A Bayesian Approach to Real Options: The Case of Distinguishing between Temporary and Permanent Shocks

STEVEN R. GRENADIER and ANDREY MALENKO*

ABSTRACT

Traditional real options models demonstrate the importance of the "option to wait" due to uncertainty over future shocks to project cash flows. However, there is often another important source of uncertainty: uncertainty over the permanence of past shocks. Adding Bayesian uncertainty over the permanence of past shocks augments the traditional option to wait with an additional "option to learn." The implied investment behavior differs significantly from that in standard models. For example, investment may occur at a time of stable or decreasing cash flows, respond sluggishly to cash flow shocks, and depend on the timing of project cash flows.

DURING THE PAST TWO DECADES the real options approach to valuation of irreversible investment opportunities has become part of the mainstream literature in financial economics. The central idea is that the opportunity to invest is equivalent to an American call option on the underlying investment project. As a consequence, the problem of optimal investment timing is analogous to the optimal exercise decision for an American option. Applications of the real options approach are now numerous.¹

One feature that is common to virtually all real options models is that the underlying uncertainty pertains only to future shocks. Since future cash flows are uncertain, there is an important opportunity cost of investing today: the value of the asset might go up so that tomorrow will be an even better time to invest. This opportunity cost is often referred to in the literature as the "option

*Grenadier and Malenko are at the Graduate School of Business, Stanford University. We thank Geert Bekaert, Jonathan Berk, Thomas Chemmanur, Barney Hartman-Glaser, Campbell Harvey (the Editor), Nadya Malenko, Robert McDonald, Paul Pfleiderer, Lukas Schmid, Ilya Strebulaev, Neng Wang, Jeffrey Zwiebel, the anonymous Associate Editor, two anonymous referees, seminar participants at BYU Marriott School and Stanford GSB, and participants at the UNC/Duke 2008 Corporate Finance Conference, 2008 EFA Meeting in Athens and the 8th Doctoral Conference at LBS for their helpful comments.

¹The early literature is well summarized in Dixit and Pindyck (1994). McDonald and Siegel (1986) provided the (now standard) setting, which was later extended to account for a time-tobuild feature (Majd and Pindyck (1987)) and strategic interactions among several option holders (Grenadier (1996, 2002), Lambrecht and Perraudin (2003), Novy-Marx (2007)). Real options modeling is used to study specific industries such as natural resources (Brennan and Schwartz (1985)) or real estate (Titman (1985), Williams (1991)). Recent developments include incorporating agency conflicts (Grenadier and Wang (2005)) and behavioral preferences (Grenadier and Wang (2007), Nishimura and Ozaki (2007)) into the standard setting. to wait." Conversely, there is also a benefit of immediate investment in that the project begins to generate cash flows immediately. At the point at which the costs and benefits of waiting an additional instant are equalized, investment is made. The implied investment strategy in this standard setting is to wait until the expected value of the asset reaches some upper threshold value and then invest.

In this paper we propose a very different kind of real options problem. While uncertainty about future shocks is important, it may not be the most important let alone only type of uncertainty faced by a firm. Indeed, another critical source of uncertainty is uncertainty about *past* shocks. This is the case when the firm observes the past shock but fails to identify its exact properties. As time passes, the firm updates its beliefs about the past shock. Therefore, unlike the standard models, the uncertainty in our setting is a combination of uncertainty over future shocks and Bayesian updating, or learning about the nature of past shocks. In this context, there is an additional benefit to waiting: the opportunity to learn more about the nature of past shocks and use this information to make a better investment decision. This gives rise to an additional trade-off between investing now and waiting, embodied by what we call the "option to learn."

Specifically, this paper focuses on the case in which uncertainty about past shocks comes from the firm's inability to distinguish between temporary and permanent shocks to the cash flow process. To gain intuition on the trade-off between investing now and waiting, consider a positive shock to the underlying cash flows of the project. In the standard setting the firm's strategy is rather simple: the firm should invest if and only if the shock is high enough so that the value of the cash flow process exceeds some threshold level. However, when the firm does not know if the shock is temporary or permanent, it may want to wait in order to learn more about the identity of the shock. Indeed, if the firm waits until tomorrow and the value of the cash flow process is still at a high level, then the past shock is more likely to represent a positive fundamental change. On the other hand, if the cash flow process decreases, then there is a greater chance that the past positive shock was simply a result of temporary, nonfundamental fluctuations.

To account for the two fundamentally different notions of uncertainty, we model the evolution of the cash flow process in the following way. First, the cash flow process is subject to standard Brownian innovations, as in traditional real options models. Second, the cash flow process is subject to jump shocks, which may be permanent or temporary. If a shock is permanent, the cash flow process is changed forever. In contrast, if a shock is temporary, the cash flow process is changed only temporarily until the shock reverts at some random point in the future. Importantly, one does not know with certainty the true nature of these jump shocks. While learning about Brownian innovations is done by simply observing the increments of the cash flow process, learning about the permanent or temporary nature of past jump shocks involves waiting to see if they revert.

Our argument is based on two building blocks that distinguish this model from the prior real options literature. The first building block is the presence of both temporary and permanent shocks to cash flows. While virtually all real options models focus solely on permanent shocks, the presence of temporary shocks is a natural feature of many real world economic environments. The focus on only permanent shocks is clearly a simplification that permits one to model the cash flow process as a geometric Brownian motion. However, as Gorbenko and Strebulaev (2010) show, assuming a geometric Brownian motion cash flow process leads to a number of undesirable empirical properties.²

The second building block is the inability of the firm to distinguish between permanent and temporary shocks. This is an especially important feature of investments in natural resources, where it is often unclear when a change in the commodity price represents a fundamental or ephemeral shift. Indeed, as the World Bank's Global Economic Prospects annual report for 2000 states:

Most important, distinguishing between temporary and permanent shocks to commodity prices can be extraordinary difficult. The swings in commodity prices can be too large and uncertain to ascertain their causes and nature. The degree of uncertainty about duration of a price shock varies. For example, market participants could see that the sharp jump in coffee prices caused by the Brazilian frost of 1994 was likely to be reversed, assuming a return to more normal weather. By contrast, most analysts assumed that the high oil prices during the mid 1970s and early 1980s would last indefinitely.³

We show that augmenting the traditional uncertainty over future shocks with Bayesian uncertainty over past shocks gives rise to a number of important and novel implications. First, in contrast to the standard real options setting in which investment occurs when cash flows rise to a constant trigger, we find that once we account for Bayesian uncertainty about past shocks, the investment trigger is a function of the firm's changing beliefs about the nature of past shocks. This distinction is due to the fact that while in standard real options models uncertainty is constant over time, in our Bayesian framework uncertainty depends on the timing of past shocks. More specifically, uncertainty is high soon after the arrival of a large shock since the firm is very unsure whether it is permanent or temporary. However, as time goes by, uncertainty about the past shock declines as the firm becomes increasingly more confident about the shock's permanence. Thus, in addition to the standard option to wait for realizations of future shocks, there is now also an option to wait in order to learn more about the nature of past shocks. We find that these two options to delay investment are additive, and appear to be of similar magnitude when past shocks are significant. Ignoring either option would lead to a similar investment timing error. Because uncertainty increases after a shock arrives, the model rationalizes a sluggish response of investment to cash flow

²For example, it implies that the volatilities of cash flow and asset growth are equal, while empirically volatilities of cash flow growth are much higher than volatilities of asset value growth. ³World Bank, 2000. *Global Economic Prospects and Developing Countries*, p. 110.

shocks. Intuitively, after a new shock arrives, the firm might find it optimal to postpone investment in order to wait and see if the past shock is indeed permanent.

A second important implication of the Bayesian model is that investment may occur in the face of stable or even decreasing cash flows. Since standard real options models imply that investment will be triggered when shocks push the underlying cash flow level up to a fixed threshold, the firm can rationally exercise the investment option only when the underlying cash flows increase. However, when the traditional real options model is extended to include the Bayesian feature, the firm may invest even when the cash flow process is stable or declining simply because it becomes sufficiently certain about the permanent nature of past shocks. While investment in the face of stable or decreasing cash flows can be obtained in other models with more than one state variable, our argument is special in that it links these features to the timing of observable shocks and thereby offers an intuitive explanation for firms (and industries) investing in markets where cash flows are stable, or even declining. For example, as we discuss in Section III, learning about the persistence of the 1973 oil price shock was likely one reason for the rapid office construction in Denver and Houston during the late 1970s and early 1980s when the underlying property cash flows were declining.

Finally, in contrast to standard real options models with uncertainty only over future shocks, in the context of Bayesian uncertainty over past shocks the exercise strategy depends on the timing structure of the project's cash flows. In standard real options models, the net present value that triggers exercise is constant, regardless of whether the project's cash flows are "front-loaded" or "back-loaded." That is, irrespective of whether the project's cash flows occur mostly up-front or mostly later in the future, exercise will be triggered when the net present value of future cash flows equals the same premium over the cost. However, with Bayesian uncertainty over past shocks, the exercise strategy depends critically on the timing structure of the project's cash flows. The greater the "front-loadedness" of the option payoff, the less important is the assessment of the relative likelihood that a shock is temporary or permanent. As a result, the option to learn is less valuable, so the firm invests at a lower trigger. Thus, while the firm might invest in a short-term project when the value of the future cash flows exceeds the investment cost by a relatively small amount, the firm might invest in a long-term project only if the value of the future cash flows exceeds the investment cost by a large premium.

While we model shocks as jumps, it should be noted that a jump specification is not necessary for the cash flow process to be a combination of temporary and permanent shocks. One could model the cash flow process as a combination of any two processes, where the first governs long-term cash flows from the project and the second is responsible for short-term deviations from the longterm mean. Uncertainty about the nature of past shocks is then equivalent to observing the total cash flow process but not its separate components. However, modeling shocks as jumps, while not being a necessary feature, allows us to develop a framework that is both tractable and intuitive. Several papers are closely related to our model. Gorbenko and Strebulaev (2010) incorporate temporary shocks into a contingent claims framework of capital structure and show that the presence of temporary shocks provides a potential explanation for several puzzles in corporate finance. Our paper differs from theirs in two important respects. First, and most importantly, our argument relies heavily on the firm's imperfect information and learning about the permanence of past shocks, while Gorbenko and Strebulaev (2010) assume perfect information about past shocks. Second, while we consider firms' investment decisions, their focus is on financing policy.

Our paper is also related to models of parameter uncertainty such as Décamps, Mariotti, and Villeneuve (2005) and Klein (2007, 2009), who solve for the optimal investment timing when the firm is uncertain about the drift parameter of the state process. These models of parameter uncertainty are based on very different notions of uncertainty and learning. First, in these models, the firm learns more about the true parameter over time. In the fullness of time, full information can be achieved. In contrast, our model deals with uncertainty about the state of the economy, which, unlike uncertainty about a parameter of the model, is evolving over time: as new shocks arrive and old ones revert, the learning problem changes dynamically, and even in the fullness of time the true nature of uncertainty is never revealed. Second, in these models of parameter uncertainty all learning is driven by evolutions in the Brownian component of uncertainty.⁴ The Brownian and Bayesian components are thus fundamentally intertwined, which results in difficulty in decomposing the two. For example, a sequence of positive Brownian innovations not only lowers the option to wait but also affects the option to learn as it implies a greater likelihood of a higher drift parameter. In our model's setting, the Brownian and Bayesian components are decoupled. The Brownian uncertainty is identical to that from standard real options models, while learning is driven by a separate shock process. Thus, our model is more appropriate for settings in which learning is driven by the evolution of economic shocks and not by the volatility of the process itself. These distinctions result in important differences in the resulting investment behavior. First, Klein (2007, 2009) shows that it may be the case that Bayesian learning about parameter uncertainty may actually lead to earlier investment than in the full-information setting. In contrast, our model of Bayesian uncertainty over past shocks always leads to delay due to the option to learn. Second, since the arrival of a shock increases uncertainty and thus the value of waiting to learn, our model implies that investment

⁴See also papers on optimal experimentation that deal with problems of optimal control under learning (e.g., Bolton and Harris (1999), Keller and Rady (1999), and Moscarini and Smith (2001)). These papers focus on active learning through experimentation, in contrast to passive learning through observation found in the present paper. Similarly to Décamps et al. (2005) and Klein (2007, 2009), these papers feature Brownian uncertainty and learning, with the Brownian motion providing the noise that makes learning necessary in the first place. In the present paper, by contrast, Brownian uncertainty is irrelevant to learning and simply drives the standard option to wait. responds sluggishly to shocks, an implication not present in models of parameter uncertainty.

