INVESTMENT TIMING UNDER INCOMPLETE INFORMATION*

by

Jean-Paul Décamps
GREMAQ-IDEI, Université de Toulouse

Thomas Mariotti
London School of Economics and Political Science

Stéphane Villeneuve
Université d'Evry Val d'Essonne

Contents:

The Suntory Centre
Suntory and Toyota International Centres for
Economics and Related Disciplines
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
Tel.: 020-7955 6698

Discussion Paper
No.TE/03/444
January 2003

---

* We are grateful to Bruno Biais, Nicole El Karouni, Christian Gollier and Damien Lamberton for thoughtful discussions and suggestions. We would also like to thank seminar participants at ESC Toulouse, Institut Henri Poincaré and Séminaire Bachelier for their comments. Financial support from STICERD is gratefully acknowledged by the second author. We remain, of course, solely responsible for the content of this paper.

Abstract

We study the decision of when to invest in an indivisible project whose value is perfectly observable but driven by a parameter that is unknown to the decision maker ex ante. This problem is equivalent to an optimal stopping problem for a bivariate Markov process. Using filtering and martingale techniques, we show that the optimal investment region is characterised by a continuous and non-decreasing boundary in the value/belief state space. This generates path-dependency in the optimal investment strategy. We further show that the decision maker always benefits from an uncertain drift relative to an 'average' drift situation. However, a local study of the investment boundary reveals that the value of the option to invest is not globally increasing with respect to the volatility of the value process.

Keywords: Real options, incomplete information, optimal stopping.
JEL Nos.: C61, D83.

Contact address: Dr Thomas Mariotti, Department of Economics, London School of Economics and Political Science, Houghton Street, London, WC2A 2AE, UK. Email: t.mariotti@lse.ac.uk

1. Introduction

Uncertainty and irreversibility have long been recognized as crucial factors of investment (Arrow and Fisher (1974), Henry (1974)). In contrast with the standard net present value rule, which implicitly requires investment expenditures to be fully recoverable, or investment opportunities to be seized on a now-or-never basis, the real option literature has emphasized the firms' ability to delay irreversible investment decisions (Dixit and Pindyck (1994)). In the presence of sunk costs, this flexibility in the timing of investments is valuable because it gives firms the option to wait for new information. Conversely, the loss of this option value at the time the firm invests generates an additional opportunity cost of investment. As a consequence, investment options are exercised significantly above the point at which expected discounted cash-flows cover the sunk investment expenditures.

In the benchmark case of a single indivisible project, the optimal investment policy of the firm can be mathematically determined as the solution of an optimal stopping problem. The prototype of this approach is the model of McDonald and Siegel (1986), in which the underlying value of the investment evolves as a geometric Brownian motion. Under this assumption, the optimal investment time can be explicitly characterized via Samuelson and McKean's (1965) celebrated smooth-fit principle. These results have been recently extended in various directions. For instance, Hu and Øksendal (1998) study an environment in which the investment cost is driven by a sum of correlated geometric Brownian motions, while Mordecki (1999) considers the case of a jump-diffusion value process. However, a common feature of these papers is their focus on complete information settings, in which investors have no uncertainty about the fundamental characteristics of investment projects.

In this paper, we analyze the problem of finding the optimal time to invest in an indivisible project whose value, while still perfectly observable, is driven by a parameter that is unknown to the decision maker ex ante. That is, there is a structural element of uncertainty besides the standard diffusion component of the value process. This captures in a simple way a variety of empirically relevant investment situations. For instance, a firm might ignore the exact growth characteristics of a market on which it contemplates investing. Alternatively, the owner of an asset who considers selling it might ignore how the willingness to pay of potential buyers will evolve in the future. By observing the evolution of the value, the decision maker can update his beliefs about the uncertain drift of the value process. This information is noisy, however, since it does not allow to distinguish perfectly between the relative contributions of the drift and diffusion components to the instantaneous variations of the project's value.

The filtering techniques of Liptser and Shiryaev (1977) allow us to re-state our problem recursively as an optimal stopping problem for a bivariate Markov process. The relevant state variables are the current value of the project and the decision maker's posterior beliefs about the unknown drift of the value process. The existence of an optimal investment strategy is then an immediate consequence of this filtering formulation. The multi-dimensionality of the state space is a key feature of our problem, that distinguishes it from related investment or learning models. It reflects the fact that the value process coincides with the observation process, so that its diffusion part is both a genuine component of the investment's payoff, and a source of noise for the identification of the drift. Unfortunately, this also prevents us to use smooth-fit techniques to characterize the optimal investment strategy. Nonetheless, the martingale approach developed by Lakner (1995) and Karatzas and Zhao (1998) in the context of portfolio optimization problems with partial observation allows us to derive some basic, yet useful analytical properties of the value of the investment option as a function of the current state variables.

The main analytical result of the paper is that the optimal investment region is characterized by a continuous and non-decreasing boundary in the value/belief state space. Thus, in contrast with standard real options models, the optimal decision rule is not described by a simple threshold for the current value of the investment, above which it becomes optimal to invest no matter the past evolution of the value. The presence of learning thus generates path-dependency, although suspension or abandonment of the project are not feasible options in our model, in contrast with Dixit's (1989). A striking feature of the optimal investment strategy is that it may be rational to invest after a drop in the investment's value. This is because such a drop brings bad news about the uncertain drift, and thus about the future evolution of the value, thereby reducing the current opportunity cost of investment. If the current value of the investment is high enough relative to his new estimate of the drift, the decision maker may give up on learning, reflecting that "a bird in the hand is worth two in the bush".

An important question is whether uncertainty about the drift actually benefits or penalizes the investor. To answer this question, we compare the value of an investment opportunity with drift uncertainty and learning with that of an investment opportunity in which the drift of the value process is known, and equal to the prior expectation of the drift in the first project. Using dynamic programming techniques, we show that the decision maker always prefer the former investment option, despite the fact that the value process only conveys a noisy signal of the drift. Hence, an investment opportunity with uncertain growth prospects always dominates one with average growth prospects. The intuition of this result is particularly easy to grasp in the case where the instantaneous variance of the observation/value process is small. Indeed, when this variance decreases, the value becomes a more accurate signal of the drift, so that learning occurs at a faster rate. The fact that the incomplete information problem is preferred by the decision maker to the average drift problem then simply reflects the fact that the value of the latter is convex with respect to the drift.

To get some intuition about the shape of the investment boundary, as well as about the wedge between the incomplete information problem and the average drift problem, we perform a local analysis for small values of the volatility of the observation/value process. We show that, as this volatility converges to zero, the loss in value arising from the need to learn about the drift vanishes, so that the value of the incomplete information problem converges to that of the complete information problem. An interesting by-product of this analysis is that, in contrast with standard models of investment under uncertainty, the value of an investment opportunity with uncertain drift is not everywhere increasing with respect to the variance of the value process. This illustrates again the duality of the value process in our model. An increase of the variance makes the decision maker's payoff upon investing more volatile, which per se has a positive impact on the value of the option to invest by increasing the incentive to delay investment. However, since the value process is also the observation process, this also reduces the speed at which the decision maker accumulates information about the uncertain drift, which tends to lower the value of the investment option by reducing the incentive to delay investment.

The paper is organized as follows. The model is described in Section 2. In Section 3, we provide the recursive formulation of our problem, and derive some basic properties of the value function. Section 4 derives the continuity of the investment region boundary. In Section 5, we compare the incomplete information problem with the average drift problem. Section 6 is devoted to a local study of the investment boundary as the volatility of the value process vanishes. Section 7 concludes.

2. An Investment Problem

2.1. The Model

Time is continuous, and labeled by $t \geq 0$. Uncertainty is modeled by a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For any stochastic process $X = \{X_t;\, t \geq 0\}$ defined on $(\Omega, \mathcal{F}, \mathbb{P})$, we denote by $\mathcal{F}^X = \{\mathcal{F}_t^X;\, t \geq 0\}$ the $\mathbb{P}$–augmentation of the filtration $\{\sigma(X_s;\, s \leq t);\, t \geq 0\}$ generated by X, and by $\mathcal{T}^X$ the set of $\mathcal{F}^X$–adapted stopping times that are $\mathbb{P}$–almost surely finite.

Payoffs. We consider an infinitely lived decision maker, whose task is to choose when to invest in a risky project. Investment is irreversible and entails a sunk cost $I > 0$. The value of the project follows a Brownian motion with constant uncertain drift μ and known variance σ,

$$dV_t = \mu\, dt + \sigma\, dW_t; \quad t \geq 0, \tag{1}$$

where $(W, \mathcal{F}^W)$ is a standard Wiener process independent of μ. While some of our results can easily be generalized to any finite number of possible values for μ, we shall hereafter assume for simplicity that μ can take only two values, 0 or 1. We denote by v the initial value of the project. The decision maker is risk neutral, and discounts future revenues and costs at a constant rate $r > 0$.

Information Structure. A key assumption of our model is that the decision maker does not know ex ante the true value of μ. We denote by $p \in [0, 1]$ his prior belief that $\mu = 1$. Ex post, the decision maker perfectly observes the value process V, but neither the drift μ, nor the evolution of W; i.e., his information at any time t is summarized by $\mathcal{F}_t^V$. It is clear from (1) that the value process conveys some information about μ. However, because of the shocks σW to the value, this information is noisy.

Statement of the Problem. At any time t prior to investment, the decision maker chooses whether to pay the sunk cost I to earn the gross profit V_t, or to delay further his investment. Since the only information available to him ex post is generated by the value process, his decision problem is to find a stopping time $\tau^* \in \mathcal{T}^V$ such that:

$$\sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[e^{-r\tau}(V_\tau - I)\right] = \mathbb{E}\left[e^{-r\tau^*}(V_{\tau^*} - I)\right]. \tag{2}$$

We shall denote this problem by $\mathcal{P}$. The objective of this paper is to characterize as fully as we can the optimal investment strategy for $\mathcal{P}$.

Remark. Note that under the formulation (1), the value of the investment project can become negative. Alternatively, one might set $I = 0$ and interpret V as the difference between the value of the project and the investment cost.

2.2. Relation to the Literature

Before proceeding with the analysis, it might be helpful to briefly contrast our model with some closely related lines of research.

Investment under Uncertainty. As in the standard real option framework, we model the investment decision as an optimal stopping problem (see, for instance, McDonald and Siegel (1986), or Dixit and Pindyck (1994, §5.1)). The distinguishing feature of our setting is that, since the decision maker can observe neither the drift μ nor the evolution of W, he has incomplete information about the dynamics of the value process V. Hence, his beliefs about μ become sensitive to information about the past evolution of the value, which in turn affects his anticipations about the future evolution of the value. This implies that the current value of the investment is not a sufficient statistics for the investment problem $\mathcal{P}$.

Sequential Statistical Testing. The problem of sequential testing of the two alternative hypotheses on the drift of the process (1) has been solved by Chernoff (1972, §17.5), assuming constant linear waiting costs. This is a pure optimal stopping problem, whose solution consists to accumulate information until an upper or lower threshold is reached by the belief process. A key feature of our model, however, is that the signal received coincides with the gross value of the investment, so that the diffusion term σW enters directly into the decision maker's payoff function. This implies that the decision maker's current beliefs about the drift of the observation process are not a sufficient statistics for the investment problem $\mathcal{P}$.

Optimal Experimentation. The optimal experimentation literature has recently investigated various sequential control problems under incomplete information and learning. As most of these papers, we follow Chernoff (1972) in parameterizing the unknown state of nature by the drift μ of the observation process V. Again, a common feature of these papers is that beliefs about μ are a sufficient statistics, either because the observation process represents a cumulative payoff (as in Jovanovic's (1979) job matching model, or in the strategic version of the multi-armed bandit problem studied by Bolton and Harris (1999), and in the monopoly/duopoly pricing models considered by Felli and

Harris (1996), Keller and Rady (1999), and Bergemann and Välimäki (1997, 2000)), or because terminal payoffs depend only on beliefs (as in the optimal control/stopping problem analyzed by Moscarini and Smith (2000)).

Learning in Financial Markets. The classical consumption/portfolio problem introduced by Merton (1971) has been recently extended to the case where the investor is uncertain about the drift parameter of the stock price process. While Lakner (1995) tackles the problem via martingale methods, Karatzas and Zhao (1998) show that the Bellman principle still applies and leads to explicit solutions for particular choices of utility functions (see also Brennan (1998)). Last, Veronesi (1999, 2000) and Alexander and Veronesi (2000) investigate standard financial puzzles when the drift of the stock dividend process is unknown and may change at random times. Equilibrium asset prices are then derived from the dynamics of investors' beliefs about the unobservable drift. Here again standard dynamic programming techniques apply.

3. A Markov Formulation of the Problem

In this section, we first derive a recursive formulation of problem $\mathcal{P}$, for which a natural Markov state variable is the pair (V, P) formed by the current value of investing in the project and the current beliefs about the drift of the value process. This allows us to prove the existence of an optimal stopping time for $\mathcal{P}$. We then use a change of measure transformation to derive some basic properties of the value function.

