

Investment Timing and Learning Externalities

Jean-Paul Décamps^{*†}

Thomas Mariotti^{‡§}

Revised: November 2003

^{*}Université de Toulouse I, GREMAQ, and IDEI, 21 Allée de Brienne, 31000 Toulouse, France.

[†]Corresponding author. Tel: +33-5-61-12-85-99. Fax: +33-5-61-12-86-37. *E-mail address:* decamps@cict.fr (J.-P. Décamps).

[‡]Department of Economics, London School of Economics and Political Science, Houghton Street, WC2A 2AE London, UK.

[§]Université de Toulouse I, CNRS-GREMAQ, and IDEI, 21 Allée de Brienne, 31000 Toulouse, France.

Abstract

We study a duopoly model of investment, in which each player learns about the quality of a common value project by observing some public background information, and possibly the experience of his rival. Investment costs are private information, and the background signal takes the form of a Poisson process conditional on the quality of the project being low. The resulting attrition game has a unique, symmetric equilibrium, which depends on initial public beliefs. We determine the impact of changes in the cost and signal distributions on investment timing, and how equilibrium is affected when a first-mover advantage is introduced.

JEL classification: C73; D82; D83

Keywords: Real Options; Duopoly; Incomplete Information; War of Attrition

1. Introduction

Uncertainty and irreversibility have long been recognized as key factors of investment (see Dixit and Pindyck [5] for a survey of this literature). Much attention has recently been devoted to the role of informational spillovers. The basic insight is that agents have a strategic incentive to delay investment in order to learn from their competitors. Rob [13] and Caplin and Leahy [2] develop models of industry equilibrium in which firms learn about the size of the market or the profitability of entry by observing past quantities and prices. Hendricks and Kovenock [8] and Chamley and Gale [3] study timing games in which the common value from exercising an investment option is correlated with players' private signals. Bolton and Harris [1] and Cripps, Keller and Rady [4] extend the classic two-armed bandit problem to a strategic setting in which players may learn from the experience of their rivals.

The purpose of this paper is to study the impact of asymmetric information about private investment costs in a duopoly model where a player's investment enhances the quality of the signal received by his rival about the return of a common value project. Our model has three distinctive features.

Private versus common values. Players' payoffs incorporate both a common and a private value components. The return of the project is assumed to be the same for both players, and independent of whom invests first. By contrast, players may have different opportunity costs of investment.

Information structure. Players have incomplete but symmetric information about the value of the investment project, but asymmetric information about their investment costs. This differs from the pure common value models of Hendricks and Kovenock [8] and Chamley and Gale [3], in which players have private signals about the common value of an investment, but not about privately borne investment costs, and from the symmetric information strategic experimentation models of Bolton and Harris [1] and Cripps, Keller and Rady [4].

Learning. There are two sources of public information in the model. Independently of players' investment decisions, a background signal provides free information about the value of the project. Moreover, once a player, acting as a leader, has sunk his investment, an additional signal is generated that may be used by the follower to optimally adjust his investment decision. This formulation is closer to the bandit problems studied by Bolton and Harris [1] and Cripps, Keller and Rady [4] than to the model of Chamley and Gale [3], in which investment by a player does not generate any information besides that revealed by the very fact that he invested.

Our objective is to study the learning externality due to the increase in the signal's quality generated by the leader's investment. The analysis focuses on how private information about investment costs is gradually revealed when new public information about the value of the investment becomes available. By delaying investment, each player tries to convince his rival that his own cost is high, and thus that his rival should invest first. The difference with

a pure common value model is that each player does not care about the information of his rival per se, but only in so far as it measures the likelihood of investing second and hence of benefitting from a better signal.

The difficulty in characterizing equilibrium in a learning model with asymmetric information is that, in general, there is no reasonably simple state variable with respect to which strategies could be measurable. Without extrinsic uncertainty, time itself would be such a state variable, as in the standard war of attrition with incomplete information. In a learning context with symmetric information about the value of the investment, it might seem reasonable to focus on strategies contingent on the players' common beliefs about the quality of the project. A natural candidate is to consider trigger strategies, as Bolton and Harris [1] do. Conditional on his cost, each player would invest as soon as his belief that the project will yield a high return reaches a certain critical value. However, for general learning processes, such as the one considered by Bolton and Harris [1], this approach is inappropriate when players have private information about their investment costs. Indeed, as opposed to time, the prospects associated to investment may then fluctuate randomly as new information becomes available. This implies that, even if one player is using a trigger strategy conditional on his cost, his rival's best response is not itself a trigger strategy. The reason is that there is not a one-to-one correspondence between what players know about the value of the project and what they know about their rival's investment cost. Note that this cannot occur when players' private signals are relative to the value of the project: in that case, there is always a one-to-one correspondence between what a player knows about his rival and how much he would obtain on average by investing.¹

To overcome this problem, we focus on a simple, though economically relevant learning process. (Cripps, Keller and Rady [4] use a similar signal structure in a strategic experimentation context.) Specifically, the investment project may be of low or high quality. If the project is of low quality, then players eventually learn this by observing failures that occur according to a Poisson process. By contrast, a high quality project never fails. Think for instance of two countries that may develop nuclear energy. Before they invest, the only information available to them is the experience of other countries. If a major nuclear incident occurs in one of these countries, this is a clear sign that the existing nuclear technology is unsafe, and that investing in it is not desirable. However, as long as no such incident has occurred, the common belief in the safety of nuclear technology will raise in both countries. An informational externality arises naturally since, if one country develops nuclear energy then, conditional on the technology being unsafe, the hazard rate of incidents will increase.

Given this signal structure, a player's decision to invest is only contingent on his cost and on the current public belief about the value of the project, which necessarily coincides with

¹This problem does not arise in a preemption model such as Lambrecht and Perraudin's [9]. The reason is that if a player is using a trigger strategy conditional on his cost, his rival is not threatened by preemption as long as the current value of the project is below its historic maximum. As a result, there exist equilibria in trigger strategies even though the value fluctuates over time.

the historic maximum of past beliefs if no failure has occurred. As a result, any equilibrium is in trigger strategies. Each player uses the flow of public information to update his beliefs about the value of the project and the investment cost of his rival. As time goes by and new upper values are reached by public beliefs, so that players become increasingly optimistic about the project, each player's belief that his rival will invest first decreases, so that they both become increasingly pessimistic about benefitting from the second-mover advantage.

While any equilibrium of the game is in trigger strategies, it is not Markovian with respect to the common belief process. Indeed, the equilibrium trigger of each player depends, besides his investment cost, on the initial belief about the value of the project. The more optimistic players are ex-ante about the quality of the project, the higher their equilibrium trigger will be. To understand this result, consider for simplicity a symmetric equilibrium. Each player's investment trigger is an increasing function of his investment cost. In particular, a player with the lowest possible cost knows that with probability one he will invest first in equilibrium. If the initial belief is above the threshold belief he would invest at if he were the only player, his equilibrium trigger is then just equal to the initial belief. As a result, an increase of the initial belief will lead each player to raise his investment trigger. In the nuclear industry example discussed above, this implies that if an initial period a latency is imposed during which investment is prohibited but during which background information can still be collected, then, conditional on the players' investment costs, investment will occur at higher levels of beliefs about the safety of nuclear technology. Interestingly, once one aggregates over the players' costs using the cost distribution, this prediction is reversed, and investment is accelerated in any interval of beliefs about the safety of nuclear technology.

This history dependence of equilibrium contrasts with the Markov perfect equilibrium that prevails in the complete information version of the game, or in the strategic experimentation models of Bolton and Harris [1] or Cripps, Keller and Rady [4]. This reflects the fact that, when players have incomplete information about each other's costs, public beliefs not only represent the information accumulated about the desirability of the investment, but also summarize players' beliefs about their opponent's cost.

Compared with the pure common value models of Hendricks and Kovenock [8] and Chamley and Gale [3], a distinctive feature of our model is that it allows to study the comparative statics of investment decisions with respect to the investment cost distribution. We show that when the distribution of cost is shifted toward a distribution with a higher hazard rate, so that the first distribution dominates the second distribution in the sense of first-order stochastic dominance, then, conditionally on their investment costs, players tend to invest later under the second than under the first distribution. This is rather intuitive: a distribution that makes relatively more likely that one's rival has a low cost provides an additional incentive to delay investment. As in the case of comparative statics with respect to the initial belief about the quality of the project, this prediction is reversed once one aggregates over players' costs, taking into account the change in distribution: in the aggregate, investment will occur at a faster rate under the second distribution than under the first.

The paper is organized as follows. Section 2 describes the model and the single-player benchmarks, as well as the equilibrium of the complete information game. In Section 3, the existence of a unique symmetric equilibrium is established, and a sufficient condition is provided under which this is the unique equilibrium. In Section 4, several comparative statics exercises are performed, and we examine how our results are affected by the introduction of a preemption motive. Section 5 concludes.

2. Model and benchmarks

2.1. The environment

Time is continuous, and labelled by $t \geq 0$. There are two players $i = 1, 2$, who can each invest in a risky project with ex-ante unknown value. The value of the project is the same for each player, and is not affected by whom invests first. Investment is irreversible and entails a sunk cost θ^i for player i . Both players are risk-neutral and discount future revenues and costs at a rate $r > 0$.