Another related paper is Miao and Wang (2007), who solve for the entrepreneur's optimal decision to exit a business in the presence of idiosyncratic nondiversifiable risk. Specifically, Miao and Wang (2007) consider an entrepreneur with incomplete information about her entrepreneurial abilities who chooses between continuing entrepreneurial activity and taking a safe job. While their paper contains Bayesian learning over entrepreneurial ability, it does not allow for Brownian uncertainty of cash flows, as does ours. In addition, in their setting there is only a single source of Bayesian uncertainty, while in our model new shocks can arrive and old shocks can revert. Finally, the focus of their paper is on risk aversion and imperfect hedging of idiosyncratic risks, while ours is on the joint impact of various sources of uncertainty on investment timing.⁵

Finally, our paper also shares the learning feature with Lambrecht and Perraudin (2003), who study competition between two firms for a single investment opportunity when information about investment costs is private. They show that, as time goes by, while the competitor has not invested yet, each firm updates its belief about the competitor's investment costs upward.⁶

The remainder of the paper is organized as follows. Section I describes a setup of the model and for purposes of providing intuition solves the most simple case of a pure jump process with only one shock. Section II focuses on incorporating learning over the past shock, developed in Section I, into the standard geometric Brownian motion setting. Section III provides a discussion of important implications of the model. Section IV generalizes the model to the case of multiple shocks. Finally, Section V concludes.

I. Model Setup and a Simple Case

In this section, we describe the setup of the model, which incorporates uncertainty over the permanence of past shocks into the real options framework. Then, to show the main intuition behind the argument, we consider the simplest possible setting as a special case. In subsequent sections we solve for more general cases.

A. The Investment Option

Consider a standard real options framework in which a risk-neutral firm contemplates irreversible investment. By paying the investment cost I, the

⁵Other papers that include learning based on a Poisson process include Malueg and Tsutsui (1997), Bergemann and Hege (1998), and Keller, Rady, and Cripps (2005). Similar to Miao and Wang (2007), but unlike the present paper, there is only Bayesian uncertainty and no Brownian uncertainty. In these models, firms learn about profitability by observing (their own and sometimes other firms') payoff outcomes.

 6 Our paper is also related to Moore and Schaller (2002), who extend the neoclassical q theory of investment by allowing for permanent and temporary shocks to interest rates, and for uncertainty about them.

firm obtains the perpetual cash flow X(t). The firm is free to invest in the project at any time it chooses. The firm has a constant discount rate r > 0.

The evolution of the cash flow process consists of two components. As usual in the literature, the cash flow process evolves according to the geometric Brownian motion with drift α and diffusion σ , where $\alpha < r$ to ensure finite values.⁷ In addition, the cash flow process is subject to jump shocks, which can be temporary or permanent. Upon the arrival of a shock, the cash flow process jumps from X(t) to $X(t)(1+\varphi)$, with $\varphi > 0.8$ A permanent shock changes the value of the cash flow process forever, while a temporary shock changes the value of the cash flow process only temporarily until it reverts at some random point in the future. Once the shock reverts, the cash flow process jumps down by the multiple of $1 + \varphi$. Importantly, the firm is unable to distinguish between permanent and temporary shocks, and thus uses Bayesian updating to assess their relative likelihoods. We assume that the arrival and reversal (of temporary shocks) follow jump processes. Specifically, a permanent shock arrives with intensity λ_1 , a temporary shock arrives with intensity λ_2 , and each outstanding temporary shock reverts with intensity $\lambda_3 > \lambda_2$. Reversal of each outstanding temporary shock is independent of arrivals of new shocks and reversals of other outstanding temporary shocks. Arrival of a new permanent (temporary) shock is independent of arrival of a new temporary (permanent) shock and reversals of outstanding temporary shocks.

Let us formally define the evolution of the cash flow process, conditional on information about the identities of the shocks. Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a filtration $(\mathcal{F}_t) = \{\mathcal{F}_t : t \ge 0\}$. The evolution of the cash flow process X(t) is given by

$$dX(t) = \alpha X(t) + \sigma X(t) dB(t) + \varphi X(t) (dM_1(t) + dM_2(t)) - \frac{\varphi X(t)}{1 + \varphi} dN(t), \quad (1)$$

where $X(0) = X_0$, and B(t), $M_1(t)$, $M_2(t)$, and N(t) are (\mathcal{F}_t) -adapted. In (1), dB(t) is the increment of a standard Wiener process, and $M_1(t)$, $M_2(t)$, and N(t) are nonexplosive counting processes corresponding to the arrival of permanent shocks, the arrival of temporary shocks, and the reversal of outstanding temporary shocks, respectively. These shocks' intensities are λ_1, λ_2 , and $k(t)\lambda_3$, respectively, where k(t) is the number of outstanding temporary shocks at time $t: k(t) = M_2(t) - N(t)$. Conditional on \mathcal{F}_t , the increments of $M_1(t), M_2(t)$, and N(t) over the next instant are independent. Defining the total number of outstanding shocks by $n(t) = M_1(t) + M_2(t) - N(t)$, we can say that the

⁷To be more precise, the upper bound on α that guarantees finite values depends on the particular specification of the model. When the number of potential permanent shocks is finite, the upper bound is r. When the number of potential permanent shocks is infinite, the upper bound is $r - \lambda_1 \varphi$, where λ_1 and φ are defined later.

⁸The model can be extended to the case of stochastic jump amplitudes, including the possibility of both positive and negative jumps, albeit at the cost of additional mathematical complexity. However, the basic intuition and results remain unchanged.

triple (X(t), n(t), k(t)) represents the state under the full-information filtration (\mathcal{F}_t).⁹

Now, let us consider the evolution of the cash flow process from the firm's point of view. The firm observes the evolution of X(t), but cannot distinguish whether jumps in X(t) are caused by permanent shocks $dM_1(t)$ or temporary shocks $dM_2(t)$. Mathematically, the information available to the firm is represented by the partial-information filtration (\mathcal{G}_t) generated by (B, M, N), where M(t) = $M_1(t) + M_2(t)$. Hence, at any time t, the firm knows the total number n(t) =M(t) - N(t) of outstanding shocks, but does not know the number $k(t) = M_2(t) -$ N(t) of outstanding temporary shocks. However, given the Bayesian learning process described in detail in the next subsection, the firm possesses posterior probabilities on the number k(t) of outstanding temporary shocks:

$$p_k(t) \equiv \mathbb{P}(k(t) = k \mid \mathcal{G}_t). \tag{2}$$

The firm therefore attaches probability $p_1(t)$ that there is exactly one outstanding temporary shock, $p_2(t)$ that there are two outstanding temporary shocks, ..., $p_{n(t)}(t)$ that all of the n(t) outstanding shocks are temporary. Thus, at time *t* the firm estimates that over the next instant (t, t + dt), a new shock arrives with probability $(\lambda_1 + \lambda_2)dt$ and an existing shock reverts with probability $\sum_{k=1}^{n(t)} p_k(t)k\lambda_3dt$.

Let S(t) denote the expected (under the firm's information set) present value of the project's underlying cash flows:

$$S(t) = \mathbb{E}\left[\int_{t}^{\infty} e^{-r(s-t)} X(s) \, ds \, | \, \mathcal{G}_t\right].$$
(3)

We can then say that the firm's investment problem is to find a stopping time $\tau^* \in (\mathcal{G}_t)$ such that

$$\sup_{\tau \in (\mathcal{G}_t)} \mathbb{E}[e^{-r\tau}(S(\tau) - I) | \mathcal{G}_t] = \mathbb{E}\left[e^{-r\tau^*}(S(\tau^*) - I) | \mathcal{G}_t\right].$$
(4)

B. The Bayesian Learning Process

Since the firm is unsure whether past shocks are permanent or temporary, an important state variable is the firm's belief about the identity of past shocks.

⁹Because the dynamics of X(t) are driven by both Brownian motion and jump processes, X(t) is a jump-diffusion process (see Chapter 11.4.1 in Shreve (2004) for a discussion) and, more generally, a semimartingale (see Chapter 2 in Protter (2004) for an overview). Mathematically, the stochastic structure of the model is similar to models of correlated default in which default of one claim causes changes in the default intensity of other claims. For example, see the primary-secondary default model in Jarrow and Yu (2001). Both in Jarrow and Yu (2001) and in our model, there is a "primary" jump process with a constant intensity, whose arrivals influence the intensity of a "secondary" jump process.

Consider moments t and t + dt for any t and infinitesimal positive dt. Depending on the history between t and t + dt, there are three cases to analyze:¹⁰

- 1. no new shocks occur or outstanding shocks reverse between t and t + dt;
- 2. a new shock occurs; and
- 3. an outstanding shock reverses.

Consider the first case, where there are no changes in the number of outstanding shocks between t and t + dt. Suppose that there are n outstanding shocks and at time t the firm assesses the probability that k of them are temporary at $p_k(t)$. Then the beliefs of the firm are defined by vector $p(t) = (p_0(t)p_1(t) \dots p_n(t))$. If k of the outstanding shocks are temporary, over a short period of time dta shock reverts back with probability $k\lambda_3dt$. Similarly, a new shock arrives with probability $(\lambda_1 + \lambda_2)dt$. Thus, conditional on k of the shocks being temporary, the probability that there are no new shocks or reversions over dt equals $1 - (\lambda_1 + \lambda_2)dt - k\lambda_3dt$. Using Bayes's rule, the posterior probability $p_k(t + dt)$ is given by

$$p_{k}(t+dt) = \frac{p_{k}(t)(1-(\lambda_{1}+\lambda_{2})dt-k\lambda_{3}dt)}{1-(\lambda_{1}+\lambda_{2})dt-\lambda_{3}\sum_{i=1}^{n}p_{i}(t)idt}.$$
(5)

Equation (5) is a direct application of Bayes's rule. The numerator is the joint probability of having k temporary outstanding jumps and observing no jumps or reversions between t and t + dt. The denominator is the sum over i = 0, 1, ... of the joint probabilities of having i temporary outstanding jumps and observing no jumps or reversions between t and t + dt.

We can rewrite (5) as

$$\frac{p_k(t+dt) - p_k(t)}{dt} = -\frac{\lambda_3 p_k(t) \left(k - \sum_{i=1}^n p_i(t)i\right)}{1 - (\lambda_1 + \lambda_2) dt - \lambda_3 \sum_{i=1}^n p_i(t) i dt}.$$
(6)

Taking the limit as $dt \rightarrow 0$,

$$\frac{dp_k(t)}{dt} = -\lambda_3 p_k(t) \left(k - \sum_{i=1}^n p_i(t)i \right).$$
(7)

The dynamics of $p_k(t)$ while there are no new shocks or reversions has two interesting properties. First, $p_k(t)$ increases over time when $k < \sum_{i=1}^{n} p_i(t)i$, and decreases over time otherwise. Intuitively, when $k < \sum_{i=1}^{n} p_i(t)i$, the likelihood of reversion conditional on having k outstanding temporary jumps is lower

¹⁰Note that the probability of observing more than one new shock or reversion between t and t + dt has the order $(dt)^2$. Since dt is infinitesimal, we can ignore these cases. The same is true for the likelihood of both a new shock and a reversion occurring at the same instant.

than the unconditional likelihood of reversion. Therefore, when the firm does not observe a new shock or a reversion, it updates its beliefs $p_k(t)$ upward. The opposite is true when $k > \sum_{i=1}^{n} p_i(t)i$. Second, the speed of learning is proportional to λ_3 . In other words, if temporary shocks are more short-term, then the firm updates its beliefs faster than if they are more long-term.

Now consider the second case. If a new shock occurs between t and t + dt, then the updated beliefs equal

$$p_k(t+dt) \equiv \hat{p}_k(p(t)) = p_{k-1}(t)\frac{\lambda_2}{\lambda_1+\lambda_2} + p_k(t)\frac{\lambda_1}{\lambda_1+\lambda_2}.$$
(8)

The intuition behind (8) is simple. After the shock, there can be k temporary shocks outstanding either if there were k - 1 temporary shocks and the new shock is temporary or if there were k temporary shocks and the new shock is permanent.