3.1. An Existence Result

The Filtering Problem. The decision maker faces a standard signal extraction problem. As μ is either 0 or 1, his beliefs about μ at any time t are summarized by the a posteriori probability $P_t = \mathbb{P}\left[\mu = 1 \mid \mathcal{F}_t^V\right]$ conditional on information available up to t. Naturally, $P_0 = p$, the a priori probability that $\mu = 1$. From Theorem 9.1 in Liptser and Shiryaev (1977), the belief process $(P, \mathcal{F}^V)$ satisfies the stochastic filtering equation:

$$dP_t = \frac{P_t\,(1 - P_t)}{\sigma}\, d\overline{W}_t; \quad t \geq 0, \tag{3}$$

where $(\overline{W}, \mathcal{F}^V)$ is a standard Wiener process relative to the decision maker's information (the so-called innovation process of filtering theory, see Liptser and Shiryaev (1977, Theorem 7.12)), that verifies:

$$d\overline{W}_t = \frac{dV_t - P_t\, dt}{\sigma}; \quad t \geq 0. \tag{4}$$

In other words, the change in beliefs dP_t is normally distributed with mean 0 and variance $P_t^2\,(1-P_t)^2/\sigma^2\,dt$. In particular, beliefs follow a martingale.

It is worth noting that 0 and 1 are absorbing barriers for the belief process. That is, if beliefs start at one of these values, they will stay constant almost surely. However, if beliefs do not start at one of these values, then both cannot be attained in finite time. These useful properties are summarized in the following lemma.

LEMMA 3.1 *For any $p \in [0,1]$, the unique solution $(P^p, \mathcal{F}^V)$ of the stochastic differential equation (3) satisfying $P_0^p = p$ lies $\mathbb{P}$–almost surely in $[0,1]$ and satisfies:*

(i) $P^0 \equiv 0$ *and* $P^1 \equiv 1$, $\mathbb{P}$*–almost surely;*

(ii) *If* $p \notin \{0,1\}$, *then* $\inf\{t \geq 0 \mid P_t^p \notin (0,1)\} = \infty$, $\mathbb{P}$*–almost surely.*

The proof is standard, and relegated to the Appendix.

The Recursive Formulation. Note that, since W is not a Brownian motion under the filtration $\mathcal{F}^V$ of the observation process, the formulation (2) of problem $\mathcal{P}$ is not recursive. What the filtering approach allows us to do is to transform $\mathcal{P}$ into a stopping time problem for a bi-dimensional Markov process. Indeed, it is immediate from (3)-(4) that the joint observation/belief process can be rewritten as:

$$d\begin{pmatrix} V_t^v \\ P_t^p \end{pmatrix} = \begin{pmatrix} P_t^p \\ 0 \end{pmatrix} dt + \begin{pmatrix} \sigma \\ P_t^p\,(1-P_t^p)\,/\sigma \end{pmatrix} d\overline{W}_t; \quad t \geq 0, \tag{5}$$

where we have made explicit the dependence of V and P upon their respective initial values v and p. Since the innovation process $\overline{W}$, unlike W, is a Brownian motion under $\mathcal{F}^V$, it is clear from (5) that the joint observation/belief process $X^{v,p} = (V^v, P^p)$ is a Markov process under $\mathcal{F}^V$. We can thus re-state $\mathcal{P}$ as the problem of finding a value function $G^* : \mathbb{R} \times [0,1] \to \mathbb{R}$ and a stopping time $\tau^* \in \mathcal{T}^V$ such that:

$$G^*(v,p) = \sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\,g(X_\tau^{v,p})\right] = \mathbb{E}\left[\mathrm{e}^{-r\tau^*}\,g(X_{\tau^*}^{v,p})\right], \tag{6}$$

for any $(v,p) \in \mathbb{R} \times [0,1]$, where $g(v,p) \equiv v - I$.

Existence. Given the recursive representation (6), we are in position to apply standard results in optimal stopping theory for Markov processes to prove the existence of a solution to problem $\mathcal{P}$. Assuming that the value function G^* is well-defined, which will be proved shortly, let $\mathcal{S}^* = \{(v,p) \in \mathbb{R} \times [0,1] \mid G^*(v,p) = g(v,p)\}$ be the coincidence set for our problem. Also, define $b^*(p) = \inf\{v \in \mathbb{R} \mid (v,p) \in \mathcal{S}^*\}$ for any given belief $p \in [0,1]$. The following result is proved in the Appendix.

PROPOSITION 3.1 *The following holds*

(i) *There exist an optimal value function G^* and an optimal stopping time $\tau^* \in \mathcal{T}^V$ solution to (6);*

(ii) *The coincidence set $\mathcal{S}^*$ is non-empty and satisfies $\mathcal{S}^* = \bigcup_{p\in[0,1]}[b^*(p), \infty)$.*

From Proposition 3.1, $\tau^* = \inf\{t \geq 0 \mid V_t^v \geq b^*(P_t^p)\}$, i.e., the optimal investment strategy consists to invest when the value crosses an upper boundary that depends on the current beliefs about μ. In what follows, b^* will be referred to as the *investment boundary function*, and $\mathcal{S}^*$ as the *optimal investment region*.

Remark. Note that the extreme points of the investment boundary corresponding respectively to $p = 0$ and $p = 1$ can be explicitly characterized by exploiting the fact that they are absorbing barriers for the belief process. Solving $\mathcal{P}$ for $p \in \{0, 1\}$ simply amounts to find an optimal stopping time for a discounted Brownian motion with constant and ex ante known drift, a problem for which a closed-form solution is available. We shall come back in Section 5 to the comparison between $\mathcal{P}$ and this standard problem. For the time being, let us just point out that $b^*(0) = I + \sigma^2/\sqrt{2r\sigma^2}$ and $b^*(1) = I + \sigma^2/(\sqrt{1 + 2r\sigma^2} - 1)$, so that in particular $b^*(1) > b^*(0)$.

3.2. Properties of the Value Function

We now derive the properties of the value function G^* of problem $\mathcal{P}$ using a Girsanov transformation introduced by Lackner (1995) and Karatzas and Zhao (1998) to study portfolio maximization problems under partial information. This transformation allows us to construct a probability measure $\mathbb{Q}$ under which μ is independent of V, thereby leading us to an alternative formulation of the value function.

A Girsanov Transformation. Since μ and W are independent by assumption, W is also a Brownian motion with respect to the enlarged filtration $\mathcal{F}^{\mu,W}$ generated by both μ and W. Let us define a probability measure $\mathbb{Q}$ by its Radon-Nikodym derivative with respect to $\mathcal{F}^{\mu,W}$,

$$\frac{d\mathbb{Q}}{d\mathbb{P}}_{|\mathcal{F}_t^{\mu,W}} = Z_t = \exp\left(-\left(\frac{\mu}{\sigma}\right) W_t - \frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2 t\right); \quad t \geq 0. \tag{7}$$

Note that by construction, $(Z, \mathcal{F}^{\mu,W})$ is a positive local martingale. Moreover, since $\mathbb{E}[Z_t \mid \mu] = 1$ for any time t, it follows that $\mathbb{E}[Z_t] = 1$ and thus that $(Z, \mathcal{F}^{\mu,W})$ is a martingale (see Karatzas and Shreve (1991, §3.5.D)). Therefore, Girsanov theorem implies that the process $(B, \mathcal{F}^{\mu,W})$ defined by:

$$B_t = W_t + \left(\frac{\mu}{\sigma}\right) t; \quad t \geq 0 \tag{8}$$

is a Brownian motion under the probability $\mathbb{Q}$. Since μ is $\mathcal{F}_0^{\mu,W}$–measurable and independent of W, and B has independent increments, it is clear that B and μ are independent under the probability $\mathbb{Q}$. It follows that B is also a standard Brownian motion with respect to its own filtration $\mathcal{F}^B$ under the probability $\mathbb{Q}$. Moreover, since $dV_t = \sigma\, dB_t$, the filtrations $\mathcal{F}^B$ and $\mathcal{F}^V$ generated by B and V are identical, as well as the sets of stopping times $\mathcal{T}^B$ and $\mathcal{T}^V$. Then, the following holds.

Proposition 3.2 *For any* $(v,p) \in \mathbb{R} \times [0,1]$,

$$G^*(v,p) = \sup_{\tau \in \mathcal{T}^B} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} \left(1 + p \left(\exp\left(\frac{B_\tau}{\sigma} - \frac{\tau}{2\sigma^2} \right) - 1 \right) \right) (v + \sigma B_\tau - I) \right]. \tag{9}$$

Proof: Define, for $i = 0, 1$:

$$H_t(i, B_t) = \exp\left(\left(\frac{i}{\sigma}\right) B_t - \frac{1}{2} \left(\frac{i}{\sigma}\right)^2 t \right); \quad t \geq 0, \tag{10}$$

where B_t is given by (8). Clearly, $(H(i,B), \mathcal{F}^B)$ is a martingale. Therefore, using the expression (7) for the Radon-Nikodym derivative of $\mathbb{Q}$ with respect to $\mathbb{P}$, we can rewrite the value function G^* as:

$$G^*(v,p) = \sup_{\tau \in \mathcal{T}^V} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} H_\tau(\mu, B_\tau)\, (v + \sigma B_\tau - I) \right],$$

for any $(v,p) \in \mathbb{R} \times [0,1]$. It follows that:

$$\begin{aligned}
G^*(v,p) &= \sup_{\tau \in \mathcal{T}^V} \mathbb{E}_{\mathbb{Q}} \left[\mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} H_\tau(\mu, B_\tau)\, (v + \sigma B_\tau - I) \mid \mathcal{F}_\tau^V \right] \right] \\
&= \sup_{\tau \in \mathcal{T}^B} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau}\, \mathbb{E}_{\mathbb{Q}} \left[H_\tau(\mu, B_\tau) \mid \mathcal{F}_\tau^B \right] (v + \sigma B_\tau - I) \right] \\
&= \sup_{\tau \in \mathcal{T}^B} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} \left(p\, H_\tau(1, B_\tau) + (1-p)\, H_\tau(0, B_\tau) \right) (v + \sigma B_\tau - I) \right] \\
&= \sup_{\tau \in \mathcal{T}^B} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} \left(1 + p \left(\exp\left(\frac{B_\tau}{\sigma} - \frac{\tau}{2\sigma^2} \right) - 1 \right) \right) (v + \sigma B_\tau - I) \right],
\end{aligned}$$

where the second inequality follows immediately from $\mathcal{F}^V = \mathcal{F}^B$, and the third from the fact that, for any $\tau \in \mathcal{T}^V$, μ and B_τ are independent under the probability $\mathbb{Q}$, which implies in particular that $\mathbb{Q}[\mu = 1 \mid \mathcal{F}_\tau^B] = p$. ■

This result admits a natural interpretation. Under the probability $\mathbb{Q}$, the value process $V = v + \sigma B$ is independent of μ, so that no learning occurs: essentially, we have transformed an incomplete information stopping problem into a complete information one. Of course, this requires a corresponding modification of the payoffs. The decision maker now maximizes the expectation of a discounted weighted average of $H(i, B)\,(v + \sigma B - I)$, $i = 0, 1$, where the weights $1 - p$ and p reflect his prior beliefs about the value of μ.

The above Girsanov transformation allows us to represent the belief process in terms of the value process. A direct application of Bayes formula yields:

$$P_t^p = \frac{p\,H_t(1, B_t)}{p\,H_t(1, B_t) + 1 - p} = \frac{p\,\exp\left(\dfrac{V_t^v - v}{\sigma^2} - \dfrac{t}{2\sigma^2}\right)}{p\,\exp\left(\dfrac{V_t^v - v}{\sigma^2} - \dfrac{t}{2\sigma^2}\right) + 1 - p}; \quad t \geq 0, \tag{11}$$

for any $(v, p) \in \mathbb{R} \times [0, 1]$, see Shiryayev (1978, §4.2.1). It is immediate to see from (11) that beliefs satisfy a non-crossing property, in the sense that $p > p'$ implies that $P_t^p > P_t^{p'}$ at any time t, $\mathbb{P}$–almost surely.

Properties of G^*. Using the characterizations provided in Propositions 3.1 and 3.2, we can now derive some basic properties of the value function. From (5) and (6), note first that, for any $(v, p) \in \mathbb{R} \times [0, 1]$,

$$G^*(v, p) = \sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(v + \int_0^\tau P_t^p\,dt + \sigma \overline{W}_\tau - I\right)\right]. \tag{12}$$

Hence, using the non-crossing property of the belief process and the fact that the optimal stopping time for $\mathcal{P}$ is $\mathbb{P}$–almost surely finite, we immediately obtain the following monotonicity result.

Corollary 3.1 *The following holds:*

(i) For any $p \in [0, 1]$, the mapping $v \mapsto G^*(v, p)$ is increasing on $\mathbb{R}$;

(ii) For any $v \in \mathbb{R}$, the mapping $p \mapsto G^*(v, p)$ is non-decreasing on $[0, 1]$.

Our next result follows immediately from (9) and the fact that the supremum of a family of linear functions is convex.

Corollary 3.2 *The following holds:*

(i) For any $p \in [0,1]$, the mapping $v \mapsto G^*(v,p)$ is convex on $\mathbb{R}$;

(ii) For any $v \in \mathbb{R}$, the mapping $p \mapsto G^*(v,p)$ is convex on $[0,1]$.

In other words, the decision maker is ready to accept risky bets on the initial value v of the project, for a fixed value of p, or on the probability p that the project has a high drift, for a fixed value of v; the latter implies that costless information about μ, in the form of a mean-preserving spread over p, always has a positive value for the decision maker. However, it does not follow from (9) that G^* is convex with respect to the pair (v,p), and thus that the decision maker would be ready to accept risky bets on both v and p simultaneously. Indeed, we shall argue in Section 6 that G^* cannot be convex on its whole domain.

