Information acquisition and strategic investment timing*

Rishabh Kirpalani and Erik Madsen†

April 30, 2018

Abstract

We study a model of strategic investment timing with costly dynamic information acquisition. Two firms investigate a nonrival investment opportunity, by expending costly effort to stochastically uncover signals about its quality. Each firm can acquire at most one signal, and signals are conditionally independent across firms. Effort and signals are private, while investment is public and provides a channel for social learning. We characterize the set of perfect Bayesian equilibria, and find that social learning can lead to both reductions in effort (the traditional free-rider effect) and delay in investment even after a positive signal is acquired. We show that investment delay can actually mitigate the inefficiency of free-riding: when signal acquisition costs aren't too high, equilibria exhibiting investment delay improve aggregate welfare or even Pareto-dominate those which don't.

1 Introduction

Consider a firm deciding when (if ever) to invest in a risky project. At any time, the firm may devote costly effort towarding assessing the project's viability. It also has the opportunity to observe whether other firms decide to invest in the project. When should the firm acquire information about the project, and how long should it wait to see what other firms do? In this paper, we study a model with costly information acquisition, social learning, and strategic investment timing to answer these questions.

*PRELIMINARY. We would like to thank Nageeb Ali, David Pearce, and Andy Skrzypacz for helpful discussions and comments.

†Department of Economics, Pennsylvania State University (Email: rishabh.kirpalani@psu.edu) and Department of Economics, New York University (Email: emadsen@nyu.edu).

One important application motivating our model is the market for venture capital funding. Venture capitalists frequently delay project evaluation and investment in order to see whether another fund invests first, a phenomenon known in the industry as "waiting for social proof." The following passage from a popular handbook for entrepreneurial fundraising succinctly captures the situation:

> Assuming that you are talking with multiple potential investors, you can generally categorize them into one of three groups: leaders, followers, and everyone else...Your goal is to find a lead VC. This is the firm that is going to put down the term sheet [and] take a leadership role in driving to a financing... [A follower] seems interested, but doesn't really step up his level of engagement. This VC seems to be hanging around, waiting to see if there's any interest in your deal.[1]

Another important application is the decision by manufacturers of whether to adopt cost-saving or quality-enhancing innovations in production processes. We model the opportunities to strategically time both evaluation of and investment in projects which arise in such markets.

In our model, two firms have the opportunity to invest in a nonrival risky project. Ex ante the project has negative expected return. Firms may exert costly effort ("prospect") at each moment in time for the chance of receiving a binary signal which is informative about the project's value. Each firm can acquire at most one signal, and signals are conditionally i.i.d. As a result, firms would benefit from sharing signals. However, they cannot observe the prospecting decisions or the signal of the other firm; they can only see whether the other firm has invested.

We fully characterize the set of perfect Bayesian equilibria of this game under an assumption that signal acquisition costs aren't too high. There are exactly three equilibria, all in pure strategies. In the unique symmetric equilibrium, each firm prospects at the maximal rate until a cutoff time, after which it quits forever if it has not seen investment by the other firm. If at any time before the cutoff a firm receives a positive signal, it invests without delay. If it receives a negative signal it never invests. There are also two asymmetric "leader-follower" equilibria. In these equilibria, one firm takes the role of a lead investor. It prospects at the maximal rate until acquiring a signal and then invests immediately if and only if the signal is positive. Meanwhile the other firm, the follower, eventually both shirks from acquiring a signal and waits to invest in the event it does obtain a positive signal.

These equilibria exhibit both the free-riding inefficiency of strategic experimentation

[1] Feld and Mendelson (2016), pg. 44-45.

models, as well as the investment delays typical of strategic investment models. In particular, a key property of the leader-follower equilibrium is that when costs are sufficiently low, with positive probability the follower both acquires a signal and waits to see what the leader does even if the signal is positive. This occurs despite the fact that investment has positive expected returns conditional on the signal. This equilibrium therefore exhibits a strong form of "waiting to see", in addition to the free-riding effect of effort reduction by the follower.

We show that when costs and time discounting aren't too high, the leader-follower equilibrium improves aggregate welfare or even Pareto-dominates the symmetric no-delay equilibrium. This reveals an important interaction between free-riding and strategic investment. While each of these sources of delay constitutes an inefficiency when studied in isolation, when taken together the ability to delay investment can actually mitigate the costs of free-riding.