Finally, considering the third case, if an outstanding shock reverses between t and t + dt, by Bayes's rule the firm's updated beliefs equal

$$p_{k}(t+dt) \equiv \tilde{p}_{k}(p(t)) = \frac{p_{k+1}(t)(k+1)}{\sum_{i=1}^{n} p_{i}(t)i}.$$
(9)

When a shock reverses, the firm learns that it was temporary. Hence, there are k outstanding temporary shocks after the reversal if and only if there were k + 1 outstanding temporary shocks before that. The joint probability of having k + 1 temporary shocks at time t and observing a reversion between t and t + dt equals $p_{k+1}(t)(k+1)\lambda_3 dt$. The probability of observing a reversion between t and t + dt conditional only on being at time t equals $\sum_{i=1}^{n} p_i(t)i\lambda_3 dt$. Dividing the former probability by the latter yields (9).

Equations (7) to (9) fully characterize the dynamics of the firm's beliefs. Notice that reversion of a shock decreases the number of outstanding shocks by one, while arrival of a new shock increases this number by one.

To keep the model both intuitive and tractable, in the early sections of the paper we focus on the case in which there is only one shock, which can be either permanent or temporary.¹¹ In this case, the firm's beliefs can be characterized by a single state variable $p(t) \equiv p_1(t)$, which is the probability that the outstanding shock is temporary. Before the arrival and after the reversion, there is no uncertainty about past shocks. At the moment t_0 of the arrival of the shock, the firm assesses the prior probability of the shock being temporary as $p(t_0) \equiv \pi_0 = \lambda_2/(\lambda_1 + \lambda_2)$. After the shock arrives, the firm continuously updates its assessment of this probability according to (7),

$$\frac{dp(t)}{dt} = -\lambda_3 p(t)(1 - p(t)),$$
(10)

¹¹Section IV considers the multi-shock setting.

which is solved by

$$p(t) = \frac{\lambda_2}{\lambda_1 e^{\lambda_3 (t-t_0)} + \lambda_2}.$$
(11)

C. A Simple Case

Before presenting the details of our analysis in subsequent sections, we consider the simplest setting to illustrate the intuition behind our results. Specifically, for the cash flow specification in (1), we set $\alpha = \sigma = 0$ and assume that there is only one shock, which can be either permanent or temporary. In this case the cash flow process is rather simple. At all times prior to the arrival of the shock, the cash flow stream is fixed at X_0 . Upon the arrival of the shock, it jumps up to $X_0(1 + \varphi)$. With a permanent shock the cash flow remains at the higher level $X_0(1 + \varphi)$ forever after, but with a temporary shock the cash flow eventually reverts to the level X_0 at some (random) point in the future.

Suppose that the shock has arrived. Let us begin by calculating some simple values. If the firm knows for sure that the shock is permanent, then the present value of the project's cash flows is $\frac{X_0(1+\varphi)}{r}$, which is simply the present value of the perpetual cash flow $X_0(1+\varphi)$. Conversely, if the firm knows for sure that the shock is temporary, then the present value of the project's cash flows is $\frac{X_0(1+\varphi+\lambda_3/r)}{r+\lambda_3}$, which is the present value of receiving a flow of $X_0(1+\varphi)$ until the shock is reversed and X_0 thereafter. Therefore, the present value of the project's underlying cash flows upon investing at the belief level p, defined by (3), is equal to

$$S(p) = (1-p)\frac{X_0(1+\varphi)}{r} + p\frac{X_0(1+\varphi+\lambda_3/r)}{r+\lambda_3}.$$
 (12)

To ensure that a nontrivial solution exists for the problem, we make the assumption that X_0 satisfies

$$\frac{rI}{1+\varphi} < X_0 < \frac{r+\lambda_3\pi_0}{1+\varphi+\frac{\lambda_3}{r}\pi_0}I.$$
(13)

Economically, the lower bound means that if the shock is known to be permanent, the net present value of investing into the project is positive, $\frac{X_0(1+\varphi)}{r} - I > 0$, or else there would never be any investment in this Bayesian setting. The upper bound is equivalent to an assumption that learning has positive value. In particular, this assumption implies that if the shock is known to be temporary, then it is not optimal to invest: $\frac{X_0(1+\varphi+\lambda_3/r)}{r+\lambda_3} - I < 0$.

Let G(p) denote the value of the option to invest, while a shock persists, where p is the current value of the belief process. By Itô's lemma for jump processes (see Shreve (2004)), the dynamics of G(p) under (\mathcal{F}_t) are given by

$$dG(p) = -\lambda_3 p(1-p)G'(p)dt + (H(X_0) - G(p))dN(t),$$
(14)

where $H(X_0)$ is the value of the option to invest after the shock reverts. In (14), the intensity of N(t) is equal to $k(t)\lambda_3$, where k(t) is the indicator of an outstanding temporary shock at time t, known under (\mathcal{F}_t) . Because the firm believes that the shock is temporary with probability p, $\mathbb{E}[dN(t) | \mathcal{G}_t] = p\lambda_3 dt$. Therefore, the expectation of the differential with respect to the information set of the firm is equal to

$$\mathbb{E}\left[\frac{dG(p)}{dt} \left| \mathcal{G}_t \right] = -\lambda_3 p(1-p)G'(p) + p\lambda_3(H(X_0) - G(p)).$$
(15)

By definition, it must be the case that $\mathbb{E}[dG(p)/dt | \mathcal{G}_t] = rG(p)$. Therefore, G(p) must satisfy the following differential equation:

$$(r + p\lambda_3)G(p) = -\lambda_3 p(1 - p)G'(p) + p\lambda_3 H(X_0).$$
(16)

By (13), the option will not be exercised after the reversion, so $H(X_0) = 0$ and

$$(r+p\lambda_3)G(p) = -\lambda_3 p(1-p)G'(p).$$
(17)

The general solution to (17) is

$$G(p) = C \cdot (1-p) \left(\frac{1}{p} - 1\right)^{\frac{r}{\lambda_3}},$$
(18)

where *C* is a constant to be determined by the appropriate boundary conditions.

The option will be exercised when the conditional probability of the shock being temporary decreases to a lower threshold. Intuitively, it is optimal for the firm to invest only when it becomes sufficiently sure that the past shock is permanent. Let \bar{p} denote the trigger at which the option is exercised. The exercise trigger \bar{p} and constant *C* are jointly determined by the following boundary conditions:

$$G(\bar{p}) = (1 - \bar{p})\frac{X_0(1 + \varphi)}{r} + \bar{p}\frac{X_0(1 + \varphi + \lambda_3/r)}{r + \lambda_3} - I,$$
(19)

$$G_p(\bar{p}) = -\frac{X_0(1+\varphi)}{r} + \frac{X_0(1+\varphi+\lambda_3/r)}{r+\lambda_3}.$$
 (20)

The first equation is the value-matching condition. It reflects the fact that upon exercise, the value of the project is equal to its net present value, $S(\bar{p}) - I$. The second equation is the smooth-pasting condition.¹² It ensures that the trigger \bar{p} maximizes the option value.

1960

 $^{^{12}}$ This condition is also known as the high-contact condition (see Krylov (1980) or Dumas (1991) for a discussion).

Combining (18) with (19) and (20) yields the optimal investment threshold \bar{p} and constant *C*:

$$\bar{p} = \frac{X_0(1+\varphi) - rI}{\lambda_3 \left(I - \frac{X_0}{r}\right)},\tag{21}$$

$$C = \frac{1}{1 - \bar{p}} \left(\frac{1}{\bar{p}} - 1\right)^{-\frac{r}{\lambda_3}} \left[(1 - \bar{p}) \frac{X_0(1 + \varphi)}{r} + \bar{p} \frac{X_0(1 + \varphi + \lambda_3/r)}{r + \lambda_3} - I \right].$$
(22)

Now we can see that the upper bound on X_0 from (13) is equivalent to assuming that $\bar{p} < \pi_0$. This restriction seems entirely reasonable since it ensures that there is some benefit to learning. Combined with the lower bound on X_0 , it guarantees that $\bar{p} \in (0, \pi_0)$. Also, note that in this simple case, the option will never be exercised prior to the arrival of the shock.¹³ This is because the value of exercising prior to the arrival of the shock is dominated by the value of waiting until the shock arrives and then exercising immediately, as shown in the Appendix (which itself is dominated by optimal exercise after the shock arrives).

This simple case highlights the key notion of Bayesian learning in a real options context. Perhaps the best way to gain intuition on the optimal investment policy is to rewrite the expression for the optimal trigger \bar{p} outlined in equation (21) as

$$X_0(1+\varphi) = \bar{p}\lambda_3 \left(I - \frac{X_0}{r}\right) + rI.$$
(23)

The intuition behind expression (23) is the trade-off between investing now versus investing a moment later if the past shock persists, where the value of waiting is explicitly an option to learn. If the firm invests now, it gets the benefit of the cash flow $X_0(1 + \varphi)$ over the next instant. This is the term on the left-hand side of the equal sign. If the firm waits a moment and invests if the past shock persists, it faces a small chance of the shock reversing, in which case the expected value gained by not investing is equal to $I - \frac{X_0}{r}$. Importantly, by waiting that additional moment, it gains the opportunity to forgo investment should the past shock prove to be temporary. The second term on the right-hand side of the equal sign, rI, is the savings from delaying the investment cost by an instant. At the optimal trigger, these two sides are exactly equal, and the firm is indifferent between investing now and waiting a moment to learn.

Figure 1 plots simulated sample paths of the cash flow process, X(t), and the corresponding firm's belief process, p(t). Before the arrival of the shock, investment is suboptimal since the project does not generate high enough cash flows. When the shock arrives, the cash flow process jumps from X_0 to $X_0(1 + \varphi)$. Despite the net present value of the project being positive, the firm finds investment suboptimal because of the option to learn more about the nature of

¹³This result will not hold in several generalized versions of the model that follow.



Figure 1. Simulation of the firm's investment strategy. The top graph shows the simulated sample path of the cash flow process, X(t), for the simple model of Section I.C. An upward step corresponds to the arrival of a shock. For this sample path, the shock arrives at time 0.67 and persists through year 3, implying that over time the firm becomes more confident that the shock is permanent. The bottom graph shows the dynamics of the firm's belief process as well as the investment threshold \bar{p} . The investment occurs when the firm's belief process falls to \bar{p} . At this point, the benefits of learning and the costs of deferring the receipt of the cash flows are precisely offset. The parameter values are r = 0.04, $\varphi = 0.2$, $\lambda_1 = 0.5$, $\lambda_2 = 1$, $\lambda_3 = 2$, I = 25, and $X_0 = 0.95$.

the past shock. As time goes by and the shock does not revert back, the firm updates its beliefs downwards. When the firm becomes sufficiently confident that the past shock is permanent, it invests. In the example in Figure 1 this happens more than 1.5 years after the arrival of the shock.

II. Model with Bayesian and Brownian Uncertainties

In the simple case in the previous section, the cash flow process X(t) is a pure jump process. Since most traditional real options models are based on Brownian uncertainty, we generalize the simple case by allowing X(t) to follow a combined jump and geometric Brownian motion process. In this case the firm faces two fundamentally different types of uncertainty. First, as in traditional models, the firm is uncertain about future shocks. This is captured by the cash flow process being subject to the Brownian innovations and unknown timing of the shock before its arrival. This gives rise to the option to wait for realizations of future shocks usually studied in the real options literature. Second, the firm is also uncertain about past shocks, as it is unable to perfectly identify whether the past shock is permanent or temporary. This gives rise to the option to learn more about the nature of past shocks. In order to highlight the interaction between these uncertainties, we assume only a single shock. Section IV extends the model to the general case of multiple shocks.

Since there is only one shock, which can be either permanent or temporary, we solve the model by backward induction. First, we consider the optimal timing of investment conditional on a temporary shock having reversed. In this case, there is neither uncertainty about past shocks nor the possibility of new jumps, so the problem is the standard real options problem studied in the literature. Second, we move one step back and consider the situation after a shock arrives, when the firm is uncertain about its nature. Finally, we consider the situation before a shock arrives.

