

Corporate Investment Over Uncertain Business Cycles

Thomas Dangl
Vienna University of Technology

Youchang Wu
University of Wisconsin-Madison

June 27, 2011

Abstract

We present a dynamic model of irreversible corporate investment over the business cycle with uncertainty about the true state of the economy. We show that the firm's optimal investment threshold, defined by the ratio of the current demand factor to the capital stock, is a concave function of the firm's posterior probability of being in an expansion. Our model replicates some important features of the capital growth rate empirically observed. Despite the strong positive skewness of capital growth rate at the firm level, the average capital growth rate across firms, as well as its first-order difference, exhibits a negative skewness. The sharp-decline-slow-recovery feature of the average capital growth rate is due to both the concavity of the investment boundary and the distribution of firms relative to the boundary, which together cause firms in aggregate to react more strongly to bad signals in good times than to good signals in bad times.

Thomas Dangl can be reached at thomas.dangl@tuwien.ac.at. Youchang Wu can be reached at ywu@bus.wisc.edu. We acknowledge a grant from University of Wisconsin-Madison Graduate School. We also appreciate helpful discussions with David Brown, Bjorn Eraker, Christopher Hennessy, Richard Green, Antonio Mello, Erwan Quintin, Mark Ready, Suresh Sundaresan, Neng Wang, Toni Whited, Wei Zhang, and comments from seminar participants at University of Vienna, Vienna University of Economics and Business.

1 Introduction

Economists have long recognized that the economic downturn is generally more abrupt and violent than the upturn. For example, Keynes (1936) made the following famous observation in his *General Theory*, “There is, however, another characteristic of what we call the trade cycle which our explanation must cover; namely, the phenomenon of the crisis – the fact that the substitution of a downward for an upward tendency often takes place suddenly and violently, whereas there is, as a rule no such sharp turning point when an upward is substituted for a downward tendency.”

Using the quarterly Compustat-CRSP merged database from 1975 to 2008, we find striking differences between the patterns of the real capital growth rate at the individual firm level and those at the aggregate level. Consistent with the previous findings of business cycle asymmetry (for example, Neftci (1985), Sichel (1993)), we find that the average capital growth rate across firms declines fast during the recession and recovers slowly during the expansion. Both its level and its first-order difference are negatively skewed, indicating that the troughs is further below the trend than peaks are above, and that the downturn is steeper than the upturn. By contrast, the real capital growth rate of individual firms exhibits strong positive skewness in levels and virtually no skewness in its first-order differences, indicating the existence of positive spikes in corporate investment and the symmetry of the upward and downward slopes. How can we reconcile these contrasting features of asymmetry in a unified framework of optimal investment? This is the research question we try to answer in the theoretical part of the paper.

We propose a dynamic model of corporate investment over the business cycle with uncertainty about the true state of the economy. We consider a representative firm facing stochastic demand shocks. The expected growth rate of the demand factor depends on the state of the economy, which shifts between a high growth state (expansion) and a low growth state (recession) at random times. The firm is risk neutral, and has an infinite time horizon. Capital can be expanded incrementally and instantaneously at any point in time

at a constant marginal cost. However, capital investment is irreversible. Furthermore, the true state of the economy is not directly observable. The firm updates its belief about the state of the economy by observing the realizations of their own operating profit and a public signal. We model the updating of the firm’s belief about the true state of the economy in a continuous-time Bayesian learning framework, and derive the firm’s optimal investment policy. Under our assumptions, the optimal investment policy is characterized by a reflecting boundary. The firm takes no action until it hits the boundary, upon which it expands its capital stock instantaneously. We estimate the key model parameters using the Compustat-CRSP merged database, and solve for the optimal investment boundary numerically.

Because of the irreversibility, the firm’s investment decision depends not only on the current demand factor, but also on the belief about future demand growth.¹ The higher expected demand growth rate in an expansion leads to a lower investment threshold, meaning that the firm is more willing to invest during an expansion. However, due to the option value of waiting associated with the uncertainty about the current state of economy, the firm’s investment threshold is not a linear function of the posterior probability of being in an expansion. Instead, this relation is concave, and the concavity increases with the information quality of the signals. The intuition is as follows. Let π_t denote the firm’s posterior probability of being in an expansion at time t . The conditional variance of its belief is then $\pi_t(1 - \pi_t)$, which is highest when $\pi_t = 0.5$, i.e., when the firm assigns equal probability to both states, and lowest when π_t equals one or zero. Consider a situation in which π_t is close to 1, i.e., the firm is almost sure that it is in the high growth state. If the firm gets a bad signal, π_t will move away from 1 and move towards 0.5. This has two effects. First, it lowers the expected growth rate of the demand. Second, it increases the uncertainty about the current state, thus generating a higher option value of waiting.

¹If investment is perfectly reversible, then capital can be adjusted to any realized future demand shock without any cost. As a result, investment depends only on the current demand factor, and not on the belief about future demand growth. In this case, decline and recovery of the capital growth rate in the aggregate are symmetric due to the law of large numbers.

These two effects together lead to a sharp increase in the firm's investment threshold, thus inducing the firm to stop investment. By contrast, if a good signal arrives when the firms are almost sure that they are in an expansion period (π_t is close to 0), the expected demand growth rate increases, but the uncertainty also increases as π_t moves away from 0. The higher uncertainty increases the option value of waiting, thus partially offsetting the positive effect of the higher expected demand on the firm's incentive to invest. Therefore, the firm only add capacity slowly. When the information quality of the signals is higher, the incentive for waiting is stronger, thus leading to a more concave investment boundary.

After deriving the optimal investment boundary, we simulate a large number firms following this investment strategy, and examine the patterns of the capital growth rate, both at the firm level and in the aggregate. Our simulated data display patterns that match the empirical patterns remarkably well: the average capital growth rate declines fast and recovers slowly; it exhibits a negative skewness both in levels and in first-differences, while the capital growth rate of individual firms has a positive skewness in levels and virtually no skewness in first-differences.

The positive skewness of the capital growth rate at the individual firm level is a natural outcome of investment irreversibility, which implies that the expansion of capacity is unconstrained, while the decline of capital stock is limited by the depreciation rate. The sharp-decline-slow-recovery feature of the average capital growth rate is due to the two properties of the optimal investment policy mentioned above. First, since the investment threshold is lower during an expansion, more firms are close to the investment boundary during an expansion than during a recession. Therefore more firms are affected by the signal arriving during an expansion. By contrast, during a recession most firms are far away from the investment boundary. They do not react to the signal due to their over-capacity. Second, since the investment threshold is concave in the posterior belief, a bad signal arriving during an expansion leads to a sharp increase in the investment threshold, while a good signal arriving during a recession only leads to a modest decline in the threshold. As a result, the change of belief has an asymmetric effect on the firms that are close to the

boundary.

Our paper is an addition to a large body of literature on investment under uncertainty (see Dixit and Pindyck (1994) for an excellent survey of this strand of literature). Our model shares a common feature of the real options approach to investment decision, i.e., uncertainty tends to increase the option value of waiting, and therefore leads to delay in investment (see for example, Bernanke (1983), McDonald and Siegel (1986) and Pindyck (1988)). However, unlike most of the models in this literature, the degree of uncertainty in our model is not exogenously given. Instead we model the evolution of the uncertainty as an endogenous process resulting from the optimal updating of the belief by the agents in the model. Alti (2003) also investigates optimal investment policy in a Bayesian learning framework. However, his model does not allow for regime shifts therefore it is unable to explain the cyclical features of investment.

Our paper is part of a recent strand of literature involving learning about unobservable regime shifts (see for example, David (1997), Veronesi (1999), Veronesi (2000), Brown (2009)). David (1997) shows that learning about productivity switches can lead to negative skewness, excessive kurtosis and predictive asymmetry in aggregate stock market returns. Veronesi (1999) presents a model in which optimal learning induces investors to overreact to bad signals in good times and underreact to good signals in bad times. Both David (1997) and Veronesi (1999) are focused on the consequence of investor learning on the empirical features of stock market returns. The main reason for the asymmetric response to signals in these models is investors' risk aversion. By contrast, our paper is focused on the optimal irreversible investment decisions of risk neutral firms, and the asymmetry in the investment behavior is entirely due to the time-varying option value of waiting.

Our paper is also related to an extensive strand of literature on business cycle asymmetry. We contribute to this literature by presenting contrasting empirical patterns of capital growth rate at the firm level and the aggregate level, and providing a theoretical explanation for these patterns. Unlike the existing theories on business cycle asymmetry, which typically rely on some sort of exogenously assumed externality, the asymmetry of

investment behavior arises in our model from an intrinsic feature of Bayesian learning in an environment with unobservable regime shifts.²

The rest of our paper is organized as follows. Section 2 introduces the setup of our model, derives the optimal learning process and the firm value dynamics. Section 3 characterizes the optimal investment policy. Section 4 compares our simulated capital growth rates with the empirical data. Section 5 concludes the paper.