The investment project can be of low or high quality. We denote by $p_0 \in (0, 1)$ the prior probability that the project is of high quality. Neither player is informed of the quality of the project. However, players can learn about it at no cost by observing the outcome of a background public signal, which may be interpreted as the information generated by, e.g., investment analyzes, consumer surveys, or the experience of other (unmodelled) players. This signal is modelled as a Poisson process B with failure rate $\lambda^B > 0$ if the project is of low quality, and zero otherwise. A failure is common knowledge when it occurs, and provides an unambiguous signal of the low quality of the investment project.

A key feature of our model is that investment decisions are public, and once a player (the leader) has invested, the remaining player (the follower) receives more information about the project than provided by the background signal. Specifically, a high quality project never fails, and generates a profit $d > 0$ per unit of time. A low quality project generates the same instantaneous profit, but may fail according to a Poisson process with failure rate $\lambda^L > 0$, after which it does not generate any profit. It is assumed that conditional on its quality being low, failures of an installed project are independent from that generated by the background signal, and are perfectly observable by the follower. Therefore, everything happens as if the follower were observing a Poisson process F with intensity $\lambda^F = \lambda^B + \lambda^L$. This increase in the quality of the signal creates a strategic incentive to delay investment, besides the pure real option effect induced by the availability of background information. Given this signal structure, a high quality project generates discounted profits $h = d/r$ and a low quality project generates expected discounted profits $l = d/(r + \lambda^L)$.

The investment costs θ^i incorporate both the value of foregone opportunities in other activities and expenditures at the date of investment. Each player is informed of his own cost, but not of the cost of his rival. At date zero, costs are independently drawn from a common knowledge distribution G with continuously differentiable density g over a compact

interval $[\underline{\theta}, \bar{\theta}] \subset \mathbb{R}_+$. Throughout the paper, we assume that $[\underline{\theta}, \bar{\theta}] \cap [l, h] \neq \emptyset$ and that $g(\theta) > 0$ for all $\theta \in (\underline{\theta}, \bar{\theta})$. Each player's cost is independent from the quality of the project and the signal processes. Thus players have incomplete but symmetric information about the quality of the project at any stage of the game.

2.2. The single-player problem

We first derive the optimal investment policies for a single player with cost $\theta \in (l, h)$ when the signal available to him is either the background or the follower's signal. Suppose that the player starts from a belief $p \in (0, 1]$ that the project is of high quality and observes signal $k \in \{B, F\}$ over a period of time dt . Then, neglecting terms of order $o(dt)$, a failure occurs with probability $\lambda^k dt$ if the project is of low quality, and the posterior belief then drops to zero. If no failure occurs over the period of time dt , then the updated belief at the end of that period is:

$$p + dp = \frac{p}{p + (1-p)(1-\lambda^k dt)},$$

by Bayes' rule. Simplifying, we get that $dp = \lambda^k p(1-p)dt$. In particular, the belief that the project is of high quality increases deterministically as long as no failure occurs. Since $\lambda^F > \lambda^B$, beliefs evolve faster under the follower's signal than under the background signal.

Before investment takes place, the value function $V^k(p, \theta)$ for the single-player investment problem under signal k satisfies the following Bellman equation:

$$\lambda^k p(1-p) \partial_1 V^k(p, \theta) = [r + (1-p)\lambda^k] V^k(p, \theta).$$

Note that the effective discount rate of the player is $r + (1-p)\lambda^k$ to reflect the possibility that a failure will occur with probability $(1-p)\lambda^k$ over the period of time dt . It is optimal to invest as soon as p exceeds some critical threshold $\pi^k(\theta)$. Let $\Delta = h - l$. Then $\pi^k(\theta)$ satisfies the value matching condition $V^k(\pi^k(\theta), \theta) = \pi^k(\theta)\Delta + l - \theta$ and the smooth pasting condition $\partial_p V^k(\pi^k(\theta), \theta) = \Delta$. Using the fact that $V^k(0, \theta) = 0$, one can find explicit solutions for $V^k(p, \theta)$ and $\pi^k(\theta)$. For each $p \in [0, 1)$, let:

$$u^k(p) = p \left(\frac{p}{1-p} \right)^{r/\lambda^k}.$$

Then, for each $p < \tilde{p}$, the quantity $D^k(p, \tilde{p}) = u^k(p)/u^k(\tilde{p})$ represents the expected discounted value of a claim to a dollar at the first date where the belief that the project is of high quality attains \tilde{p} given a current value p , if this ever happens. Then, for any $p < \pi^k(\theta)$,

$$V^k(p, \theta) = D^k(p, \pi^k(\theta))[\pi^k(\theta)\Delta + l - \theta],$$

where the optimal investment trigger is given by:

$$\pi^k(\theta) = \frac{(\theta - l)(1 + r/\lambda^k)}{\theta - l + \Delta r/\lambda^k}.$$

For each $\theta \in (l, h)$, $\pi^k(\theta)$ is strictly greater than the Marshallian trigger $(\theta - l)/\Delta$ at which the net present value of investing is equal to zero. As in standard real option models, this reflects the value of waiting for information before investing. Similarly, the fact that $\pi^B(\theta) < \pi^F(\theta)$ and $V^F(p, \theta) > V^B(p, \theta)$ for each $p \in (0, \pi^F(\theta))$ reflects that the follower's signal is more informative than the background signal.

2.3. The complete information game

As a benchmark, it is helpful to consider a complete information version of the investment game in which players know each others' investment costs. A backward induction argument similar to the one used by Ghemawat and Nalebuff [7] to characterize exit decisions from an industry allows to derive the following result.

Proposition 1 *If $l < \theta^1 < \theta^2 < h$, there exists a unique subgame-perfect equilibrium. If the current belief that the project is of high quality is p and no player has invested yet, then:*

- (i) *If $p \geq \pi^F(\theta^2)$, both players simultaneously invests;*
- (ii) *If $\pi^B(\theta^1) \leq p < \pi^F(\theta^2)$, player 1 invests, and player 2 does not invest;*
- (iii) *If $p < \pi^B(\theta^1)$, no player invests.*

If the current belief that the project is of high quality is p and player i has invested, player j invests if and only if $p \geq \pi^F(\theta^j)$.

Since $\pi^B(\theta^2) > \pi^B(\theta^1)$ and $\pi^F(\theta^2) > \pi^F(\theta^1)$, player 2 can credibly threaten player 1 to always invest after him. Since he knows that he will never benefit from the high-quality signal, player 1 plays in equilibrium according to the single-player strategy $\pi^B(\theta^1)$. Thus the only reason why concession by player 1 may not occur at date zero is that $p_0 < \pi^B(\theta^1)$ —a pure real option effect. (Note that whenever there is no cost asymmetry, the game also admits mixed strategy equilibria.) It is worth noting that this equilibrium is inefficient. When taking his investment decision, the leader does not internalize the positive informational externality he exerts on the follower.

3. Equilibrium analysis

3.1. The incomplete information game

Since the optimal behavior of a follower is obviously given by the Markov policy π^F , we may focus without loss of generality on the following game of timing: (i) The only signal is the background signal B ; (ii) The game is over as soon as one player invests; (iii) If player i invests at a time where the current belief is p , he receives $p\Delta + l - \theta^i$, and player j receives $V^F(p, \theta^j)$. A pure strategy for player i specifies whether or not he should invest at time t , for each $t \geq 0$ and each history of the game up to time t , given his cost θ^i . Our equilibrium concept is perfect Bayesian equilibrium (PBE).

Since the only payoff relevant variable for the players, besides their own investment costs, is described by their common belief p about the underlying quality of the project, a natural question is whether there is a loss of generality in describing the players' behavior on the equilibrium path by investment triggers as in the benchmark cases previously analyzed. It is not difficult to show that, given a prior p_0 , any PBE is associated to measurable functions $\pi^i : [\underline{\theta}, \bar{\theta}] \rightarrow [p_0, 1] \cup \{+\infty\}$, $i = 1, 2$, such that, conditional on investing first, player i does so at the first time p exceeds $\pi^i(\theta^i)$, if this ever happens. (Note that one may have $\bar{\theta} > h$, in which case some types of the players never find optimal to invest.) The intuition for this result is simple. The Poisson specification of the background signal implies that, conditional on no failure and no investment occurring up to a certain time, all histories of the game up to this time are identical. Hence in any pure strategy PBE, each player knows at which time, if ever, he will invest if he invests first.

This does not however imply that a PBE is Markovian with respect to the belief process $\{p_t^B\}$ generated by the background signal. Quite on the contrary, the trigger functions π^i , $i = 1, 2$, depend on the initial belief p_0 . The reason for this is that since players have asymmetric information about their investment costs, the relevant Markov state variable (in the sense of Maskin and Tirole [11]) embodies players' beliefs about their opponent's type, as well as their beliefs about the quality of the project. A convenient feature of the Poisson specification of the background signal is that, conditional on no failure occurring, the belief process $\{p_t^B\}$ is monotone. Thus at any time before a failure occurs or an investment decision is taken, there is a one-to-one correspondence between the expected payoff that a player can get by investing and what he knows about his opponent.

3.2. Statement of the problem

Given a prior belief p_0 , and a trigger function π^j for player j , let $V(p_0, p^i, \pi^j, \theta^i)$ be the expected discounted payoff of player i with cost θ^i if he plans to invest at $p^i \geq p_0$,

$$V(p_0, p^i, \pi^j, \theta^i) = \int_{\{p^i > \pi^j(\theta^j)\}} D^B(p_0, \pi^j(\theta^j)) V^F(\pi^j(\theta^j), \theta^i) dG(\theta^j) + D^B(p_0, p^i) (p^i \Delta + l - \theta^i) \int_{\{\pi^j(\theta^j) \geq p^i\}} dG(\theta^j). \quad (1)$$

The first term on the right-hand side of (1) is i 's expected gain if j invests before p^i is reached by the belief process. The second term is i 's expected gain otherwise. A necessary and sufficient condition for a pair of trigger functions π^i , $i = 1, 2$, to be induced by a PBE is:

$$\pi^i(\theta^i) \in \operatorname{argmax}_{p^i \in [p_0, 1] \cup \{+\infty\}} V(p_0, p^i, \pi^j, \theta^i), \quad (2)$$

for all $i, j = 1, 2$ and $\theta \in [\underline{\theta}, \bar{\theta}]$. First, we show that the system (2) has a unique symmetric solution. We then provide a condition that ensures that all the solutions of (2) are symmetric.