A. Optimal Investment after a Shock Reverses

First, consider the situation after a temporary shock reverses. In this case there is no uncertainty about past shocks, so the only underlying state variable is X(t), which follows

$$dX(t) = \alpha X(t) dt + \sigma X(t) dB(t).$$
(24)

Let H(X) denote the value of the investment option after the shock reverses, where X is the current value of the cash flow process. Using standard arguments (e.g., Dixit and Pindyck (1994)), in the range before investment, H(X)must solve the following differential equation:

$$rH = \alpha XH_X + \frac{1}{2}\sigma^2 X^2 H_{XX}.$$
 (25)

Equation (25) must be solved subject to the following boundary conditions:

$$H(X^*) = \frac{X^*}{r - \alpha} - I, \qquad (26)$$

$$H'(X^*) = \frac{1}{r - \alpha},\tag{27}$$

where X^* is the investment threshold. The first boundary condition is the value-matching condition, which states that at the exercise time the value of the option is equal to the net present value of the project. The second boundary condition is the smooth-pasting condition, which guarantees that the exercise strategy is chosen optimally. The last boundary condition is H(0) = 0, reflecting the fact that X(t) = 0 is the absorbing barrier for the cash flow process.¹⁴

¹⁴This boundary condition applies to all of the valuation equations. However, to avoid repetition, we do not list it in the future valuation equations.

Solving (25) subject to (26) and (27) yields the investment threshold X^* ,

$$X^* = \frac{\beta}{\beta - 1} (r - \alpha) I, \qquad (28)$$

where β is the positive root of the fundamental quadratic equation $\frac{1}{2}\sigma^2\beta(\beta-1) + \alpha\beta - r = 0$:

$$\beta = \frac{1}{\sigma^2} \left[-\left(\alpha - \frac{\sigma^2}{2}\right) + \sqrt{\left(\alpha - \frac{\sigma^2}{2}\right)^2 + 2r\sigma^2} \right] > 1.$$
 (29)

The corresponding value of the investment option H(X) is then given by

$$H(X) = \begin{cases} \left(\frac{X}{X^*}\right)^{\beta} \left(\frac{X^*}{r-\alpha} - I\right), & \text{if } X < X^*, \\ \frac{X}{r-\alpha} - I, & \text{otherwise.} \end{cases}$$
(30)

B. Optimal Investment while a Shock Persists

Consider the situation in which there is an outstanding shock that can be either permanent or temporary. In this case, analogous to (1), the evolution of X(t) under (\mathcal{F}_t) is given by

$$dX(t) = \alpha X(t) dt + \sigma X(t) dB(t) - \frac{\varphi X(t)}{1 + \varphi} dN(t), \qquad (31)$$

where the intensity of N(t) is equal to $k(t)\lambda_3$, and where k(t) is the indicator for an outstanding temporary shock at time t, known under (\mathcal{F}_t) but unknown by the firm.

We begin by calculating some simple values. If the firm knows for sure that the shock is permanent, then the present value of the project's cash flows is $\frac{1}{r-\alpha}X(t)$. If the firm knows for sure that the shock is temporary, then the present value of the project's cash flows is $\frac{1+\lambda_3/((r-\alpha)(1+\varphi))}{r-\alpha+\lambda_3}X(t)$. Therefore, the expected present value of the project's underlying cash flows, defined by (3), is equal to

$$S(X(t), p(t)) = \left[(1 - p(t))\frac{1}{r - \alpha} + p(t)\frac{1 + \frac{\lambda_3}{(r - \alpha)(1 + \varphi)}}{r - \alpha + \lambda_3} \right] X(t).$$
(32)

Let G(X, p) denote the value of the option to invest while the shock persists, where X and p are the current values of the cash flow and belief processes, respectively. By Itô's lemma for two-dimensional processes with jumps (see Theorem 11.5.4 in Shreve (2004)), in the range of (X, p) at which the option is

1964

not exercised, the dynamics of G(X, p) under (\mathcal{F}_t) satisfy

$$dG = \left(\alpha XG_X + \frac{\sigma^2}{2}X^2G_{XX} - \lambda_3 p(1-p)G_p\right)dt + \sigma XG_X dB(t) + \left(H\left(\frac{X}{1+\varphi}\right) - G\right)dN(t),$$
(33)

where H(X) is the value of the option when no more jumps can occur, given by (30). Because the firm believes that the shock is temporary with probability p, $\mathbb{E}[dN(t) | \mathcal{G}_t] = p\lambda_3 dt$. Hence, the instantaneous conditional expected change in G(X, p) is equal to

$$\mathbb{E}\left[\frac{dG}{dt}\left|\mathcal{G}_{t}\right]=\alpha XG_{X}+\frac{\sigma^{2}}{2}X^{2}G_{XX}-\lambda_{3}p(1-p)G_{p}+p\lambda_{3}\left(H\left(\frac{X}{1+\varphi}\right)-G\right).$$
 (34)

Thus, analogous to equation (16), in the range of (X, p) at which the option is not exercised, G(X, p) must satisfy the following partial differential equation:

$$(r+p\lambda_3)G = \alpha XG_X + \frac{\sigma^2}{2}X^2G_{XX} - \lambda_3 p(1-p)G_p + p\lambda_3 H\left(\frac{X}{1+\varphi}\right).$$
 (35)

Exercise will be triggered when cash flows rise to an upper trigger, which is itself a function of the firm's beliefs. Let $\bar{X}(p)$ denote the exercise trigger function. Conjecture that $\bar{X}(p) < (1 + \varphi)X^*$ for all $p \in (0, \pi_0]$.¹⁵ Equation (35) is solved subject to the following value-matching and smooth-pasting conditions:

$$G(\bar{X}(p), p) = \left[(1-p)\frac{1}{r-\alpha} + p\frac{1 + \frac{\lambda_3}{(r-\alpha)(1+\varphi)}}{r-\alpha+\lambda_3} \right] \bar{X}(p) - I, \quad (36)$$

$$G_{X}(\bar{X}(p), p) = \left[(1-p)\frac{1}{r-\alpha} + p\frac{1 + \frac{\lambda_{3}}{(r-\alpha)(1+\varphi)}}{r-\alpha+\lambda_{3}} \right],$$
 (37)

$$G_p(\bar{X}(p), p) = \left[-\frac{1}{r-\alpha} + \frac{1 + \frac{\lambda_3}{(r-\alpha)(1+\varphi)}}{r-\alpha+\lambda_3} \right] \bar{X}(p).$$
(38)

The value-matching condition (36) captures the fact that at the time of investment the value of the investment option equals the expected payoff from immediate investment $S(\bar{X}(p), p) - I$, while the smooth-pasting conditions (37)

¹⁵This conjecture is verified in the Appendix.

and (38) guarantee that the trigger function is chosen optimally. In the Internet Appendix we use the results in Peskir and Shiryaev (2006) to show that the boundary $\bar{X}(p)$ indeed satisfies the smooth-pasting conditions (37) and (38).¹⁶

Evaluating (35) at X(p), plugging in the boundary conditions (36) to (38), and simplifying, we get the following equation for the optimal trigger function $\overline{X}(p)$:¹⁷

$$\bar{X}(p) = p\lambda_3 \left[\left(\frac{\bar{X}(p)}{(1+\varphi)X^*} \right)^{\beta} \left(\frac{X^*}{r-\alpha} - I \right) - \left(\frac{\bar{X}(p)}{(1+\varphi)(r-\alpha)} - I \right) \right]$$

$$+ rI + \frac{\sigma^2}{2} \bar{X}(p)^2 G_{XX}(\bar{X}(p), p),$$
(39)

where $G_{XX}(\bar{X}(p), p)$ is the limit of $G_{XX}(X, p)$ from the left, $\lim_{X \to \bar{X}(p)^-} G_{XX}(X, p)$. In the Internet Appendix we demonstrate that there does indeed exist a solution $\bar{X}(p)$ to the fixed-point problem in equation (39) for all $p \in (0, \pi_0]$.

Equation (39) provides intuition on the additive forces of the option to wait and the option to learn. If the firm chooses to invest now over waiting an instant, it receives the benefits of the current cash flow, which is the term on the left-hand side of (39). If the firm waits a moment, it faces a small chance of the shock reverting, in which case the benefits of not investing are equal to $(\frac{\bar{X}(p)}{(1+\varphi)X^*})^{\beta}(\frac{X^*}{r-\alpha} - I) - (\frac{\bar{X}(p)}{(1+\varphi)(r-\alpha)} - I)$. Thus, the first term on the righthand side of (39) corresponds to the value of waiting to learn more about the nature of the past shock. The second term on the right-hand side of (39), rI, is simply the savings from delaying the investment cost by an instant. Finally, the third term is a new convexity term, $\frac{\sigma^2}{2}\bar{X}(p)^2G_{XX}(\bar{X}(p), p)$, which represents the traditional option to wait for future realizations of Brownian uncertainty.¹⁸ In the sense that the right-hand side of (39) is the sum of the three components, Brownian uncertainty is additive to Bayesian uncertainty. In other words, the addition of Brownian uncertainty to the model increases the investment trigger by a further component due to the option value of waiting for the evolution of Brownian uncertainty over future cash flows.

We find that the trigger function X(p) is increasing in p in all of our numerical results. The intuition for this monotonicity result is as follows. An increase in p makes the option to learn more valuable, in that waiting an instant longer provides insurance against a greater likelihood of reversal. We clearly see this

¹⁶An Internet Appendix for this article is available online in the "Supplements and Datasets" section at http://www.afajof.org/supplements.asp.

¹⁷In general, $\bar{X}(p)$ and G(X, p) can be computed numerically. In the Internet Appendix we also derive the closed-form solution for the special case of $\sigma = 0$.

¹⁸In order to ensure optimality of the exercise trigger, $\bar{X}(p)$, G(X, p) must be convex at $\bar{X}(p)$. To see this, note that for a given p, G(X, p) > h(p)X - I for all $X < \bar{X}(p)$, where $h(p) = [(1 - p)\frac{1}{r-\alpha} + p\frac{1 + \frac{\lambda_3}{(r-\alpha)(1+\phi)}}{r-\alpha+\lambda_3}]$. From the value-matching condition, at $\bar{X}(p)$, $G(\bar{X}(p), p) = h(p)\bar{X}(p) - I$, and from the smooth-pasting condition, $G_X(\bar{X}(p), p) = h(p)$. Thus, at $\bar{X}(p)$ it must be the case that $G_{XX}(\bar{X}(p), p) > p$

0.



Figure 2. Investment trigger functions for different values of the Brownian volatility parameter. The graph plots the investment trigger function $\bar{X}(p)$ for different values of the Brownian volatility parameter σ . The bottom curve corresponds to the case of pure Bayesian uncertainty ($\sigma = 0$). The middle and top curves correspond to the case of both Bayesian and Brownian uncertainties ($\sigma = 0.05$ and $\sigma = 0.10$, respectively). As a result, the change of the trigger along each curve is due to the impact of Bayesian uncertainty, while the upward shift of the whole trigger function is due to the impact of Brownian uncertainty. The parameter values are r = 0.04, $\alpha = 0.02$, $\varphi = 0.2$, $\lambda_1 = 0.5$, $\lambda_2 = 1$, $\lambda_3 = 2$, and I = 25.

effect for the case of $\sigma = 0$, which is solved in closed form in the Internet Appendix, in which there is only the option to learn at work. Thus, in order for $\bar{X}(p)$ to be decreasing over some range, it would have to be the case that the option to wait due to Brownian uncertainty decreases substantially as p increases, so much so that it counteracts the increase in the option to learn. However, our numerical results show that the option to wait (the convexity term) is relatively independent of p. This relative independence between the option to learn and the option to wait is not surprising, as in our model learning and Brownian innovations are separable.

The quantitative effects of the addition of Brownian uncertainty to the model are illustrated in Figure 2, which shows the investment trigger function $\bar{X}(p)$ for different values of σ .¹⁹ Importantly, while σ is constant, p is falling over

¹⁹To compute the trigger functions we used a variation of the least squares method developed by Longstaff and Schwartz (2003). The procedure is outlined in the Internet Appendix.

time, so that the Bayesian effect is dynamic and declining over time as learning accumulates. Both Bayesian and Brownian uncertainties lead to a significant increase in the investment trigger. If there were no Brownian or Bayesian uncertainty and all shocks were permanent, the trigger would equal one. Now consider the case in which we add only Bayesian uncertainty, corresponding to the bottom curve where $\sigma = 0$. For the case of p = 2/3, which equals the value of π_0 for our parameter specification, we find that the impact of pure Bayesian uncertainty increases the investment threshold by 12.6%, which is labeled the "Bayesian effect" in Figure 2. Now consider the additional impact of Brownian uncertainty. The middle and top curves correspond to the cases of $\sigma = 0.05$ and $\sigma = 0.10$, respectively. The addition of Brownian uncertainty leads to an additional increase in the threshold by 5.6% for the case of $\sigma = 0.05$ and by 20% for $\sigma = 0.10$, which are labeled the "Brownian effects" in Figure 2.²⁰ As discussed above, the addition of Brownian uncertainty does not undo the effect that Bayesian learning has on the optimal exercise rule. Indeed, the shape of the trigger function X(p) does not change much as the Brownian volatility parameter σ increases. For any σ , X(p) is an increasing and concave function of p.