The last result of this section concerns the continuity of the value function with respect to the initial conditions (v,p) and the variance σ of the observation/value process. In the latter case, we denote the value function by G^*_σ instead of G^*.

Corollary 3.3 *The following holds:*

(i) The mapping $(v,p) \mapsto G^*(v,p)$ is continuous on $\mathbb{R} \times [0,1]$;

(ii) The mapping $\sigma \mapsto G^*_\sigma(v,p)$ is continuous on $\mathbb{R}_{++}$.

The proof relies on the two formulations of the value function given in Propositions 3.1 and 3.2. Details are provided in the Appendix.

4. The Optimal Investment Region

Since $\mathcal{P}$ is intrinsically a bi-dimensional problem, the standard partial differential equation approach to optimal stopping is of little help, since no closed form solution for G^* is available. In particular, it is not clear whether or not the usual smooth pasting condition holds along the investment boundary (see Shiryayev (1978, §3.8.1) for a discussion of this point). Instead of focusing on the optimal value function, a task we shall return to in Section 5, we first determine some properties of the investment boundary function b^*. We then discuss some qualitative features of the optimal investment strategy.

4.1. Properties of the Investment Boundary Function

Our main findings about b^* are summarized in the following result.

THEOREM 4.1 *The investment boundary function* $b^* : [0,1] \to \mathbb{R}_{++}$ *is continuous and non-decreasing on* $[0,1]$.

The proof of Theorem 4.1 is divided into two steps. The monotonicity and left-continuity of b^* follow from standard arguments (see for instance Villeneuve (1999, Proposition 3.2)).

LEMMA 4.1 *The investment boundary function* b^* *is non-decreasing on* $[0,1]$.

PROOF: By Proposition 3.1, $G^*(v,p) = v - I$ for any $v \geq b^*(p)$. Moreover, the mapping $p \mapsto G^*(v,p)$ is non-decreasing according to Corollary 3.1. It follows that $G^*(v,p_0) = v - I$ for any $p_0 \leq p$ and $v \geq b^*(p)$, which implies that $b^*(p_0) \leq b^*(p)$. ■

The monotonicity of the investment boundary function with respect to beliefs captures the intuitive idea that, for any given current value of the project, the more confident the decision maker is that the drift of the value process is high, the more he is willing to 'experiment', i.e., to delay his investment. This generalizes to our incomplete information setting the standard result that the optimal investment trigger for a project whose value follows a Brownian motion with a constant and known drift is increasing in the value of this drift, see Section 5.1.

LEMMA 4.2 *The investment boundary function* b^* *is left-continuous on* $(0,1]$.

PROOF: Let $\{p_n\}$ be a non-decreasing sequence in $[0,1]$ converging to $p \in (0,1]$. From Lemma 4.1, the sequence $\{b^*(p_n)\}$ is non-decreasing and upper bounded by $b^*(p)$, and therefore converges to a limit $b^*(p^-)$. By definition, $G^*(b^*(p_n),p_n) = b^*(p_n) - I$ for any $n \in \mathbb{N}$. By Corollary 3.2, G^* is continuous, so that $G^*(b^*(p^-),p) = b^*(p^-) - I$. Hence $b^*(p) \leq b^*(p^-)$, and thus $b^*(p) = \lim_{n\to\infty} b^*(p_n)$, which implies the result. ■

The proof that b^* is right-continuous is a bit more involved. We first need the following lemma.

LEMMA 4.3 *For all* $(p_0,x) \in [0,1) \times \mathbb{R}_{++}$, *and for each* $p \in [p_0,1]$,

$$\sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x + \int_0^\tau (P_t^p - P_t^{p_0})\,dt\right)\right] \leq \frac{p-p_0}{r} + \sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x - \frac{P_\tau^p - P_\tau^{p_0}}{r}\right)\right].$$

PROOF: Fix $p_0 \in [0,1)$. The result is obvious for $p = p_0$. Consider some $p \in (p_0, 1)$. The non-crossing property of the belief process P implies that the difference $P^p - P^{p_0}$ remains positive $\mathbb{P}$–almost surely. Hence, for any $\tau \in \mathcal{T}^V$, we obtain that:

$$\begin{aligned}\mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x + \int_0^\tau (P_t^p - P_t^{p_0})\,dt\right)\right] &\leq \mathbb{E}\left[\mathrm{e}^{-r\tau}x + \int_0^\tau \mathrm{e}^{-rt}\,(P_t^p - P_t^{p_0})\,dt\right] \\ &= \frac{p - p_0}{r} + \mathbb{E}\left[\mathrm{e}^{-r\tau}x - \int_\tau^\infty \mathrm{e}^{-rt}\,(P_t^p - P_t^{p_0})\,dt\right],\end{aligned}$$

where the equality follows from the monotone convergence theorem. Next,

$$\begin{aligned}\mathbb{E}\left[\int_\tau^\infty \mathrm{e}^{-rt}\,(P_t^p - P_t^{p_0})\,dt\right] &= \mathbb{E}\left[\mathbb{E}\left[\int_\tau^\infty \mathrm{e}^{-rt}\,(P_t^p - P_t^{p_0})\,dt \mid \mathcal{F}_\tau^V\right]\right] \\ &= \mathbb{E}\left[\mathrm{e}^{-r\tau}\,\mathbb{E}\left[\int_0^\infty \mathrm{e}^{-rt}\,(P_{t+\tau}^p - P_{t+\tau}^{p_0})\,dt \mid \mathcal{F}_\tau^V\right]\right] \\ &= \mathbb{E}\left[\mathrm{e}^{-r\tau}\int_0^\infty \mathrm{e}^{-rt}\,\mathbb{E}\left[(P_{t+\tau}^p - P_{t+\tau}^{p_0}) \mid \mathcal{F}_\tau^V\right]dt\right] \\ &= \mathbb{E}\left[\mathrm{e}^{-r\tau}\,\frac{P_\tau^p - P_\tau^{p_0}}{r}\right],\end{aligned}$$

where the third equality follows from the monotone convergence theorem and the fourth from the strong Markov property together with the fact that $P^p - P^{p_0}$ is a martingale. The result then follows immediately from the previous inequality. ■

Our next result is a simple consequence of the fact that the belief process P^{p_0} is locally Lipschitzian with respect to its initial condition $p_0 \in (0,1)$.

LEMMA 4.4 *For any $(p_0, x) \in [0,1) \times \mathbb{R}_{++}$, there exists $\eta \in (0, 1 - p_0)$ such that:*

$$\sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x + \int_0^\tau (P_t^p - P_t^{p_0})\,dt\right)\right] = x \tag{13}$$

for all $p \in [p_0, p_0 + \eta)$.

PROOF: Fix $p_0 \in [0,1)$. The result is obvious for $p = p_0$. Consider some $p \in (p_0, 1)$. Since the supremum in (13) is greater or equal than x, it follows from Lemma 4.3 that we need only to prove that there exists $\eta \in (0, 1 - p_0)$ such that:

$$\sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x - \frac{P_\tau^p - P_\tau^{p_0}}{r}\right)\right] = x - \frac{p - p_0}{r} \tag{14}$$

for any $p \in (p_0, p_0 + \eta)$. Suppose first that $p_0 \in (0,1)$, and let $p \in (0,1)$. Then, from (11), the process $P^p - P^{p_0}$ can be written as $(p - p_0)\, f(H(1,B), p, p_0)$, where B is given by (8), $H(1,B)$ is defined as in (10), and for any $h \in \mathbb{R}_+$,

$$f(h,p,p_0) = \frac{h}{pp_0\,(h-1)^2 + (p+p_0)\,(h-1) + 1}. \tag{15}$$

It is easy to check from (15) that the mapping $h \mapsto f(h,p,p_0)$ reaches a maximum on $\mathbb{R}_+$ at $h(p,p_0) = \sqrt{(1-p)(1-p_0)/pp_0}$, and that the function $p \mapsto f(h(p,p_0),p,p_0)$ is bounded above by some positive constant $C(p_0)$ in a neighborhood of p_0. Thus,

$$\sup_{t \geq 0} |P_t^p - P_t^{p_0}| \leq C(p_0)\,|p - p_0|$$

for all p in a neighborhood of p_0, $\mathbb{P}$–almost surely. Since $r, x > 0$ and $H(1,B)$ is non-negative, there exists $\eta \in (0, 1 - p_0)$ such that $x - (P^p - P^{p_0})/r$ is a positive martingale whenever $p \in (p_0, p_0 + \eta)$. Hence, by the optional sampling theorem,

$$\begin{aligned}\mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x - \frac{P_\tau^p - P_\tau^{p_0}}{r}\right)\right] &\leq \mathbb{E}\left[x - \frac{P_\tau^p - P_\tau^{p_0}}{r}\right] \\ &= x - \frac{p - p_0}{r}\end{aligned}$$

for all $\tau \in \mathcal{T}^V$ and all $p \in (p_0, p_0 + \eta)$. Since the supremum in (14) is greater than $x - (p - p_0)/r$, the result follows. Suppose now that $p_0 = 0$ and $p \in (0,1)$. Then, from Lemma 2.1, we need only to consider the problem:

$$G^\dagger(p) = \sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(x - \frac{P_\tau^p}{r}\right)\right]. \tag{16}$$

If $rx \geq 1$, $G^\dagger(p) = x - p/r$ since $x - P^p/r$ is then a positive martingale. If $rx < 1$, a standard computation (see, e.g., Bolton and Harris (1999)) yields that a solution to (16) is given by $\tau^\dagger = \inf\{t \geq 0 \mid P_t^p \leq p^\dagger\}$, where $p^\dagger = (\gamma - 1)\,rx/(\gamma - 2rx + 1) > 0$ with $\gamma = \sqrt{1 + 8r\sigma^2}$. Hence $G^\dagger(p) = x - p/r$ if $p \leq p^\dagger$, which implies the result. ■

We are now ready to complete the proof of Theorem 4.1.

LEMMA 4.5 *The investment boundary function b^* is right-continuous on $[0,1)$.*

PROOF: Fix $p_0 \in [0,1)$, and suppose by way of contradiction that $\lim_{p \downarrow p_0} b^*(p) = b^*(p_0^+) > b^*(p_0)$. Fix some $v \in (b^*(p_0), b^*(p_0^+))$. For any $p \in (p_0, 1]$, let $T^*_{v,p}$ be the optimal stopping time for our problem starting at (v, p). Then, by (6) and (12),

$$\begin{aligned} G^*(v,p) &= \mathbb{E}\left[\mathrm{e}^{-rT^*_{v,p}}\left(v + \int_0^{T^*_{v,p}} P_t^p \, dt + \sigma \overline{W}_{T^*_{v,p}} - I\right)\right] \\ &\leq b^*(p_0) - I + \mathbb{E}\left[\mathrm{e}^{-rT^*_{v,p}}\left(v - b^*(p_0) + \int_0^{T^*_{v,p}} (P_t^p - P_t^{p_0}) \, dt\right)\right]. \end{aligned} \tag{17}$$

where we have used the fact that $G^*(b^*(p_0), p_0) = b^*(p_0) - I$. Since $v - b^*(p_0) > 0$, Lemma 4.4 implies that there exists $\eta \in (0, 1 - p_0)$ such that, for each $p \in (p_0, p_0 + \eta)$,

$$\mathbb{E}\left[\mathrm{e}^{-rT^*_{v,p}}\left(v - b^*(p_0) + \int_0^{T^*_{v,p}} (P_t^p - P_t^{p_0}) \, dt\right)\right] \leq v - b^*(p_0).$$

Hence, from (17), $G^*(v,p) \leq v - I$ for any $p \in (p_0, p_0 + \eta)$ and, since the reverse inequality always hold, $G^*(v,p) = v - I$, so $b^*(p) \leq v$ by definition of b^*. As b^* is non-decreasing by Lemma 4.1, it follows that $v \geq b^*(p_0^+)$ for any $p \in (p_0, p_0 + \eta)$. This contradicts the fact that, by assumption, $v < b^*(p_0^+)$. Hence the result. ■

4.2. Some Qualitative Features of the Investment Strategy

We begin with the following simple observation. Since $b^*(1) > b^*(0)$ and b^* is continuous on $[0,1]$ by Theorem 4.1, the optimal strategy τ^* for $\mathcal{P}$ is not a trigger strategy relative to the process V, i.e., a stopping time of the form $T_{\overline{b}} = \inf\left\{t \geq 0 \mid V_t \geq \overline{b}\right\}$ for some threshold $\overline{b}$. In contrast with the predictions of standard models of irreversible investment under uncertainty (see Dixit and Pindyck (1994)), the value of the project at the time of the investment does not therefore necessarily coincide with the maximum historic value. As pointed out in the Introduction, a rational investor may even optimally decide to invest after a drop of the value. The model thus allows some form of ex post regret. For instance, the owner of a house who sells it at a discount might regret not having accepted an earlier high quote because he then anticipated a sustained boom of the housing market.

This path-dependency reflects the fact that the innovations of the belief process are positively correlated with the fluctuations of the value, as is easily seen from (5). Overall, the fluctuations of the value have two opposite effects on the investment decision. A positive (resp. negative) innovation in V has a direct positive (resp. negative) impact on the payoff from investing immediately, and thus on the attractiveness of the investment for the investor. On the other hand, a positive (resp. negative) innovation in V indirectly increases (resp. decreases) the subjective probability that the unknown drift μ has a high value, thus providing good (resp. bad) news about the likelihood of further future increases of V. Ceteris paribus, this indirect effect makes the investor more (resp. less) willing to delay his investment.