Our results depart from existing work on strategic experimentation in several ways. (See the literature review below for an overview of related work.) We find that adding strategic investment timing reduces the multiplicity of equilibria common in the literature. Our model yields exactly one symmetric and one asymmetric equilibrium, up to relabeling of players, which can be sharply characterized. In contrast, the typical experimentation model suffers from a large number of equilibria, even with only two players. As a result, predictions for welfare and free riding are ambiguous. We also find patterns of effort sharing which contrast with common outcomes in the literature. Our symmetric equilibrium exhibits a bang-bang structure with both firms exerting maximal effort until abandoning the project, whereas interior effort is typical of symmetric equilibria in the literature. And in our asymmetric equilibrium, if the follower exerts effort at all prior to investment by the leader, it does so only early in the game. In the literature, followers typically shirk initially and only jump in later on.

The rest of the paper is organized as follows. Section 1.1 briefly surveys related literature. Section 2 describes the model. Section 3 characterizes the set of perfect Bayesian equilibria of the model. Section 4 performs the welfare comparison betwen the different equilibria. Section 5 concludes.

1.1 Related literature

Our paper builds on two strands of the literature. The first strand studies collective experimentation by multiple agents.[2] In this setting, agents engage in social learning by observing

[2]See Horner and Skrzypacz (2017) for an excellent survey of these models.

the outcome of experimentation by other agents. Experimentation is typically modeled via a Poisson bandit framework. Agents have the opportunity to repeatedly pull the arm of a slot machine with unknown average payout, and must decide whether to abandon the machine by learning from past pulls of the arm. The profitability of the slot machine is partially or perfectly correlated across players, opening a channel for social learning. Bolton and Harris (1999), Keller, Rady, and Cripps (2005), and Keller and Rady (2010, 2015) adopt this framework under the assumption that the actions of each agent and payoffs from the slot machine are publicly observable. Bonatti and Hörner (2011, 2017) also use this formulation, but with private actions, so that only payoffs are observed.[3]

This literature focuses on the information spillovers of each agent's information acquisition to other agents' experimentation. A universal effect is the free-rider problem, where costly experimentation by one agent acts as a substitute for experimentation by others. Depending on the learning process, there may also be encouragement effects, where good news acquired by one agent spurs other agents to experiment more. These papers abstract from the endogenous timing of investment, as agents need never commit irreversibly to pulling the arm of the slot machine forever.

Another literature studies the timing of irreversible investment in risky projects when multiple players learn about the project privately but invest publicly. This literature takes seriously the strategic nature of the investment timing choice when agents learn from other agents' actions. Papers in this literature include Chamley and Gale (1994), Chari and Kehoe (2004), Rosenberg, Solan, and Vieille (2007), and Murto and Välimäki (2011, 2013). All of these papers feature exogenous information arrival, and so abstract from the choice of when and whether to experiment. A recent paper by Aghamolla and Hashimoto (2018) endogenize the precision of the time-zero signal in the model studied by Chamley and Gale (1994) and Murto and Välimäki (2013), but do not allow agents to acquire any further information over the course of the game.

Our paper lies at the intersection of these two literatures. We build a model which includes both endogenous information acquisition as well as the strategic timing of investment. These features also link our paper with several others which incorporate related assumptions. Ali (2018) studies the consequences of endogenous information acquisition in the context of a classic herding model, where players act in a pre-specified order. And Frick and Ishii (2016) build a model of investment timing, where investment boosts the arrival rate of public signals rather than signaling an agent's private information.

[3] Bonatti and Hörner (2011) also incorporate a payoff externality, in that a successful pull of the arm yields not only an informational benefit but also an immediate payout to all players in the game.

2 The model

Two firms have the opportunity to invest in a nonrival risky project of unknown quality. The project has underlying type θ and is either Good ($\theta = G$) or Bad ($\theta = B$). The payoff of the project is R if $\theta = G$, and 0 otherwise, with $R > 1$. Each firm is risk-neutral with discount rate $r > 0$. The project is indivisible, investment in the project is irreversible, and project outcomes are not observed until the end of the game. Each firm is free to invest in the project at any time $t \in \mathbb{R}_+$.