C. Optimal Investment Prior to the Arrival of a Shock

To complete the solution, consider the value of the investment option before the shock arrives. In this case, the only underlying uncertainty concerns the future shocks to X(t). Specifically, X(t) evolves according to

$$dX(t) = \alpha X(t) dt + \sigma X(t) dB(t) + \varphi X(t) (dM_1(t) + dM_2(t)), \tag{40}$$

where $M_1(t)$ and $M_2(t)$ are arrival processes of permanent and temporary shocks, respectively.

Denote the option value by F(X). Prior to the investment, F(X) solves²¹

$$(r + \lambda_1 + \lambda_2)F(X) = \alpha XF'(X) + \frac{1}{2}\sigma^2 X^2 F''(X) + (\lambda_1 + \lambda_2)G(X(1 + \varphi), \pi_0), \quad (41)$$

where G(X, p) is the value of the investment option while the shock persists. Because the option has only one state variable, the optimal trigger will be a fixed value, \hat{X} .

Since in this region Bayesian learning plays no role, we relegate the derivation of \hat{X} to the Appendix. The investment trigger prior to the arrival of the shock is the unique solution to

$$\hat{X} = \frac{\gamma_1}{\gamma_1 - 1} \frac{(r - \alpha + \lambda_1 + \lambda_2)rI}{r + \lambda_1 + \lambda_2} + \tilde{A}\hat{X}^{\gamma_2}, \tag{42}$$

 20 Note that this implies that the total instantaneous volatility of the cash flow process is higher than σ , since it also includes volatility due to the jumps.

²¹As before, equation (40) describes the dynamics of X(t) under (\mathcal{F}_t) . However, in terms of valuation, it is immaterial, because $\mathbb{E}[dM_1(t) + dM_2(t) | \mathcal{F}_t] = \mathbb{E}[dM_1(t) + dM_2(t) | \mathcal{G}_t] = (\lambda_1 + \lambda_2)dt$.

1968

where \tilde{A} is defined in the Appendix, and where γ_1 and γ_2 are the positive and negative roots of the fundamental quadratic equation $\frac{1}{2}\sigma^2\gamma(\gamma-1) + \alpha\gamma = r + \lambda_1 + \lambda_2$. The corresponding value of the investment option F(X) is shown in the Appendix.

We can now fully summarize the optimal investment strategy in the following proposition:

PROPOSITION 1: The optimal investment strategy for the model with both Bayesian and Brownian uncertainty is:

- 1. If the shock does not occur until X(t) reaches \hat{X} , then it is optimal for the firm to invest when $X(t) = \hat{X}$;
- 2. If the shock occurs before X(t) reaches \hat{X} , and if at the time of the shock X(t) is above $\bar{X}(\pi_0)/(1+\varphi)$, then it is optimal for the firm to invest immediately after the shock;
- 3. If the shock occurs before X(t) reaches \hat{X} , and if at the time of the shock X(t) is below $\bar{X}(\pi_0)/(1+\varphi)$, then it is optimal for the firm to invest at the first time when $X(t) = \bar{X}(p(t))$;
- 4. If the shock occurs before X(t) reaches \hat{X} , and if at the time of the shock X(t) is below $\bar{X}(\pi_0)/(1+\varphi)$ but it reverts back before X(t) reaches $\bar{X}(p(t))$, then it is optimal to invest at the first time when X(t) reaches X^* after the shock reverts.

D. Discussion

Proposition 1 demonstrates that there are four different scenarios for the firm's investment timing. Under the first scenario, the cash flow process X(t) increases up to \hat{X} before the shock occurs. There is no learning prior to the arrival of the shock, so the investment trigger is constant. One simulated sample path that satisfies this scenario is shown in the upper left corner of Figure 3.

Under the other three possible scenarios, the shock occurs before the cash flow process reaches \hat{X} , so the firm does not invest prior to the arrival of the shock. After the shock arrives, the investment trigger is a function of the firm's beliefs about the past shock as given by equation (39). When the firm observes that the shock does not revert back, it lowers its assessment of the probability that the past shock was temporary. Since immediate investment is more attractive when the past shock is permanent, the investment trigger $\bar{X}(p(t))$ decreases over time.

Under the second scenario, the value of the cash flow process immediately after the shock overshoots $\bar{X}(\pi_0)$. In this case, the firm invests immediately after the shock arrives. A sample path satisfying this scenario is shown in the upper right corner of Figure 3.

Under the third scenario, the cash flow process X(t) reaches the investment trigger $\overline{X}(p(t))$ after a period of learning but before the shock reverts. A simulated sample path illustrating this scenario is shown in the lower left corner



Figure 3. Simulations of the firm's investment strategies. The figure shows the simulated paths of the cash flow process X(t) (plotted in solid lines) and the corresponding investment triggers (plotted in dashed lines) for the four different scenarios specified in Proposition 1. The optimal exercise strategy is to invest at the first time when X(t) reaches the investment trigger for the first time. The parameter values are r = 0.04, $\alpha = 0.02$, $\sigma = 0.10$, $\varphi = 0.2$, $\lambda_1 = 0.5$, $\lambda_2 = 1$, $\lambda_3 = 2$, I = 25, and X(0) = 1.

of Figure 3. This graph illustrates several interesting properties of investment in our Bayesian setting such as the dependence on the timing of past shocks and the sluggish response of investment to shocks. We discuss these and other implications in more detail in the next section.

Finally, under the fourth scenario, the past temporary shock reverts back before the firm invests. If this happens, the problem becomes standard. After the reversal, the investment trigger is constant at the level X^* . The firm invests at the first time when the cash flow process reaches X^* . A simulated sample path that describes the fourth scenario is shown in the lower right corner of Figure 3.

III. Model Implications

In this section we analyze important implications of the model.

A. Bayesian Uncertainty and the Option to Learn

In traditional real options models focusing on uncertainty about future shocks only, investment takes place when cash flows rise to a constant trigger. This is due to the fact that uncertainty is entirely driven by the Brownian component, which is constant over time. This trigger embodies the notion of the traditional option to wait as investment is triggered when the real option is sufficiently in the money.

When we add Bayesian uncertainty over past shocks, the investment trigger is no longer constant, but rather is a function of the firm's beliefs about the nature of past shocks. More specifically, the trigger is high soon after the arrival of a shock, but then goes down as time goes by. Thus, in the Bayesian framework, uncertainty is no longer constant, but instead declines over time as the firm becomes increasingly more confident in past shock's permanence. Hence, in addition to the standard option to wait for future shocks, we now also have an option to learn more about the nature of past shocks. We see in equation (39) that the investment trigger exceeds the Marshallian trigger of rIdue to two forces: the option to learn (waiting until a moment later to see if a past shock reverses) and the option to wait (and benefit from cash flows rising over the next moment). As demonstrated in Figure 2, these forces are additive and appear to be of similar magnitudes. Ignoring either option would lead to a similar investment timing error.

An important implication of Bayesian uncertainty over past shocks is the sluggishness in responding to positive cash flow shocks. In traditional real options models, in which uncertainty is about future shocks only, investment responds immediately to past cash flow shocks. In contrast, uncertainty about past shocks rationalizes the sluggish response of investment to shocks. Specifically, the firm might find it optimal to wait an additional period to see if the past positive shock is indeed permanent. Given this, the firm is now less likely to invest soon after the arrival of a positive shock, when the option to learn is most valuable. As time goes by, however, uncertainty is reduced, and the firm is willing to invest at a lower level of the cash flow process.²²

The slow response of investment to cash flow shocks generated by the option to learn is consistent with the slow response of economic variables such as investment, labor demand, and prices to shocks (e.g., Caballero and Engel (2003)). Another stylized fact that is consistent with Bayesian learning about past shocks is that the response of investment to shocks is time-dependent. In particular, aggregate investment is shown to be more responsive to shocks if the period preceding expansion is longer (Caballero, Engel, and Haltiwanger

 $^{^{22}}$ See Bloom (2009) for evidence that uncertainty shocks such as the 1973 oil price shock have large aggregate effects due to more "cautious" behavior by economic agents.

(1995), Bachmann, Caballero, and Engel (2008)). Our argument appears to be consistent with this evidence since firms are likely to believe that positive shocks following the period of expansion are more likely to be permanent, so the option to learn is less valuable in this case. The result that uncertainty about the identity of the shock can lead to a sluggish response of investment to shocks is very general and can hold in different environments with shocks of different types. For example, Moore and Schaller (2002) show in simulations that uncertainty about the persistence of interest rate shocks can lead to sluggishness of investment in the context of the neoclassical q theory.

The magnitude of the slow response of investment to cash flow shocks depends on two key parameters: the size of the shock φ and the speed of learning λ_3 . Papers that calibrate models of investment decisions usually assume that temporary shocks are relatively persistent. For example, Thomas (2002) and Gourio and Kashyap (2007) assume that a productivity shock follows an AR(1) process with coefficient 0.9225. In the context of our model this would mean that learning about the shock's type would be rather slow. This suggests that incorporating uncertainty about whether shocks are permanent or temporary into models of aggregate investment can have important quantitative effects.

B. Investment in the Face of Stable or Declining Cash Flows

Another interesting property of our model is that investment may occur in the face of stable or even decreasing cash flows. In standard real options models, investment will be triggered when shocks increase the cash flow process up to a constant upper threshold. As a result, the firm can rationally exercise the investment option only at times of increasing cash flows. In contrast, when the traditional real options model is extended to include Bayesian uncertainty over the nature of past shocks, the firm may invest at times of stable or even declining cash flows simply because it becomes sufficiently confident in the permanence of the past shocks.

As an illustration of this property, consider a particular sample path of cash flows and the investment response presented in Figure 4. We see that investment is made at a time of steady cash flows. At the moment of investment, it was the reduction in Bayesian uncertainty rather than the increase in current cash flows that triggered investment. This is made possible by the monotonically decreasing trigger function, $\bar{X}(p(t))$.

Empirically, we certainly see examples of firms (and industries) investing in markets where cash flows are stable, or even declining. As detailed in Grenadier (1996), during the late 1970s and early 1980s several U.S. cities saw explosive growth in office building development in the face of rapidly increasing office vacancy rates. Consider the cases of the Denver and Houston office markets. Over the 30-year period from 1960 to 1990, over half of all office construction was completed in a 4-year interval: 1982 to 1985. Since the average length of time between the initiation and completion of office space construction is 2.5 years, this investment was likely initiated over the 1979 to 1983 period. Throughout this period, office vacancies in these two cities were around 30%,



Figure 4. Investment in the face of stable cash flows: Illustration. The figure shows the simulated path of the cash flow process X(t) (plotted in the solid line) and the corresponding investment trigger (plotted in the dashed line). The optimal exercise strategy is to invest at the first time when X(t) reaches the investment trigger for the first time. The parameter values are r = 0.04, $\alpha = 0.02$, $\sigma = 0.10$, $\varphi = 0.2$, $\lambda_1 = 0.5$, $\lambda_2 = 1$, $\lambda_3 = 2$, I = 25, and X(0) = 1.

considerably above previous levels. Notably, however, these two cities (as well as most of the cities experiencing unprecedented office growth during this period) were oil-patch cities where developers likely concluded (incorrectly) that high oil prices in the late 1970s and early 1980s would last indefinitely (see the quote provided in the introduction). Our model explains this period of rapid construction in times of declining cash flows by suggesting that the firms became very sure that high oil prices, and hence high rents, in these oil cities would last forever.

Figure 4 also makes evident that accounting for uncertainty about past shocks may lead to a failure of the "record-setting news principle," which states that investment occurs only at instants in which the value of the cash flow process reaches a new peak.²³ In Figure 4 we see that investment occurs at a cash flow level that is more than 5% lower than its previous maximum. Thus,

 $^{^{23}}$ The record-setting news principle, a term credited to Maxwell Stinchcombe, is usually formulated in terms of asset values, rather than cash flows (see, for example, Boyarchenko (2004)). Because the model is defined in cash flows, we focus on the cash flow analogue of the record-setting

even when current cash flows are not at their all-time maximum, the decline in uncertainty about past shocks can trigger investment.