5. Comparison with the Constant Drift Model

In this section, we provide a comparison between problem $\mathcal{P}$ and the standard problem of finding an optimal investment time for a project whose value follows a Brownian motion with constant and ex ante known drift. In particular, we want to determine whether a rational investor would choose to exchange an investment project with unknown drift μ equal to 1 with probability p and to 0 otherwise, with an 'average drift' project with a known drift equal to p.

5.1. The Average Drift Problem

Suppose that the decision maker has also the opportunity to invest, at the same cost I, into an alternative project whose value is observable and follows a Brownian motion with constant and known drift $p \in [0,1]$, and known variance σ,

$$d\widehat{V}_t = p\,dt + \sigma\,dW_t; \quad t \geq 0. \tag{18}$$

We assume that the two investment projects with values V and $\widehat{V}$ are mutually exclusive and that the decision maker must make up his mind at date zero about the project he might later invest in. If he chooses the project with value $\widehat{V}$, we can state his decision problem as of finding a value function $\widehat{G}(\cdot,p) : \mathbb{R} \to \mathbb{R}$ and a stopping time $\widehat{\tau} \in \mathcal{T}^W$ such that, for any $v \in \mathbb{R}$,

$$\widehat{G}(v,p) = \sup_{\tau \in \mathcal{T}^W} \mathbb{E}\,[\mathrm{e}^{-r\tau}\,(\widehat{V}_t^v - I)], \tag{19}$$

where v refers to the initial value of the project. We shall call this problem $\widehat{\mathcal{P}}$, or the *average drift problem*. Indeed, $\widehat{\mathcal{P}}$ only differs from $\mathcal{P}$ in that in the former problem, the drift of the value is known and equal to p, whereas in the latter, it is unknown and p is to be interpreted as the prior belief that $\mu = 1$.

It is well-known that the optimal investment strategy for $\widehat{\mathcal{P}}$ is a trigger strategy, i.e., a stopping time of the form $\widehat{T}_{\overline{b}} = \inf\{t \geq 0 \mid \widehat{V}_t^v \geq \overline{b}\}$, for some $\overline{b} > I$. Specifically, let $f(p) = \sqrt{p^2 + 2r\sigma^2} - p$ and $\widehat{b}(p) = I + \sigma^2/f(p)$. We have the following result.

LEMMA 5.1 *$\widehat{T}_{\widehat{b}(p)}$ is an optimal stopping time for $\widehat{\mathcal{P}}$.*

This standardly results from the fact that the Laplace transform of $\widehat{T}_{\overline{b}}$ is given by $\mathbb{E}\left[\mathrm{e}^{-r\widehat{T}_{\overline{b}}}\right] = \exp\left((v - \overline{b})\, f(p)/\sigma^2\right)$, see Karatzas and Shreve (1991, §3.5.C). The value function for $\widehat{\mathcal{P}}$ can then be written as:

$$\widehat{G}(v,p) = \begin{cases} \exp\left(-1 + \dfrac{f(p)}{\sigma^2}\,(v - I)\right) \dfrac{\sigma^2}{f(p)} & \text{if} \quad v < I + \dfrac{\sigma^2}{f(p)} \\ \\ v - I & \text{if} \quad v \geq I + \dfrac{\sigma^2}{f(p)} \end{cases} . \tag{20}$$

There is no a priori obvious relationship between G^* and $\widehat{G}$, except of course when $p \in \{0,1\}$, i.e., the decision maker in $\mathcal{P}$ knows with certainty the true value of μ ex ante, in which case they coincide. With a slight abuse of terminology, we will refer to $\widehat{b}$ as the investment boundary function for the average drift problem.

5.2. *The Comparison Result*

Since the objective function in (19) is linear in the drift p of the value process, the value function $\widehat{G}$ for the average drift problem is convex in p. This implies that a risk neutral investor is ready to exchange the option to invest in the average drift problem for an option to invest in a project with value $V_t = v + \mu\, t + \sigma\, W_t$ and uncertain drift $\mu \in \{0,1\}$ drawn according to $\mathbb{P}[\mu = 1] = p$, provided he is informed of the value of μ immediately after its realization, and can thus take his investment decision under complete information about μ. In problem $\mathcal{P}$, the decision maker's information structure is coarser, since he has only access to an imperfect learning technology—namely, the observation of V. Nevertheless, we have the following result.

THEOREM 5.1 *For any $(v,p) \in \mathbb{R} \times [0,1]$, $G^*(v,p) \geq \widehat{G}(v,p)$.*

That is, the value of the option to invest in the project with an uncertain drift is higher than that of investing in the project with an average drift, despite the fact that, in the former case, the decision maker has to learn about μ before taking his investment decision. Geometrically, the investment boundary b^* for the incomplete information problem $\mathcal{P}$ is to the right of the investment boundary $\widehat{b}$ for the average drift problem $\widehat{\mathcal{P}}$, as illustrated on Figure 1 below.

The proof of Theorem 5.1 relies on standard discrete time approximations of problems $\mathcal{P}$ and $\widehat{\mathcal{P}}$. We shall need the following notation. Let Ψ be the space of continuous functions $\psi : \mathbb{R} \times [0,1] \to \mathbb{R}$ such that the family $\{\psi(V_\tau^v, P_\tau^p) \mid \tau \in \mathcal{T}^V\}$ is uniformly integrable. Similarly, let $\widehat{\Psi}$ be the space of continuous functions $\widehat{\psi} : \mathbb{R} \to \mathbb{R}$ such that the family $\{\widehat{\psi}(\widehat{V}_\tau^v) \mid \tau \in \mathcal{T}^W\}$ is uniformly integrable. For any fixed $\delta > 0$, interpreted as the duration of a discrete time period, define the following operators acting respectively on Ψ and $\widehat{\Psi}$ by:

$$Q_\delta(\psi)(v,p) = \max\left\{\psi(v,p), \mathbb{E}\left[\mathrm{e}^{-r\delta}\, \psi(V_\delta^v, P_\delta^p)\right]\right\} \tag{21}$$

and:

$$\widehat{Q}_\delta(\widehat{\psi})(v) = \max\left\{\widehat{\psi}(v), \mathbb{E}\left[e^{-r\delta}\, \widehat{\psi}(\widehat{V}_\delta^v)\right]\right\}. \tag{22}$$

It is clear from these definitions that Q_δ (resp. $\widehat{Q}_\delta$) maps Ψ (resp. $\widehat{\Psi}$) into itself. Therefore, for any $\psi \in \Psi$, $\widehat{\psi} \in \widehat{\Psi}$ and $n \in \mathbb{N}$, we can define recursively the iterates $Q_\delta^{n+1}(\psi) = Q_\delta(Q_\delta^n(\psi))$ and $\widehat{Q}_\delta^{n+1}(\widehat{\psi}) = \widehat{Q}_\delta(\widehat{Q}_\delta^n(\widehat{\psi}))$. Note also that Q_δ is monotone, i.e., if $\psi, \widetilde{\psi} \in \Psi$ and $\psi \leq \widetilde{\psi}$, then $Q_\delta(\psi) \leq Q_\delta(\widetilde{\psi})$, and that $\widehat{Q}_\delta$ preserves convexity by the convexity of the maximum operator.

For any $v \in \mathbb{R}$, let $\widehat{\psi}(v) = v - I$, and let $\pi : \mathbb{R} \times [0,1] \to \mathbb{R}$ denote the projection of elements of $\mathbb{R} \times [0,1]$ on their first coordinate. It is immediate to check that $\widehat{\psi} \in \widehat{\Psi}$ and $\widehat{\psi} \circ \pi \in \Psi$. Our next result follows immediately from the characterization of the value function of an optimal stopping problem as the smallest excessive majorant of the reward function (see Shiryayev, 1978, §3.2.2, Lemma 3, and §3.3.1, Theorem 1).

Lemma 5.2 *For any* $(v,p) \in \mathbb{R} \times [0,1]$,

(i) $\lim_{\delta \to 0} \lim_{n \to \infty} Q_\delta^n(\widehat{\psi} \circ \pi)(v,p) = G^*(v,p)$;

(ii) $\lim_{\delta \to 0} \lim_{n \to \infty} \widehat{Q}_\delta^n(\widehat{\psi})(v) = \widehat{G}(v,p)$.

The interpretation of this approximation result is clear. For any fixed $\delta > 0$, the limit with respect to n yields the value of the infinite horizon, discrete time problem, where the investor is constrained to stopping times with range in $\{n\delta \mid n \in \mathbb{N}\}$. Letting then δ go to zero yields the value of the continuous time problem. Given Lemma 5.2, Theorem 5.1 is a direct consequence of the following lemma.

LEMMA 5.3 *For any* $(\delta, n) \in \mathbb{R}_{++} \times \mathbb{N}$, $Q_\delta^n(\widehat{\psi} \circ \pi) \geq \widehat{Q}_\delta^n(\widehat{\psi}) \circ \pi$.

PROOF: We proceed by induction. First, for any $(v, p) \in \mathbb{R} \times [0, 1]$,

$$\begin{aligned} Q_\delta(\widehat{\psi} \circ \pi)(v, p) &= \max\left\{\widehat{\psi}(v), \mathbb{E}\left[\mathrm{e}^{-r\delta}\, \widehat{\psi}(v + \mu\delta + \sigma W_\delta)\right]\right\} \\ &\geq \max\left\{\widehat{\psi}(v), \mathbb{E}\left[\mathrm{e}^{-r\delta}\, \widehat{\psi}(v + p\delta + \sigma W_\delta)\right]\right\} \\ &= (\widehat{Q}_\delta(\widehat{\psi}) \circ \pi)(v, p), \end{aligned}$$

where the inequality follows from the convexity of $\widehat{\psi}$ and the independence between μ and W, together with Jensen inequality. Next, suppose that $Q_\delta^n(\widehat{\psi} \circ \pi) \geq \widehat{Q}_\delta^n(\widehat{\psi}) \circ \pi$ for some $n \in \mathbb{N}$. Then, for any $(v, p) \in \mathbb{R} \times [0, 1]$,

$$\begin{aligned} Q_\delta^{n+1}(\widehat{\psi} \circ \pi)(v, p) &= Q_\delta(Q_\delta^n(\widehat{\psi} \circ \pi))(v, p) \\ &\geq Q_\delta(\widehat{Q}_\delta^n(\widehat{\psi}) \circ \pi)(v, p) \\ &\geq (\widehat{Q}_\delta(\widehat{Q}_\delta^n(\widehat{\psi})) \circ \pi)(v, p) \\ &= (\widehat{Q}_\delta^{n+1}(\widehat{\psi}) \circ \pi)(v, p), \end{aligned}$$

where the first inequality follows from the induction hypothesis and the monotonicity of Q_δ, and the second inequality from the first part of the proof together with the fact that $\widehat{Q}_\delta$ preserves convexity. Hence the result. ∎

5.3. A Remark on Trigger Strategies

In the previous subsection, we have used tools from dynamic programming theory to compare the incomplete information problem $\mathcal{P}$ with the average drift problem $\widehat{\mathcal{P}}$. Since the optimal stopping time in $\widehat{\mathcal{P}}$ is a trigger strategy, the reader might wonder whether a more direct approach would not consist to restrict the strategy space in $\mathcal{P}$ to trigger strategies conditional on the current value V of the project in order to compare the value of this constrained problem to that of $\widehat{\mathcal{P}}$. It turns out that this approach is not conclusive; it is nevertheless instructive to understand why.

To do so, let $\overline{\mathcal{T}}^V$ be the set of trigger stopping times for V, i.e., the subset of $\mathcal{T}^V$ composed of stopping times of the form $T_{\overline{b}} = \inf\{t \geq 0 \mid V_t \geq \overline{b}\}$ for $\overline{b} \in \mathbb{R}$, and consider the following constrained problem:

$$\overline{G}(v,p) = \sup_{\tau \in \overline{\mathcal{T}}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(V_\tau^v - I\right)\right]. \tag{23}$$

Using (9), a standard computation based on Girsanov theorem and on the formula for the Laplace transform of a trigger stopping time for a Brownian motion with drift yields the following result.

Lemma 5.4 *For any* $(v,p) \in \mathbb{R} \times [0,1]$,

$$\overline{G}(v,p) = \max_{\overline{b} \geq v} \Gamma_{\overline{b}}(v,p), \tag{24}$$

where, for any $\overline{b} \in \mathbb{R}$:

$$\Gamma_{\overline{b}}(v,p) = \left(p \exp\left(\frac{v-\overline{b}}{\sigma^2}\, f(1)\right) + (1-p) \exp\left(\frac{v-\overline{b}}{\sigma^2}\, f(0)\right)\right)(\overline{b} - I). \tag{25}$$

The expression (25) for the maximand in (24) admits a natural interpretation. When choosing an optimal trigger $\overline{b}$ in $\mathcal{P}$, the decision maker maximizes a weighted average of the payoffs he would get from playing $\overline{b}$ if he knew the true value of μ, where the weights $1-p$ and p reflect his prior beliefs about μ.