2 The Model

2.1 Setup

Our model builds on the model of Guo, Miao, and Morellec (2005). Consider a representative firm with an infinite time horizon. Time is continuous, the firm is risk neutral, and investment is irreversible. The firm faces a stochastic demand factor, which drives its cash flow. Depending on the state of the economy, the expected growth rate of the demand factor shifts between two different values at random times. A key difference between our model and that of Guo, Miao, and Morellec (2005) is that they assume the regime shift to be perfectly observable, while we assume that the firm can only infer the true regime from noisy signals.

Cash Flows: The operating income (before depreciation) of the firm is assumed to be

²Examples of theories based on externalities in the production process include Gale (1996) and Acemoglu and Scott (2003). Gale (1996) shows the dependence of a firm's profit on the general level of economic activity leads to strategic delays in investment during a recession, which prolongs the process of a recovery. Acemoglu and Scott (2003) show the dependence of a firm's investment costs on its past activity slows down the upturn and amplifies shocks in the trough of a business cycle. Explanations based on informational externalities include Chamley and Gale (1994), Chalkley and Lee (1998), Veldkamp (2005), and Nieuwerburgh and Veldkamp (2006). A general theme of this strand of literature is that a firm's investment generates information for other firms. In Chamley and Gale (1994), the recovery is delayed because firms wait in order to learn from the actions of the others. In the other three papers, the information value of public signals is higher in good times, either because lower fractions of projects are undertaken by noise traders, or because more production generates more information. As a result, firms react to public signals more quickly in good times.

given by a linearly homogeneous function $f : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfying:

$$f(x_t, k_t) = \frac{1}{1 - \alpha} x_t^\alpha k_t^{1 - \alpha}, \quad (1)$$

where $(k_t)_{t \geq 0}$ represents the process of the firm's capital stock, $(x_t)_{t \geq 0}$ represents the process of a demand factor.³ Assuming that the firm's output is nonstorable, equation (1) can be interpreted as the profit of either a price taking firm with a decreasing returns to scale technology, or a monopolist facing a constant returns to scale technology and a constant elasticity demand curve (see Abel and Eberly (1996) and Morellec (2001)).

Demand Shocks: Assume that the demand factor for the firm, x_t , evolves according to the following stochastic differential equation:

$$\frac{dx_t}{x_t} = \mu_t dt + \sigma_x dW_{x_t}, \quad x_0 > 0, \quad (2)$$

where W_{x_t} is a standard Wiener process. The volatility σ_x is a known constant. The expected growth rate of the demand factor, μ_t , is determined by the macroeconomic condition. It is low ($\mu_t = \mu_l$) in a recession period, and high ($\mu_t = \mu_h$) in an expansion period. Within a given state of the economy, the demand factor follows a standard geometric Brownian motion.

The macroeconomic condition switches between expansions and recessions at random times. Correspondingly, μ_t switches randomly between μ_h and μ_l . More specifically, we assume that μ_t is driven by a continuous time Markov jump process with the following transition probabilities:

$$P(\Delta t) = \mathbb{I} + \begin{pmatrix} -\lambda_{h,l} & +\lambda_{h,l} \\ +\lambda_{l,h} & -\lambda_{l,h} \end{pmatrix} (\Delta t + o(\Delta t)), \quad (3)$$

³More generally, x_t can be interpreted as an index of business conditions, which is positively related to the strength of the demand for the firm's product or the firm's productivity, and negatively related to the firm's cost of factors other than capital.

where \mathbb{I} is the identity matrix and $\lambda_{h,l}$ and $\lambda_{l,h}$ are the constant instantaneous intensities of transition from μ_h to μ_l and vice versa.⁴ The magnitudes of $\lambda_{h,l}$ and $\lambda_{l,h}$ determine the persistence of the expansion and the recession, respectively. The lower the transition intensity, the higher is the persistence.

Investment and Depreciation: The firm can add capital incrementally and instantaneously at any point in time at the constant marginal cost of one. Installed capital cannot be disinvested but depreciates at a constant rate of ξ . It is straightforward to extend the model to allow for costly reversibility, i.e., disinvestment at a certain cost. As long as there is some degree of irreversibility, the economic mechanisms we model, which lead to asymmetric investment behavior over the business cycle, will still stay.

Sudden Termination: Firms face the risk that their productive capital becomes obsolete. We model such an event as an exogenous random shock that leads to a sudden termination of the firm's operation, and assume that it occurs with a constant probability per unit of time. More specifically, we describe the dynamics of this destructive shock by a jump process q_t with the following jump probability:

$$P_q\{\Delta q(\Delta t) = 1\} = \lambda_q(\Delta t + o(\Delta t)), \quad q_0 = 0, \quad (4)$$

with λ_q as the constant jump intensity.

Information: It is assumed that the firm is able to observe its own realized demand shock x_t through the realized operating profit f_t , but not its expected growth rate μ_t . In other words, the true state of the economy that determines μ_t is a hidden process. This implies that W_{xt} is not observable as well. The firm's problem is that when observing a certain increase or decrease in the cash flow, it is not immediately clear which portion of this observed change comes from the currently prevailing growth rate μ_t and which portion comes from the noise terms W_{xt} . However, the distinction between these possible sources is of core relevance for the firm's investment decision. Since investment is irreversible, it

⁴The limitation to only two states is not essential from a technical point of view and can easily be relaxed to a finite number of states.

depends not only on the current demand factor, but also on expectations about its future growth rate.

In reality, the firm's own profit is clearly not the only source of information about the state of the economy. To summarize the other sources of information available to the firm, we assume that the firm can also observe a public signal s_t about the current state of the economy. The signal evolves according to the following stochastic differential equation:

$$\frac{ds_t}{s_t} = \mu_t dt + \sigma_s dW_{st}, \quad s_0 > 0. \quad (5)$$

where W_{st} is a standard Wiener process, σ_s is the constant volatility of the signal process, which is publicly known. We also assume the instantaneous correlation between the two Wiener processes, W_{st} and W_{xt} , to be a known parameter, $\rho \in [-1, 1]$. Note that the drift terms in equations (2) and (5) are identical. This assumption is made for simplicity. Our main results remain unchanged if we allow for imperfect correlation between the expected growth rates of x_t and s_t .⁵

Allowing the firm to observe an additional signal other than the firm's own profit not only makes our model more realistic, but also allows us to distinguish between the effect of the volatility of its own cash flow and the effect of the quality of information about the state of the economy. The instantaneous volatility of the signal, σ_s , measures the noisiness of the signal. It characterizes (inversely) the accuracy of the publicly available data about the economy. By varying σ_s without changing the volatility of the x_t , we can analyze the effect of information quality on the firm's investment. If observing the gradual realization of the cash flow process is the only source of information, then obviously the effects of information quality and cash flow volatility are coupled, because the information about the current growth rate is more accurate when the instantaneous volatility of the innovation term, σ_x , is lower.

To summarize the information structure of our model, let \mathcal{F}_t be the canonical non-

⁵See Brown (2009) for an example of learning when the drift terms of individual firms are imperfectly correlated.

decreasing filtration jointly created by the firm's own demand factor process x_t and the public signal s_t . Our assumptions about observability state that μ_t , dW_{xt} , and dW_{st} are not measurable with respect to the information set \mathcal{F}_t .

2.2 Bayesian Learning About the State of the Economy

Observing the firm's demand factor x_t and the signal s_t over the time interval $[0, t)$, an agent can continuously update her belief about the state of the economy in which the firm is currently operating. To formalize the rational learning rule, we denote by π_t the probability that $\mu_t = \mu_h$ conditional on \mathcal{F}_t and a prior π_0 . Consequently, the rational belief about demand growth rate is given by $\pi_t\mu_h + (1 - \pi_t)\mu_l$. Unexpected changes in x_t , which means deviations from the expected changes, give rise to an update of the belief. The same is true for unexpected changes in s_t .

Learning under this circumstance is a standard nonlinear filtering problem and can be characterized by the following proposition:

Proposition 1. *The optimal updating of the belief satisfies the following stochastic differential equation:*

$$d\pi_t = [-\pi_t\lambda_{h,l} + (1 - \pi_t)\lambda_{l,h}]dt + (\mu_h - \mu_l)\pi_t(1 - \pi_t)\mathbf{1}'(\Phi')^{-1}d\mathbf{W}_t^{\mathcal{F}}, \quad (6)$$

where $\mathbf{W}_t^{\mathcal{F}}$ is a two-dimensional independent Wiener process with respect to \mathcal{F}_t defined as

$$d\mathbf{W}_t^{\mathcal{F}} \equiv \begin{pmatrix} dW_{xt}^{\mathcal{F}} \\ dW_{st}^{\mathcal{F}} \end{pmatrix} \equiv \Phi^{-1} \begin{pmatrix} \frac{dx_t}{x_t} - E(\mu_t|\mathcal{F}_t)dt \\ \frac{ds_t}{s_t} - E(\mu_t|\mathcal{F}_t)dt \end{pmatrix}, \quad (7)$$

Φ is a 2×2 matrix satisfying

$$\Phi\Phi' = \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_s \\ \rho\sigma_x\sigma_s & \sigma_s^2 \end{pmatrix},$$

and $\mathbf{1}$ is a two-dimensional column vector with both elements equal to 1.