C. Investment Strategy and the Timing Structure of the Project's Cash Flows

In the standard real options literature, there is a simple equivalence between options that pay off a stream of cash flows and those that pay off an identical lump sum value. Thus, the timing of cash flows, per se, does not matter for the firm's investment strategy.²⁴ Ultimately, the firm chooses a trigger such that the value of future cash flows at exercise provides the needed option premium to induce investment. For example, suppose that the project's cash flows are driven solely by Brownian uncertainty about future shocks, and the maturity structure of these cash flows is relatively short. Suppose that for this option on a short-term project the firm chooses to exercise when the value of the future cash flows exceeds the investment cost by 50%, a premium consistent with the value of the option to wait. Then it is also true that the optimal exercise policy for an option on a project with a long-term cash flow structure but the same other characteristics will be one where investment is triggered when the value of future cash flows exceeds the investment cost by 50%. This equivalence relation is previously shown by Dixit and Pindyck (1994). In Chapter 5, these authors consider the optimal exercise rule for options that pay off with a lump sum value of V(t). Then, in Chapter 6, they perform an analogous analysis for options that pay off with a perpetual cash flow of P(t), with present value identical to the lump sum value V(t). They show that the optimal exercise rules are identical.

However, in the context of valuations that are driven by uncertainty over past shocks, the timing of cash flows can be quite important in determining the investment strategy. The greater is the front-loadedness of the project's cash flows, the less important is the assessment of the relative likelihood that a shock is temporary or permanent. As a result, the option to learn is less valuable, so the firm invests at a lower trigger. Thus, while the firm might invest in a short-term project when the value of the future cash flows exceeds the investment cost by 50%, the firm might only invest in a longer-term project if the value of the future cash flows exceeds the investment cost by 100%. This is because the option to learn is fundamentally linked to the cash flow structure of the project.

In this section we consider a simple parametrization of the front-loadedness of the option payoff, ranging from payoffs that are equivalent to a one-time lump sum payment to payoffs that are equivalent to perpetual cash flows.

news principle. The two definitions are equivalent if the cash flow process is the only state variable and the asset value is a monotone function of the current value of the cash flow process.

²⁴When we refer to the standard real options literature, we refer to models in which there is an initial lump-sum cost followed by a single lump-sum cash flow or a sequence of positive cash flows, and in which investment is irreversible. However, in a broader class of real options, in which costs could be paid over time and partial reversibility is possible, the timing of the project's cash flows could matter for investment.

Consider the simplest case studied in Section I.C, but with one alteration. Assume now that if an option is exercised at time τ , it provides a stream of payments $(1 + \frac{w}{r})e^{-w(t-\tau)}X(t), t \geq \tau$. The parameter $w \in [0, +\infty)$ captures the degree of front-loadedness of the project's cash flows. Projects with low values of w are relatively back-loaded: much of their cash flows are generated long after the exercise time. High values of w mean that the project is relatively front-loaded, with most cash flows coming close to the exercise time. The particular parametrization was chosen so as to make the present value of cash flows from the immediate exercise of the project in the no-shock case independent of w: $\int_{\tau}^{\infty} X(t)(1 + \frac{w}{r})e^{-w(t-\tau)}e^{-r(t-\tau)}dt = \frac{X(\tau)}{r}$. Of course, other reasonable parameterizations are possible. This specification of cash flows captures two cases widely used in the real options literature. First, when w = 0, the model reduces to the one studied in Section I.C. In this case, the project pays a perpetual flow of X(t) upon exercise. Second, if $w \to \infty$, payments from the project converge to a one-time lumpy payment of $\frac{X(\tau)}{r}$ at the time of exercise τ .

Similar to (13), to ensure that the project has a potentially positive net present value and that there is positive value to learning, we make the assumption that X_0 satisfies

$$\frac{rI}{1+\varphi} < X_0 < \frac{r+\lambda_3\pi_0}{1+\varphi+\frac{\lambda_3}{r}\pi_0+\frac{\lambda_3}{r}\pi_0\frac{w\varphi}{(r+\lambda_3+w)}}I.$$
(43)

Compared to (13), (43) puts the same lower bound and a more restrictive upper bound. As previously, these bounds guarantee that the solution to the investment timing problem is nontrivial.

As in Section I.C, let the value of the option while the shock persists be denoted by G(p). Over the range of p over which the option is not exercised, the standard argument implies

$$(r+p\lambda_3)G(p) = -G'(p)\lambda_3p(1-p).$$
(44)

This equation has the general solution

$$G(p) = C_w (1-p) \left(\frac{1}{p} - 1\right)^{\frac{r}{\lambda_3}},$$
(45)

where C_w is a constant. Note that this general solution coincides with the one obtained in Section I.C. However, because of the more general cash flow timing assumption, the boundary conditions are now different.

Because the payoff from the project if the current shock is temporary is affected by the parameter w, the value-matching condition at the exercise trigger \bar{p}_w is now

$$G(\bar{p}_w) = (1 - \bar{p}_w) \frac{X_0(1 + \varphi)}{r} + \bar{p}_w \frac{X_0 \left(1 + \varphi + \frac{\lambda_3}{r} + \frac{w(1 + \varphi)}{r}\right)}{r + \lambda_3 + w} - I.$$
(46)

Note that although the present value of the firm's cash flows in the case of the permanent shock does not depend on w, in the case of the temporary shock it depends on w positively. Intuitively, a more front-loaded project allows the firm to capture more of the temporary high cash flows than a more back-loaded project.

As in Section I.C, the exercise trigger is chosen to maximize the value of the option (or equivalently, to satisfy the smooth-pasting condition), giving the optimal trigger value

$$\bar{p}_{w} = \frac{X_{0}(1+\varphi) - rI}{\lambda_{3} \left(I - \frac{X_{0}}{r} - \frac{X_{0}w\varphi}{r(r+\lambda_{3}+w)}\right)}.$$
(47)

Note that when w = 0, then $\bar{p}_w = \bar{p}$, which is the trigger we obtained in Section I.C. As in Section I.C, it is straightforward to show that the option will never be exercised prior to the arrival of the shock.

Consider how the parameter of front-loadedness affects the trigger value:

$$\frac{\partial \bar{p}_w}{\partial w} = \frac{r\varphi(r+\lambda_3)}{\lambda_3} \frac{X_0(1+\varphi) - Ir}{[X_0(r+\lambda_3+w(1+\varphi)) - Ir(r+\lambda_3+w)]^2} X_0 > 0.$$
(48)

We therefore find that the greater the front-loadedness, the earlier the timing of investment.

Now consider how the parameter of front-loadedness affects the value at exercise. The option premium at the time of exercise is

$$\frac{G(\bar{p}_w)}{I} = -1 + \frac{1}{I} \left(1 - \frac{X_0(1+\varphi) - rI}{\lambda_3 \left(I - \frac{X_0}{r} - \frac{X_0 w\varphi}{r(r+\lambda_3 + w)}\right)} \right) \frac{X_0(1+\varphi)}{r} + \frac{1}{I} \frac{X_0(1+\varphi) - rI}{\lambda_3 \left(I - \frac{X_0}{r} - \frac{X_0 w\varphi}{r(r+\lambda_3 + w)}\right)} \frac{X_0(1+\varphi) + \frac{\lambda_3}{r} + \frac{w(1+\varphi)}{r}}{r+\lambda_3 + w}.$$
(49)

Differentiating with respect to w gives us

$$\frac{\partial \left(\frac{G(\bar{p}_w)}{I}\right)}{\partial w} = -\frac{X_0\varphi(X_0(1+\varphi)-rI)^2}{I((X_0-rI)(r+\lambda_3+w)+X_0\varphi w)^2} < 0.$$
(50)

Thus, the greater the front-loadedness of the project, the lower the option premium necessary to induce exercise.

Intuitively, if the project's cash flows are very front-loaded, then knowing the identity of the shock is not very important for the firm. In this case the firm gets a large part of the project cash flows very soon after the exercise date. As

a result, the learning option is not very valuable, so the firm invests at a lower premium. On the other hand, if the project's cash flows are very back-loaded, then it is very important for the firm to be sure that past shocks are permanent. As a result, the learning option is very valuable and the firm invests at a higher premium.

IV. A Model with Multiple Shocks

In this section, we generalize the model of Section II for the case of multiple shocks. Specifically, we assume that at any time t there can be any number $n(t) = 0, 1, \ldots, \bar{N}$ of shocks outstanding. If $n(t) < \bar{N}$, the evolution of X(t) under (\mathcal{F}_t) is given by (1), where the probability of observing a reversal over the next instant depends on the number k(t) of outstanding temporary shocks. In particular, if n(t) = 0, there are no outstanding shocks, so the intensity of the reversal process N(t) is zero. If $n(t) = \bar{N}$, no new shocks can occur, so the evolution of X(t) under (\mathcal{F}_t) is given by (1) with the intensities of the shock processes $M_1(t)$ and $M_2(t)$ being zero. When the maximum number of outstanding shocks \bar{N} increases to infinity, the model approaches the case of an unlimited number of shocks, each of which can be permanent or temporary.

Under the information set of the firm, at each time *t* the state can be described by a triple (X(t), n(t), p(t)), where X(t) is the value of the cash flow process, n(t)is the number of outstanding shocks, and $p(t) = (p_0(t) \ p_1(t) \dots \ p_{n(t)}(t))'$ is the vector of the firm's beliefs. The Bayesian learning process p(t) is derived in Section I.B.

As before, we begin by calculating the present value of the project's underlying cash flows, defined by (3). It is equal to

$$S(X(t), n(t), p(t)) = a_0^{n(t)} X(t) + \sum_{i=1}^{n(t)} p_i(t) \left(a_i^{n(t)} - a_0^{n(t)} \right) X(t),$$
(51)

where the constants $a_0^{n(t)}, a_1^{n(t)}, \ldots$ are defined in the Internet Appendix. Intuitively, if the firm knew for sure that there are exactly *k* outstanding temporary shocks out of n(t), then the present value of the project's cash flows would be linear in X(t) and equal to $a_k^{n(t)}X(t)$ for some constant $a_k^{n(t)}$. Taking the expectation with respect to the beliefs of the firm yields (51).

Let G(X, n, p) denote the value of the option to invest, where X and p are the current values of the cash flow and belief processes, respectively, and n is the current number of outstanding shocks. The optimal investment decision can be described by a trigger function $\overline{X}(n, p)$. In the Internet Appendix we provide the details of the determination of $\overline{X}(n, p)$ and G(X, n, p), which is done using the same argument as in Section II.B. In particular, we show that $\overline{X}(n, p)$ can

be expressed as

$$\begin{split} \bar{X}(n,p) &= \left(\lambda_3 \sum_{k=1}^n p_k k\right) \Biggl[G\Biggl(\frac{\bar{X}(n,p)}{1+\varphi}, n-1, \tilde{p}(p) \Biggr) \\ &+ I - S\left(\frac{\bar{X}(n,p)}{1+\varphi}, n-1, \tilde{p}(p) \Biggr) \Biggr] \\ &+ rI + \frac{\sigma^2}{2} \bar{X}(n,p)^2 G_{XX}(\bar{X}(n,p),n,p), \end{split}$$
(52)

where $\tilde{p}(p)$ is the vector of beliefs updated after the reversion, defined by (9).

Given the similarity between (52) and (39), it becomes clear how the multishock case generalizes the single-shock case. As before, the intuition behind the equation for the optimal trigger is the trade-off between investing now versus investing a moment later if all of the past shocks persist, where the value of waiting is explicitly the option to learn. If the firm invests now, it gets the benefit of the cash flow $\bar{X}(n, p)$ over the next instant. This is the term on the left-hand side of the equal sign. If the firm waits a moment, it faces a small chance of one of the shocks reversing, in which case the expected value gained by waiting is equal to the first term on the right-hand side. Specifically, $G(\frac{\bar{X}(n,p)}{1+\varphi}, n-1, \tilde{p}(p)) +$ $I - S(\frac{\bar{X}(n,p)}{1+\varphi}, n-1, \tilde{p}(p))$ is the value gained by not investing should an existing shock reverse over the next instant, and $\lambda_3 \sum_{k=1}^{n} p_k k$ is the likelihood of this occurring. The second term on the right-hand side of the equal sign, rI, is the savings from delaying the investment cost by an instant. The third term is the convexity term, which represents the value of waiting due to Brownian

uncertainty. At the optimal trigger, X(n, p), the two sides of equation (52) are exactly equal and the firm is indifferent between investing now and waiting a moment. We can now summarize the solution to the optimal investment timing prob-

We can now summarize the solution to the optimal investment timing problem in this multi-shock setting.