It is clear from (25) that the choice of an optimal trigger in (24) depends on the the initial value v of the project as well as on the prior belief p about the quality of the project. Thus this choice is not time-consistent: a rational investor constrained to trigger strategies would like to revise his optimal trigger strategy as the value changes and new information about μ becomes available.

Formulas (24)-(25) does not allow to derive a closed-form solution for the optimal trigger $\overline{b}(v,p)$ conditional on the initial state (v,p). However, one can unambiguously compare $\overline{b}(v,p)$ with the optimal trigger $\widehat{b}(p)$ for the average drift problem, at least for some values of v and p. Specifically, let $\widetilde{b}(p) = I + \sigma^2/(pf(1) + (1-p)f(0))$. The following result is proved in the Appendix.

Proposition 5.1 *For any* $p \in (0,1)$ *and* $v \geq \widetilde{b}(p)$, $\overline{b}(v,p) = v$.

Since f is convex and positive, $\widetilde{b}(p) < \widehat{b}(p)$ whenever $p \in (0,1)$. Hence, for any such p and $v \in [\widetilde{b}(p), \widehat{b}(p))$, $\overline{b}(v,p) < \widehat{b}(p)$ and thus $\overline{G}(v,p) = v - I < \widehat{G}(v,p)$. For these values of (v,p), the decision maker, when constrained to use trigger strategies, invests as if the drift was constant and equal to $f^{-1}(\sigma^2/(v-I)) < p$, making him just ready to invest immediately. This reflects an implicit risk-aversion due to the additional uncertainty generated by the randomness of μ.

Clearly, Proposition 5.1 does not us allow to compare $\mathcal{P}$ and $\widehat{\mathcal{P}}$. In a sense, this is not surprising: the restriction to trigger strategies in $\mathcal{P}$ essentially amounts to deprive the decision maker from the benefits of learning about μ.

6. A Local Study of the Investment Boundary

While the comparison result of Section 5 provides us with a useful lower bound for the value function G^* of problem $\mathcal{P}$, it does not yield any information about the shape of the investment boundary b^*, nor about the wedge between the incomplete information problem and the average drift problem. We now address these and related questions in the case where the volatility σ of the observation/value process is small.

6.1. A Strict Comparison Result

In the remaining of the paper, we systematically index the value processes V and $\widehat{V}$, the value functions G^* and $\widehat{G}$, and the investment boundary functions b^* and $\widehat{b}$ by the volatility parameter σ. Our objective in this section is to compare G^*_σ and $\widehat{G}_\sigma$, as well as the investment boundaries b^*_σ and $\widehat{b}_\sigma$ for small values of σ. Specifically, consider the open domain $\mathcal{D} = (I, I + 1/r) \times (0,1)$. Our discussion will be based on the following strict comparison result.

Theorem 6.1 *For any compact subset $K \subset \mathcal{D}$, there exists $\sigma_K > 0$ such that for any $\sigma \in (0, \sigma_K]$, $G^*_\sigma(v,p) > \widehat{G}_\sigma(v,p)$ for any $(v,p) \in K$.*

The proof of Theorem 6.1 is based on three uniform convergence lemmas. First, we study the properties of the solution to the average drift problem $\widehat{\mathcal{P}}$ in the neighborhood of $\sigma = 0$. Next, we introduce an auxiliary problem $\widetilde{\mathcal{P}}$ which we show to be uniformly equivalent to $\mathcal{P}$ in the neighborhood of $\sigma = 0$. We conclude by proving that $\widetilde{\mathcal{P}}$ yields a strictly higher value than $\widehat{\mathcal{P}}$ to the decision maker.

Local Study of $\widehat{\mathcal{P}}$. First, we study the behavior of $\widehat{G}_\sigma$ when the variance σ of $\widehat{V}_\sigma$ becomes arbitrarily small. When $\sigma = 0$, solving $\widehat{\mathcal{P}}$ simply amounts to find a maximum of $t \mapsto \mathrm{e}^{-rt}(v + pt - I)$, yielding $\widehat{t}_0(v,p) = \max\{0, (I-v)/p + 1/r\}$ whenever $p > 0$ and $\widehat{t}_0(v,0) = 0$ whenever $v > I$. It is then straightforward to check that:

$$\widehat{G}_0(v,p) = \begin{cases} \exp\left(-1 + \dfrac{r}{p}(v - I)\right)\dfrac{p}{r} & \text{if} \quad v < I + \dfrac{p}{r} \\ \\ v - I & \text{if} \quad v \geq I + \dfrac{p}{r} \end{cases} \tag{26}$$

for any $(v,p) \in \mathcal{D}$. We then have the following uniform convergence result.

Lemma 6.1 $\lim_{\sigma \to 0} \widehat{G}_\sigma(v,p) - \widehat{G}_0(v,p) = 0$ *uniformly on any compact subset of* $\mathcal{D}$.

Proof: Pointwise convergence follows immediately from (20) and (26) together with the fact that $\lim_{\sigma \to 0} \sigma^2 / f_\sigma(p) = p/r$ by L'Hôpital rule. Next, for any $p \in [0,1]$, the quantity $\sigma^2/f_\sigma(p)$ is an increasing function of σ. Since $\widehat{G}_0$ is continuous and the mapping $x \mapsto \exp(-1 + (v-I)/x)$ is increasing on $[v - I, \infty)$ for any $v \in \mathbb{R}$, the result follows immediately from Dini's theorem. ■

A key observation is that, for any initial value v, $\widehat{G}_0(v,p)$ is convex in p. In particular, $p\,\widehat{G}_0(v,1) + (1-p)\,\widehat{G}_0(v,0) > \widehat{G}_0(v,p)$ for any $(v,p) \in \mathcal{D}$.

An Auxiliary Problem. To compare G^*_σ and $\widehat{G}_\sigma$, it will be helpful to consider an auxiliary optimal stopping problem that differs from $\widetilde{\mathcal{P}}$ only in that the Gaussian component of the value is omitted in the decision maker's payoff:

$$\sup_{\tau \in \mathcal{T}^{V_\sigma}} \mathbb{E}\left[\mathrm{e}^{-r\tau}(v + \mu\tau - I)\right]. \tag{27}$$

We shall call this problem $\widetilde{\mathcal{P}}$. For any v, and for any prior belief p that $\mu = 1$, we denote by $\widetilde{G}_\sigma(v,p)$ the value of the supremum in (27). Note that the information structure is the same in $\mathcal{P}$ and $\widetilde{\mathcal{P}}$. This leads immediately to the following result.

Lemma 6.2 $\lim_{\sigma \to 0} G^*_\sigma(v,p) - \widetilde{G}_\sigma(v,p) = 0$ *uniformly on* $\mathbb{R} \times [0,1]$.

Proof: From (6) and (27), we have, for each $(v,p) \in \mathbb{R} \times [0,1]$,

$$\begin{aligned} \left| G^*_\sigma(v,p) - \widetilde{G}_\sigma(v,p) \right| &\leq \sigma \sup_{\tau \in \mathcal{T}^{V_\sigma}} \mathbb{E}\left[\mathrm{e}^{-r\tau}|W_\tau|\right] \\ &\leq \sigma \sup_{\tau \in \mathcal{T}^{W}} \mathbb{E}\left[\mathrm{e}^{-r\tau}|W_\tau|\right], \end{aligned} \tag{28}$$

where the second inequality takes advantage from the inclusion $\mathcal{F}^{V_\sigma} \subset \mathcal{F}^{\mu,W}$ and the independence of μ and W. For any $a > 0$, let $T_{(-a,a)} = \inf\{t \geq 0 \mid |W_t| \geq a\}$. From standard optimal stopping theory,

$$\sup_{\tau \in \mathcal{T}^W} \mathbb{E}\left[\mathrm{e}^{-r\tau}|W_\tau|\right] = \sup_{a>0} a\,\mathbb{E}\left[\mathrm{e}^{-rT_{(-a,a)}}\right] = \frac{a^*}{\cosh\left(\sqrt{2r}\,a^*\right)},$$

where a^* is the unique positive solution of $\sqrt{2r}\,a \tanh\left(\sqrt{2r}\,a\right) = 1$ (see Karatzas and Shreve (1991, §2.8.C)). The result follows then immediately from (28). ■

It is worth noting that, given a prior belief p that $\mu = 1$, it is possible to secure the payoff $\widehat{G}_0(v,p)$ in $\mathcal{P}$ or $\widetilde{\mathcal{P}}$ by mimicking the optimal strategy in $\widehat{\mathcal{P}}$ when $\sigma = 0$, i.e., by delaying investment by the deterministic time $\widehat{t}_0(v,p)$. This follows immediately from the linearity of the payoff functions in (6) and (27) with respect to μ.

Local Comparison Between $\widehat{P}$ and $\widetilde{P}$. We now prove that, as σ goes to zero, $\widetilde{G}_\sigma$ converges uniformly to a function that is a strict upper bound for $\widehat{G}_0$ on $\mathcal{D}$, which implies Theorem 6.1 given the uniform convergence results of Lemmas 6.1 and 6.2. Specifically, we have the following result.

LEMMA 6.3 $\lim_{\sigma\to 0} \widetilde{G}_\sigma(v,p) = p\,\widehat{G}_0(v,1) + (1-p)\,\widehat{G}_0(v,0)$ *uniformly on any compact subset of $\mathcal{D}$.*

PROOF: Note first that, for any $\sigma > 0$ and for any prior belief $p \in [0,1]$ that $\mu = 1$, we have, from (27) and the inclusion $\mathcal{F}^{V_\sigma} \subset \mathcal{F}^{\mu,W}$,

$$\begin{aligned}\widetilde{G}_\sigma(v,p) &\leq \sup_{\tau \in \mathcal{T}^{\mu,W}} \mathbb{E}\left[\mathrm{e}^{-r\tau}(v + \mu\tau - I)\right] \\ &= p\,\widehat{G}_0(v,1) + (1-p)\,\widehat{G}_0(v,0).\end{aligned}$$

Thus, $\limsup_{\sigma\to 0} \widetilde{G}_\sigma(v,p) \leq p\,\widehat{G}_0(v,1) + (1-p)\,\widehat{G}_0(v,0)$ for any $(v,p) \in \mathbb{R} \times [0,1]$. Conversely, consider the following strategy in $\widetilde{\mathcal{P}}$. First, wait for a deterministic time $\varepsilon \in (0, \widehat{t}_0(v,p))$. Next, delay further investment by $\widehat{t}_0\left(v + \varepsilon P^p_\varepsilon, P^p_\varepsilon\right)$, i.e., the amount of time that is optimal in $\widetilde{\mathcal{P}}$ in the state $(v + \varepsilon P^p_\varepsilon, P^p_\varepsilon)$. We have:

$$\begin{aligned}\widetilde{G}_\sigma(v,p) &\geq \mathbb{E}\left[\mathrm{e}^{-r\left(\varepsilon + \widehat{t}_0(v+\varepsilon P^p_\varepsilon, P^p_\varepsilon)\right)}\left(v + \mu\left(\varepsilon + \widehat{t}_0\left(v + \varepsilon P^p_\varepsilon, P^p_\varepsilon\right)\right) - I\right)\right] \\ &= \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\mathbb{E}\left[\mathrm{e}^{-r\widehat{t}_0(v+\varepsilon P^p_\varepsilon, P^p_\varepsilon)}\left(v + \mu\,\varepsilon + \mu\,\widehat{t}_0\left(v + \varepsilon P^p_\varepsilon, P^p_\varepsilon\right) - I\right) \mid \mathcal{F}^{V_\sigma}_\varepsilon\right]\right] \quad (29) \\ &= \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \varepsilon P^p_\varepsilon, P^p_\varepsilon\right)\right],\end{aligned}$$

where the last equality follows from the definition of $\widehat{G}_0$. Rewriting (11) and making the dependence of P^p on σ explicit, we get:

$$P^p_\varepsilon(\sigma) = \frac{p\,\exp\left(\dfrac{\sigma W_\varepsilon + \left(\mu - \frac{1}{2}\right)\varepsilon}{\sigma^2}\right)}{p\,\exp\left(\dfrac{\sigma W_\varepsilon + \left(\mu - \frac{1}{2}\right)\varepsilon}{\sigma^2}\right) + 1 - p},$$

from which we get that $\lim_{\sigma\to 0} P^p_\varepsilon(\sigma) = \mu$, $\mathbb{P}$–almost surely. As $\widehat{G}_0\left(v + \varepsilon P^p_\varepsilon(\sigma), P^p_\varepsilon(\sigma)\right)$ is positive and upper bounded by $\widehat{G}_0\left(v+\varepsilon, 1\right)$ for any $\sigma > 0$, it follows from (29) and Lebesgue's dominated convergence theorem that:

$$\begin{aligned}\liminf_{\sigma\to 0}\, \widetilde{G}(v,p) \;&\geq\; \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \mu\,\varepsilon, \mu\right)\right] \\ &=\; \mathrm{e}^{-r\varepsilon}\left(p\,\widehat{G}_0\left(v+\varepsilon, 1\right) + (1-p)\,\widehat{G}_0\left(v, 0\right)\right).\end{aligned}$$