Both firms begin with common prior belief π_0 that the project is Good. Each firm can additionally exert costly effort to search for an informative signal about the project's type, an activity we will refer to as *prospecting*. A signal, when it arrives, is binary with $s \in \{H, L\}$, i.e. High and Low, and is distributed as $\Pr(s = H \mid \theta = G) = q^H$ and $\Pr(s = L \mid \theta = B) = q^L$ with $q^H, q^L > 1/2$. For a given belief $\mu \in [0, 1]$ that $\theta = G$, let

$$h(\mu) \equiv q^H \mu + (1 - q^L)(1 - \mu)$$

be the total probability that an arriving signal is High, and similarly

$$l(\mu) \equiv (1 - q^H)\mu + q^L(1 - \mu) = 1 - h(\mu)$$

be the total probability that an arriving signal is Low. The values $h(\mu)$ and $l(\mu)$ are the transition probabilities that a firm's posterior belief jumps up or down upon receiving a signal.

Each firm can obtain at most one signal, and firms observe conditionally IID signals. We will denote the posterior beliefs induced by one or more signals as follows: π_+ and π_{++} are the posteriors induced by one and two High signals, respectively; similarly π_- and π_{--} are the posteriors induced by one and two Low signals. Finally, π_{+-} is the posterior induced by one High and one Low signal. (Exchangeability implies that posterior beliefs are independent of the order of receipt of signals.) Given that High signals are more likely when the state is Good, and conversely for Low signals when the state is Bad, $\pi_{++} > \pi_+ > \pi_0, \pi_{+-} > \pi_- > \pi_{--}$. Note that in general $\pi_{+-} \neq \pi_0$, except in the special case when $q^H = q^L$. If $q^H > q^L$ then $\pi_{+-} < \pi_0$, and if $q^H < q^L$ then $\pi_{+-} > \pi_0$.

Assumption 1. $\pi_0 < 1/R < \pi_+$.

Under this assumption, investment in the project is ex ante unprofitable, but becomes profitable conditional on observation of a High signal.[4] Note that $1/R < \pi_+$ holds so long as q^H is sufficiently large, i.e. a High signal is sufficiently correlated with a Good state.

Assumption 2. *$\pi_{+-} < 1/R$.*

This assumption is satisfied so long as q^L is not too much smaller than q^H. Under this assumption, even after observing a High signal making investment profitable, observation of a Low signal would push beliefs back below the breakeven threshold. Without this assumption no equilibrium would exhibit "wait and see" behavior, since the optimality of investment following receipt of a High signal would not depend on the information obtained by the other firm. (Note that $\pi_{+-} < 1/R$ does not inevitably imply waiting to see, and indeed we will construct an equilibrium in which such behavior does not arise.)

Prospecting is a dynamic process unfolding in continuous time. At each instant dt, firm i's signal arrives with probability $\lambda\, dt$ when firm i exerts effort $C(\lambda)\, dt$. We will maintain the assumption of a linear cost structure:

$$C(\lambda) = \begin{cases} c\lambda, & \lambda \in [0, \overline{\lambda}], \\ \infty, & \lambda \in (\overline{\lambda}, \infty) \end{cases}$$

for some constant marginal cost $c > 0$ and maximum prospecting rate $\overline{\lambda}$, both of which are symmetric across firms. Conditional on prospecting rates, signal arrival times are independent across firms and independent of the type of the project.

Firms cannot observe each other's signals or prospecting intensities, nor can they observe whether another firm has received a signal. There are also no communication channels between firms. However, all investment decisions are public, introducing a channel for social learning. Our goal is to characterize all perfect Bayesian equilibria of this game.

2.1 Strategies, payoffs, and beliefs

For each $i \in \{1, 2\}$, let s^i be the process tracking what signal, if any, firm i has received at each moment in time. That is, $s^i_t \in \{\emptyset, H, L\}$ for each t, with $s^i_0 = \emptyset$ and s^i jumping at most once to either H or L at the time a signal is received. We will use $\nu^i = \inf\{t \;:\; s^i_t \neq \emptyset\}$ to denote the first time firm i receives a signal. Also let $\mathbb{F}^i$ be the filtration generated by s^i and a randomization device privately observed by i, with the latter allowing for mixed strategies.