It should be noted that, for each i and θ^i , $\pi^i(\theta^i)$ must be at least as large as $p_0 \vee \pi^B(\theta^i)$, and no larger than $p_0 \vee \pi^F(\theta^i)$.² Indeed, since there are no strategic interactions besides informational ones, each player must be ready to wait at least as much as if he were the only player. Symmetrically, there is no point for a player in investing at a higher belief than his optimal trigger as a follower, unless p_0 is larger than the latter. Denoting by ϕ^k , $k = B, F$, the inverse of π^k on $[0, 1]$, we will henceforth assume that $\Theta(p_0) = (\phi^F(p_0) \vee \underline{\theta}, h \wedge \bar{\theta}) \neq \emptyset$ to ensure that the problem is not trivial.

3.3. Separating properties of equilibria

It is clear that any PBE involves pooling for $\theta \leq \phi^F(p_0)$ and $\theta \geq h$. It follows that if $\underline{\theta} < \phi^F(p_0)$ (respectively $\bar{\theta} > h$), the distribution of investment triggers has an atom at p_0 (respectively at $+\infty$). The main result of this subsection is that any PBE of the investment game is separating on the relevant set of types $\Theta(p_0)$. Specifically, we have the following result.

Lemma 1 *In any PBE, the equilibrium trigger functions π^i , $i = 1, 2$, are strictly increasing on $\Theta(p_0)$.*

The proof of this result goes through a series of steps. The first consists to show that the equilibrium trigger functions are non-decreasing on $\Theta(p_0)$. This follows from a standard revealed preference argument. The bulk of the proof consists to show that in any PBE, the equilibrium distributions of investment triggers have no atoms, i.e., $\int_{\{\pi^i(\theta^i)=p\}} dG(\theta^i) = 0$ for each $i = 1, 2$ and $p \in (p_0, 1)$.

This part of the proof is more involved than in other models of wars of attrition (Riley [12], Fudenberg and Tirole [6]), because it is not a priori clear that if player j 's trigger distribution has an atom at $p \in (p_0, 1)$, all types of player i will be better-off waiting that the belief process reaches p before investing. In particular, if j 's hazard rate of investment is high in a small interval $[p - \varepsilon, p]$, we cannot exclude that some types of player i with an opportunity cost θ^i arbitrarily close to, but smaller than $\phi^F(p)$, would be ready to wait almost up to the point where p is attained by the belief process before investing. If it were the case, and thus a positive measure of types of player i were investing in any interval $[p - \varepsilon, p]$, the standard argument used to show that the types of player j investing at p would be better-off investing slightly before p breaks down.

To address this problem, we first show that if p is an atom in the distribution of j 's investment trigger, no type θ^i of player i such that $\theta^i > \phi^F(p)$ will invest in an interval $[p - \varepsilon, p]$ for $\varepsilon > 0$ close enough to zero. Now consider the case of $\theta^i \leq \phi^F(p)$. Intuitively, if there is an atom at p in the distribution of player j 's investment triggers, then a set of types $\theta^i \leq \phi^F(p)$ of player i of positive measure must invest in any interval of the form $[p - \varepsilon, p]$, or some types of player j who invest at p would have an incentive to deviate. In

²For any real numbers x and y , $x \vee y$ denotes the maximum of x and y , and $x \wedge y$ the minimum of x and y . By convention, $\pi^B(\theta) = \pi^F(\theta) = 0$ if $\theta < l$, and $\pi^B(\theta) = \pi^F(\theta) = +\infty$ if $\theta > h$.

particular, the marginal type $\phi^F(p)$ of player i must invest exactly at p . But whatever the hazard rate of investment of player j in a left neighborhood of p , it is not optimal for player i with cost $\phi^F(p)$ to wait up to p before investing. Technically, this is because the smooth pasting condition for the follower's problem at $\theta = \phi^F(p)$ implies that, as $\tilde{p} \uparrow p$, the relative gain of investing immediately at \tilde{p} becomes infinite relative to the marginal gain obtained by delaying investment. This contradiction establishes the result. It follows that in any PBE, the inverses $\phi^i = (\pi^i)^{-1}$ of the trigger functions π^i , $i = 1, 2$, are single-valued and continuous on $\pi^i(\Theta(p_0))$.

3.4. The symmetric equilibrium

In this section, we establish the existence of a unique symmetric PBE. First, we argue that if such an equilibrium exists, the common equilibrium trigger function π must be continuous on the relevant set of types $\Theta(p_0)$. This follows from the fact that, in any PBE, if the equilibrium trigger functions π^i , $i = 1, 2$, are discontinuous, the corresponding gaps in the players' distributions of investment triggers cannot overlap. Since the intuition is straightforward, we only sketch the argument here. If this did not hold, then there would exist at least one player i and one discontinuity point θ of π^i , such that almost surely no investment by player j takes place in an interval of beliefs I with upper bound $\pi^i(\theta^+)$. Since clearly $\pi^i(\theta^+) > \pi^B(\theta)$, one can choose the lower bound of I to be strictly greater than $\pi^B(\theta)$. But then, there is no point for player i to invest just beyond $\pi^i(\theta^+)$ rather than at the lower bound of I whenever his investment cost is just above θ . By doing so, he only marginally increases the probability of being a follower—as there cannot be an atom at $\pi^i(\theta^+)$ in the distribution of player j 's investment triggers—while foregoing current benefits since $\pi^B(\theta) < \pi^i(\theta^+)$. It follows that in a symmetric PBE, $\pi(\Theta(p_0))$ is an interval.

The following result gives the necessary first-order conditions that must be satisfied in a symmetric PBE.

Lemma 2 *In a symmetric PBE, the inverse ϕ of the common equilibrium trigger function π is differentiable on $\pi(\Theta(p_0))$, and:*

$$\partial\phi(p) = f(p, \phi(p)) \tag{3}$$

for all $p \in \pi(\Theta(p_0))$, where, by definition:

$$f(p, \theta) = \left[\frac{1 - G(\theta)}{g(\theta)} \right] \left[\frac{\partial_1^- D^B(p, p) (p\Delta + l - \theta) - \Delta}{V^F(p, \theta) - p\Delta - l + \theta} \right] \tag{4}$$

on the domain $\mathcal{D} = \{(p, \theta) \in (0, 1) \times (l, h) \cap (\underline{\theta}, \bar{\theta}) \mid \pi^B(\theta) < p < \pi^F(\theta)\}$.

Intuitively, the first-order condition (3) states that a player i with investment cost $\phi(p) \in \Theta(p_0)$ cannot increase his payoff by investing at $p + dp$ instead of p . Heuristically, this formula can be derived as follows. Note first that the probability that player j 's investment trigger

lies in $[p, p + dp)$ is $\partial\phi(p) g(\phi(p)) dp$. Hence, a marginal change $dp > 0$ by player i yields him an expected gain:

$$[V^F(p, \phi(p)) - p\Delta - l + \phi(p)] \partial\phi(p) g(\phi(p)) dp + o(dp). \quad (5)$$

This change however costs him $p\Delta + l - \phi(p) - D^B(p, p + dp) [(p + dp)\Delta + l - \phi(p)]$ if player j 's investment trigger lies in $[p + dp, 1] \cup \{+\infty\}$, which has probability $1 - G(\phi(p + dp))$. The expected incremental cost is therefore equal to:

$$\{-\Delta - \partial_2^+ D^B(p, p) [p\Delta + l - \phi(p)]\} [1 - G(\phi(p))] dp + o(dp). \quad (6)$$

Equating (5) and (6), and taking advantage of $\partial_1^- D^B(p, p) = -\partial_2^+ D^B(p, p)$ yields (3). It should be noted that the smooth pasting condition for the single-player problem with signal B implies that $\lim_{\theta \uparrow \phi^B(p)} f(p, \theta) = 0$, and that the value matching condition for the single-player problem with signal F implies that $\lim_{\theta \downarrow \phi^F(p)} f(p, \theta) = +\infty$. It follows that a solution to (3) entering the domain \mathcal{D} will remain in this domain. As a result, the equilibrium trigger function π in a symmetric PBE lies strictly between π^B and π^F on $\Theta(p_0)$. Intuitively, since the identity of the player with the lowest investment cost is not common knowledge ex-ante, each player wants to delay his decision in order to convince his rival that the latter has indeed the lowest cost. Hence, uncertainty about private investment costs creates further inefficient delays on top of those generated by the presence of strategic interactions.