Proof. See Theorem 9.1 in Liptser and Shiryaev (2001) for the basic filtering equation, and equation (A1) in Veronesi (2000) for an extension to the vector case. Equation (6) is obtained by applying the vector version of the theorem. The independence between $dW_{xt}^{\mathcal{F}}$ and $dW_{st}^{\mathcal{F}}$ can be established by noting

$$dW_{xt}^{\mathcal{F}}dW_{st}^{\mathcal{F}} = (1, 0) \Phi^{-1}(\Phi\Phi'dt)(\Phi^{-1})'(0, 1)' = 0.$$

□

The diffusion process π_t is bounded between 0 ($\mu_t = \mu_l$, almost sure) and 1 ($\mu_t = \mu_h$, almost sure). The drift term in equation (6) indicates that in the absence of information shocks, there is a tendency for the belief to revert towards the unconditional mean,

$$\bar{\pi} = \frac{\lambda_{l,h}}{\lambda_{h,l} + \lambda_{l,h}}, \quad (8)$$

which satisfies $[-\pi_t\lambda_{h,l} + (1 - \pi_t)\lambda_{l,h}] = 0$. Therefore, the impact of any particular information shock decays gradually over time.

The diffusion term in equation (6) characterizes the response of the belief to unexpected changes in the realized demand factor x_t and the signal s_t . $E(\mu_t|\mathcal{F}_t)dt$ represents the best forecast of $\frac{dx_t}{x_t}$ and $\frac{ds_t}{s_t}$ conditional on the information set \mathcal{F}_t , $d\mathbf{W}_t^{\mathcal{F}}$ represents the standardized forecast errors. It is straightforward to see from the equation that the belief is more sensitive to the forecast errors (i) the greater the difference ($\mu_h - \mu_l$) between the two possible scenarios; (ii) the higher uncertainty about the state of the economy, i.e., the larger the conditional variance of the belief, $\pi_t(1 - \pi_t)$. These results are quite intuitive. When the growth rates in the two states do not differ much, or when the agents are very sure that they are in one of these two states (i.e., π_t is close to one or zero), the unexpected changes in the signals do not have much impact on the belief.

An alternative formulation of the optimal updating rule, equation (A.1) in Appendix

A.1, provides some further insights to the learning process. It suggests that the optimal learning from these two jointly normally distributed signals can proceed in two steps. One first forms a minimum variance “portfolio” of both signals, and then does the updating based on this compound signal. When the realized value of this compound signal is higher than expected, the posterior belief π_t is adjusted upward. Conversely, when the realized value is lower than expected, π_t is adjusted downward. The information quality of this compounded signal is measured by the inverse of its variance, $\frac{1}{\sigma^2} = \mathbf{1}'(\Phi\Phi')^{-1}\mathbf{1}$. When this quality is high, the response of the firm’s belief to the compound signal is more pronounced. Furthermore, by the nature of the minimum variance portfolio, the optimal weighting vector \mathbf{w} assigns more weight to the signal with a lower variance, indicating that agents pays more attention to the signal that has less noise.

The fact that $\mathbf{W}_t^{\mathcal{F}}$ is a Wiener process with respect to the filtration \mathcal{F}_t means that all information about the state of the economy available at time t is already contained in the current belief π_t , or in other words, if there were some information on the gradual unfolding of information in the future, this information could be used to immediately update the current belief. Therefore, the belief π follows a \mathcal{F}_t -Markov process. Alternatively, one can view $\mathbf{W}_t^{\mathcal{F}}$ as the *perceived innovation* of the informative signals, i.e., the disturbance that drives the deviation from the expected drift.

Dynamics with respect to \mathcal{F} : The formulation of the dynamics of x_t and s_t in equations (5) and (2) is not suitable for further analysis because the drift and innovation terms are not measurable with respect to \mathcal{F}_t . A reformulation of these dynamics in terms of unexpected changes with respect to \mathcal{F}_t solves this problem. Using $dW_{xt}^{\mathcal{F}}$ and $dW_{st}^{\mathcal{F}}$ defined by equation (7), we can rewrite the joint dynamics of the individual demand factor x_t and the public signal s_t as

$$\begin{pmatrix} dx_t/x_t \\ ds_t/s_t \end{pmatrix} = [\pi_t\mu_h + (1 - \pi_t)\mu_l] \begin{pmatrix} 1 \\ 1 \end{pmatrix} dt + \Phi \begin{pmatrix} dW_{xt}^{\mathcal{F}} \\ dW_{st}^{\mathcal{F}} \end{pmatrix}, \quad (9)$$

where π_t is updated as stated in equation (6).

2.3 Firm Value Dynamics

From the previous section we know that the Bayesian belief π_t follows a \mathcal{F} -Markov process. Thus, when starting from a certain prior, all information about demand growth that can be extracted from the entire observed paths of the demand factor x and the signal s is contained in the current belief π_t . Therefore the firm's value is fully determined by the current capital stock k_t , the current demand factor x_t , and the current belief about the state of the economy π_t . Thus, the value function V can be written as $V(k_t, x_t, \pi_t)$. This value function represents the present value of the operating profit flow under the optimal investment policy.

Since there is no fixed adjustment cost, the representative firm's optimal investment policy can be characterized by a reflecting boundary, which splits the state space into an investment region and no-investment region. The firm remains inactive in the interior area of the no-investment region, and increases its capital instantaneously by an infinitesimal amount dk whenever it hits the boundary.

Let r denote the instantaneous riskless rate of interest. We have the following proposition about the dynamics of the average value of capital, i.e., Tobin's average Q :

Proposition 2. *In the no-investment region, the value of the firm can be written as $V(k, x, \pi) = kv(h, \pi)$, where $h \equiv \frac{x}{k}$, and $v \equiv \frac{V}{k}$ represents Tobin's average Q . Furthermore, in this region, the average value of capital, v , has to satisfy the following partial differential equation (Hamilton-Jacoby-Bellman equation):*

$$\begin{aligned}
 (r + \xi + \lambda_q)v &= \frac{h^\alpha}{1 - \alpha} + h(\pi\mu_h + (1 - \pi)\mu_l + \xi) \frac{\partial v}{\partial h} + \frac{1}{2} h^2 \sigma_x^2 \frac{\partial^2 v}{\partial h^2} \\
 &+ (-\pi\lambda_{h,l} + (1 - \pi)\lambda_{l,h}) \frac{\partial v}{\partial \pi} \\
 &+ \frac{1}{2} \pi^2 (1 - \pi)^2 \frac{(\mu_h - \mu_l)^2}{\sigma^2} \frac{\partial^2 v}{\partial \pi^2} + h\pi(1 - \pi)(\mu_h - \mu_l) \frac{\partial^2 v}{\partial h \partial \pi},
 \end{aligned} \tag{10}$$

where σ^2 is defined in (A.2).

Proof. A proof is provided in the Appendix A.2. □

From Proposition 2 it follows that the average value of capital, v , depends only on the ratio of the demand factor to the amount of capital installed, and the belief about the current state of the economy.⁶

The partial differential equation (10) has to be solved under proper boundary conditions. We now specify these conditions.

2.4 Boundary Conditions

Since marginal productivity of capital is diminishing for low levels of demand, the no-investment region is associated with low values of $h = x/k$. When demand becomes high relative to the magnitude of invested capital, marginal productivity of capital increases and, thus, at a certain threshold of h the firm will have an incentive to invest additional capital. This critical threshold h^* is certainly a function of π , because the value of a newly invested unit of capital crucially depends on the growth prospects of the firm's operating profit. At this threshold, every positive shock on h is offset by an appropriate increase in capital. As a result, the firm will never enter the interior area of the investment region. That is why this threshold is called a reflecting boundary.⁷

Since the marginal cost of adding one unit of capital is one (i.e., the capital is measured in the same unit as the output), a rational valuation of the firm implies that at the investment boundary, the firm value has to satisfy the following value matching condition:

$$V(x, k, \pi) = V(x, k + dk, \pi) - dk,$$

which can be written in derivative form as

$$\lim_{k \rightarrow x/h^*} \frac{\partial V}{\partial k} = 1.$$

⁶This homogeneity property simplifies the solution of the valuation equation since it reduces the dimensionality of the problem. It follows from the fact that operating profit itself shows the homogeneity property $f(x_t, k_t) = k f(x_t/k_t, 1)$ as well as the fact that there are no fixed costs of investment.

⁷The model can be extended to the case with fixed adjustment costs, in which case the investment will be lumpy. However, the main message from our paper is unaffected.