[4] The case $\pi_+ < 1/R$ is uninteresting, as the unique equilibrium involves no prospecting and no investment by either firm. To see this, note that any firm investing first must have posterior beliefs weakly below π_+, meaning investment would be unprofitable. Hence no firm ever invests, and so never acquires a signal.

Definition 1. *A strategy σ^i for firm $i \in \{1,2\}$ is a tuple $\sigma^i = (\lambda^i(T), \iota^i(T))_{T\in\{\emptyset\}\cup\mathbb{R}_+}$, where each $\lambda^i(T)$ is a $[0,\overline{\lambda}]$-valued $\mathbb{F}^i$-adapted process and each $\iota^i(T)$ is a $\{0,1\}$-valued $\mathbb{F}^i$-adapted process.*

A strategy σ^i consists of a prospecting process λ^i and an investment decision process ι^i which may condition on the timing of any past investment by the other firm. So long as firm i has not observed investment by firm $-i$, it prospects at rate $\lambda^i(\emptyset)_t$ until a signal is observed. And it invests at time $\tau^i(\emptyset) = \inf\{t \ : \ \iota^i(\emptyset) = 1\}$. After i has observed $-i$ invest at some time T, the firm prospects at rate $\lambda^i(T)_t$ until a signal is observed, and it invests at time $\tau^i(T) = \inf\{t \geq T \ : \ \iota^i(T) = 1\}$. This construction allows for the possibility that firm i, upon observing investment by firm $-i$, immediately follows and invests "afterward at the same time".[5] In particular, consider a strategy and state of the world in which firm 1 invests at time T. It will be important to allow for strategies for firm 2 under which $\tau^2(\emptyset) > T$, so that firm 2 would not invest at time T on its own, but under which $\tau^2(T) = T$, so that investment by firm 1 spurs firm 2 to act immediately.

For each firm, λ^i and ι^i are adapted to the history of its own signal as well as its randomization device. Because prospecting does not occur after a signal has arrived, the conditioning of λ^i on the signal history is redundant; however, it is important that the firm be allowed to condition on the randomization device to allow for mixed strategies. Additionally, ι^i contains information beyond what is necessary to construct the investment time τ^i. This is because we will be interested in characterizing perfect Bayesian equilibria, which require a notion of optimality off the equilibrium path. Supposing that firm i has deviated and failed to invest at time τ^i, then at time $t > \tau^i$ firm i's strategy induces the continuation investment time $\widetilde{\tau}^i = \inf\{t' \geq t \ : \ \iota^i = 1\}$. This allows for the important possibility that a firm who initially finds investment profitable may eventually become pessimistic and prefer not to invest immediately.

Fix a strategy profile $\sigma = (\sigma^1, \sigma^2)$. Firm i's expected payoff under σ is then

$$V^i(\sigma) = \mathbb{E}\left[(R\,\mathbf{1}_{\{\theta=G\}} - 1)e^{-r\tau^i(\sigma)} - c\int_0^{\min\{\nu^i,\tau^i(\sigma)\}} e^{-rt}\lambda^i(\sigma)_t\,dt\right].$$

[5]In this respect, we follow the construction of strategy profiles used by Murto and Valimaki (2011), who model "exit waves" of firms who follow others out of the market with no delay. This model timing is necessary in continuous time to ensure existence of best replies. Otherwise a firm observing another investing/exiting might want to follow "as soon as possible", meaning any strategy of delaying a finite amount of time could be improved upon by delaying a bit less.

where

$$\lambda^i(\sigma)_t = \begin{cases} \lambda^i(\emptyset)_t, & t < \tau^{-i}(\emptyset), \\ \lambda^i(\tau^{-i}(\emptyset))_t, & t \geq \tau^{-i}(\emptyset) \end{cases}$$

and

$$\tau^i(\sigma) = \begin{cases} \tau^i(\emptyset), & \tau^i(\emptyset) \leq \tau^{-i}(\emptyset), \\ \tau^i(\tau^{-i}(\emptyset)), & \tau^i(\emptyset) > \tau^{-i}(\emptyset). \end{cases}$$

The first term in $V^i(\sigma)$ is the discounted payoff from investing in the project at time $\tau^i(\sigma)$. The second term is the cumulative discounted cost of prospecting according to $\lambda^i(\sigma)$. The upper limit of integration reflects the fact that prospecting stops whenever either a signal arrives (at time ν^i) or the firm invests (at time $\tau^i(\sigma)$).