It is straightforward to check that any function ϕ satisfying the first-order condition (3) also satisfies the second-order condition for problem (2). Indeed, let $p \in \pi(\Theta(p_0))$. With obvious notation, let us rewrite (3) as:

$$\frac{g(\phi(p))}{1 - G(\phi(p))} \partial\phi(p) = \varphi(p, \phi(p)).$$

ϕ is strictly increasing on $\pi(\Theta(p_0))$ as $\pi = \phi^{-1}$ is continuous. Moreover, for each $\tilde{p} \in (0, 1)$, the mapping $\varphi(\tilde{p}, \cdot)$ is strictly decreasing on $(\phi^F(p), \phi^B(p))$. It follows that:

$$\frac{g(\phi(\tilde{p}))}{1 - G(\phi(\tilde{p}))} \partial\phi(\tilde{p}) > \varphi(\tilde{p}, \phi(p)) \quad (7)$$

if and only if $\tilde{p} < p$. In other words, a player with cost $\phi(p)$ optimally invests at p , because the expected incremental payoff of a marginal delay of investment is everywhere positive between $\pi^B(\phi(p))$ and p and everywhere negative between p and $\pi^F(\phi(p))$. Therefore, the incentive compatibility conditions are globally satisfied.

We are now ready to state and prove the main result of this section.

Proposition 2 *For each $p_0 \in (0, 1)$, there exists a unique symmetric PBE.*

The proof proceeds by establishing the existence of a unique solution to (3) satisfying appropriate boundary conditions. We distinguish two cases. First, if $\pi^F(\underline{\theta}) > p_0$, then the

boundary condition is $\pi(\underline{\theta}) = p_0 \vee \pi^B(\underline{\theta})$, since a player with cost $\underline{\theta}$ always invests first and has thus no incentive to delay investment for beliefs above $\pi^B(\underline{\theta})$. Second, if $\pi^F(\underline{\theta}) \leq p_0$, then the boundary condition is $\pi(\phi^F(p_0)) = p_0$, since a player with cost $\theta \leq \phi^F(p_0)$ has no incentive to delay investment for beliefs above $\pi^F(\theta) \leq p_0$. In the first case, the differential equation (3) satisfies a Lipschitz condition at $p_0 \vee \pi^B(\underline{\theta})$ as $\pi(\theta) < \pi^F(\underline{\theta})$. In the second case, the Lipschitz condition no longer holds at p_0 . An ad hoc argument is then needed to establish the existence of a solution to (3) satisfying the equilibrium boundary conditions.

The symmetric PBE is unique in the sense that two symmetric PBEs can only differ off the equilibrium path. If $\bar{\theta} \geq h$, then for any time t , there is a positive probability that the first investment will take place later than t , conditional on no failure occurring: in the terminology of Fudenberg and Tirole [6], there is “perpetual selection.” In this case, beliefs can be computed by Bayes’ rule at any date. If on the contrary $\bar{\theta} < h$, there is “finite selection:” there exists a time T such that the probability that the first investment will occur later than T is zero. One must then specify appropriate off-the-equilibrium beliefs for $t \geq T$.

3.5. Uniqueness of equilibrium

So far, we have focused on the symmetric PBE of the investment game. We now provide a sufficient condition for uniqueness of equilibrium. This condition states that with positive (but arbitrarily small) probability, each player’s cost is large enough to discourage him to invest whatever his beliefs about the profitability of the project, i.e., $\bar{\theta} > h$. This assumption plays the same role in our model as Assumption A(ii) in Fudenberg and Tirole [6]. In a duopoly model of exit, they argue that there is no a priori reason to impose a natural monopoly assumption as in the standard war of attrition. In the present context of investment decisions, the analogue of this point is that one cannot exclude that players with prohibitively high costs would abstain to invest even if they knew that the project is of high quality.

Proposition 3 *If $\bar{\theta} > h$, there exists a unique PBE.*

The proof of this result goes through a series of steps. First, one shows that the symmetric PBE is the unique PBE in which the equilibrium trigger functions π^i , $i = 1, 2$, are continuous. This is done by proving that the inverse functions ϕ^i , $i = 1, 2$, satisfy a system of differential equations analogous to (3) and by using an argument similar to the one yielding the uniqueness of the symmetric PBE. Intuitively, the assumption $\bar{\theta} > h$ provides a boundary condition at $p = 1$ which pins down a unique solution to the system of differential equations satisfied by ϕ^i , $i = 1, 2$. Next, we consider the case of asymmetric equilibria. By the above argument, such equilibria if they exist must exhibit discontinuities. A key lemma is that if $\theta \in \Theta(p_0)$ is a discontinuity point of π^j , then the types θ^i of player i who invest within $[p^j(\theta^-), p^j(\theta^-)]$ do so at $\pi^B(\theta^i)$ —recall that gaps in the players’ equilibrium distribution of triggers cannot overlap. One then concludes the argument by showing that if the trigger function π^j has a discontinuity at $\theta \in \Theta(p_0)$, then $\pi^j > \pi^i$ on $[\theta, h)$, and by applying the same boundary condition as the one used to rule out asymmetric continuous PBE. In the remaining of the

paper, we shall focus on the case where $\bar{\theta} > h$ and thus a unique PBE exists.

4. Comments and extensions

4.1. History dependence

As discussed in Subsection 3.1, it would be misleading to describe the PBE derived above as a Markov equilibrium in which the players' strategies are only contingent on their current beliefs about the quality of the project. Indeed, the equilibrium analysis has made clear that players' investment triggers depend on their initial belief p_0 that the project is of high quality, and more precisely on the position of p_0 with respect to $\pi^F(\underline{\theta})$. For instance, if $\pi^F(\underline{\theta}) < p_0$, the equilibrium distribution of entry triggers has an atom at p_0 , since there is a positive measure of each player's types for which the additional information provided by becoming a follower has no value at this initial level of belief. If $p_0 > \pi^B(\underline{\theta})$, a small increase dp_0 increases the PBE entry trigger of type $\underline{\theta}$ by dp_0 . Besides, since the PBE trigger function π is increasing, the probability that a player with cost $\theta > \underline{\theta}$ invests second is $G(\theta)$. Hence, in order to maintain a constant probability of investing second, the equilibrium entry trigger of each type $\theta > \underline{\theta}$ must increase. In that sense, the PBE of the game is history dependent.

This result contrasts with those obtained under complete information. Indeed, the unique subgame-perfect equilibrium derived in Proposition 1 is truly Markovian with respect to current beliefs about the quality of the project, since players' strategies do not depend on their initial assessment p_0 . The only reason why concession by the low-cost player might not occur at date zero is the real option effect induced by the availability of the background signal. When players have complete information about each other's characteristics, public beliefs only represents the information accumulated about the desirability of the investment. By contrast, when investment costs are private information, public beliefs indirectly perform the dual task of summarizing each player's beliefs about the cost of his rival.

One can strengthen this insight by studying how the marginal incentives to invest at a belief level p vary as a function of p_0 . We focus on the hazard rate of investment at a belief p , i.e., the probability that each player will invest in an interval of beliefs $[p, p + dp]$ conditional on no investment having occurred at p . Specifically, consider two PBEs with trigger functions π and $\tilde{\pi}$ corresponding to two initial beliefs p_0 and \tilde{p}_0 . If $\tilde{p}_0 > p_0$, one has $\phi > \tilde{\phi}$ on $(\tilde{p}_0, 1)$. It is then easy to check from (3) that the equilibrium hazard rates of investment of each player at a belief $p > \tilde{p}_0$ satisfy:

$$\frac{g(\tilde{\phi}(p))}{1 - G(\tilde{\phi}(p))} \partial \tilde{\phi}(p) > \frac{g(\phi(p))}{1 - G(\phi(p))} \partial \phi(p).$$

On average, an increase in the initial belief p_0 about the quality of the project raises the likelihood that each player will invest in $[p, p + dp]$ given that no player has invested yet at p . In other terms, each player, unconditionally on his type, becomes less cautious in his investment decision, and investment is accelerated at any level of beliefs. This illustrates the

fact that a slight perturbation of the initial conditions of the problem may have long-lasting consequences.

4.2. Comparative statics

We now perform three types of comparative statics exercises, with respect to the cost distribution, the quality of the signals, and the investment payoffs. The proofs are simple consequences of the characterization of the unique PBE of the investment game by the differential equation (3), and are therefore omitted.

Changes in the cost distribution. Let G and \tilde{G} be two cost distributions with densities g and \tilde{g} . For simplicity, let G and \tilde{G} be defined on the same support $[\underline{\theta}, \bar{\theta}]$ with $p_0 < \pi^B(\underline{\theta})$ and $\bar{\theta} > h$. Let $H = g/(1 - G)$ and $\tilde{H} = \tilde{g}/(1 - \tilde{G})$ be the corresponding hazard rates. In analogy with the auction literature (Maskin and Riley [10]), we compare the associated PBE trigger functions π and $\tilde{\pi}$ and the distributions $G \circ \phi$ and $\tilde{G} \circ \tilde{\phi}$ of investment triggers under the assumption that $\tilde{H} > H$ on $(\underline{\theta}, \bar{\theta})$. This implies in particular that G dominates \tilde{G} in the sense of first-order stochastic dominance.

Corollary 1 *If $\tilde{H} > H$ on $(\underline{\theta}, \bar{\theta})$, then $\tilde{\pi} > \pi$ on $(\underline{\theta}, h)$ and the hazard rate of $\tilde{G} \circ \tilde{\phi}$ is strictly greater than the hazard rate of $G \circ \phi$ on $(p_0, 1)$.*

The intuition of this result is that, if one shifts the distribution of costs from G to \tilde{G} , one makes it more likely that players have low rather than high costs. If one holds the strategy of player i fixed, then, for any $p \in (p_0, 1)$, the likelihood that investment by player i will occur in the interval of beliefs $[p, p + dp]$ given that no one has invested yet at p is higher under the distribution \tilde{G} than under G . The fact that player i thus invests at a faster rate gives player j an additional incentive to delay investment at any belief $p \in (p_0, 1)$. Hence a shift in the cost distribution that decreases the likelihood of high costs may actually incite each player to increase his investment trigger. However, this mixed strategy effect holds conditional on a particular value of a player's cost, but not when one aggregates over them. The second part of Corollary 1 shows that even if each player has, conditional on any particular value of his cost, greater incentive to wait before investing under \tilde{G} than under G , the resulting distribution $\tilde{G} \circ \tilde{\phi}$ will be more biased towards lower values of the investment triggers than $G \circ \phi$, and investment will occur at a faster rate.