Since $\widehat{G}_0$ is continuous, we get that $\liminf_{\sigma\to 0}\, \widetilde{G}(v,p) \geq p\,\widehat{G}_0(v,1) + (1-p)\,\widehat{G}_0(v,0)$ by letting ε go to 0. Since the reverse inequality holds for the lim sup, pointwise convergence follows. To prove that this convergence is uniform, note first from (3) that, for any $\sigma > 0$, we can rewrite the dynamics of $P^p(\sigma)$ as:

$$dP^p_t(\sigma) = \frac{P^p_t(\sigma)\left(1 - P^p_t(\sigma)\right)\left(\mu - P^p_t(\sigma)\right)}{\sigma^2}\,dt + \frac{P^p_t(\sigma)\left(1 - P^p_t(\sigma)\right)}{\sigma}\,dW_t; \quad t \geq 0.$$

Therefore, from the time-change theorem for diffusion processes (see Øksendal (1995, Theorem 8.11)), $P^p_t(\sigma)$ coincides in law with $P^p_{t/\sigma^2}(1)$ for any $t \geq 0$ and $\sigma > 0$. It follows that, for any $\varepsilon \in (0, \widehat{t}_0(v,p))$ and $\sigma > \widetilde{\sigma} > 0$,

$$\begin{aligned}\mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \varepsilon P^p_\varepsilon(\widetilde{\sigma}), P^p_\varepsilon(\widetilde{\sigma})\right)\right] \;&=\; \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \varepsilon P^p_{\varepsilon/\widetilde{\sigma}^2}(1), P^p_{\varepsilon/\widetilde{\sigma}^2}(1)\right)\right] \\ &=\; \mathbb{E}\left[\mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \varepsilon P^p_{\varepsilon/\widetilde{\sigma}^2}(1), P^p_{\varepsilon/\widetilde{\sigma}^2}(1)\right) \mid \mathcal{F}^{V_1}_{\varepsilon/\sigma^2}\right]\right] \\ &\geq\; \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \varepsilon P^p_{\varepsilon/\sigma^2}(1), P^p_{\varepsilon/\sigma^2}(1)\right)\right] \\ &=\; \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v + \varepsilon P^p_\varepsilon(\sigma), P^p_\varepsilon(\sigma)\right)\right],\end{aligned}$$

where the first and last equalities follow from the above time-change argument, and the inequality from the fact that $P^p(1)$ is a martingale and $\widehat{G}_0$ is convex as the supremum of linear functions of (v,p), together with Jensen inequality. Hence, the mapping $\sigma \mapsto \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v+\varepsilon P^p_\varepsilon(\sigma), P^p_\varepsilon(\sigma)\right)\right]$ is decreasing. Thus, by Dini's theorem,

$$\lim_{\sigma\to 0}\mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v+\varepsilon P^p_\varepsilon(\sigma), P^p_\varepsilon(\sigma)\right)\right] = \mathbb{E}\left[\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0\left(v+\mu\,\varepsilon,\mu\right)\right] \tag{30}$$

uniformly on every compact of $\mathcal{D}$. Note that this holds for every $\varepsilon \in (0,\widehat{t}_0(v,p))$. From (26), the right-hand side of (30) is equal to $p\,\widehat{G}_0(v,1) + (1-p)\,\mathrm{e}^{-r\varepsilon}\,\widehat{G}_0(v,0)$ whenever $\varepsilon \in (0,\widehat{t}_0(v,p))$. As ε goes to 0, this converges uniformly to $p\,\widehat{G}_0(v,1)+(1-p)\,\widehat{G}_0(v,0)$ on every compact of $\mathcal{D}$. The result follows. ■

This result relies on the following intuition. As already mentioned, the decision maker can always secure the payoff $\widehat{G}_0(v,p)$ in $\widetilde{\mathcal{P}}$. In fact, he can do strictly better by first waiting for a deterministic time interval $\varepsilon > 0$ and only then playing the optimal strategy in $\widehat{\mathcal{P}}$ for $\sigma = 0$, conditional on the expected value of $(v+\mu\,\varepsilon,\mu)$ at time ε. The delay ε corresponds to a learning phase, during which the decision maker accumulates information about μ before taking his decision. As σ goes to zero, the accuracy of the value process as a signal of μ becomes infinite, so this learning phase can be made arbitrarily short. In the limit, everything happens as if the decision maker knew exactly the value of μ at date zero, which implies pointwise convergence.

The fact that convergence is uniform in (v,p) on compact subsets of $\mathcal{D}$ follows from a time-change argument. For a fixed duration ε of the learning phase, reducing the variance of the observation process effectively amounts to increasing the duration of the learning phase while keeping the variance constant. Since beliefs follow a martingale, this generates a mean-preserving spread in the decision maker's beliefs. As his expected gain at the end of the learning phase is convex in $(v+\varepsilon\,P_\varepsilon, P_\varepsilon)$, the payoff from this investment strategy increases as σ decreases to zero, which implies the result. Theorem 6.1 is then an immediate consequence of Lemmas 6.1, 6.2, and 6.3, together with the fact that $p\,\widehat{G}_0(v,1)+(1-p)\,\widehat{G}_0(v,0) > \widehat{G}_0(v,p)$ for any $(v,p)\in\mathcal{D}$.

6.2. Comments and Interpretation

The previous results allow us to study the qualitative properties of the solution to $\mathcal{P}$ for small values of the variance σ of the observation/value process.

Three Investment Problems. An immediate consequence of Lemmas 6.1, 6.2, and 6.3 is that $\lim_{\sigma\to 0} G^*_\sigma(v,p) - (p\,\widehat{G}_\sigma(v,1)+(1-p)\,\widehat{G}_\sigma(v,0)) = 0$ uniformly on any compact subset of $\mathcal{D}$. This means that, as σ goes to zero, the value of the incomplete information problem $\mathcal{P}$ converges uniformly to the value of an investment problem where the drift μ of the value process is first selected according to a lottery on 0 and 1 with respective probabilities $1-p$ and p, and immediately revealed to the decision maker, who then takes his investment decision under complete information about μ. In other terms, the loss in value arising from the need to learn about μ in $\mathcal{P}$ vanishes as the variance of the observation process converges to 0. The fact that the incomplete information problem $\mathcal{P}$ is strictly preferred by the decision maker to the average drift problem $\widehat{\mathcal{P}}$ for small values of σ simply reflects the fact that the value function $\widehat{G}_\sigma$ of $\widehat{\mathcal{P}}$ is strictly convex on $\mathcal{D}$ with respect to p and that learning about μ in $\mathcal{P}$ is fast when σ is small.

Local Comparison of $\widehat{b}_\sigma$ and b^*_σ. From Lemma 5.1, the optimal investment strategy in the average drift problem $\widehat{\mathcal{P}}$ with drift p consists to delay investment until the value $\widehat{V}$ hits the threshold $\widehat{b}_\sigma(p)$. It is easy to check that, as σ goes to zero, $\widehat{b}_\sigma$ converges monotonically from above to the mapping $p \mapsto I + p/r$. This implies that any pair (v, p) in the triangle $\mathcal{D}' = \{(v, p) \in \mathcal{D} \mid v > I + p/r\}$ satisfies $\widehat{G}_\sigma(v, p) = v - I$ for σ close enough to 0. It follows then from Theorem 6.1 that $G^*_\sigma(v, p) > v - I$, i.e., that (v, p) belongs to the continuation region of $\mathcal{P}$ for σ close enough to 0. Moreover, this reasoning can be made uniform on compact subsets of $\mathcal{D}$.

Corollary 6.1 *For any compact subset $K \subset \mathcal{D}'$, there exists $\sigma_K > 0$ such that for any $\sigma \in (0, \sigma_K]$, $G^*_\sigma(v, p) > v - I = \widehat{G}_\sigma(v, p)$ for any $(v, p) \in K$.*

We can actually characterize exactly the asymptotic behavior of the investment boundary function b^*_σ as σ converges to 0.

Corollary 6.2 *The following holds:*

(i) $\lim_{\sigma \to 0} b^*_\sigma(0) = I$;

(ii) For any $p \in (0, 1]$, $\lim_{\sigma \to 0} b^*_\sigma(p) = I + 1/r$.

That is, the optimal investment boundary function b^*_σ converges pointwise to the discontinuous function $b^*_0(p) = I + \chi_{(0,1]}(p)/r$ as σ goes to 0. The proof simply consists to apply Corollary 6.1 to an increasing sequence $\{K_n\}$ of compact subsets of $\mathcal{D}'$ such that $\bigcup_{n=0}^{\infty} K_n = \mathcal{D}'$. To interpret this result, consider the limit problem arising from $\mathcal{P}$ when the Gaussian component is omitted both in the decision maker's information and in his payoff:

$$G_0^*(v,p) = \sup_{\tau \in \mathcal{T}^{V_0}} \mathbb{E}\left[\mathrm{e}^{-r\tau}\,(v + \mu\tau - I)\right]. \tag{31}$$

Since there is no noise in the signal V_0, all learning about μ takes place at date zero, as the belief process jumps instantaneously to one of its absorbing barriers, 0 (with probability $1-p$), or 1 (with probability p). If $v \geq I$ then, in the first case, it is optimal to invest immediately, while in the second, it is optimal to wait until the value reaches the threshold $I + 1/r$. Thus b_0^* can be roughly interpreted as the investment boundary for the limit problem (31). Note that $G_0^*(v,p) = p\,\widehat{G}_\sigma(v,1) + (1-p)\,\widehat{G}_\sigma(v,0)$, in accordance with Lemma 6.3. The area $\mathcal{D}'$ represents the discrepancy between the incomplete information problem and the average drift problem as σ converges to 0.

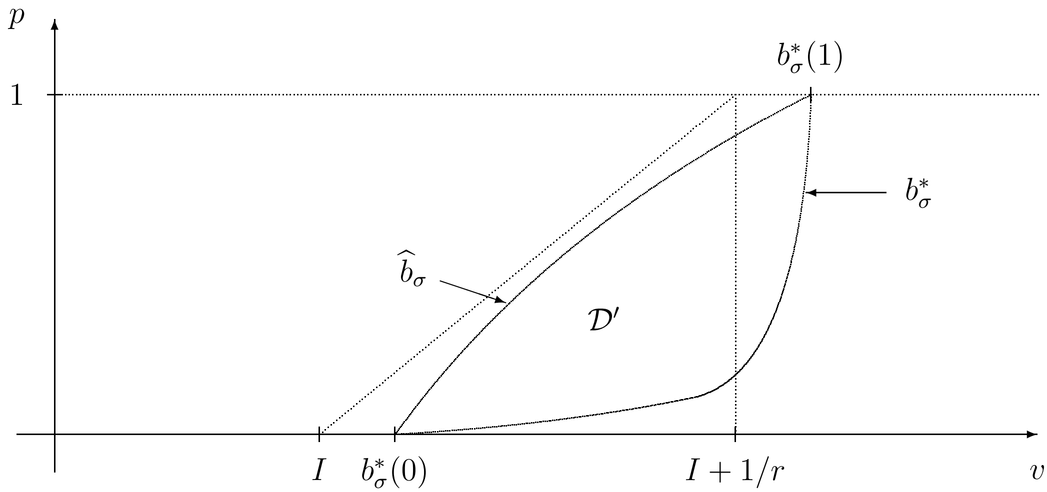

FIGURE 1: Local Comparison Between $\mathcal{P}$ and $\widehat{\mathcal{P}}$

Our results are illustrated on Figure 1. While $\widehat{b}_\sigma$ is unambiguously convex, the concave shape of b_σ^* is only meant to be suggestive. Note, however, that in virtue of Corollary 6.2, b_σ^* cannot be globally convex on $[0,1]$ when σ is small. This implies in particular that the value function G_σ^* is not globally convex with respect to the value/belief pair. Indeed, if it were, then, for any two points (v,p) and $(\widetilde{v},\widetilde{p})$ on the investment boundary b_σ^*, and for any convex combination (v_λ, p_λ) of these points, we would have $G_\sigma^*(v_\lambda, p_\lambda) \leq \lambda\, G_\sigma^*(v,p) + (1-\lambda)\, G_\sigma^*(\widetilde{v},\widetilde{p}) = v_\lambda - I$, so that $G_\sigma^*(v_\lambda, p_\lambda) = v_\lambda - I$ and (v_λ, p_λ) would belong to the investment region as well. But this can only hold if b_σ^* is globally convex on $[0,1]$, a contradiction.

A Non-Monotonocity Result. The convergence of the investment boundary b^*_σ to the discontinuous function b^*_0 as σ goes to zero has a striking consequence. Indeed, consider some $(v, p) \in \mathcal{D}'$ such that, for some $\sigma > 0$, $G^*_\sigma(v, p) = v - I$, i.e., it is optimal to invest in the state (v, p) when the observation/value process has variance σ. It is clearly possible to find a triple (v, p, σ) satisfying this condition, since b^*_σ is continuous with respect to p and $b^*_\sigma(0)$ is close to I for small enough σ. But since $v < I + 1/r$ and $p > 0$, Corollary 6.2 implies that, for all $\widetilde{\sigma} < \sigma$ that are close enough to 0, $b^*_{\widetilde{\sigma}}(p) > v$ and thus $G^*_{\widetilde{\sigma}}(v, p) > v - I$. In other terms, the value of problem $\mathcal{P}$ can be a decreasing function of the volatility σ of the value process, at least locally.