Given a strategy profile, it will be useful to define each firm's posterior beliefs about θ at each moment in time conditional on their private signal and lack of investment by the other firm.[6] We will let $\mu^i(t) = \Pr(\theta = G \mid s^i_t = \emptyset, \tau^{-i}(\emptyset) \geq t)$ be firm i's posterior beliefs at time t conditional on having seen no signal so far, with $\mu^i_+(t) = \Pr(\theta = G \mid s^i_t = H, \tau^{-i}(\emptyset) \geq t)$ and $\mu^i_-(t) = \Pr(\theta = G \mid s^i_t = L, \tau^{-i}(\emptyset) \geq t)$ similarly representing firm i's beliefs conditional on having observed a High and Low signal, respectively. By Bayes' rule, these three probabilites are linked via the identities $\mu^i_+(t) = q^H \mu^i(t)/h(\mu^i(t))$ and $\mu^i_-(t) = (1 - q^L)\mu^i(t)/l(\mu^i(t))$.

Note that if $\tau^{-i}(\emptyset) = t$, then firm i is not able to observe this fact until after making his own initial investment decision, so its beliefs at time t cannot condition on this fact. Hence the appropriate conditioning for "no investment by firm $-i$ up to time t" is the event $\tau^{-i}(\emptyset) \geq t$.

2.2 Belief updating under social learning

In this subsection we discuss the impact of observing investment on a firm's beliefs about the quality of the project.

Fix a firm i and a prospecting strategy λ^{-i} for each firm $-i$, and suppose that firm $-i$ invests immediately upon receiving a High signal and does not invest otherwise. Let $\Omega^{-i}(t)$ be the (ex ante) probability that $-i$ has received no signal by time t. Then

$$\Omega^{-i}(t) = \exp\left(-\int_0^t \lambda^{-i}(s)\, ds\right),$$

[6]Lemma 2 shows that these beliefs are always uniquely pinned down by Bayes' rule in any PBE strategy profile.

and so by Bayes' rule

$$\mu^i(t) = \frac{\Omega^{-i}(t)\pi_0 + (1 - \Omega^{-i}(t))l(\pi_0)\pi_-}{\Omega^{-i}(t) + (1 - \Omega^{-i}(t))l(\pi_0)}.$$

Note that in stark contrast to Poisson bandit models, each firm's posterior beliefs about the state are independent of their own history of prospecting. Lack of signal acquisition in our model does not signal anything, positive or negative, about the true project state; no news truly is no news until a signal arrives. Nonetheless, if the other firm is prospecting and investing, then each firm's beliefs $\mu^i(t)$ *do* deteriorate over time due to the negative inference from continued lack of investment by the other firm. More precisely, continued inaction by firm $-i$ leads firm i to infer that $-i$ is likely dormant due to receipt of a Low signal, rather than due to a long string of bad luck leading to no signal.

Of course, the rate at which beliefs deteriorate is an endogenous property of a particular equilibrium, and in particular will dependent on whether each firm expects the other to be prospecting and investing, or shirking and waiting. This linkage of beliefs about the state and about the (unobserved) strategy of one's opponent will play a crucial role in the construction of equilibria in this setting.

2.3 Autarky

Consider a single firm prospecting and investing on its own, or equivalently with the social learning channel of our model shut down. We will refer to this benchmark setting as the *autarky case*. So long as the firm has acquired no signal, every continuation game is isomorphic to the original game. Thus an optimal prospecting strategy is stationary, except in the knife-edge case in which the firm just breaks even in expectation by prospecting. This behavior is very different from the cutoff strategies which are optimal when learning from Poisson bandits, and is driven by the very different learning dynamics, as mentioned above. Also note that once the firm has acquired a signal, no further information may be acquired and the firm faces a simple choice of whether to invest or abandon the project once and for all.