This condition characterizes an important feature of the investment boundary: at the boundary, the marginal value of capital, i.e., Tobin's marginal Q , is always equal to the marginal cost of adding capital, which is constant in our model. Using the homogeneity feature of the value function, we can then rewrite this boundary condition as

$$\lim_{h \rightarrow h^*} \frac{\partial v(h, \pi)}{\partial h} = \frac{v(h^*, \pi) - 1}{h^*}. \quad (11)$$

Note that the value matching condition (11) is derived from a rational valuation argument, and is thus a necessary condition that holds even if the investment boundary $h^*(\pi)$ is not optimal. The optimality of the endogenously determined boundary requires the smoothness of the marginal value of capital at the boundary. This implies the following super contact (or smooth pasting) conditions (see Dumas (2001)): at the boundary,

$$\frac{\partial^2 V}{\partial k \partial x} = 0, \quad \frac{\partial^2 V}{\partial k^2} = 0, \quad \frac{\partial^2 V}{\partial k \partial \pi} = 0.$$

These translate into the following boundary conditions for $v(h, \pi)$:

$$\lim_{h \rightarrow h^*} \frac{\partial^2 v(h, \pi)}{\partial h^2} = 0, \quad \frac{\partial v(h, \pi)}{\partial \pi} + h \frac{\partial^2 v(h, \pi)}{\partial h \partial \pi} = 0. \quad (12)$$

3 Optimal Investment Policy

Our model does not have an analytical solution, so we solve it numerically. In this section, we first describe our procedure to calibrate the model parameters, and then discuss the main properties of the numerically derived optimal investment policy.

3.1 Calibration

Our model consists of eleven parameters: α in the operating profit function; regime switching intensities $\lambda_{h,l}$ and $\lambda_{l,h}$; expected demand growth rates in the good and bad states μ_h

and μ_t ; sudden termination rate λ_q ; instantaneous volatilities of the demand and of the public signal, σ_x and σ_s ; depreciation rate ξ ; risk free rate r ; and the correlation between the Wiener Processes driving the demand factor x_t and the signal s_t , ρ .

We first estimate the parameters of the Markov switching process that characterizes the evolution of the expected growth rate μ_t of the demand factor x_t . Note that from equation (1), we can back out the demand factor x_t using the firm's operating profit $f(x_t, k_t)$ and capital stock k_t :

$$x_t = \left[\frac{(1 - \alpha)f(x_t, k_t)}{k_t^{1-\alpha}} \right]^{1/\alpha}. \quad (13)$$

Using the annual Compustat-CRSP merged database from 1950 to 2008, we compute the x_t for each individual firm in each year. We measure $f(x_t, k_t)$ by operating income before depreciation (OIBDP), and measure k_t by the average net property, plant and equipment (PPENT) at the beginning and end of the year. We set the operating profit parameter α equal to 0.53, following Guo, Miao, and Morellec (2005). To adjust for inflation, the current-dollar operating income is converted into the year 2005 dollar term using the GDP deflator, while the current-dollar PPENT is converted into the year 2005 dollar term using a fixed asset deflator constructed based on the U.S. Bureau of Economic Analysis (BEA) data.⁸

We calculate the continuously-compounded annual growth rate of the demand factor for each firm by taking the first-order difference of the natural logarithm of x_t estimated using equation (13). We exclude the growth rates that are below the 1% or above the 99% quantile of the whole sample. We then calculate the cross-sectional average of the annual growth rates in each year. This gives us an a time series of 57 observations. We apply the EM algorithm described in Chapter 22 of Hamilton (1994) to estimate the parameters

⁸BEA publishes both the nominal value (at historical costs) and a real quantity index (relative to the year 2005) of the net stock of the private nonresidential fixed assets in the U.S.. We construct the fixed asset deflator series based on the following relation:

$$\text{Quantity index for year } i = \frac{\frac{\text{Nominal fixed asset in year } i}{\text{Fix asset deflator for year } i} \times 100}{\text{Nominal fixed asset in 2005}} \times 100.$$

of the hidden Markov chain: $\lambda_{h,l}$, $\lambda_{l,h}$, μ_h , μ_l . The results are summarized in Table 1, together with the values of other parameters.⁹

[Table 1 about here.]

Our estimates of the transition intensities imply that the expected durations of the good and bad states are 3.89 years ($=1/0.2573$) and 1.58 years ($=1/0.6326$), respectively. This is well in line with the earlier estimates of Hamilton (1989) using the quarterly GDP data. According to his estimation, the expected durations of expansions and recessions are 10.5 quarters and 4.1 quarters, respectively. We plot the time series of the estimated posterior belief π_t in Figure 1, along with the historical expansion and recession periods identified by National Bureau of Economic Research (NBER). One can see that the π_t series matches the NBER-designated business cycles very well. The recession periods designated by NBER are always associated with π_t less than 0.5, while the expansion periods are associated with π_t higher than 0.5 except for the years 1952 and 1967. Our results also indicate that expected growth rates in the good and bad states are quite different: 0.1060 vs. -0.1085. This reflects the volatile nature of the demand factor, which highlights the importance of the business cycle on firms' profitability and investment decisions.

[Figure 1 about here.]

The other parameter values are chosen as follows. We set the sudden termination rate, λ_q , equal to 0.0604, which is the average annual rate at which firms drop out of the CRSP-Compustat database from 1966 to 2008. The depreciation rate ξ and risk free rate r are set equal to 0.1 and 0.05, respectively. The instantaneous volatilities of individual demand factor and of the signal, σ_x and σ_s , are set equal to 0.20, while the instantaneous correlation of these two processes is set equal to 0.05.

⁹Our procedure also estimates the instantaneous volatility of the average demand growth rate to be 0.0917. Since the average growth rate has lower volatility than individual growth rates do, this can be viewed as a lower bound for σ_x .

3.2 Optimal Investment Boundary

In this subsection we determine the investment boundary $h^*(\pi)$ at which the firm will find it optimal to invest additional capital. As long as x/k is below this threshold (i.e., invested capital is high relative to the observed demand), marginal value of capital, i.e., marginal Q , is below 1. Thus the firm does not invest, and the firm value is characterized by the solution of equation (10). At the investment boundary, marginal Q equals one and, thus, the firm instantaneously invests additional capital. Since the Hamilton-Jacoby-Bellman equation (10) is a partial differential equation, it is solved numerically on the estimated parameter set listed in Table 1. We apply the approach derived in Nelson and Ramaswamy (1990) to map the dynamics of the belief π onto a recombining tree, and then use a two-dimensional tree (as outlined in Boyle, Evnine, and Gibbs (1989)) to jointly determine the firm value and the optimal investment boundary.

[Figure 2 about here.]

Panel (a) of Figure 2 illustrates the shape of the optimal investment boundary $h^*(\pi)$. The solid, concave curve constitutes the optimal investment threshold for the baseline parameterization given in Table 1. When the firm believes that the economy is in an expansion ($\pi \approx 1$), it invests earlier, i.e., at lower values of $h = x/k$, than when it assigns a high probability of a recession ($\pi \approx 0$). The difference between the investment thresholds in these two cases is due to the irreversibility of investment. If investment is perfectly reversible, i.e., if the firm can adjust its capital stock to any level that is appropriate for the realized demand factor without any cost and delay, then the firm's optimal capacity choice depends only on the current value of the demand factor, and not on the belief about its future growth rate. In this case, the optimal investment boundary will be a vertical line, since it is identical for all values of π .

A notable feature of the optimal investment boundary is that it is concave in the belief π , indicating that the firm's investment decision is relatively insensitive to the change in the belief π when π is low. This results from the interaction of the expected growth rate

and the uncertainty about the growth rate. Starting from π close to zero, i.e., a situation in which the firm is almost sure of being in the low growth rate state, a positive signal increases the expected growth rate of future demand. At the same time, it also increases the uncertainty about the current state (captured by the conditional variance $\pi(1 - \pi)$), and therefore increases the option value of waiting. As a result, the firm is reluctant to increase its capital stock immediately. On the contrary, if a bad signal is received in a good state (i.e., when π is high), both the resulting lower expected growth rate and higher uncertainty diminish the firm's incentive to invest. Therefore the investment threshold increases sharply. This asymmetric response of the firm's investment decision to the signal is an important reason for the asymmetry of capital growth rate over business cycles.

The dashed curves illustrate investment boundaries for higher information quality, i.e., low σ_s of the common signal s_t . They indicate that better information quality increases the concavity of the boundary. Other things equal, the higher the precision of the information, the stronger is the response of the belief to the signal, as we discuss in Subsection 2.2. This leads to a higher volatility of the belief, which in turn generates a higher option value of waiting. Higher information quality thus exacerbates the concavity of the boundary.

As the noisiness of the signal, σ_s , goes to zero, the optimal investment policy in our model converges to the policy under perfect information derived analytically by Guo, Miao, and Morellec (2005). It is characterized by two distinct threshold values, a lower one for the good state and a higher one for the bad state of the economy. The two dashed horizontal lines depict these two distinct thresholds using the results of Guo, Miao, and Morellec (2005).¹⁰

Comparing the investment boundaries under different information qualities, we find an interesting relation between information quality and investment policy. Higher information quality leads to a lower investment boundary as long as the uncertainty about the true state is low, but it leads to a higher investment boundary when π is in the intermediate range, i.e., when the uncertainty is high. Since investment is irreversible, the marginal

¹⁰We adapt their solutions to account for the sudden termination of the firm introduced in our model.

value of installed capital depends on future investment policy. If the information quality is high, future investment decisions are more informed. This increases the marginal value of capital today, thus lowering the investment threshold when the uncertainty about the current state is low. However, if the uncertainty about the current state is high (π close to 0.5), then the higher option value of waiting associated with the high quality signal dominates this effect, thus leading to a higher investment threshold. In the extreme case of σ_s approaching zero, the option value of waiting is so high that the firm behaves very conservatively, i.e., it always invests as if it were in a recession unless it is almost sure of being in an expansion. As a result, the investment boundary becomes a vertical line for π below one and a horizontal line as π approaches one.