Changes in the quality of public signals. We now examine the sensitivity of the equilibrium to increases in the quality of public signals. As above, we denote by π and $\tilde{\pi}$ the PBE trigger functions associated with different intensities λ^B and $\tilde{\lambda}^B$ for the background signal, or different intensities λ^L and $\tilde{\lambda}^L$ for the signal generated by the leader, and we assume that $p_0 < \pi^B(\underline{\theta})$.

Corollary 2 *If $\tilde{\lambda}^B > \lambda^B$ or $\tilde{\lambda}^L > \lambda^L$, then $\tilde{\pi} > \pi$ on $(\underline{\theta}, h)$.*

The impact of an increase in the quality of the background signal is twofold. First, it affects directly the investment decision of a player with cost $\underline{\theta}$, who in any case invests first in

equilibrium. An increase in the background signal's intensity from λ^B to $\tilde{\lambda}^B$ leads him to change his investment trigger from $\pi^B(\underline{\theta})$ to $\tilde{\pi}^B(\underline{\theta}) > \pi^B(\underline{\theta})$. This pure real option effect modifies the initial condition for (3), which shifts upward the whole PBE trigger function. The second impact is strategic: by reducing the cost of waiting to become a follower, it increases the marginal incentives of each player with cost $\theta > \underline{\theta}$ to delay investment.

An increase in the quality of the signal generated by the leader's investment affects equilibrium play in two ways. First, it increases the value of investing second in order to benefit from a more informative signal. Second, it decreases the expected discounted payoff $l = d/(r + \lambda^L)$ of a low quality project by increasing its failure rate. This in turn increases the payoff differential Δ between a high and a low quality project, thus providing an additional incentive to delay investment.

Changes in payoffs. Last, we analyze the sensitivity of the equilibrium to changes in the payoff structure. We compare the PBE trigger payoff functions π and $\tilde{\pi}$ corresponding respectively to minimal values l and $\tilde{l} > l$ of the project, holding $\Delta = h - l = \tilde{h} - \tilde{l}$ constant. Intuitively, this corresponds to an investment subsidy $s = \tilde{l} - l$, received at the time of investment. We assume that $p_0 < \tilde{\pi}^B(\underline{\theta})$, so that the increase in value is not so drastic as to imply immediate investment by a player with cost $\underline{\theta}$. An obvious impact of such a policy is to incite players with costs $\theta \in [h, h + s]$ to invest, provided they become sufficiently optimistic about the quality of the project. For the remaining types of the players, we have the following result.

Corollary 3 *If $\tilde{l} > l$ and Δ remains constant, then $\pi > \tilde{\pi}$ on $(\underline{\theta}, h)$.*

As for an increase in the quality of the background signal, the impact of a subsidy at the time of investment is twofold. First, it affects directly the behavior of a player with cost $\underline{\theta}$, who will change his trigger to $\tilde{\pi}^B(\underline{\theta}) < \pi^B(\underline{\theta})$. Second, by reducing the relative cost of becoming a leader, it decreases the marginal incentives of each player with cost $\theta > \underline{\theta}$ to delay investment.

4.3. Preemption versus war of attrition

We have so far assumed that the value of the project is the same for each player and independent of the timing of investment. We now study how our results are affected when a preemption motive is introduced besides the pure informational externality.

The only change in the model is that the instantaneous profit for the player who invests first is αd instead of d , where $\alpha > 1$. Thus the follower benefits from a more informative signal, but his project is less valuable than the leader's. The optimal investment policy of a follower is the same as in the basic model. The optimal investment trigger for a single player with cost θ is $\pi_\alpha^B(\theta) = \pi^B(\theta/\alpha) < \pi^B(\theta)$. In addition, we denote by $\pi_\alpha^P(\theta)$ the belief about the quality of the project at which a player with cost θ is exactly indifferent between preempting his rival or being preempted by him. Formally, this is the solution to:

$$V^F(p, \theta) = \alpha(p\Delta + l) - \theta. \quad (8)$$

It is easy to check that $\pi_\alpha^P(\theta)$ is well-defined and unique for any $\theta \in [\underline{\theta} \vee \alpha l, \bar{\theta} \wedge \alpha h]$, and that it lies between the Marshallian trigger $(\theta/\alpha - l)/\Delta$ at which the net present value of investing as a leader is equal to zero, and the follower's trigger $\pi^F(\theta)$. Note that as a player with cost $\theta \geq h$ effectively loses any opportunity to invest if he is preempted, $\pi_\alpha^P(\theta)$ is equal to the Marshallian trigger for any $\theta \geq h$. By continuity, one can verify that $\pi_\alpha^P(\theta) < \pi_\alpha^B(\theta)$ if θ is close enough to h . This reflects the fact that the preemption motive is prevalent for firms with high costs.

Equilibrium analysis. As in the basic model, any PBE of this modified game is characterized by a pair of trigger functions π_α^i , $i = 1, 2$. It is easy to check that for each i and θ^i , $\pi_\alpha^i(\theta^i)$ must be at least as large as $p_0 \vee [\pi_\alpha^B(\theta^i) \wedge \pi_\alpha^P(\theta^i)]$, and no larger than $p_0 \vee [\pi_\alpha^B(\theta^i) \vee \pi_\alpha^P(\theta^i)]$. It follows immediately that if π_α^B and π_α^P intersect at some $\theta \in [\underline{\theta}, \bar{\theta}]$, then $\pi_\alpha^i(\theta) = \pi_\alpha^B(\theta) = \pi_\alpha^P(\theta)$ in any PBE. (To simplify the analysis, we assume that $p_0 < \pi_\alpha^B(\underline{\theta}) \vee \pi_\alpha^P(\underline{\theta})$, so that a player with cost $\underline{\theta}$ will not invest at time zero.) By contrast, if $\pi_\alpha^B(\theta) \neq \pi_\alpha^P(\theta)$, then by continuity there exists a maximal open interval $I_\theta \subset [\underline{\theta}, \bar{\theta}]$ containing θ such that $\pi_\alpha^B(\theta) - \pi_\alpha^P(\theta)$ and $\pi_\alpha^B(\tilde{\theta}) - \pi_\alpha^P(\tilde{\theta})$ have the same sign for any $\tilde{\theta} \in I_\theta$. If $\tilde{\theta} \in I_\theta$, two players with costs θ and $\tilde{\theta}$ face the same type of incentives. Two cases must therefore be distinguished.

If $\pi_\alpha^B(\theta) < \pi_\alpha^P(\theta)$, the attrition motive dominates in I_θ . The analysis parallels that of Section 3, except that π_α^P now plays the role of π^F . In a symmetric PBE, the inverse ϕ_α of the common equilibrium trigger function π_α is differentiable on $\pi_\alpha(I_\theta)$, and:

$$\partial \phi_\alpha(p) = f_\alpha(p, \phi_\alpha(p)) \quad (9)$$

for all $p \in \pi_\alpha(I_\theta)$, where, by definition:

$$f_\alpha(p, \theta) = \left[\frac{1 - G(\theta)}{g(\theta)} \right] \left\{ \frac{\partial_1^- D^B(p, p) [\alpha(p\Delta + l) - \theta] - \alpha\Delta}{V^F(p, \theta) - \alpha(\Delta p + l) + \theta} \right\} \quad (10)$$

on the domain $\mathcal{D}_\alpha = \{(p, \theta) \in (0, 1) \times (\underline{\theta}, \bar{\theta}) \mid \pi_\alpha^P(\theta) \wedge \pi_\alpha^B(\theta) < p < \pi_\alpha^P(\theta) \vee \pi_\alpha^B(\theta)\}$. The differential equation (9) simply expresses the first-order condition for a player with cost $\phi_\alpha(p)$. A key observation is that the boundary condition for (9) is an initial condition as in the pure attrition model:

$$\phi_\alpha(\pi_\alpha^B(\inf I_\theta)) = \inf I_\theta.$$

If $\pi_\alpha^B(\theta) > \pi_\alpha^P(\theta)$, the preemption motive dominates in I_θ . The derivative of ϕ_α is again given by (9) but, in contrast with the attrition case, the boundary condition for (9) is a terminal condition, reflecting the threat of preemption:

$$\phi_\alpha(\pi_\alpha^P(\sup I_\theta)) = \sup I_\theta.$$

If in particular $\bar{\theta} = h$ then, whenever the posterior belief that the project is of high quality approaches $\pi_\alpha^P(\bar{\theta}) = (h/\alpha - l)/\Delta$ and no one has invested yet, the hazard rate of being preempted per unit of time explodes to infinity. The terminal condition expresses then the fact that at the upper end of the cost distribution, the option value of waiting disappears, and

players' equilibrium triggers become arbitrarily close to their Marshallian triggers. Following the same steps as in Section 3, one can establish that the investment game has a unique symmetric PBE.

