen, L., Sargent, T., and Williams, N. (2001). Robustness and pricing with uncertain growth. *Review of Financial Studies*, 15(2):363–404.

Christiano, L., Ilut, C., Motto, R., and Rostagno, M. (2008). Monetary policy and stock market boom-bust cycles. Working Paper Series 955, European Central Bank.

Clark, J. M. (1935). *Strategic Factors in Business Cycles*. National Bureau for Economic Research, New York.

Comin, D. (2009). On the integration of growth and business cycles. *Empirica*, 36(2):165–176.

Comin, D. and Mulani, S. (2009). A theory of growth and volatility at the aggregate and firm level. *Journal of Monetary Economics*, 56(8):1023 – 1042.

Den Haan, W. J. and Kaltenbrunner, G. (2009). Anticipated growth and business cycles in matching models. *Journal of Monetary Economics*, 56(3):309–327.

Dixit, A. and Pindyck, R. S. (1994). Investment under uncertainty. *Princeton University Press, Princeton, NY*.

Eusepi, S. and Preston, B. (2011). Expectations, learning and business cycle fluctuations. *American Economic Review*, forthcoming.

Evans, G., Honkapohja, S., and Romer, P. (1998). Growth cycles. *American Economic Review*, 88:495–515.

Fisher, J. D. M. (2006). The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy*, 114(3):413–451.

Flodén, M. (2007). Vintage capital and expectations driven business cycles. *Centre for Economic Policy Research discussion paper*, (6113).

Greenwood, J., Hercowitz, Z., and Huffman, G. (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 78:402–217.

Greenwood, J., Hercowitz, Z., and Krusell, P. (2000). The role of investment specific technological change in the business cycle. *European Economic Review*, 44:91–115.

Greenwood, J. and Yorukoglu, M. (1997). 1974. *Carnegie-Rochester Conference Series on Public Policy*, 46:49–95.

Gunn, C. and Johri, A. (2011). News and knowledge capital. *Review of Economic Dynamics*, 14(1):92–101.

Guo, S. (2008). News shocks, expectation driven business cycles and financial market frictions. Manuscript, Concordia University.

Hansen, G. D. (1985). Indivisible labor and the business cycle. *Journal of Monetary Economics*, 16:309–327.

Helpman, E. and Trajtenberg, M. (1994). A time to sow and a time to reap: Growth based on general purpose technologies. *NBER Working Paper*, (4854).

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

Jovanovic, B. and Lach, S. (1989). Entry, exit, and diffusion with learning by doing. *American Economic Review*, 79(4):690–99.

Justiniano, A. and Primiceri, G. E. (2008). The time-varying volatility of macroeconomic fluctuations. *The American Economic Review*, 98(3):604 – 641.

Karnizova, L. (2010). The spirit of capitalism and expectation-driven business cycles. *Journal of Monetary Economics*, 57(6):739 – 752.

Kasa, K. (2000). Forecasting the forecasts of others in the frequency domain. *Review of Economic Dynamics*, 3(4):726–756.

Keiichiro, K., Tomoyuki, N., and Masaru, I. (2007). Collateral constraint and news-driven cycles. Discussion papers 07013, Research Institute of Economy, Trade and Industry (RIETI).

King, R. G., Plosser, C., and Rebelo, S. T. (1988). Production, growth, and business cycles: I the basic neoclassical model. *Journal of Monetary Economics*, 21:195–232.

Kobayashi, K. and Nutahara, K. (2010). Nominal rigidities, news-driven business cycles, and monetary policy. *The B.E. Journal of Macroeconomics*, 10(1).

Lorenzoni, G. (2009). A theory of demand shocks. *The American Economic Review*, 99(5):2050–2084.

Pigou, A. C. (1926). *Industrial Fluctuations*. Macmillan. London.

Sakellaris, P. and Wilson, D. (2004). Quantifying embodied technological change. *Review of Economic Dynamics*, 7:1–26.

Van Nieuwerburgh, S. and Veldkamp, L. (2006). Learning asymmetries in real business cycles. *Journal of Monetary Economics*, 53:753–772.