To better illustrate the intuition of the optimal investment policy, we present an alternative plot of the investment boundary in Panel (b) of Figure 2. Instead of plotting $h = x/k$ as a function of π , we plot the inverse, $1/h = k/x$, as a function of π . This boundary can be interpreted as the firm's optimal normalized (by demand factor x) capacity under different beliefs about the current state of the economy. It is a lower bound for the firm's normalized capital stock. A firm increases its capital whenever it hits this boundary from above (as k/x declines) or from left (as π increases). As we can see, the investment boundaries defined this way are convex in π , indicating that optimal capacity increases very slowly as π moves from zero toward the middle range, but declines very fast as π decreases from the peak value of one.

Note that while the shape of the investment boundary determines how sensitive the firm's optimal capacity is to the change in the belief, it alone does not determine the speed of capital adjustment after a regime shift, because the latter also depends on how quickly the firm learns about the true state of the economy, i.e., how quickly π converges to zero or one. When the information is very precise, the investment boundary $h^*(\pi)$ is very concave, but since firms can form a precise belief about the true state of the economy very fast, π will only be in the intermediate range around $\pi = 0.5$ for a relatively short period of time. In this case, firms can react to a good signal in bad times quickly, even though the

investment boundary is highly concave.

4 Capital Growth Rate

We now examine the dynamics of capital growth rate under the optimal investment policy derived in Section 3. We simulate a sample of 1000 firms that follow their optimal investment policy over the time horizon of 100 years during which the economic conditions follow the Markov process (3).¹¹ Each firm observes its own demand factor x_t following (2), and the common signal process s_t . When a firm is hit by a sudden termination, its path is terminated. At the same moment, a new and optimally capitalized firm is introduced which has the same belief as the terminated firm. On the set of simulated data we analyze the time series properties of the capital growth rate of individual firms as well as the average capital growth rate across firms.

We compare our simulated data with the empirical data of the capital growth rate. For the empirical analysis, we use the quarterly Compustat-CRSP merged database from 1975 to 2008.¹² We exclude all financial firms (SIC codes between 6000 and 6999), utilities (SIC codes between 4900-4999), government entities (SIC codes greater or equal to 9000) and firms/quarters with missing net capital stock (PPE) data. We end up with a sample with on average 3986 firms in each quarter. We use the book value of net property, plant and equipment (PPENT) as the measure of a firm's capital stock. The nominal values of net PPE are converted into the year 2005 dollar using the fixed asset deflator, as explained in footnote 8. Since this series is only available at the annual frequency, we construct the quarterly deflator by linear interpolation.

We measure a firm's capital growth by the continuously-compounded growth rate of its net PPE. Unlike the simple growth rate, which is inherently asymmetric since it cannot go

¹¹We drop the first 20 years of this generated data set such that the investment patterns are independent of initial conditions.

¹²Before 1975, the quarterly data on firms' capital stock is only sparsely available. We have also analyzed the empirical pattern of capital growth rate using the annual Compustat-CRSP merged database from 1950 to 2008, the results are similar.

below -100% due to the non-negativeness of capital stock, the continuously compounded capital growth rate can, in principle, be symmetric. The growth rate of net capital measures net investment, i.e., gross investment minus depreciation, as a percentage of the beginning-of-the-period net capital. This measure of corporate investment activity is broader than the capital expenditure item in Compustat. The latter does not reflect a firm’s capital accumulation through merger or acquisition, or capital reduction due to asset sales, and is only widely available at quarterly frequency since 1985.

4.1 Average Capital Growth Rate over Business Cycles

[Figure 3 about here.]

Figure 3 the simulated, as well as the empirically estimated, equal weighted average capital growth rate over a full business cycle (starting with the recession). Time zero represents the beginning of the expansion period. The black dashed line represents the average capital growth rate under the base case parameterization specified in Table 1 with $\sigma_s = 20\%$. The red dashed line shows the average capital growth rate under high information quality, $\sigma_s = 0.5\%$, with other parameter values unchanged. The solid line represents the equal weighted quarterly capital growth rate calculated using the Compustat-CRSP merged database from 1975 to 2008. For the empirical curve, the time zero is the troughs of various business cycles dated by the National Bureau of Economic Research.

Consistent with the Keynes’s observation of business cycle asymmetry, the decline of the empirical capital growth rate during the recession is steeper than its recovery in the expansion. Our simulation under the baseline parameterization ($\sigma_s = 20\%$) replicates this sharpe-decline-slow-recovery pattern, although the trough of the simulated data appears to be significantly deeper. The simulation under high information quality ($\sigma_s = 0.5\%$) features an even more dramatic decline at the beginning of the recession, as well as an almost equally sharp and immediate rebound once the recession is over.

There are two reasons for the asymmetry of decline and recovery in our model. First,

since the investment threshold ($h^*(\pi)$) is lower while the demand growth rate is higher in the expansion, there are generally a larger number of firms close to the investment boundary at the end of an expansion than at the end of a recession. These are the marginal firms that react to the changes in beliefs. Firms that are far away from the investment boundary may not invest even if π goes from zero to one, since their capital stock is too high even relative to the low threshold of an expansion period. As a result, when a bad signal comes during an expansion, the number of reacting firms is higher than when a good signal comes during a recession. This explains why the decline appears to be slightly more dramatic than the recovery even when information is almost perfect.

Second, as we explain in Section 3.2, the effect of uncertainty about the true state of economy is asymmetric over business cycles. During an expansion, in which π is close to one, a decrease in π triggered by a bad signal leads to a sharp decline in a firm's optimal capacity, as it not only lowers the expected demand growth rate, but also increases the option value of waiting. By contrast, when a good signal arrives during a recession, its positive impact on expected growth rate is partially offset by the higher option value of waiting, leaving the optimal capacity largely unaffected. Therefore, even for the firms that are close the investment boundary, the capital growth rate increases only slowly.

The simulation results presented in Figure 3 suggest that these two effects together can generate to a large extent the asymmetry of capital growth rate observed in the data. Nevertheless, the recession in our simulated data is much deeper than the empirical counterpart and the recovery is faster. This has to do with the fact that our model is abstract from any frictions in the financial market. In the real world, when the economy just comes out of a recession, it is usually difficult for many firms to raise capital. Such financial constraints further reduce firms' capital growth rate at the initial stage of expansion.

4.2 Skewness of the Capital Growth Rate

To further gauge the empirical plausibility of our model, we examine the asymmetry of the capital growth rate, both at the individual firm level and in the aggregate. Following Sichel (1993), we distinguish between two types of asymmetry: level asymmetry (depth) and slope asymmetry (steepness). Level asymmetry refers to the characteristic that troughs are further below trend than peaks are above, while slope asymmetry refers to the characteristic that contractions are steeper than expansions. We use the skewness of the capital growth rate time series to measure level asymmetry; and use the skewness of the first-order difference of the capital growth rate to measure the slope asymmetry.

To capture the asymmetries of the capital growth rate over different time horizons, we examine the n -quarter capital growth rate for n varying from 1 and 16, following Nieuwerburgh and Veldkamp (2006).¹³ To measure the asymmetries at the firm level, we first calculate the skewness of the n -quarter capital growth rate and its first-order difference for each individual firm, and then average them across all the firms that have at least 20 observations. To account for potential data errors, and to limit the impact of some extreme outliers, we exclude the capital growth rates that are below the 1% or above the 99% quantile of the whole sample.¹⁴ For the asymmetries at the aggregate level, we first calculate the average of the n -quarter capital growth rates across firms. This gives us a total of $n=16$ time series of average capital growth rate. The skewnesses of these time series are then used to measure the level asymmetry in the aggregate. The slope asymmetry in the aggregate is measured by the skewness of the first-order difference of the average capital growth rate.

[Figure 4 about here.]

¹³Specifically, the n -quarter capital growth rate is defined as

$$g_{t,t-n} \equiv \ln(PPE_t) - \ln(PPE_{t-n}),$$

where n varies from 1 to 16.

¹⁴Our results remain largely unchanged if we do not exclude any outliers.

Panel (a) of Figure 4 shows the skewness of the capital growth rate and its first-order difference over various time horizons (between 1 to 16 quarters) at the individual firm level. The solid curves are the empirical estimates based on the capital growth rates of the firms in the quarterly Compustat-CRSP merged database from 1975 to 2008. The dashed curves show the skewness calculated for our simulated data. As we can see from the graph, our model replicates an important feature of the capital growth rate at the individual firm level: it exhibits strong positive skewness in levels and almost zero skewness in first differences. For both the empirical and simulated data, the positive skewness in levels is highest at the one quarter horizon, and decrease steadily as the time horizon lengthens, but remains positive even at the 16-quarter interval. This is consistent with the fact that disinvestment is more costly than investment. Capital is built up for specific uses, it may be very difficult to use it for other purposes. Therefore firms may be able to expand their capital stock in a large quantity, but unable to reduce it dramatically. A noticeable difference between the empirical and the simulated results is that the positive skewness of the simulated data is higher in magnitude. This is because we assume complete irreversibility of investment in our model, while investment in the real world is likely to be partially reversible. The skewness of the first difference of is almost negligible in both the empirical and the simulated data, indicating that at the firm level, increases and decreases in the capital growth rate are symmetric in magnitude.