Given our assumptions on posterior beliefs following signal acquisition, the firm invests immediately if it receives a High signal, and abandons the project if it receives a Low signal. The HJB equation characterizing the firm's autarkic value function V prior to obtaining a signal is therefore

$$rV = \sup_{\lambda \in [0,\overline{\lambda}]} \left\{ \lambda(-c + h(\pi_0)(\pi_+ R - 1) - V) \right\}.$$

Lemma 1. *The firm's expected profits under autarky are*

$$V = \frac{\overline{\lambda}}{\overline{\lambda} + r} \max\{h(\pi_0)(\pi_+ R - 1) - c, 0\}.$$

If $h(\pi_0)(\pi_+ - R) > c$, the unique optimal prospecting strategy under autarky is $\lambda_t = \overline{\lambda}$ for all time. If $h(\pi_0)(\pi_+ - R) < c$, the unique optimal prospecting strategy under autarky is $\lambda_t = 0$ for all time. In the knife-edge case $h(\pi_0)(\pi_+ - R) = c$, every prospecting strategy is optimal.

The firm's behavior under autarky depends critically on the sign of $h(\pi_0)(\pi_+ R - 1) - c$, which weighs the average profits from signal acquisition against the costs of acquiring the signal. This profit threshold may be equivalently written as a threshold in initial beliefs about the state prior to signal acquisition. Given initial beliefs μ about the state, receipt of a High signal updates beliefs to $\mu_+ = q^H \mu / h(\mu)$. It is easy to check that $h(\mu)(\mu_+ R - 1) - c$ is an increasing function of μ. In fact, it may be written

$$h(\mu)(\mu_+ R - 1) - c = K(\mu - \pi^A),$$

where $K \equiv q^H(R-1) + (1 - q^L) > 0$ and

$$\pi^A \equiv \frac{c + (1 - q^L)}{q^H(R-1) + (1 - q^L)}.$$

We will refer to π^A as the *autarky threshold*, as for prior beliefs below π^A an autarkic firm optimally abandons the project without prospecting, while for beliefs above π^A it prospects to acquire a signal.

Assumption 3. *c is sufficiently small that $\pi_0 > \pi^A$.*

Note that as c approaches 0, π^A approaches the prior belief $\underline{\mu}$ for which $\underline{\mu}_+ R - 1 = 0$. Since by assumption $\pi_+ R - 1 > 0$, it must be that $\pi_0 > \pi^A$ for sufficiently small c. This assumption corresponds to the interesting case in which prospecting is possible in equilibrium. If $\pi_0 < \pi^A$, then the unique equilibrium outcome has both firms abandon the project immediately.

3 Equilibrium analysis

In this section we show that our model has exactly three perfect Bayesian equilibria. One equilibrium is symmetric and exhibits no delay in investment, but features a time threshold past which both firms stop prospecting and abandon investigation of the project if they haven't already observed the other firm invest. The other two have a "leader-follower" structure, and differ only in the firm playing each role. One firm, the leader, never abandons the project or delays investment if it observes a High signal. The other, the follower, eventually both stops prospecting and delays investing following receipt of a High signal until the leader acts first.

The development of this section proceeds by backward induction. We progressively identify information sets whose continuation strategies in equilibrium can be pinned down, and then roll back the game tree to further characterize equilibrium behavior at earlier information sets. Ultimately this will yield a complete description of equilibrium behavior in every information set of the game.

3.1 Preliminaries

We begin by identifying several information sets in which equilibrium behavior can be directly characterized.

Definition 2. *A firm's strategy is* regular *if:*

- *Investment never occurs after receipt of a Low signal,*
- *Investment without a signal occurs only in histories in which the other firm has invested.*

Regular strategies treat receipt of a Low signal as a terminal node of the game, with the firm abandoning the project at that point. They also ignore the opportunity to invest until some information about the project's value has been received, either by receiving a signal or observing investment.

Lemma 2. *In any perfect Bayesian equilibrium, each firm's strategy is regular.*

The intuition for this lemma is very simple. Once a Low signal has been received, a firm's posterior beliefs about the state can never rise above π_{+-}, no matter its inference about the other firm's information. Then as $\pi_{+-} < 1/R$ by assumption, it must never invest once in possession of a Low signal. Given this, observing lack of investment by the other firm must be weakly bad news about the state, so a firm's beliefs absent a signal or observation of investment cannot rise above π_0 in equilibrium. As $\pi_0 < 1/R$, investment can't be optimal at this point in the game either.