PROPOSITION 2: The optimal investment strategy for the model outlined in this section is to invest when X(t) exceeds $\overline{X}(n(t), p(t))$ for the first time.

To provide intuition for the multi-shock extension, Figure 5 plots a simulated path of the cash flow process X(t) along with the corresponding investment trigger $\bar{X}(t) \equiv \bar{X}(n(t), p(t))$. Due to computational complexity, for this numerical solution we consider the case of $\bar{N} = 2$ and $\alpha = \sigma = 0$. The numerical procedure is outlined in the Internet Appendix. In this example, each shock increases the value of the cash flows by 10% ($\varphi = 0.1$). A new permanent shock occurs, on average, every 3 years ($\lambda_1 = 0.33$). A new temporary shock occurs, on average, every year ($\lambda_2 = 1$) and takes 4 months to revert ($\lambda_3 = 3$). The dynamics of the cash flow process are straightforward. In the 5-year period there are four arrivals of new shocks, corresponding to upward jumps in X(t). Two out of four shocks revert, corresponding to downward jumps in X(t). The dynamics of the investment boundary are more interesting. Prior to the arrival of



Figure 5. Simulation of the firm's investment strategy. The top graph shows the simulated sample path of the cash flow process, X(t), for the model of Section IV. Each of the four upward steps corresponds to the arrival of a new shock. Similarly, each of the two downward steps corresponds to the reversal of an outstanding shock. The bottom graph adds the investment trigger $\bar{X}(t) = \bar{X}(n(t), p(t))$ to the simulated sample path from the top graph. The optimal investment strategy is to invest at the first time when X(t) crosses the investment trigger. The parameter values are r = 0.04, $\alpha = 0$, $\sigma = 0$, $\varphi = 0.1$, $\lambda_1 = 0.33$, $\lambda_2 = 1$, $\lambda_3 = 3$, I = 25, $\bar{N} = 2$, and X(0) = 0.92.

the first shock, there is no uncertainty about past shocks. Since there is no Bayesian updating at this time, the trigger is constant at $\bar{X}(0, 1) = rI$. After the first shock arrives, the trigger jumps up to $\bar{X}(1, 1 - \pi_0, \pi_0)$ as the firm becomes unsure if the outstanding shock is permanent or temporary. As time goes by and the cash flow process does not revert back, the firm believes that the first shock is more likely to be permanent. As a result, the investment trigger decreases. This intuition underlies the whole dynamics of the trigger in Figure 5. An arrival of a new shock leads to an upward jump in the trigger due to an increase in uncertainty. When the cash flow process is stable, the trigger goes down as the firm updates its beliefs about the past shocks. When an

outstanding shock reverses, the firm learns that one of the outstanding shocks is temporary. Hence, reversal leads to a downward jump in the trigger. The investment occurs when X(t) exceeds the dynamic investment trigger $\bar{X}(t)$ for the first time. In the example in Figure 5, this happens at t = 2.1, when the firm becomes sufficiently sure that the outstanding shock is permanent.

By allowing for multiple shocks, the model of this section illuminates the nature of the Bayesian investment problem beyond that of the simpler one-shock model. Most notably, the exercise trigger, $\bar{X}(n(t), p(t))$, is a multi-dimensional function, which demonstrates its dependence on all persisting past shocks. Thus, as illustrated in Figure 5, the investment trigger jumps by a discrete amount at all moments when either a new jump occurs or an existing shock reverses. In between such jumps, the trigger declines as uncertainty about past shocks is reduced.

V. Conclusion

This paper studies the optimal timing of investment in the presence of uncertainty about both future and past shocks. We augment the standard Brownian uncertainty driving traditional real options models with additional Bayesian uncertainty over distinguishing between the temporary and permanent nature of past cash flow shocks. As a result, the evolving uncertainty is no longer constant and is driven by Bayesian updating, or learning. We show that this gives rise to two real options. In addition to the "waiting" option as in traditional literature, the firm now has a "learning" option: the option to learn more about the nature of past shocks. We solve for the optimal investment rule and show that the implied investment behavior differs significantly from that predicted by standard real options models. In particular, investment often responds sluggishly to cash flow shocks, investment may occur at a time of stable or even decreasing cash flows, and the optimal investment timing depends critically on the timing structure of the project's cash flows.

Several further extensions of the model would prove interesting. First, the model could be extended to a multiple-firm industry equilibrium in the manner of Grenadier (2002). Of particular interest would be the fact that firms need to update their beliefs about not only their own past shocks, but also those of their competitors. Thus, while competition is known to erode the option to wait, the benefits from waiting to learn could mitigate such erosion. Second, the structure of the jump component could be made richer by permitting more than just either permanent or temporary shocks. For example, shocks could have a greater variety of (unobserved) types, such as *n* types with reversion intensities $\lambda_1^{reversion}, \lambda_2^{reversion}, \ldots, \lambda_n^{reversion}$.

Appendix: Proofs

Proof that in the simple case of Section I.C, the option will never be exercised prior to the arrival of the shock.

First, we demonstrate that the strategy of exercising prior to the arrival of the shock is dominated by the strategy of investing immediately upon the arrival of the shock. If the firm invests prior to the arrival of the shock, it pays the cost I and obtains the cash flow X_0 until the arrival of the shock. Upon the arrival of the shock, it obtains the additional cash flow φX_0 until the time that the shock reverses, if ever. If the firm instead waits until the shock arrives and then invests immediately, it gets no cash flow before the shock arrives, pays the cost I when the shock arrives, and obtains the cash flow φX_0 until the time that the shock reverses, if ever.

Let *NB* denote the net benefit of investing prior to the shock versus waiting until the shock arrives. Let τ denote the random arrival time of the shock. Then we can write *NB* as

$$NB = \mathbb{E}\left[\int_0^\tau X_0 e^{-rt} dt\right] + (\mathbb{E}[e^{-r\tau}]I - I) = \left(\frac{X_0}{r} - I\right)(1 - \mathbb{E}[e^{-r\tau}]), \quad (A1)$$

where NB is equal to the value of the cash flow X_0 until the arrival of the shock (the first term), plus the cost of paying I immediately rather than at the time of the shock's arrival (the second term).

The upper bound on X_0 in (13) implies

$$X_0 - rI < \frac{r + \lambda_3 \pi_0 - r - r\varphi - \lambda_3 \pi_0}{1 + \varphi + \frac{\lambda_3}{r} \pi_0}I = -\frac{r\varphi}{1 + \varphi + \frac{\lambda_3}{r} \pi_0}I < 0.$$
(A2)

Thus, NB < 0.

We have proven that it is always better to wait until the shock arrives and invest immediately at that time than to invest prior to the shock. In addition, once the shock arrives, it is always optimal to wait further until p(t) falls to the threshold \bar{p} . Thus, investing prior to the arrival of the shock is dominated by the strategy of waiting until the shock arrives and then waiting until p(t) falls to the threshold \bar{p} .

Proof of Section II.B conjecture that $\bar{X}(p) < (1+\varphi)X^*$.

From the proof of the existence of a boundary $\overline{X}(p)$ in the Internet Appendix, it follows that it is always optimal to exercise the option whenever $X > (1 + \varphi)X^*$. Therefore, for any $p \in (0, \pi_0]$, $\overline{X}(p)$ must be below $(1 + \varphi)X^*$.

Derivation of the investment trigger \hat{X} and the investment option value F(X) for the model with Brownian uncertainty in Section II.C.

If the investment option is exercised prior to the arrival of a shock, the firm receives

$$S(X) - I = \frac{X}{r - \alpha} + \frac{\varphi X}{r - \alpha + \lambda_1 + \lambda_2} \left(\frac{\lambda_1}{r - \alpha} + \frac{\lambda_2}{r - \alpha + \lambda_3} \right) - I.$$
(A3)

Conjecture that it is strictly optimal to wait when X(t) is below $\overline{X}(\pi_0)/(1 + \varphi)$, that is, $\hat{X} \geq \overline{X}(\pi_0)/(1 + \varphi)$. This is quite intuitive, as it implies that if it is not optimal to invest at an instant prior to the shock, then it would not be optimal to invest if the jump occurs and immediately reverts. The conjecture is verified later. It follows that the investment option value F(X) is equal to either $F_L(X)$ or $F_H(X)$, corresponding to the lower and the upper regions, respectively. By Itô's lemma, F(X) satisfies the following differential equations:

• in the region $X < \overline{X}(\pi_0)/(1+\varphi)$,

$$(r + \lambda_1 + \lambda_2)F_L(X) = \alpha X F'_L(X) + \frac{1}{2}\sigma^2 X^2 F''_L(X) + (\lambda_1 + \lambda_2)G((1 + \varphi)X, \pi_0);$$
(A4)

• in the region $\bar{X}(\pi_0)/(1+\varphi) < X < \hat{X}$,

$$(r + \lambda_1 + \lambda_2)F_H(X) = \alpha X F'_H(X) + \left(\lambda_1 \frac{1+\varphi}{r-\alpha} + \lambda_2 \frac{1+\varphi + \frac{\lambda_3}{r-\alpha}}{r-\alpha + \lambda_3}\right) X - (\lambda_1 + \lambda_2)I.$$
(A5)

Equations (A4) and (A5) differ due to the implied investment behavior at the moment of the arrival of a shock. In the lower region the arrival of a shock does not induce immediate investment, while in the higher region it does. Equations (A4) and (A5) are solved subject to the following boundary conditions:

$$F_H(\hat{X}) = \frac{\hat{X}}{r-\alpha} + \frac{\varphi \hat{X}}{r-\alpha+\lambda_1+\lambda_2} \left(\frac{\lambda_1}{r-\alpha} + \frac{\lambda_2}{r-\alpha+\lambda_3}\right) - I, \quad (A6)$$

$$F'_{H}(\hat{X}) = \frac{1}{r-\alpha} + \frac{\varphi}{r-\alpha+\lambda_{1}+\lambda_{2}} \left(\frac{\lambda_{1}}{r-\alpha} + \frac{\lambda_{2}}{r-\alpha+\lambda_{3}}\right), \tag{A7}$$

$$\lim_{X\uparrow \bar{X}(\pi_0)/(1+\varphi)} F_L(X) = \lim_{X\downarrow \bar{X}(\pi_0)/(1+\varphi)} F_H(X),$$
(A8)

$$\lim_{X\uparrow\bar{X}(\pi_0)/(1+\varphi)} F'_L(X) = \lim_{X\downarrow\bar{X}(\pi_0)/(1+\varphi)} F'_H(X),\tag{A9}$$

$$\lim_{X \to 0} F_L(X) = 0. \tag{A10}$$

As before, the value-matching condition (A6) imposes equality at the exercise point between the value of the option and the net present value of the project's cash flows, while the smooth-pasting condition (A7) ensures that the exercise point is chosen optimally. Conditions (A8) and (A9) guarantee that the value of the investment option is continuous and smooth. Finally, (A10) is a boundary condition that reflects the fact that X = 0 is the absorbing barrier for the cash flow process.