Panel (b) presents the skewness calculated for the average capital growth rate across firms and its first-order difference. Here again the solid lines represents the empirical estimates and dashed lines are for the simulated data. The patterns are dramatically different from those displayed in Panel (a). Consistent with the macroeconomic evidence for business cycle asymmetry in the previous literature, empirical capital growth rate exhibits both level asymmetry and slope asymmetry. The average capital growth rate over time horizons up to 3 years is negatively skewed. At the quarterly interval, the skewness is -0.70. It then increases slowly as the time horizon gets longer and approaches zero at the 12-quarter interval, after which it becomes positive. This indicates that troughs are further

below trend then peaks are above for time horizons up to 3 years. The skewness of the first-order difference of the average capital growth rate is non-monotonic in the time horizon. It starts at -0.29 at the quarterly interval, decreases to -0.83 at the three-quarter interval, then slowly increases afterwards, and decreases again at around the 11-quarter interval. The negative skewness of the first-differenced average capital growth rate indicates that the aggregate investment tends to increase slowly and decreases sharply, lending support to the observation of Keynes (1936) that recessions arrive violently while recoveries only occur slowly.

Our simulated average capital growth rate exhibits patterns that accords remarkably well with those observed in the empirical data. Its level is negatively skewed over short horizon, and gradually becomes positively skewed at longer horizon. Furthermore, its first-order difference is also negatively skewed, and the skewness is “tilde”-shaped over investment horizons from 1 to 16 quarters.

Figure 4 suggests that our model is able to reconcile with two seemingly conflicting patterns in the capital growth rate: the negative skewness of the average capital growth rate, both in levels and in first-differences, and the positive skewness at the firm level. The positive skewness of the capital growth rate at the firm level is a natural outcome of investment irreversibility, which prevents the capital stock from shrinking faster than depreciation. The level asymmetry of the average capital growth rate has to do with the fact that expansions usually last longer than recessions do, which implies that the distribution of short-run growth rates is mainly populated by observations from the expansion period. The relatively small number of observations from the recession period then form a long tail in the left side of the distribution, leading to a negative skewness. When the capital growth rates are summed up for multiple quarters, the successive investments over a long expansion period generate outliers in right tail of the distribution, thus resulting in a positive skewness for the average capital growth rate over a longer horizon. The slope asymmetry of the average capital growth rate reflects the sharpe-decline-slow-recovery feature of the business cycle. As we explain in the previous subsection, this occurs for two

reasons. First, there are more firms close to investment boundary during an expansion than during a recession, therefore more firms are affected when a bad signal arrives during an expansion. Second, for the firms that are close to the investment boundary, a bad signal arriving during an expansion leads to a stronger reaction because it not only reduces the expected demand growth rate, but also increases the option value of waiting, while a good signal arriving during a bad time has a lower impact because its positive effect on expected demand growth is partially offset by the increased option value of waiting.

5 Conclusion

In this paper we present a dynamic model of irreversible corporate investment over the business cycle with uncertainty about the true state of the economy. We consider a representative firm facing stochastic demand shocks. The expected growth rate of the demand factor depends on the state of the economy, which shifts between an expansion and a recession at random times. Because of the irreversibility, the firm's investment decision depends not only on the current demand factor, but also on the belief about the current state of economy. We show that the firm's optimal investment threshold, defined by the ratio of the current demand factor to the capital stock, is a concave function of the firm's posterior probability of being in the expansion period of the business cycle.

Our model replicates some important features of the capital growth rate empirically observed. Despite the strong positive skewness of capital growth rate at the firm level, the average capital growth rate across firms, as well as its first-order difference, exhibits a negative skewness. The positive skewness of the individual firm's capital growth rate is a natural outcome of investment irreversibility, which limits the speed of capital shrinkage of an individual firm. The sharp-decline-slow-recovery feature of the average capital growth rate is due to both the concavity of the investment boundary and the distribution of firms relative to the boundary, which together cause firms in aggregate to react more strongly to bad signals in good times than to good signals in bad times. As a result, the average

capital growth rate across firms increases gradually during an expansion but drops sharply at the beginning of a recession.

An important limitation of our model is that it abstracts from any frictions in the financial markets. In the real world, when the economy just comes out of a recession, it is usually difficult for firms to raise capital. Such financial constraints tend to further delay firms' investment at the initial stage of expansion. Incorporating financial frictions over business cycles into the firm's investment decision is a fruitful venue for future research.

Appendix

A.1 An Alternative Formulation of the Optimal Updating Rule

Equation (6) can be rewritten as:

$$d\pi_t = [-\pi_t\lambda_{h,l} + (1 - \pi_t)\lambda_{l,h}]dt + \frac{(\mu_h - \mu_l)\pi_t(1 - \pi_t)}{\sigma^2} \mathbf{w} \begin{pmatrix} \frac{dx_t}{x_t} - E_t(\mu_t|\mathcal{F}_t)dt \\ \frac{ds_t}{s_t} - E_t(\mu_t|\mathcal{F}_t)dt \end{pmatrix} \quad (\text{A.1})$$

where

$$\sigma^2 \equiv \frac{1}{\mathbf{1}'(\Phi\Phi')^{-1}\mathbf{1}}, \quad (\text{A.2})$$

$$\mathbf{w} \equiv \frac{\mathbf{1}'(\Phi\Phi')^{-1}}{\mathbf{1}'(\Phi\Phi')^{-1}\mathbf{1}} = \begin{pmatrix} \frac{\sigma_s^2 - \rho\sigma_x\sigma_s}{\sigma_x^2 + \sigma_s^2 - 2\rho\sigma_x\sigma_s} & \frac{\sigma_x^2 - \rho\sigma_x\sigma_s}{\sigma_x^2 + \sigma_s^2 - 2\rho\sigma_x\sigma_s} \end{pmatrix}. \quad (\text{A.3})$$

Note that $\Phi\Phi'$ is simply the instantaneous variance-covariance matrix of $\frac{dx_t}{x_t}$ and $\frac{ds_t}{s_t}$. Readers familiar with the classic mean-variance portfolio analysis will immediately recognize that σ^2 is the minimum instantaneous variance that can be obtained using all possible linear combinations of $\frac{dx_t}{x_t}$ and $\frac{ds_t}{s_t}$, while \mathbf{w} is a vector that specifies the weights of each individual signal in the minimum variance combination. This formulation thus reveals an important feature of the Bayesian learning process. When there are multiple jointly normally distributed signals, the agent can form a minimum variance “portfolio” of all the available signals, and base the learning on this compound signal. When the realized value

of this compound signal is higher than expected, the posterior belief π_t is adjusted upward. Conversely, when the realized value is lower than expected, π_t is adjusted downward.

The standard deviation of this optimally constructed compound signal, σ , measures the noisiness of the overall information content of all the signals. The learning equation (6) thus implies that the response to the forecasting is more pronounced when the precision, $\frac{1}{\sigma}$, of the signals is high. Furthermore, by the nature of the minimum variance portfolio, the optimal weighting vector \mathbf{w} assigns more weight to the signal with a lower variance, indicating that agents pay more attention to the signal that has less noise. In particular, when $\sigma_s = \rho\sigma_x < \sigma_x$, the optimal weight of $\frac{dx_t}{x_t}$ in the compound signal is zero. Learning is entirely based on the signal $\frac{ds_t}{s_t}$. More surprisingly, when $\sigma_s < \rho\sigma_x$, the optimal weight of $\frac{dx_t}{x_t}$ is even negative. This implies that when the realized value of the more precise external signal s_t is just as expected, while the firm's own demand factor x_t is higher than expected, agents will adjust their belief π_t downward. The intuition is as follows. Since the shocks to the two signals are highly correlated, a higher-than-expected realized value of $\frac{dx_t}{x_t}$ suggests that it is very likely that $\frac{ds_t}{s_t}$ has also received a positive shock, i.e., it is above its true mean. This therefore suggests that the current belief of the mean, which is right at the realized value of $\frac{ds_t}{s_t}$, is too high and should be downward updated.