Lemma 3. *Fix any history in which firm i has obtained no signal and has observed firm $-i$ invest at some point in the past. Then in any perfect Bayesian equilibrium:*

- *If* $\pi_+ R - 1 > \frac{\overline{\lambda}}{\overline{\lambda}+r}(h(\pi_+)(\pi_{++}R - 1) - c)$, *firm* i *invests immediately,*
- *If* $\pi_+ R - 1 < \frac{\overline{\lambda}}{\overline{\lambda}+r}(h(\pi_+)(\pi_{++}R - 1) - c)$, *firm* i *prospects at rate* $\overline{\lambda}$ *until obtaining its own signal, then invests immediately iff that signal is High.*

(In the knife-edge case, there are many optimal continuation strategies. Any strategy which either prospects at the maximum rate or invests immediately at each moment in time will be optimal.) This lemma characterizes equilibrium continuation play after one firm has exited the game by investing. Supposing the remaining firm has not yet obtained a signal, this exit reduces the game to an autarkic problem with beliefs π_+.[7]

The solution to this problem depends on the cost and maximum rate of prospecting, similar to our analysis of the autarky case. If an additional signal can be obtained quickly and at relatively low cost, the firm will do so. Otherwise, the firm will simply follow the lead of the first mover and invest immediately. Recall that in contrast to Poisson bandit models, when a signal is worth obtaining the firm never quits prospecting until it does so.

Definition 3. *Signals are* complements *if* $\pi_+ R - 1 \leq \frac{\overline{\lambda}}{\overline{\lambda}+r}(h(\pi_+)(\pi_{++}R - 1) - c)$. *Otherwise they are* substitutes.

In particular, when either $\overline{\lambda}/r$ or c is sufficiently large, signals are substitutes, and when both are small, they are complements. Whether signals are complements or substitutes will turn out to have important implications for the structure of the equilibrium set.

It will be useful for the work that follows to define the upper bound

$$\overline{c} \equiv h(\pi_+)(\pi_{++}R - 1) - (\pi_+ R - 1)$$

on costs. The condition $c \leq \overline{c}$ says roughly that signals would be complements if the firm could prospect for a signal arbitrarily quickly. It is always satisfied under complementary signals, and is also satisfied when signals are substitutes so long as at least one of c or R is not too large. Our main characterization of the structure of equilibrium will assume that $c \leq \overline{c}$.[8]

[7] There is a technical subtlety with the construction of posterior beliefs after such a history, since all such histories are of measure 0. We discuss this issue further in the Appendix, and choose a definition of perfect Bayesian equilibrium for our setting ensuring that firm i's posterior beliefs assign probability 1 to firm $-i$ having observed a High signal.

[8] We can show that when this bound is violated, there continues to exist a unique symmetric equilibrium. In contrast to the structure for low c, it features interior declining effort and no abandonment of the project.

3.2 Investment delay in equilibrium

In principle, the optimal strategy for a firm after acquiring a High signal could be quite complicated - it could invest immediately; wait forever to see if the other firm invests; wait for some time interval and then invest if it sees no activity; or even randomize over its waiting time. In this section we establish a key lemma showing that without loss of generality, all optimal investment strategies may be assumed to take a very simple cutoff rule form, so long as one's opponent employs a regular strategy. Before the cutoff time the firm invests immediately upon obtaining a High signal, and afterward the firm waits forever absent investment by the other firm.

Lemma 4. *Suppose firm $-i$ uses a regular strategy σ^{-i}. Then there exists a cutoff time $t_i^* \in \mathbb{R}_+ \cup \{\infty\}$ and a best reply σ^i for firm i such that, at any time t and history in which firm i has obtained a High signal and seen no investment, under σ^i firm i invests immediately if $t < t_i^*$ and waits until firm $-i$ invests if $t > t_i^*$.*

Further, any other best reply $\widetilde{\sigma}^i$ to σ^{-i} induces the same distribution over investment times as σ^i, regardless of the strategy chosen by firm $-i$.