The general solutions to equations (A4) and (A5) are given by

$$F_L(X) = C_1 X^{\gamma_1} + C_2 X^{\gamma_2} + \frac{2(\lambda_1 + \lambda_2)}{\sigma^2(\gamma_1 - \gamma_2)} (\Gamma_2(X) - \Gamma_1(X)),$$
(A11)

$$\begin{aligned} F_H(X) &= A_1 X^{\gamma_1} + A_2 X^{\gamma_2} \\ &+ \frac{(\lambda_1 + \lambda_2)((1+\varphi)(r-\alpha) + \lambda_3) + \lambda_1 \lambda_3 \varphi}{(r-\alpha + \lambda_1 + \lambda_2)(r-\alpha)(r-\alpha + \lambda_3)} X \\ &- \frac{\lambda_1 + \lambda_2}{r + \lambda_1 + \lambda_2} I, \end{aligned}$$
(A12)

where

$$\Gamma_i(X) = X^{\gamma_i} \int_0^X \frac{G((1+\varphi)x, \pi_0)}{x^{\gamma_i+1}} dx, i \in \{1, 2\}.$$
 (A13)

We have five boundary conditions (A6) to (A10) to determine the four unknown constants $(A_1, A_2, C_1, \text{ and } C_2)$ and the investment trigger \hat{X} . Consider condition (A10). It implies

$$\lim_{X \to 0} C_2 X^{\gamma_2} + \frac{2(\lambda_1 + \lambda_2)}{\sigma^2(\gamma_1 - \gamma_2)} \left(\lim_{X \to 0} \Gamma_2(X) - \lim_{X \to 0} \Gamma_1(X) \right) = 0.$$
(A14)

Because $\gamma_1>0, \lim_{X\to 0}\Gamma_1(X)=0$. Also, by l'Hôpital's rule,

$$\begin{split} \lim_{X \to 0} \Gamma_2(X) &= \lim_{X \to 0} X^{\gamma_2} \int_0^X \frac{G((1+\varphi)x, \pi_0)}{x^{\gamma_2 + 1}} dx \\ &= \lim_{X \to 0} \frac{G((1+\varphi)X, \pi_0)X^{-\gamma_2 - 1}}{-\gamma_2 X^{-\gamma_2 - 1}} \\ &= -\frac{1}{\gamma_2} \lim_{X \to 0} G((1+\varphi)X, \pi_0) = 0. \end{split}$$
(A15)

Therefore, condition (A10) implies $C_2 = 0$. Using (A11) and (A12) and combining equations (A8) and (A9), we get

$$A_{2} = \left(\frac{\bar{X}(\pi_{0})}{1+\varphi}\right)^{-\gamma_{2}} \times \left[\frac{1-\gamma_{1}}{\gamma_{1}-\gamma_{2}}\frac{(\lambda_{1}+\lambda_{2})((1+\varphi)(r-\alpha)+\lambda_{3})+\lambda_{1}\lambda_{3}\varphi}{(r-\alpha+\lambda_{1}+\lambda_{2})(r-\alpha)(r-\alpha+\lambda_{3})}\frac{\bar{X}(\pi_{0})}{1+\varphi} + \frac{\gamma_{1}}{\gamma_{1}-\gamma_{2}}\frac{\lambda_{1}+\lambda_{2}}{r+\lambda_{1}+\lambda_{2}}I + \frac{2(\lambda_{1}+\lambda_{2})}{\sigma^{2}(\gamma_{1}-\gamma_{2})}\Gamma_{2}\left(\frac{\bar{X}(\pi_{0})}{1+\varphi}\right)\right].$$
(A16)

Combining equations (A6) and (A7), we get

$$\hat{X} = \frac{\gamma_1}{\gamma_1 - 1} \frac{(r - \alpha + \lambda_1 + \lambda_2)rI}{r + \lambda_1 + \lambda_2} + \tilde{A}\hat{X}^{\gamma_2}, \tag{A17}$$

where $\tilde{A} = A_2(r - \alpha + \lambda_1 + \lambda_2)(\gamma_1 - \gamma_2)/(\gamma_1 - 1).$

1983

The corresponding value of the investment opportunity is

$$F(X) = \begin{cases} F(X) = \\ C_1 X^{\prime_1} + \frac{2(\lambda_1 + \lambda_2)}{\sigma^2(\gamma_1 - \gamma_2)} (\Gamma_2(X) - \Gamma_1(X)), & X \le \frac{\bar{X}(\pi_0)}{1 + \varphi} \\ A_1 X^{\prime_1} + A_2 X^{\prime_2} + \frac{(\lambda_1 + \lambda_2)((1 + \varphi)(r - \alpha) + \lambda_3) + \lambda_1 \lambda_3 \varphi}{(r - \alpha + \lambda_1 + \lambda_2)(r - \alpha)(r - \alpha + \lambda_3)} X & \frac{\bar{X}(\pi_0)}{1 + \varphi} \le X \le \hat{X} \\ - \frac{\lambda_1 + \lambda_2}{r + \lambda_1 + \lambda_2} I, \\ \frac{X}{r - \alpha} + \frac{\varphi X}{r - \alpha + \lambda_1 + \lambda_2} \left(\frac{\lambda_1}{r - \alpha} + \frac{\lambda_2}{r - \alpha + \lambda_3}\right) - I, & X \ge \hat{X}, \end{cases}$$
(A18)

where A_2 is given by (A16), and A_1 and C_1 satisfy

$$A_1 = \hat{X}^{-\gamma_1} \left(\frac{\hat{X}}{r - \alpha + \lambda_1 + \lambda_2} - \frac{rI}{r + \lambda_1 + \lambda_2} - A_2 \hat{X}^{\gamma_2} \right), \tag{A19}$$

$$C_{1} = A_{1} + A_{2} \left(\frac{\bar{X}(\pi_{0})}{1+\varphi} \right)^{\gamma_{2}-\gamma_{1}} + \left(\frac{\bar{X}(\pi_{0})}{1+\varphi} \right)^{-\gamma_{1}} \\ \times \left(\frac{(\lambda_{1}+\lambda_{2})((1+\varphi)(r-\alpha)+\lambda_{3})+\lambda_{1}\lambda_{3}\varphi}{(r-\alpha+\lambda_{1}+\lambda_{2})(r-\alpha)(r-\alpha+\lambda_{3})} \frac{\bar{X}(\pi_{0})}{1+\varphi} \right) \\ - \frac{\lambda_{1}+\lambda_{2}}{r+\lambda_{1}+\lambda_{2}} I \\ + \frac{2(\lambda_{1}+\lambda_{2})}{\sigma^{2}(\gamma_{1}-\gamma_{2})} \left(\Gamma_{1} \left(\frac{\bar{X}(\pi_{0})}{1+\varphi} \right) - \Gamma_{2} \left(\frac{\bar{X}(\pi_{0})}{1+\varphi} \right) \right) \right).$$
(A20)

We now confirm the conjecture that $\hat{X} \geq \bar{X}(\pi_0)/(1+\varphi)$. By contradiction, suppose that it is optimal to invest at some trigger X' below $\bar{X}(\pi_0)/(1+\varphi)$. Let us modify the problem in the following way. Suppose that the upward jump occurs immediately, that is, $\lambda_1 + \lambda_2 = +\infty$ with π_0 being unchanged. For each sample path, the project in the modified problem yields the same cash flows as the project in the original problem with the difference that the range of extra cash flows generated by the upward shock occurs earlier. Hence, investment in the modified problem occurs earlier than in the original problem. In particular, since it was optimal to invest at X' in the original problem the jump occurs immediately, the value of the investment option is $G(X(1+\varphi), \pi_0)$. However, it is strictly optimal to wait when $X(1+\varphi) < \bar{X}(\pi_0)$. Therefore, it cannot be optimal to invest at any such X'. This implies that, indeed, it is strictly optimal to wait for any X below $\bar{X}(\pi_0)/(1+\varphi)$.

REFERENCES

- Bachmann, Ruediger, Ricardo J. Caballero, and Eduardo M.R.A. Engel, 2008, Aggregate implications of lumpy investment: New evidence and a DSGE model, NBER Working paper 12336, University of Michigan.
- Bergemann, Dirk, and Ulrich Hege, 1998, Venture capital financing, moral hazard, and learning, Journal of Banking and Finance 22, 703–735.
- Bloom, Nicholas, 2009, The impact of uncertainty shocks, Econometrica 77, 623-685.
- Bolton, Patrick, and Christopher Harris, 1999, Strategic experimentation, *Econometrica* 67, 349–374.
- Boyarchenko, Svetlana, 2004, Irreversible decisions and record-setting news principle, American Economic Review 94, 557–568.
- Brennan, Michael J., and Eduardo S. Schwartz, 1985, Evaluating natural resource investments, Journal of Business 58, 135–157.
- Caballero, Ricardo J., and Eduardo M.R.A. Engel, 2003, Adjustment is much slower than you think, NBER Working paper 9898, Massachusetts Institute of Technology.
- Caballero, Ricardo J., Eduardo M.R.A. Engel, and John C. Haltiwanger, 1995, Plant-level adjustment and aggregate investment dynamics, *Brookings Papers on Economic Activity* 2, 1–54.
- Décamps, Jean-Paul, Thomas Mariotti, and Stephane Villeneuve, 2005, Investment timing under incomplete information, *Mathematics of Operations Research* 30, 472–500.
- Dixit, Avinash K., and Robert S. Pindyck, 1994, Investment under Uncertainty (Princeton University Press, Princeton, NJ).
- Dumas, Bernard, 1991, Super contact and related optimality conditions, Journal of Economic Dynamics and Control 15, 675–685.
- Gorbenko, Alexander S., and Ilya A. Strebulaev, 2010, Temporary versus permanent shocks: Explaining corporate financial policies, *Review of Financial Studies* 23, 2591–2647.
- Gourio, Francois, and Anil K. Kashyap, 2007, Investment spikes: New facts and a general equilibrium exploration, *Journal of Monetary Economics* 54, 1–22.
- Grenadier, Steven R., 1996, The strategic exercise of options: Development cascades and overbuilding in real estate markets, *Journal of Finance* 56, 1653–1679.
- Grenadier, Steven R., 2002, Option exercise games: An application to the equilibrium investment strategies of firms, *Review of Financial Studies* 15, 691–721.
- Grenadier, Steven R., and Neng Wang, 2005, Investment timing, agency, and information, Journal of Financial Economics 75, 493–533.
- Grenadier, Steven R., and Neng Wang, 2007, Investment under uncertainty and time-inconsistent preferences, *Journal of Financial Economics* 84, 2–39.
- Jarrow, Robert A., and Fan Yu, 2001, Counterparty risk and the pricing of defaultable securities, Journal of Finance 56, 1765–1799.
- Keller, Godfrey, and Sven Rady, 1999, Optimal experimentation in a changing environment, *Review* of *Economic Studies* 66, 475–507.
- Keller, Godfrey, Sven Rady, and Martin Cripps, 2005, Strategic experimentation with exponential bandits, *Econometrica* 73, 39–68.
- Klein, Manuel, 2007, Irreversible investment under incomplete information, Working paper, IN-SEAD.
- Klein, Manuel, 2009, Comment on "Investment timing under incomplete information", Mathematics of Operations Research 35, 249–254.
- Krylov, Nicolai V., 1980, Controlled Diffusion Processes (Springer, Berlin).
- Lambrecht, Bart M., and William Perraudin, 2003, Real options and preemption under incomplete information, Journal of Economic Dynamics and Control 27, 619–643.
- Longstaff, Francis A., and Eduardo S. Schwartz, 2003, Valuing American options by simulation: A simple least-squares approach, *Review of Financial Studies* 14, 113–147.
- Majd, Saman, and Robert S. Pindyck, 1987, Time to build, option value, and investment decisions, Journal of Financial Economics 18, 7–27.
- Malueg, David A., and Shunichi O. Tsutsui, 1997, Dynamic R&D competition with learning, Rand Journal of Economics 28, 751–772.

- McDonald, Robert, and David Siegel, 1986, The value of waiting to invest, *Quarterly Journal of Economics* 101, 707–728.
- Miao, Jianjun, and Neng Wang, 2007, Experimentation under uninsurable idiosyncratic risk: An application to entrepreneurial survival, Working paper, Boston University.
- Moore, Bartholomew, and Huntley Schaller, 2002, Persistent and transitory shocks, learning, and investment dynamics, *Journal of Money, Credit, and Banking* 34, 650–677.
- Moscarini, Giuseppe, and Lones Smith, 2001, The optimal level of experimentation, *Econometrica* 69, 1629–1644.
- Nishimura, Kiyohiko G., and Hiroyuki Ozaki, 2007, Irreversible investment and Knightian uncertainty, *Journal of Economic Theory* 136, 668–694.
- Novy-Marx, Robert, 2007, An equilibrium model of investment under uncertainty, *Review of Financial Studies* 20, 1461–1502.
- Peskir, Goran, and Albert Shiryaev, 2006, *Optimal Stopping and Free-Boundary Problems* (Lectures in Mathematics, ETH Zurich, Birkhauser).
- Protter, Philip E., 2004, *Stochastic Integration and Differential Equations* (Springer-Verlag, New York, NY).
- Shreve, Steven E., 2004, Stochastic Calculus for Finance II: Continuous-Time Models (Springer-Verlag, New York, NY).
- Thomas, Julia K., 2002, Is lumpy investment relevant for the business cycle? *Journal of Political Economy* 110, 508–534.
- Titman, Sheridan, 1985, Urban land prices under uncertainty, American Economic Review 75, 505–514.
- Williams, Joseph T., 1991, Real estate development as an option, Journal of Real Estate Finance and Economics 4, 191–208.
- World Bank, 2000, Global Economic Prospects and Developing Countries.