A.2 Proof of Proposition 2

Consider the firm value V as a claim on the firm's operating profit as a function of its current demand factor x , the installed capital k and the belief about the current state of the economy π , i.e., $V = V(x, k, \pi)$. For the risk neutral decision maker the value function must satisfy the following Hamilton-Jacobi-Bellman equation:

$$\begin{aligned} rV(x, k, \pi)dt &= f(x, k) dt + E(dV(x, k, \pi)|\mathcal{F}_t) \\ &= f(x, k) dt + (1 - \lambda_q dt)E[dV(x, k, \pi)|\mathcal{F}_t, dq_t = 0] \\ &\quad + \lambda_q dt E[dV(x, k, \pi)|\mathcal{F}_t, dq_t = 1]. \end{aligned}$$

The expected change in the value function $E(dV)$ is split into two mutually exclusive branches, one conditional on the fact that sudden termination does not occur during dt and the second conditional on the fact that sudden termination occurs during dt . Sudden termination implies that $E[dV(x, k, \pi)|\mathcal{F}_t, dq_t = 1] = -V$. Further, it will be shown below that $E[dV(x, k, \pi)|\mathcal{F}_t, dq_t = 0]$ is of order dt , thus, ignoring terms of $o(dt)$ leads to the Hamilton-Jacobi-Belman Equation

$$(r + \lambda_q)V(x, k, \pi)dt = f(x, k) dt + E[dV(x, k, \pi)|\mathcal{F}_t, dq_t = 0].$$

The expectation of dV is to be determined by Itô's Lemma using the \mathcal{F} -dynamics of the state variables x , k and π . When applying Itô's Lemma to the value function, note that since $dW_{xt}^{\mathcal{F}}$ and $dW_{st}^{\mathcal{F}}$ are uncorrelated, we have

$$\begin{aligned} (d\pi)^2 &= d\pi(d\pi)' \\ &= [\pi(1 - \pi)(\mu_h - \mu_l)]^2 [\mathbf{1}'(\Phi')^{-1}d\mathbf{W}_t^{\mathcal{F}}] [\mathbf{1}'(\Phi')^{-1}d\mathbf{W}_t^{\mathcal{F}}]' \\ &= [\pi(1 - \pi)(\mu_h - \mu_l)]^2 \frac{1}{\sigma^2} dt, \\ dx d\pi &= dx(d\pi)' \\ &= x\pi(1 - \pi)(\mu_h - \mu_l) [(1, 0) \Phi d\mathbf{W}_t^{\mathcal{F}}] [\mathbf{1}'(\Phi')^{-1}d\mathbf{W}_t^{\mathcal{F}}]' \\ &= x\pi(1 - \pi)(\mu_h - \mu_l) dt. \end{aligned}$$

Therefore, in the no-investment region, in which $dk = -\xi k dt$, we have

$$\begin{aligned} E[dV(x, k, \pi)|\mathcal{F}_t, dq_t = 0] &= \left[-\frac{\partial V}{\partial k} \xi k + \frac{\partial V}{\partial x} x(\pi\mu_h + (1 - \pi)\mu_l) \right. \\ &\quad + \frac{1}{2} \frac{\partial^2 V}{\partial x^2} x^2 \sigma_x^2 + \frac{\partial V}{\partial \pi} (-\pi\lambda_{h,l} + (1 - \pi)\lambda_{l,h}) \\ &\quad \left. + \frac{1}{2} \frac{\partial^2 V}{\partial \pi^2} \pi^2 (1 - \pi)^2 (\mu_h - \mu_l)^2 \frac{1}{\sigma^2} + \frac{\partial^2 V}{\partial x \partial \pi} x\pi(1 - \pi)(\mu_h - \mu_l) \right] dt. \end{aligned}$$

Substituting the above expression into the Hamilton-Jacobi-Belman equation and drop-

ping dt from both sides of the equation yields

$$\begin{aligned}
(r + \lambda_q) V &= \frac{1}{1 - \alpha} x^\alpha k^{(1-\alpha)} - \frac{\partial V}{\partial k} \xi k + x(\pi \mu_h + (1 - \pi) \mu_l) \frac{\partial V}{\partial x} + \frac{1}{2} x^2 \sigma_x^2 \frac{\partial^2 V}{\partial x^2} \\
&+ (-\pi \lambda_{h,l} + (1 - \pi) \lambda_{l,h}) \frac{\partial V}{\partial \pi} \\
&+ \frac{1}{2} \pi^2 (1 - \pi)^2 \frac{(\mu_h - \mu_l)^2}{\sigma^2} \frac{\partial^2 V}{\partial \pi^2} + x \pi (1 - \pi) (\mu_h - \mu_l) \frac{\partial^2 V}{\partial x \partial \pi}
\end{aligned} \tag{A.4}$$

The last part of the proof shows that writing $V(x, k, \pi)$ as $V = kv(h, \pi)$ with $h = \frac{x}{k}$ gives equation (10). This is shown by substituting the following partial derivatives into equation (A.4):

$$\begin{aligned}
\frac{\partial V}{\partial k} &= v(h, \pi) - h \frac{\partial v(h, \pi)}{\partial h}, \\
\frac{\partial V}{\partial x} &= \frac{\partial v(h, \pi)}{\partial h}, \\
\frac{\partial^2 V}{\partial x^2} &= \frac{1}{k} \frac{\partial^2 v(h, \pi)}{\partial h^2}, \\
\frac{\partial V}{\partial \pi} &= k \frac{\partial v(h, \pi)}{\partial \pi}, \\
\frac{\partial^2 V}{\partial \pi^2} &= k \frac{\partial^2 v(h, \pi)}{\partial \pi^2}, \\
\frac{\partial^2 V}{\partial x \partial \pi} &= \frac{\partial^2 v(h, \pi)}{\partial x \partial \pi}.
\end{aligned}$$

We have therefore proved that $V(x, k, \pi) = kv(h, \pi)$ satisfies equation (A.4) as long as $v(h, \pi)$ satisfies equation (10).

References

- Abel, Andrew B., and Janice C. Eberly, 1996, Optimal Investment with Costly Reversibility, *Review of Economic Studies* 63, 581–593.
- Acemoglu, Daron, and Andrew Scott, 2003, Asymmetric Business Cycles: Theory and Time-Series Evidence, *Journal of Monetary Economics* 40, 501–533.
- Alti, Aydogan, 2003, How Sensitive Is Investment to Cash Flow When Financing Is Frictionless?, *Journal of Finance* 58, 707–722.
- Bernanke, Ben, 1983, Irreversibility, Uncertainty, and Cyclical Investment, *Quarterly Journal of Economics* 8, 85–106.
- Boyle, Phelim P., Jeremy Evnine, and Steven Gibbs, 1989, Numerical Evaluation of Multivariate Contingent Claims, *Review of Financial Studies* 2, 241–250.
- Brown, David, 2009, The Systematic and Non-Systematic Variations in the Systematic Risk of Stock Returns, Working paper, University of Wisconsin-Madison.
- Chalkley, Martin, and In Ho Lee, 1998, Learning and Asymmetric Business Cycles, *Review of Economic Dynamics* 1, 623–645.
- Chamley, Christophe, and Douglas Gale, 1994, Information Revelation and Strategic Delay in a Model of Investment, *Econometrica* 62, 1065–1085.
- David, Alexander, 1997, Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility, *Journal of Financial and Quantitative Analysis* 32, 427–462.
- Dixit, Avinash, and Robert Pindyck, 1994, *Investment Under Uncertainty*. (Princeton University Press).
- Dumas, Bernard, 2001, Super contact and related optimality conditions, *Journal of Economic Dynamics and Control* 15, 675–685.

- Falk, Barry, 1986, Further Evidence on the Asymmetric Behavior of Economic Time Series over the Business Cycle, *Journal of Political Economy* pp. 1096–1109.
- Gale, Douglas, 1996, Delay and Cycles, *Review of Economic Studies* 63, 169–198.
- Guo, Xin, Jianjun Miao, and Erwan Morellec, 2005, Irreversible Investment with Regime Shifts, *Journal of Economic Theory* 132, 37–59.
- Hamilton, James, 1989, A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle, *Econometrica* 57, 357–384.
- Hamilton, James, 1994, *Time Series Analysis*. (Princeton University Press).
- Keynes, John M., 1936, *The General Theory of Employment, Interest, and Money*. (Macmillan, London).
- Liptser, Robert. S., and Albert. N. Shiryaev, 2001, *Statistics of Random Processes I, General Theory*. (Springer Berlin).
- McDonald, Robert, and Daniel Siegel, 1986, The Value of Waiting to Invest, *Quarterly Journal of Economics* pp. 707–728.
- Morellec, Erwan, 2001, Asset Liquidity, Capital Structure and Secured Debt, *Journal of Financial Economics* 61, 173–206.
- Neftci, Salih N., 1985, Are Economic Time Series Asymmetric over the Business Cycle, *Journal of Political Economy* 92, 307–328.
- Nelson, Daniel B., and Krishna Ramaswamy, 1990, Simple Binomial Processes as Diffusion Approximations in Financial Models, *Review of Financial Studies* 3, 393–430.
- Nieuwerburgh, Stijn Van, and Laura Veldkamp, 2006, Learning Asymmetries in Real Business Cycles, *Journal of Monetary Economics* 53, 753–772.

- Pindyck, Robert, 1988, Irreversible Investment, Capacity Choice, and the Value of the Firm, *American Economic Review* pp. 969–985.
- Ramsey, James B., and Philip Rothman, 1996, Time Irreversibility and Business Cycle Asymmetry, *Journal of Money, Credit and Banking* 28, 1–21.
- Sichel, Daniel E., 1993, Business Cycle Asymmetry: A Deeper Look, *Economic Inquiry* 31, 224–236.
- Veldkamp, Laura, 2005, Slow Boom, Sudden Crash, *Journal of Economic Theory* 124, 230–257.
- Veronesi, Pietro, 1999, Stock Market Overreaction on Bad News in Good Times: A Rational Expectations Equilibrium Model, *Review of Financial Studies* 12, 975–1007.
- Veronesi, Pietro, 2000, How does Information Quality Affect Stock Returns?, *Journal of Finance* 55, 807–837.