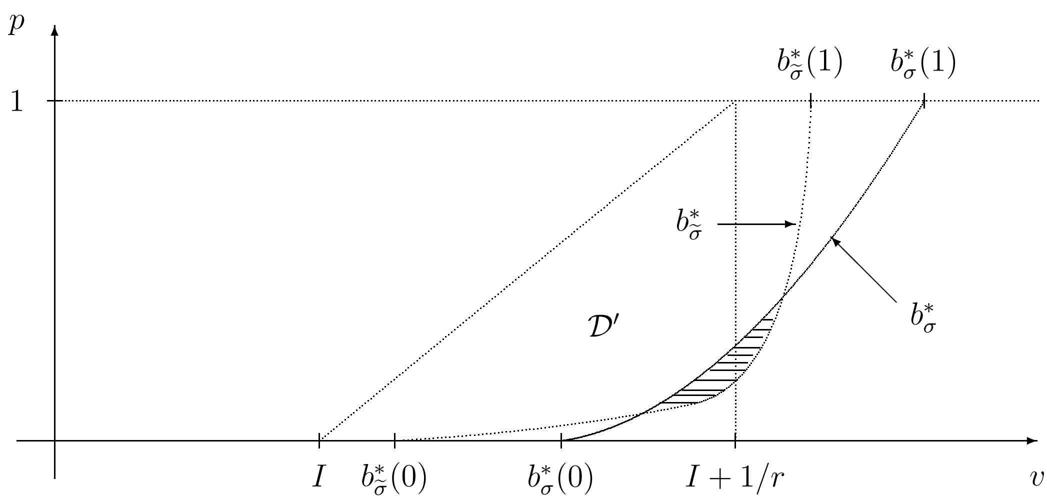

FIGURE 2: Non-Monotonocity with Respect to σ

Our findings are illustrated on Figure 2. Here, $\widetilde{\sigma} < \sigma$, and, on the hatched zone, $G^*_{\widetilde{\sigma}}(v, p) > G^*_\sigma(v, p) = v - I$. Again, the exact shapes of b^*_σ and $b^*_{\widetilde{\sigma}}$, as well as the fact that they cross only twice, are only meant to be suggestive.

This non-monotonicity result contrasts sharply with the predictions of standard real option models (see Dixit and Pindyck (1994, §5.4)). There, a greater uncertainty (in the sense of a higher σ) typically increases the value of a firm's investment opportunity, and increases the critical value at which investment takes place by raising the opportunity cost of exercising the option to invest. A similar result also holds for the average drift problem, as $\widehat{b}_\sigma$ and therefore $\widehat{G}_\sigma$ are clearly increasing in σ. The intuition is that when the volatility increases, the decision maker can achieve a higher exposition to upside realizations of the value by increasing his investment trigger, while being protected from downside risk.

In the incomplete information problem, by contrast, the decision maker is not protected from downside risk, since the investment boundary is not flat in p. An increase in σ has thus an ambiguous effect on the value of the option to invest, because of the interplay between two opposite effects. On the one hand, an increase in σ raises the volatility of the decision maker's payoff at the time of the investment, which tends to increase the value of the option to invest; one might call this standard effect the *real option effect* by analogy with the complete information case. On the other hand, a raise in σ decreases the volatility of the belief process, which tends to impede learning about μ, and thus to reduce the opportunity cost to exert the option to invest. One might call this countervailing effect the *inefficient learning effect*. On the hatched zone in Figure 2, the inefficient learning effect clearly dominates. In this zone, a reduction in σ is likely to delay investment. This casts some doubt on the effectiveness of policies aiming at promoting investment by reducing the level of uncertainty, at least when such a reduction in uncertainty facilitates learning.

Overall, the impact of an increase in σ on the value of the option to invest depends on which of the real option and the inefficient learning effects dominates. It is difficult to map precisely the parameter space in terms of this distinction. Note however that, since $b^*_\sigma(0)$ and $b^*_\sigma(1)$ are both increasing functions of σ and $b^*_\sigma(p)$ is a continuous non-decreasing function of p, $b^*_\sigma(p)$ must at least be locally increasing in σ in neighborhoods of 0 and 1, as well as $G^*_\sigma(v, p)$ for v close enough to $b^*_\sigma(p)$ (see Figure 2 for an illustration of this effect). Intuitively, if the decision maker is already fairly confident in his estimate of μ, an increase in σ will only have a marginal impact on the efficiency of learning, since his beliefs are unlikely to change very fast anyway. The increased volatility of his payoff make him however willing to delay his investment further, thereby increasing the value of his option to invest. In that case, the real option effect compensates for the decreased efficiency of learning.

7. Concluding Remarks

This paper has focused on the qualitative properties of the optimal decision to invest in a project whose value is observable but driven by a parameter that is unknown to the decision maker ex ante. We have shown that the optimal investment strategy is characterized by a continuous and non-decreasing boundary in the value/belief state space. The presence of learning implies that the optimal investment strategy is path-dependent. In particular, the value of the project at the time of the investment does not necessarily coincide with its historic maximum.

We have shown that the decision maker always benefit from being uncertain about the drift of the value process. That is, he prefers the option to invest in a project with unknown drift to that of investing in a project with a constant drift equal to the prior expectation of the drift in the first option. Thus one might expect the value of claims on structurally uncertain assets—e.g., in an emerging sector in which future growth prospects are uncertain—to be higher than that of claims on assets in more traditional sectors with otherwise identical risk characteristics.

A significant point of departure with the standard real option model is that the value of the option to invest is not everywhere increasing with respect to the volatility of the value process. Thus, while drift uncertainty always benefit a risk neutral investor, non-structural uncertainty might prove harmful. As we argued, this non-monotonicity can be interpreted in terms of two countervailing effects: the real option effect and the inefficient learning effect.

Appendix

Proof of Lemma 3.1: Note that for any $t \geq 0$, $P_t^0 = \int_0^t P_s^0 (1 - P_s^0)\, d\overline{W}_s$ satisfies:

$$\mathbb{E}\left[(P_t^0)^2\right] = \int_0^t \mathbb{E}\left[(P_s^0)^2(1 - P_s^0)^2\right] ds \leq \int_0^t \mathbb{E}\left[(P_s^0)^2\right] ds.$$

Thus, by Gronwall's lemma, we obtain that $\mathbb{E}\left[(P_t^0)^2\right] = 0$, and therefore $P_t^0 = 0$, $\mathbb{P}$–almost surely. Since $1 - P_t^1 = \int_0^t P_s^0 (1 - P_s^0)\, d\overline{W}_s$, this argument also implies that $P_t^1 = 1$, $\mathbb{P}$–almost surely. Part (ii) is a direct application of Feller's test for explosions (Karatzas and Shreve (1991, Theorem 5.5.29). ■

Proof of Proposition 3.1: Note first that, since immediate stopping is always a feasible strategy, the supremum in (6) must be non-negative. Next, from (5) and the definition of g, it follows that, for all $(v,p) \in \mathbb{R} \times [0,1]$,

$$\begin{aligned} G^*(v,p) &= \sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(v + \sigma\overline{W}_\tau + \int_0^\tau P_t^p\, dt - I\right)\right] \\ &\leq \sup_{\tau \in \mathcal{T}^V} \mathbb{E}\left[\mathrm{e}^{-r\tau}\left(v + \tau + \sigma\overline{W}_\tau - I\right)\right] \qquad (32) \\ &= \mathbb{E}\left[\mathrm{e}^{-r\tau^{\max}}\left(v + \tau^{\max} + \sigma\overline{W}_{\tau^{\max}} - I\right)\right], \end{aligned}$$

where $\tau^{\max} = \inf\{t \geq 0 \mid v + t + \sigma\overline{W}_t \geq b^{\max}\}$ for some $b^{\max} > 0$ that can be explicitly computed given that $\overline{W}$ is a Brownian motion under $\mathcal{F}^V$, see Lemma 5.1. It is easy to check that $b^{\max} = b^*(1)$ and that the right-hand side of (32) coincides with $G^*(v,1)$. In particular, G^* is well-defined. Let $\tau^* = \inf\{t \geq 0 \mid X_t^{v,p} \notin \mathcal{C}^*\}$, where $\mathcal{C}^* = \{(v,p) \in \mathbb{R} \times [0,1] \mid G^*(v,p) > g(v,p)\}$ is the continuation region for our problem. Since the family of random variables $\{\mathrm{e}^{-r\tau}\overline{W}_\tau^- \mid \tau \in \mathcal{T}^V\}$ is uniformly integrable, a sufficient condition for:

$$G^*(v,p) = \mathbb{E}\left[\mathrm{e}^{-r\tau^*} g(X_{\tau^*}^{v,p})\right] = \mathbb{E}\left[\mathrm{e}^{-r\tau^*}\left(v + \sigma\overline{W}_{\tau^*} + \int_0^{\tau^*} P_t^p\, dt - I\right)\right]$$

is that τ^* be finite, $\mathbb{P}$–almost surely (see Øksendall (1995, Theorem 10.9, and the remark p. 195)). To prove this, note first that $\tau^* \leq T_{b^*(1)}^v = \inf\{t \geq 0 \mid v + \sigma\overline{W}_t + \int_0^t P_u^p\, du \geq b^*(1)\}$, $\mathbb{P}$–almost surely. Indeed, if not, then with positive $\mathbb{P}$–probability, $G^*\big(b^*(1), P^p_{T^v_{b^*(1)}}\big) > b^*(1) - I = G^*(b^*(1),1)$ by definition of τ^*, which contradicts (32). Since $T_{b^*(1)}^v \leq \inf\{t \geq 0 \mid v + \sigma\overline{W}_t \geq b^*(1)\}$ which is $\mathbb{P}$–almost surely finite, part (i) follows and $\mathcal{S}^* = \mathbb{R} \times [0,1] \setminus \mathcal{C}^* \neq \emptyset$. Last, suppose that $(v,p) \in \mathcal{S}^*$, so that $G^*(v,p) = v - I$. For any $h > 0$, discounting implies that $G^*(v+h,p) \leq G^*(v,p) + h = v + h - I$. Since the reverse inequality always hold as immediate stopping is always a feasible strategy, we obtain $G^*(v+h,p) = v + h - I$, hence $(v+h,p) \in \mathcal{S}^*$, which implies the second half of (ii). ■

Proof of Corollary 3.3: First, note from (6) that, for any $(v, v_0, p) \in \mathbb{R}^2 \times [0,1]$,

$$|G^*(v,p) - G^*(v_0,p)| \leq |v - v_0|. \tag{33}$$

Next, from (9), we have, for any $(v, p, p_0) \in \mathbb{R} \times [0,1]^2$,

$$\begin{aligned}
(p - p_0)\, \mathbb{E}_{\mathbb{Q}} &\left[\mathrm{e}^{-rT^*_{v,p_0}} \left(v + \sigma B_{T^*_{v,p_0}} - I \right) \left(\exp\left(\frac{B_{T^*_{v,p_0}}}{\sigma} - \frac{T^*_{v,p_0}}{2\sigma^2} \right) - 1 \right) \right] \\
&\leq\ G^*(v,p) - G^*(v,p_0) \\
&\leq\ (p - p_0)\, \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-rT^*_{v,p}} \left(v + \sigma B_{T^*_{v,p}} - I \right) \left(\exp\left(\frac{B_{T^*_{v,p}}}{\sigma} - \frac{T^*_{v,p}}{2\sigma^2} \right) - 1 \right) \right].
\end{aligned}$$

By Corollary 3.1, the mapping $p \mapsto G^*(v,p)$ is non-decreasing. Hence:

$$\begin{aligned}
\left| \frac{G^*(v,p) - G^*(v,p_0)}{p - p_0} \right| &\leq \max_{\widetilde{p} \in \{p, p_0\}} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-rT^*_{v,\widetilde{p}}} \left(v + \sigma B_{T^*_{v,\widetilde{p}}} - I \right) \left(\exp\left(\frac{B_{T^*_{v,\widetilde{p}}}}{\sigma} - \frac{T^*_{v,\widetilde{p}}}{2\sigma^2} \right) - 1 \right) \right] \\
&\leq \sup_{\tau \in \mathcal{T}^V} \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} \left(v + \sigma B_\tau - I \right) \left(\exp\left(\frac{B_\tau}{\sigma} - \frac{\tau}{2\sigma^2} \right) + 1 \right) \right] \qquad (34) \\
&\leq G^*(v,0) + G^*(v,1),
\end{aligned}$$

where the second inequality follows from the fact that both $v + \sigma B_{T^*_{v,p}} - I$ and $v + \sigma B_{T^*_{v,p_0}} - I$ must be non-negative, $\mathbb{P}$–almost surely by (9), and the third from (9) again, applied respectively to $p = 1$ and $p = 0$. Using the two uniform upper bounds (33) and (34), we obtain that for any $(v, v_0, p, p_0) \in \mathbb{R}^2 \times [0,1]^2$,

$$\begin{aligned}
|G^*(v,p) - G^*(v_0,p_0)| &\leq |G^*(v,p) - G^*(v_0,p)| + |G^*(v_0,p) - G^*(v_0,p_0)| \\
&\leq |v - v_0| + \left(G^*(v_0,0) + G^*(v_0,1) \right) |p - p_0|,
\end{aligned}$$

which implies (i). To prove (ii), define, for any $t \geq 0$, $N_t = \exp\left(B_t/\sigma - t/2\sigma^2 \right) - \exp\left(B_t/\sigma_0 - t/2\sigma_0^2 \right)$, and note from (9) that, for all $(v, p, \sigma, \sigma_0) \in \mathbb{R} \times [0,1] \times \mathbb{R}^2_{++}$,

$$\begin{aligned}
\left| G^*_\sigma(v,p) - G^*_{\sigma_0}(v,p) \right| &\leq \sup_{\tau \in \mathcal{T}^B} \left| (\sigma - \sigma_0)\, \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} \left(p \exp\left(\frac{B_\tau}{\sigma} - \frac{\tau}{2\sigma^2} \right) + 1 - p \right) B_\tau \right] \right| \\
&\quad + p \sup_{\tau \in \mathcal{T}^B} \left| \mathbb{E}_{\mathbb{Q}} \left[\mathrm{e}^{-r\tau} N_\tau (v + \sigma_0 B_\tau - I) \right] \right|.
\end{aligned} \tag{35}$$