The case $\lambda^L = +\infty$. As (8) admits no explicit solution, it is in general difficult to determine the sign of $\pi_\alpha^B(\theta) - \pi_\alpha^P(\theta)$ and hence the incentives of a player with cost θ . However, in the limit case where λ^L tends to infinity, and thus the follower's signal becomes perfectly informative about the quality of the project, it is possible to determine explicitly $\pi_\alpha^P(\theta)$. Specifically, when λ^L tends to infinity, the payoff $l = d/(r + \lambda^L)$ of a low quality project converges to zero as such a project fails almost immediately after investment, the follower's value $V^F(p, \theta)$ converges to $p(h - \theta)$, and (8) becomes $p(h - \theta) = p\alpha h - \theta$, so that:

$$\pi_\alpha^P(\theta) = \frac{\theta}{(\alpha - 1)h + \theta}.$$

It is then easy to check that, if the preemption effect is not too large, $\alpha < 1 + r/\lambda^B$, there exists a single cutoff value of the cost, $\theta^P = h[1 + (1 - \alpha)\lambda^B/r]$, such that $\pi_\alpha^B(\theta) < \pi_\alpha^P(\theta)$ if and only if $\theta < \theta^P$. Thus the attrition motive dominates for low values of the costs, while the preemption motive dominates for high values of the costs.

Things are more complex whenever λ^L is large, but finite. Indeed, one can show that there is an additional region of low values for the costs close to $\alpha l = \alpha d/(r + \lambda^L)$ such that the preemption motive dominates the attrition motive.³ The intuition is as follows. Consider hypothetically a player with cost $\theta \in (l, \alpha l)$. Then for this player, preempting gives a minimum sure payoff of $\alpha l - \theta > 0$ and thus dominates strictly being a follower at any level of belief, including, and this is the key point, $p = 0$. This implies that the preemption motive dominates in a right neighborhood of αl . Numerical simulations suggest however that this additional region tends to be extremely small, and vanishes as λ^L tends to infinity. Overall, the case of large λ^L is quantitatively very similar to the limit case of infinite λ^L .

5. Concluding remarks

In this paper, we have built a strategic investment model in which the information available to players is partly exogenous, through the background signal, and partly endogenous, through the investment decision of the leader. We have focused on the problem of private provision of information as a public good, and on the impact of cost uncertainty on the incentives of players to provide this information by investing. When the identity of the player with the smallest investment cost is unknown ex-ante, inefficient delays occur in equilibrium because each player tries to convince his rival that the latter has the lowest investment cost and should therefore bear the whole burden of experimentation. In equilibrium, the first player to invest does so at a time where the marginal value of public background information per

³This can be checked as follows. First, note that by construction, $\pi_\alpha^B(\alpha l) = \pi_\alpha^P(\alpha l) = 0$. One can then verify from (8) that $\partial\pi_\alpha^P(\alpha l^+) = 1/(\alpha\Delta) < (1 + \lambda^B/r)/(\alpha\Delta) = \partial\pi_\alpha^B(\alpha l^+)$.

se is less than its marginal cost in terms of delayed investment. This implies that failing to take into account the potential, or indirect information generated by the players' investments themselves, leads to underestimate the value of public information. Our analysis has shown that the common belief about the quality of the project is not a sufficient statistics for the investment decisions of the players. In particular, an increase in the players' initial degree of optimism has long-lasting consequences in the form of a higher investment rate at any level of beliefs. For more general learning processes, this form of history dependence would lead to more complex investment dynamics. This, and related questions, must await for future work.

Acknowledgements

We are grateful to Stéphane Auray, Martin Cripps, Jocelyn Donze, Godfrey Keller, Bart Lambrecht, Sébastien Pouget, Jean Tirole, Stéphane Villeneuve, Helen Weeds, Lucy White, as well as an anonymous referee, for thoughtful comments and suggestions. We would also like to thank seminar participants at CIRANO, GREMAQ, HEC, Institut Henri Poincaré, LSE, Nuffield College, UCL, as well as at the CEPR Conference on Real Options and at the Third International Conference on Real Options. Financial support from CNRS and STICERD is gratefully acknowledged by the authors. We remain of course solely responsible for the content of this paper.

Appendix

Proof of Proposition 1. (i) and (iii) are immediate. Suppose that $p \in [\pi^F(\theta^1), \pi^F(\theta^2))$ and that no player has invested yet. Then it is a dominant strategy for player 1 to invest. Since $V^F(p, \theta^2) > p\Delta + l - \theta^2$, the best reply of player 2 is to abstain. The most pessimistic assumption that player 2 can make about the behavior of player 1 is that he will not invest until $\pi^F(\theta^1)$ is attained, if this ever happens. Now, define a sequence $\{\pi_n\}$ as follows. First, let $\pi_0 = \pi^F(\theta^1)$. Next, for each $n \in \mathbb{N}$:

(i) If $\pi_n \leq \pi^B(\theta^2)$, $\pi_{n+1} = \pi_n$;

(ii) If for each $p \in [0, \pi_n]$, $D^B(p, \pi_n) V^F(\pi_n, \theta^2) > p\Delta + l - \theta^2$, $\pi_{n+1} = 0$;

(iii) Otherwise, $\pi_{n+1} = \sup\{p \in [0, \pi_n] \mid p\Delta + l - \theta^2 \geq D^B(p, \pi_n) V^F(\pi_n, \theta^2)\}$.

Obviously, $\pi_{n+1} \leq \pi_n$ as $\pi_0 < \pi^F(\theta^2)$. Note that if there exists some $n \in \mathbb{N}$ such that $\pi_n \leq \pi^B(\theta^2)$, then player 2's best reply at any $p \in (0, \pi_0)$ is to wait until player 1 invests, which therefore occurs at $\pi^B(\theta^1)$. Otherwise, if no player has invested at the time where π_{n+1} is attained, if this ever happens, player 1 can make no credible threat to wait until π_n to invest. Hence, if $\pi_{n+1} \geq \pi^B(\theta^1)$, player 1 should invest immediately, and player 2 should therefore abstain. To complete the proof, it is thus sufficient to show that for some $n \in \mathbb{N}$, $\pi_n \leq \pi^B(\theta^2)$. Suppose the contrary. Then, since the sequence $\{\pi_n\}$ is bounded below by $\pi^B(\theta^2)$ and decreasing, it must converge to a limit $\pi_\infty \geq \pi^B(\theta^2)$. By continuity, (iii) implies that $\pi_\infty \Delta + l - \theta^2 = V^F(\pi_\infty, \theta^2)$. This contradicts $\pi_\infty < \pi^F(\theta^2)$. \square

Proof of Lemma 1. Fix some PBE and associated trigger functions π^i , $i = 1, 2$. That π^i is non-decreasing for each i follows from a standard revealed preference argument. Now, if $p \in (p_0, 1)$ is an atom in the distribution of player j 's investment trigger, it is easy to check that for some $\varepsilon > 0$, $[p - \varepsilon, p] \cap \{\pi^i(\theta^i) \mid \theta^i > \phi^F(p)\} = \emptyset$, which in turn implies that $\pi^i(\phi^F(p)) = p$. Rearranging the incentive constraint $V(p_0, \tilde{p}, \pi^j, \phi^F(p)) \leq V(p_0, p, \pi^j, \phi^F(p))$ for player i with cost $\phi^F(p)$, we get:

$$\frac{\tilde{p}\Delta + l - \phi^F(p) - D^B(\tilde{p}, p) [p\Delta + l - \phi^F(p)]}{D^F(\tilde{p}, p) [p\Delta + l - \phi^F(p)] - \tilde{p}\Delta - l + \phi^F(p)} \leq \frac{\int_{\{p > \pi^j(\theta^j) \geq \tilde{p}\}} dG(\theta^j)}{\int_{\{\pi^j(\theta^j) \geq \tilde{p}\}} dG(\theta^j)} \quad (11)$$

for $\tilde{p} < p$. We prove that the left-hand side of (11) goes to $+\infty$ as $\tilde{p} \uparrow p$, thus obtaining a contradiction since the right-hand side is bounded above by one. Using L'Hôpital's rule and the smooth pasting condition $\Delta - \partial_1^- D^F(p, p) [p\Delta + l - \phi^F(p)] = 0$ for the follower's problem, we obtain:

$$\begin{aligned} \lim_{\tilde{p} \uparrow p} \frac{\tilde{p}\Delta + l - \phi^F(p) - D^B(\tilde{p}, p) [p\Delta + l - \phi^F(p)]}{\tilde{p} - p} &= \partial_1^- [D^F(p, p) - D^B(p, p)] [p\Delta + l - \phi^F(p)] \\ &= \frac{r/\lambda^F - r/\lambda^B}{p(1-p)} [p\Delta + l - \phi^F(p)], \end{aligned}$$

which is strictly negative since $\lambda^F > \lambda^B$ and $p > [\phi^F(p) - l]/\Delta$. Using again the smooth pasting condition for the follower's problem, we obtain:

$$\lim_{\tilde{p} \uparrow p} \frac{D^F(\tilde{p}, p) [p\Delta + l - \phi^F(p)] - \tilde{p}\Delta - l + \phi^F(p)}{\tilde{p} - p} = 0^-.$$

Taking the quotient of these two limits establishes the result. \square

Proof of Lemma 2. By Lemma 1, π has a continuous inverse ϕ on $\pi(\Theta(p_0))$. Let $p \in \pi(\Theta(p_0))$, and consider a player with cost $\phi(p)$. Then, for all $\tilde{p} \in (p, \pi^F(\phi(p)))$,

$$p\Delta + l - \phi(p) \geq \int_p^{\tilde{p}} D^B(p, q) V^F(q, \phi(p)) \frac{dG(\phi(q))}{1 - G(\phi(p))} + \frac{1 - G(\phi(\tilde{p}))}{1 - G(\phi(p))} D^B(p, \tilde{p}) [\tilde{p}\Delta + l - \phi(p)].$$