Figure 1: Posterior beliefs of being in the good state

This figure shows the time series of the posterior belief π estimated from the merged Compustat-CRSP annual database using the EM algorithm. The narrow bars correspond to the recession periods designated by NBER.





Figure 2: Optimal investment threshold for different levels of information quality

This figure depicts the optimal investment boundary for different level of information quality. Panel (a) plots $h = x/k$ as a function of the belief π . Panel (b) $1/h = k/x$ as a function of the belief π . The solid curves depict the investment threshold computed for the base case parameterization given in Table 1. The dashed curves correspond to lower volatilities of the signal process. The two dashed horizontal lines in each panel are the two investment thresholds under perfect information.

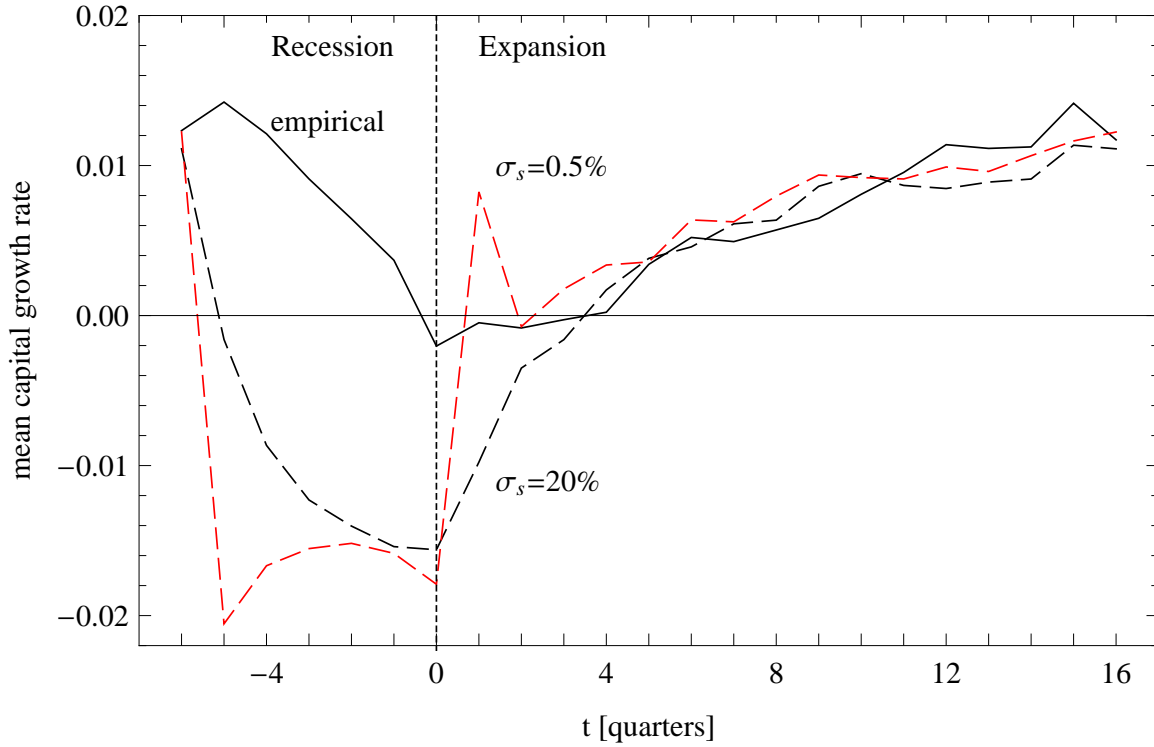
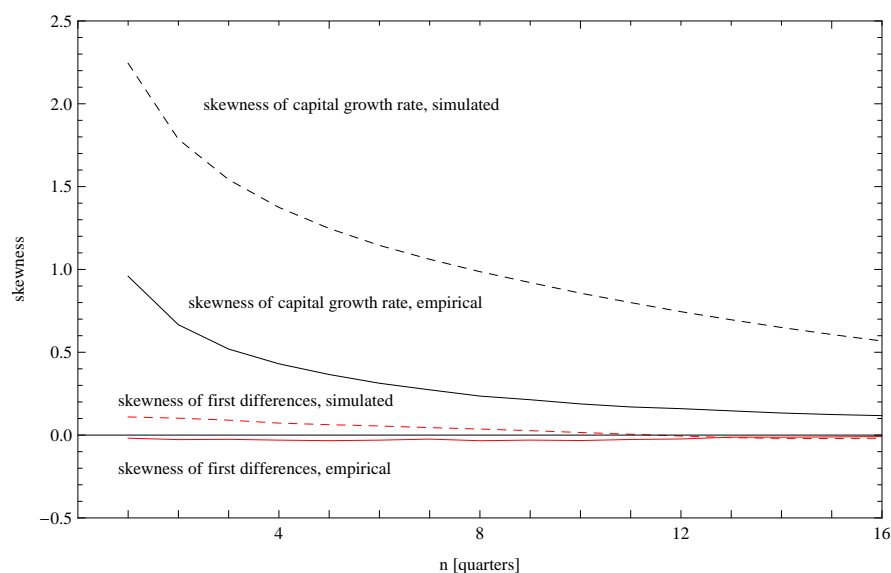
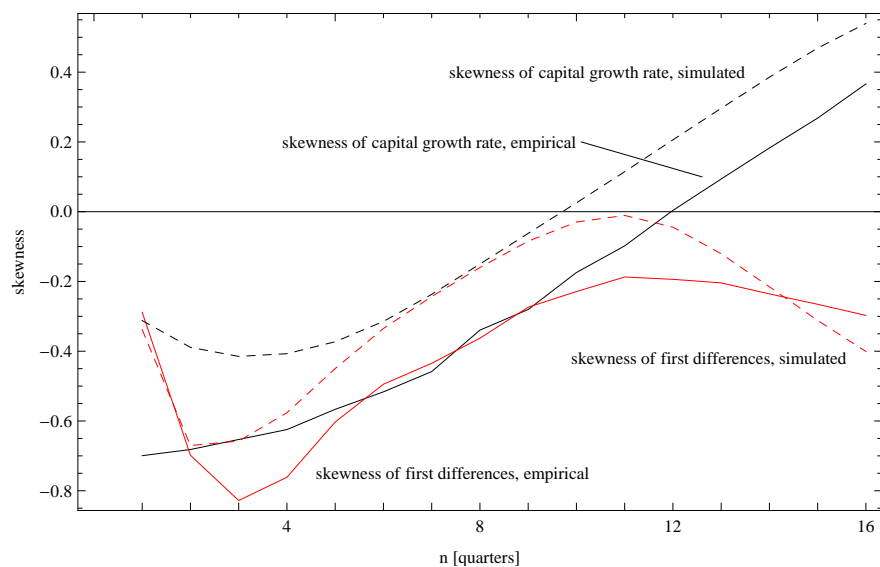


Figure 3: Mean capital growth rate over the business cycle

This graph shows the simulated, as well as the empirically estimated, equal weighted average capital growth rate over a full business cycle. The simulated firms are assumed to follow the optimal investment policy numerically derived Section 3.2. Capital growth rates are first averaged across firms and then averaged across simulated business cycles over the 80-year period of time. The solid line represents the equal weighted quarterly real capital growth rate around the NBER dated business cycle troughs calculated using the Compustat-CRSP merged database from 1975 to 2008.



(a) The skewness of the individual capital growth rate and its first-order difference



(b) The skewness of the average capital growth rate and its first-order differences

Figure 4: Skewness of capital growth rates: empirical estimates vs. simulated results

This figure shows the skewness of the capital growth rate as well as the skewness of the first-order differences of the capital growth rate over various time horizons (between 1 to 16 quarters). Panel (a) shows the skewness at the individual firm level. Panel (b) shows the skewness of the capital growth rate average across firms. The solid curves correspond to the empirical estimates based on the real capital growth rates of the firms in the quarterly Compustat-CRSP merged database from 1975 to 2008. The dashed curves correspond to results from 1000 simulated firms following the optimal investment policy derived from our model.

Table 1: **Parameter values**

This table summarizes the parameter values for our base case scenario. The parameters governing the Markov switching process, $\lambda_{h,l}$, $\lambda_{l,h}$, μ_h , μ_l , are estimated from the Compustat-CRSP database using the EM algorithm.

notation	economic meaning	value
α	operating profit parameter	0.53
μ_h	expected growth rate in good state	0.1060
μ_l	expected growth rate in bad state	-0.1085
$\lambda_{h,l}$	transition intensity from good state to bad state	0.2573
$\lambda_{l,h}$	transition intensity from bad state to good state	0.6327
λ_q	termination rate	0.0604
σ_s	instantaneous volatility of s_t	0.20
σ_x	instantaneous volatility of x_t	0.20
ξ	depreciation rate	0.1
r	risk free rate	0.05
ρ	correlation between dW_{st} and dW_{xt}	0.05