Define a probability measure $\mathbb{Q}^\sigma$ by its Radon-Nikodym derivative $d\mathbb{Q}^\sigma / d\mathbb{Q}_{|\mathcal{F}^B_t} = \exp\left(B_t/\sigma - t/2\sigma^2 \right)$ with respect to $\mathcal{F}^B$. A direct application of Girsanov theorem shows that the first term on the right-hand side of (35) is bounded above by $|\sigma-\sigma_0|\left(\sup_{\tau\in\mathcal{T}^{B^\sigma}}\mathbb{E}_{\mathbb{Q}^\sigma}\left[\mathrm{e}^{-r\tau}|B^\sigma_\tau|\right]+\sup_{t\geq 0}\mathrm{e}^{-rt}t\right)$, where B^σ is a standard Brownian motion under $\mathbb{Q}^\sigma$. This in turn clearly converges to 0 as σ converges to σ_0 since $\sup_{\tau\in\mathcal{T}^{B^\sigma}}\mathbb{E}_{\mathbb{Q}^\sigma}\left[\mathrm{e}^{-r\tau}|B^\sigma_\tau|\right]$ is finite and independent of σ, see the proof of Lemma 5.4. Using Jensen's inequality and defining a probability measure $\mathbb{Q}^{\sigma_0}$ and a Brownian motion B^{σ_0} under $\mathbb{Q}^{\sigma_0}$ by analogy with $\mathbb{Q}^\sigma$ and B^σ, a similar argument implies that, up to division by p, the second term on the right-hand side of (35) is no greater than $\sup_{\tau\in\mathcal{T}^{B^\sigma}}\mathbb{E}_{\mathbb{Q}^\sigma}\left[\mathrm{e}^{-r\tau}\left|v+\sigma_0/\sigma\,\tau+\sigma_0B_\tau-I\right|\right]+\sup_{\tau\in\mathcal{T}^{B^{\sigma_0}}}\mathbb{E}_{\mathbb{Q}^{\sigma_0}}\left[\mathrm{e}^{-r\tau}\left|v+\tau+\sigma_0B_\tau-I\right|\right]$, which can be shown to be finite along the same lines as in the proof of Lemma 5.4. The bounded convergence theorem implies that for any $\tau\in\mathcal{T}^B$,

$$\lim_{T\to\infty}\sup_{\tau\in\mathcal{T}^B}\left|\mathbb{E}_{\mathbb{Q}}\left[\mathrm{e}^{-r\tau}N_\tau(v+\sigma_0B_\tau-I)\,\chi_{\{\tau>T\}}\right]\right|=0. \tag{36}$$

Next, for any $(\tau,T)\in\mathcal{T}^B\times\mathbb{R}$, we obtain, from Jensen and Cauchy-Schwartz inequalities:

$$\begin{aligned}\left|\mathbb{E}\left[\mathrm{e}^{-r\tau}N_\tau(v+\sigma_0B_\tau-I)\,\chi_{\{\tau\leq T\}}\right]\right| &\leq \mathbb{E}\left[\mathrm{e}^{-r\tau}\left|N_\tau(v+\sigma_0B_\tau-I)\,\chi_{\{\tau\leq T\}}\right|\right]\\ &\leq \sqrt{\mathbb{E}\left[N_\tau^2\,\chi_{\{\tau\leq T\}}\right]\mathbb{E}\left[(v+\sigma_0B_\tau-I)^2\,\chi_{\{\tau\leq T\}}\right]} \qquad (37)\\ &\leq 2\sqrt{\mathbb{E}\left[N_T^2\right]\left((v-I)^2+4\sqrt{T}\,|v-I|\,\sigma_0+4T\sigma_0^2\right)},\end{aligned}$$

where the last step follows from $\mathbb{E}\left[\sup_{t\leq T}B_t\right]\leq\sqrt{\mathbb{E}\left[\left(\sup_{t\leq T}B_t\right)^2\right]}\leq\sqrt{\mathbb{E}\left[\sup_{t\leq T}B_t^2\right]}$ together with Doob inequality applied to the martingales B and N. It is easy to check from the definition of N that $N_T=\int_0^T\left((1/\sigma-1/\sigma_0)\exp\left(B_t/\sigma-t/2\sigma^2\right)+N_t/\sigma_0\right)dB_t$. From the Itô isometry and Gronwall inequality, we get, after some straightforward computations:

$$\mathbb{E}\left[N_T^2\right]\leq\frac{\sigma}{\sigma_0}\,\sigma\left(\frac{1}{\sigma}-\frac{1}{\sigma_0}\right)\left((\sigma+\sigma_0)\left(\mathrm{e}^{T/\sigma^2}-1\right)-2\,\sigma_0\left(\mathrm{e}^{T/\sigma\sigma_0}-1\right)\right)\mathrm{e}^{T/\sigma_0^2} \tag{38}$$

Let $\varepsilon>0$. By (36), there exists $T\in\mathbb{R}$ such that $\sup_{\tau\in\mathcal{T}^B}\left|\mathbb{E}_{\mathbb{Q}}\left[\mathrm{e}^{-r\tau}N_\tau(v+\sigma_0B_\tau-I)\,\chi_{\{\tau>T\}}\right]\right|<\varepsilon/2$. Similarly, from (37) and (38), $\sup_{\tau\in\mathcal{T}^B}\left|\mathbb{E}\left[\mathrm{e}^{-r\tau}N_\tau(v+\sigma_0B_\tau-I)\,\chi_{\{\tau\leq T\}}\right]\right|<\varepsilon/2$ for σ close enough to σ_0. It follows that the second term on the right-hand side of (35) converges to 0 as σ converges to σ_0, which concludes the proof of (ii). ■

Proof of Proposition 5.1: For any $(v,p)\in\mathbb{R}\times[0,1]$ and for any $\overline{b}\geq v$, one gets, using (25) and the definitions of $\widehat{b}$ and f:

$$\frac{\partial\Gamma_{\overline{b}}(v,p)}{\partial\overline{b}}=\frac{pf(1)}{\sigma^2}\exp\left(\frac{v-\overline{b}}{\sigma^2}\,f(1)\right)(\widehat{b}(1)-\overline{b})+\frac{(1-p)f(0)}{\sigma^2}\exp\left(\frac{v-\overline{b}}{\sigma^2}\,f(0)\right)(\widehat{b}(0)-\overline{b}). \tag{39}$$

It follows that $\operatorname{argmax}_{\overline{b}\geq v}\Gamma_{\overline{b}}(v,p)=\{v\}$ whenever $v\geq\widehat{b}(1)$, and that $\operatorname{argmax}_{\overline{b}\geq v}\Gamma_{\overline{b}}(v,p)\subset[\widehat{b}(0),\widehat{b}(1))$ otherwise. From now on, we focus on the latter case. Let us first show that $\overline{b}\mapsto\Gamma_{\overline{b}}(v,p)$ is quasi-concave on $[v, \widehat{b}(1)]$. If $\partial\Gamma_{\overline{b}}(v,p)/\partial\overline{b} < 0$ whenever $\overline{b} \in [v, \widehat{b}(1)]$, the result is immediate and $\operatorname{argmax}_{\overline{b}\geq v} \Gamma_{\overline{b}}(v,p) = \{v\}$. Otherwise, let $\overline{b} \in [v, \widehat{b}(1)]$ such that $\partial\Gamma_{\overline{b}}(v,p)/\partial\overline{b} = 0$. Then, from (39),

$$\begin{aligned} \frac{\partial^2\Gamma_{\overline{b}}(v,p)}{\partial\overline{b}^2} &\propto \left(\widehat{b}(1) - \overline{b}\right)\left(\overline{b} - \widehat{b}(0)\right)\frac{f(0)-f(1)}{\sigma^2} - \left(\widehat{b}(1) - \widehat{b}(0)\right) \\ &\leq \left(\widehat{b}(1) - \widehat{b}(0)\right)\left(\frac{\left(\widehat{b}(1) - \widehat{b}(0)\right)\left(f(0)-f(1)\right)}{4\sigma^2} - 1\right). \end{aligned} \tag{40}$$

A direct computation reveals that the right-hand side of (40) is negative for all $(r,\sigma) \in \mathbb{R}^2_{++}$, which implies the strict quasi-concavity of $\overline{b} \mapsto \Gamma_{\overline{b}}(v,p)$ on $[v, \widehat{b}(1)]$. By (39), for any $v \in [\widetilde{b}(p), \widehat{b}(1))$,

$$\left.\frac{\partial\Gamma_{\overline{b}}(v,p)}{\partial\overline{b}}\right|_{\overline{b}=v} = 1 + \frac{I-v}{\widetilde{b}(p)-I} < 0,$$

so that $\partial\Gamma_{\overline{b}}(v,p)/\partial\overline{b} < 0$ for each $\overline{b} \in [v, \widehat{b}(1))$ by strict quasi-concavity of $\overline{b} \mapsto \Gamma_{\overline{b}}(v,p)$ on $[v, \widehat{b}(1)]$. Therefore $\operatorname{argmax}_{\overline{b}\geq v} \Gamma_{\overline{b}}(v,p) = \{v\}$ in that case as well. ■

References

Alexander, D., and P. Veronesi (2000): "Option Prices with Uncertain Fundamentals: Theory and Evidence on the Dynamics of Implied Volatilities," working paper, Graduate Business School, University of Chicago.

Arrow, K., and A.C. Fisher (1974): "Environmental Preservation, Uncertainty, and Irreversibility," *Quarterly Journal of Economics*, 88, 312-319.

Bergemann, D., and J. Välimäki (1997): "Market Diffusion with Two-Sided Learning," *Rand Journal of Economics*, 28, 773-795.

——— (2000): "Experimentation in Markets," *Review of Economic Studies*, 67, 213-234.

Bolton, P., and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67, 349-374.

Brennan, M.J. (1998): "The Role of Learning in Dynamic Portfolio Decision," *European Finance Review*, 1, 295-306.

Chernoff, H. (1972): *Sequential Analysis and Optimal Design*, Philadelphia: Society for Industrial and Applied Mathematics.

Dixit, A.K. (1989): "Entry and Exit Decisions Under Uncertainty," *Journal of Political Economy*, 97, 620-638.

Dixit, A.K., and R.S. Pindyck (1994): *Investment under Uncertainty*, Princeton: Princeton University Press.

Felli, L., and C. Harris (1996): "Learning, Wage Dynamics, and Firm-Specific Human Capital," *Journal of Political Economy*, 104, 838-868.

Henry, C. (1974): "Investment Decisions Under Uncertainty: The Irreversibility Effect," *American Economic Review*, 64, 89-104.

Hu, Y., and B. Øksendal (1998): "Optimal Time to Invest when the Price Processes are Geometric Brownian Motions," *Finance and Stochastics*, 2, 295-310.

Jovanovic, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972-990.

Karatzas, I., and S. Shreve (1991): *Brownian Motion and Stochastic Calculus*, Berlin, Heidelberg, New York: Springer-Verlag.

Karatzas, I., and X. Zhao (1998): "Bayesian Adaptive Portfolio Optimization," working paper, Department of Mathematics, Columbia University.

Keller, G., and S. Rady (1999): "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, 66, 475-507.

Lackner, P. (1995): "Utility Maximization with Partial Information," *Stochastics Processes and their Applications*, 56, 247-273.

Liptser, R.S, and A.N. Shiryayev (1977): *Statistics of Random Processes I*, Berlin, Heidelberg, New York: Springer-Verlag.

McDonald, R., and D. Siegel (1986): "The Value of Waiting to Invest," *Quarterly Journal of Economics*, 101, 707-727.

McKean, H.P. (1965): "A Free Boundary Problem for the Heat Equation Arising from a Problem of Mathematical Economics," *Industrial Management Review*, 60, 32-39.

Merton, R. (1971): "Optimum Consumption and Portfolio Rules in a Continuous Time Model," *Journal of Economic Theory*, 3, 373-413.

Mordecki, E. (1999): "Optimal Stopping for a Diffusion with Jumps," *Finance and Stochastics*, 3, 227-236.

Moscarini, G., and L. Smith (2000): "The Optimal Level of Experimentation," forthcoming, *Econometrica.*

Øksendal, B. (1995): *Stochastic Differential Equations*, Berlin, Heidelberg, New York: Springer-Verlag.

Samuelson, P.A. (1965): "Rational Theory of Warrant Pricing," *Industrial Management Review*, 60, 13-32.

Shiryayev, A.N. (1978): *Optimal Stopping Rules*, Berlin, Heidelberg, New York: Springer Verlag.

Veronesi, P. (1999): "Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12, 975-1007.

——————— (2000): "How Does Information Quality Affect Stock Returns," *Journal of Finance*, 55, 807-837.

Villeneuve, S. (1999): "Exercise Regions of American Options on Several Assets," *Finance and Stochastics*, 3, 295-322.