Taking advantage from the fact that $D^B(p, \cdot) V^F(\cdot, \phi(p))$ is strictly decreasing on $(p, \pi^F(\phi(p)))$ and rearranging, we obtain:

$$[1 - G(\phi(\tilde{p}))] \frac{p\Delta + l - \phi(p) - D^B(p, \tilde{p}) [\tilde{p}\Delta + l - \phi(p)]}{V^F(p, \phi(p)) - p\Delta - l + \phi(p)} \geq G(\phi(\tilde{p})) - G(\phi(p)) \geq 0.$$

Dividing both sides by $\tilde{p} - p$ and letting $\tilde{p} \downarrow p$, we conclude that:

$$0 \leq \bar{\partial}^+ G(\phi(p)) = \limsup_{\tilde{p} \downarrow p} \frac{G(\phi(\tilde{p})) - G(\phi(p))}{\tilde{p} - p} \leq g(\phi(p)) f(p, \phi(p)),$$

and a similar inequality holds for any \tilde{p} in a neighborhood of p . From the proof of Lemma 1, $\tilde{p} < \pi^F(\phi(\tilde{p}))$ for each such \tilde{p} . It follows that $\bar{\partial}^+ G \circ \phi$ is bounded in a neighborhood of p , and since p is arbitrary, on any closed interval $I \subset \pi(\Theta(p_0))$. In particular, $G \circ \phi$ is Lipschitzian, hence absolutely continuous on any such I . Since the density g is bounded away from zero, ϕ is absolutely continuous, hence almost everywhere differentiable on $\pi(\Theta(p_0))$. Equation (3) follows then from differentiating $V(p_0, \tilde{p}, \pi, \phi(p))$ with respect to \tilde{p} and requiring that the resulting derivative, when it exists, be equal to zero at p . To conclude the proof, note that $f(\cdot, \phi(\cdot))$ is continuous, and that ϕ is absolutely continuous. Since ϕ satisfies (3) almost everywhere, the fundamental theorem of calculus implies that it is in fact everywhere differentiable on $\pi(\Theta(p_0))$, and its derivative is given everywhere by (3). \square

Proof of Proposition 2. By the Cauchy-Lipschitz theorem, for any $(p, \theta) \in \mathcal{D}$, there exists a unique solution $\phi_{p, \theta}$ on \mathcal{D} to (3) passing through (p, θ) . A standard extension argument based on the behavior of f at the upper and lower frontiers of \mathcal{D} implies that if $I_{p, \theta}$ is the maximal open interval $I \subset (0, 1)$ such that for each $q \in I$, $(q, \phi_{p, \theta}(q)) \in \mathcal{D}$, then $\sup_{q \in I_{p, \theta}} \phi_{p, \theta}(q) = h \wedge \bar{\theta}$. If $\pi^F(\underline{\theta}) > p_0$, the result follows readily from this observation, since the Cauchy-Lipschitz theorem can be applied in a neighborhood of $p_0 \vee \pi^B(\underline{\theta})$. Suppose now that $\pi^F(\underline{\theta}) \leq p_0$. Let $\{\vartheta_n\}$ be a decreasing sequence in $(\phi^F(p_0), \phi^B(p_0))$ converging to $\phi^F(p_0)$. For each $n \in \mathbb{N}$, there exists a unique solution $\phi_n = \phi_{p_0, \vartheta_n}$ on \mathcal{D} to (3) starting at (p_0, ϑ_n) , and defined on a maximal interval $I_n = I_{p_0, \vartheta_n}$. If $n \geq m$, then $\sup I_n \geq \sup I_m$ and $\phi_n(p) \leq \phi_m(p)$ for each $p \in [p_0, \sup I_m)$. Note that if $\bar{\theta} < h$, the functions $\{\phi_n\}$ are not defined on the same set. To remedy this, we take p_∞ to be the limit of $\{\sup I_n\}$ as $n \rightarrow +\infty$, which is well-defined since the sequence is increasing and bounded above by $\pi^F(\bar{\theta})$. Then, for each $n \in \mathbb{N}$, we extend ϕ_n by continuity by setting $\phi_n(p) = \bar{\theta}$ for each $p \in [\sup I_n, p_\infty]$. Since the sequence $\{\phi_n\}$ is bounded below by ϕ^F , it admits a pointwise limit ϕ on $[p_0, p_\infty]$. We now show that ϕ is the required solution. Since ϕ obviously satisfies the boundary condition at p_0 , we must only show that ϕ is the unique solution to (3) starting at $(p_0, \phi^F(p_0))$.

First, we show that for each $p \in (p_0, p_\infty)$, $\phi(p) > \phi^F(p)$. Suppose the contrary holds for some p . Let $\delta > 0$. Since $f(q, \theta)$ goes to $+\infty$ as $(q, \theta) \rightarrow (p, \phi^F(p))$ in \mathcal{D} and $\partial\phi^F$ is bounded, there exists an open ball $B \subset (0, 1) \times (l, h)$ centered at $(p, \phi^F(p))$ such that $f(q, \theta) \geq \partial\phi^F(q) + \delta$ for each $(q, \theta) \in B \cap \mathcal{D}$. If $\{\phi_n(p)\}$ tends to $\phi^F(p)$ as $n \rightarrow \infty$, then for each $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that $\phi_N(p) - \phi^F(p) < \varepsilon/2$. Since ϕ_N and ϕ^F are continuous, and the $\{\phi_n\}$ form by construction a decreasing sequence of functions, there exists a neighborhood $U \subset (p_0, p_\infty)$ of p such that $\phi_n(q) - \phi^F(q) \leq \varepsilon$ for each $q \in U$ and $n \geq N$. Hence, by choosing ε small enough, we can ensure that for each $q \in U$ and $n \geq N$, $(q, \phi_n(q)) \in B$ and thus $\partial\phi_n(q) \geq \partial\phi^F(q) + \delta$. Integrating from $u_0 = \inf U$, it follows that for each $n > N$, $\phi_n(p) - \phi_n(u_0) \geq \phi^F(p) - \phi^F(u_0) + \delta(p - u_0)$. As $\phi_n(u_0) > \phi^F(u_0)$ for each $n \geq N$, it follows that $\{\phi_n(p)\}$ cannot converge to $\phi^F(p)$, a contradiction.

Next, we prove that ϕ satisfies (3) on (p_0, p_∞) . For each $(n, p) \in \mathbb{N} \times (p_0, 1)$, $\phi_n(p) > \phi(p)$ and $(p, \phi(p)) \in \mathcal{D}$. Since $f \in \mathcal{C}^1(\mathcal{D}, \mathbb{R}_+)$, for each $p \in (p_0, p_\infty)$ there exists a neighborhood $V \subset (p_0, p_\infty)$ of p and a real number $M > 0$ such that for each $(n, q) \in \mathbb{N} \times V$, $\partial\phi_n(q) = f(q, \phi_n(q)) \leq M$ and therefore $|\phi_n(q) - \phi_n(r)| \leq M|q - r|$ for each $(n, q, r) \in \mathbb{N} \times V \times V$. This implies that the sequence $\{\phi_n\}$ is equicontinuous on V , and thus converges uniformly on V to its pointwise limit ϕ by Ascoli's theorem. Since f is locally Lipschitzian with respect to θ on \mathcal{D} , there exists a real number $\tilde{M} > 0$ such that for each $(n, p) \in \mathbb{N} \times V$, $|f(p, \phi_n(p)) - f(p, \phi(p))| \leq \tilde{M}|\phi_n(p) - \phi(p)|$, so that the sequence $\{f(\cdot, \phi_n(\cdot))\}$ converges uniformly to $f(\cdot, \phi(\cdot))$ on V . Letting $v_0 = \inf V$, we can therefore take the limits as $n \rightarrow \infty$ on both sides of $\phi_n(q) = \phi_n(v_0) + \int_{v_0}^q f(r, \phi_n(r)) dr$ for each $q \in V$, which yields the desired result.

We now show that ϕ is the unique solution to (3) satisfying the equilibrium boundary conditions. Supposing the contrary, let ϕ and $\tilde{\phi}$ be two solutions to (3) starting at $(p_0, \phi^F(p_0))$, and such that $\tilde{\phi}(\bar{p}) > \phi(\bar{p})$ for some $\bar{p} \in (p_0, 1)$. Since $\tilde{\phi}$ is increasing and continuous on (p_0, \bar{p}) , there exists exactly one $\underline{p} \in (p_0, \bar{p})$ such that $\tilde{\phi}(\underline{p}) = \phi(\underline{p})$. For each $(q, \theta) \in \mathcal{D}$, let $\varphi(q, \theta) = f(q, \theta)g(\theta)/[1 - G(\theta)]$. Consider now the ratio $R(p, \phi, \tilde{\phi}) = [1 - G(\phi(p))]/[1 - G(\tilde{\phi}(p))]$. Integrating (3) respectively for ϕ and $\tilde{\phi}$ on (p, \bar{p}) and (p, \underline{p}) and subtracting yields:

$$\ln(R(p, \phi, \tilde{\phi})) = \int_p^{\underline{p}} [\varphi(q, \phi(q)) - \varphi(q, \tilde{\phi}(q))] dq + \int_{\underline{p}}^{\bar{p}} \varphi(q, \phi(q)) dq \quad (12)$$

for each $p \in (p_0, \underline{p})$. Note that $\tilde{\phi}(\bar{p}) > \phi(\bar{p})$ implies that $\tilde{\phi}(q) > \phi(q)$ for each $q \in (p_0, \underline{p})$. It is easy to check that the mapping $\theta \mapsto \varphi(q, \theta)$ is decreasing on $(\phi^F(q), \phi^B(q))$ for each $q \in (0, 1)$. Hence, for each $p \in (p_0, \underline{p}]$, $R(p, \phi, \tilde{\phi}) > 1$, and by (12) it is decreasing in p . Therefore $R(p, \phi, \tilde{\phi})$ cannot converge to one as $p \downarrow p_0$, a contradiction.

To conclude, observe that (7) implies that for any $\theta \in \Theta(p_0)$, the equilibrium trigger $\pi(\theta)$ remains optimal off the equilibrium path for beliefs in $\pi(\Theta(p_0))$. If $\bar{\theta} \geq h$, then the description of the equilibrium is complete. If $\bar{\theta} < h$, then one must describe the equilibrium behavior on $(\pi(\bar{\theta}), \pi^F(\bar{\theta}))$. We shall assume that player 1 will immediately invest, while player 2 will invest at $\phi^F(\theta^2)$. This is consistent with each player believing that his opponent is of type $\bar{\theta}$. It is not difficult to see that any other system of beliefs about costs defined over this region of beliefs about v would not alter the behavior of players on the equilibrium path. \square

Proof of Proposition 3. We need only to show that there are no discontinuous PBE. It is easy to

show that if $\theta \in \Theta(p_0)$ is a discontinuity point of π^i , then π^j coincides with π^B on $\phi^B([\pi^i(\theta^-), \pi^i(\theta^+)])$. The bulk of the proof consists to show that for each $i \in I$, if $\theta \in \Theta(p_0)$ is a discontinuity point of player j 's investment trigger, then π^i is continuous and $\pi^j > \pi^i$ on $[\theta, h)$. Suppose instead that π^i has a discontinuity at some $\theta' \in [\theta, h)$. Then π^j coincides with π^B on $\phi^B([\pi^i(\theta'^-), \pi^i(\theta'^+)])$. Since π^j cannot jump downward at θ' , $\tilde{\theta} = \inf\{\theta' \mid \pi^i \text{ has a discontinuity at } \theta'\} > \theta$. Note that by construction, we must have $\pi^i < \pi^j$ on $[\theta, \tilde{\theta})$ since π^i has no discontinuity on this interval. Hence $\lim_{\theta \uparrow \tilde{\theta}} \pi^i(\theta) = \pi^j(\tilde{\theta}) = \pi^B(\tilde{\theta})$. Consider now the ratio $R(p, \phi^j, \phi^i) = [1 - G(\phi^j(p))]/[1 - G(\phi^i(p))]$, extending ϕ^j by requiring it to be constant where π^j is discontinuous. Since $\pi^i < \pi^j$ on $[\theta, \tilde{\theta})$, $R(\cdot, \phi^j, \phi^i) > 1$ on $[\pi^j(\theta^+), \pi^B(\tilde{\theta}))$. We now prove that $R(p, \phi^j, \phi^i)$ is increasing in p on $(\pi^j(\theta^+), \pi^B(\tilde{\theta}))$, thereby contradicting the fact that it must converges to one as $p \uparrow \pi^B(\tilde{\theta})$. Three cases must be distinguished. Consider first a value of $p \in [\pi^-, \pi^+)$ where $[\pi^-, \pi^+] \subset (\pi^j(\theta^+), \pi^B(\tilde{\theta}))$ is an interval of discontinuity of π^j . Then $\bar{\partial}^+ R(p, \phi^j, \phi^i) \geq 0$ since $\phi^i = \phi^B$ and ϕ^j is constant on any interval $[p, \tilde{p}]$ for \tilde{p} close enough to p . Consider next a value of $p \in (\pi^j(\theta^+), \pi^B(\tilde{\theta}))$ such that the derivative of ϕ^j exists and is strictly positive at p . The analogue of the first-order condition (3) is then satisfied by both ϕ^i and ϕ^j at p , from which it follows that $\partial R(p, \phi^j, \phi^i) = R(p, \phi^j, \phi^i) [\varphi(p, \phi^j(p)) - \varphi(p, \phi^i(p))] \geq 0$ since the mapping $\varphi(q, \cdot)$ is decreasing on $(\phi^F(q), \phi^B(q))$ for each $q \in (0, 1)$. Last, consider a value of p such that p is the upper bound of an interval of discontinuity of π^j . From the formula $f \bar{\partial}^+ g + g \bar{\partial}^+ f \geq \bar{\partial}^+(fg)$ which is valid for any nonnegative and continuous functions f and g , we get $R(p, \phi^j, \phi^i) \bar{\partial}^+[1 - G(\phi^i(p))] + [1 - G(\phi^i(p))] \bar{\partial}^+ R(p, \phi^j, \phi^i) \geq \bar{\partial}^+[1 - G(\phi^j(p))]$. A sufficient condition for $\bar{\partial}^+ R(p, \phi^j, \phi^i) \geq 0$ is that $\bar{\partial}^+[1 - G(\phi^j(p))] = 0$. Since p is the upper bound of an interval of discontinuity for player j , one must have $\pi^i(\phi^B(p)) = p$. Proceeding as in the proof of Lemma 2, we have, for any $\tilde{p} \in (p, \pi^F(\phi^B(p)))$,

$$[1 - G(\phi^j(\tilde{p}))] \frac{p\Delta + l - \phi^B(p) - D^B(p, \tilde{p}) [\tilde{p}\Delta + l - \phi^B(p)]}{V^F(p, \phi^B(p)) - p\Delta - l + \phi^B(p)} \geq G(\phi^j(\tilde{p})) - G(\phi^j(p)) \geq 0.$$

Dividing both sides of this inequality by $\tilde{p} - p$ and letting $\tilde{p} \downarrow p$, we conclude that:

$$\begin{aligned} \underline{\partial}^+ G(\phi^j(p)) &= \liminf_{\tilde{p} \downarrow p} \frac{G(\phi^j(\tilde{p})) - G(\phi^j(p))}{\tilde{p} - p} \\ &\leq \frac{1 - G(\phi^j(p))}{V^F(p, \phi^B(p)) - p\Delta - l + \phi^B(p)} \\ &\quad \times \liminf_{\tilde{p} \downarrow p} \frac{p\Delta + l - \phi^B(p) - D^B(p, \tilde{p}) [\tilde{p}\Delta + l - \phi^B(p)]}{\tilde{p} - p}, \end{aligned}$$

which is equal to zero from L'Hôpital's rule and the smooth pasting condition for the single-player problem with signal B , together with the fact that $\partial_1^- D^B(p, p) = -\partial_2^+ D^B(p, p)$. Since ϕ^j is non-decreasing, it follows that $\bar{\partial}^+[1 - G(\phi^j(p))] = \underline{\partial}^+ G(\phi^j(p)) = 0$, as claimed. Overall, $\bar{\partial}^+ R(p, \phi^j, \phi^i) \geq 0$ everywhere on $(\pi^j(\theta^+), \pi^B(\tilde{\theta}))$, hence $R(\cdot, \phi^j, \phi^i)$ is increasing on this interval, which yields the required contradiction. An identical argument, but this time applied to h instead of $\tilde{\theta}$ implies that when $\tilde{\theta} > h$ and therefore $G(h) = G(\phi^i(1)) = G(\phi^j(1)) < 1$, $R(p, \phi^j, \phi^i)$ is increasing and strictly greater than one over $[\pi^j(\theta^+), 1]$, and hence cannot converge to one as $p \uparrow 1$. \square

References

- [1] P. Bolton, C. Harris, Strategic experimentation, *Econometrica* 67 (1999) 349-374.
- [2] A. Caplin, J. Leahy, Sectorial shocks, learning and aggregate fluctuations, *Rev. Econ. Stud.* 60 (1993) 777-794.
- [3] C. Chamley, D. Gale, Information revelation and strategic delay in a model of investment, *Econometrica* 62 (1994) 1065-1085.
- [4] M. Cripps, G. Keller, S. Rady, Strategic experimentation with exponential bandits, Mimeo, University of Munich, 2003.
- [5] A.K. Dixit, R.S. Pindyck, *Investment under Uncertainty*, Princeton University Press, Princeton, 1994.
- [6] D. Fudenberg, J. Tirole, A theory of exit in duopoly, *Econometrica* 54 (1986) 943-960.
- [7] P. Ghemawat, B. Nalebuff, Exit, *RAND J. Econ.* 16 (1985) 184-194.
- [8] K. Hendricks, D. Kovenock, Asymmetric information, information externalities, and efficiency: The case of oil exploration, *RAND J. Econ.* 20 (1989) 164-182.
- [9] B. Lambrecht, W. Perraudin, Real options and preemption under incomplete information, *J. Econ. Dynam. Control* 27 (2003) 619-643.
- [10] E. Maskin, J. Riley, Asymmetric auctions, *Rev. Econ. Stud.* 67 (2000) 413-438.
- [11] E. Maskin, J. Tirole, Markov perfect equilibrium, I: Observable actions, *J. Econ. Theory* 100 (2001) 191-219.
- [12] J. Riley, Strong evolutionary equilibrium and the war of attrition, *J. Theor. Biology* 82 (1980) 383-400.
- [13] R. Rob, Learning and capacity expansion under demand uncertainty, *Rev. Econ. Stud.* 58 (1991) 655-675.