This lemma allows us to summarize each firm's investment strategy in any equilibrium as a cutoff rule $t_i^* \in \mathbb{R}_+ \cup \{\infty\}$ such that firm i invests immmediately upon obtaining a high signal if $t < t_i^*$, and otherwise waits until firm $-i$ invests. Note that in some cases, there may be other best replies which do not take this form.[9] However, the lemma ensures that all other best replies differ only off-path, regardless of the strategy actually chosen by firm $-i$. So we do not rule out any equilibria by restricting attention to cutoff strategies.

In particular, this lemma rules out randomization over investment timing. Such strategies require the firm to be indifferent between investing at different times absent action by the other firm. But that lack of action is always weakly bad news about the state, regardless of the other firm's prospecting strategy. And given that delaying investment imposes a discounting cost, it can never be the case that a firm is indifferent between investing now and in the future absent good news. So randomization can never be part of an equilibrium investment strategy.

We next characterize the possible equilibrium values of each t_i^*. To do this, we construct a critical belief threshold at which a firm optimally switches between investing and waiting, supposing that its opponent is prospecting and investing as rapidly as possible. Suppose

[9] The one case in which this can occur is if $\mu_+^i(t)$ eventually drops to $1/R$ and then stays fixed there forever. In this case any investment policy is optimal once beliefs have dropped to $1/R$. However, as the proof of the lemma establishes, no best reply by firm i ever leaves it in possession of a High signal in such a history.

that a firm has current posterior beliefs $\mu \in [\pi_-, \pi_0]$ about the state, following a history in which it has no signal and has not seen investment. If the firm then receives a High signal, let $\Delta(\mu)$ be the difference in continuation payoffs between waiting for the other firm to invest and investing immediately.

Lemma 5. *Δ is a continuous, strictly decreasing function of μ, and $\Delta(\pi_-) > 0$. Also,*

$$\Delta(\pi_0) = \frac{\overline{\lambda}}{\overline{\lambda} + r} h(\pi_+)(\pi_{++}R - 1) - (\pi_+ R - 1).$$

In particular, $\Delta(\pi_0) > 0$ whenever signals are complements.

This lemma shows that the value of waiting relative to investing rises the lower are beliefs at the time of receipt of a High signal. Intuitively, the discounting cost of waiting to learn more about the state drops as current beliefs about the payoff of the project become more pessimistic. The lemma establishes that for sufficiently low beliefs waiting is superior to investing, and when signals are complements waiting is more profitable at all beliefs.

Now, define a belief threshold $\mu^* \in (\pi_-, \pi_0]$ to be the belief at which the continuation payoffs of investing and waiting are equalized:

$$\mu^* \equiv \begin{cases} \pi_0, & \Delta(\pi_0) \geq 0, \\ \Delta^{-1}(0), & \Delta(\pi_0) < 0. \end{cases}$$

(If waiting is always superior to investing, then by convention we set μ^* equal to time-zero beliefs.) When mapped onto a corresponding time at which these beliefs are reached, μ^* will pin down one of two possible cutoff times supportable in equilibrium.

To find the corresponding cutoff time, let $\mu^{\overline{\lambda}}(t)$ be a firm's beliefs at time t supposing it has received no signal and seen no investment, when its opponent prospects at rate $\overline{\lambda}$ and invests immediately forever. Using the general updating formula derived in section 2.2, these beliefs may be explicitly written

$$\mu^{\overline{\lambda}}(t) \equiv \frac{\pi_0 e^{-\overline{\lambda}t} + (1 - e^{-\overline{\lambda}t}) l(\pi_0)\pi_-}{e^{-\overline{\lambda}t} + (1 - e^{-\overline{\lambda}t}) l(\pi_0)}.$$

This function is continuous and strictly decreasing in t, with $\mu^{\overline{\lambda}}(0) = \pi_0$ and $\mu^{\overline{\lambda}}(\infty) = \pi_-$. Then the time threshold $t^{**} \equiv (\mu^{\overline{\lambda}})^{-1}(\mu^*)$ is well-defined. Keep in mind that Δ, μ^*, and t^{**} are all independent of the particular equilibrium being played. They are rather fixed features of the game in which equilibria are constructed.

Note that if $\Delta(\pi_0) \geq 0$, then $t^{**} = 0$. In particular, this is always true when signals are complements, and is true when signals are substitutes if the reason for the substitutability is high prospecting costs rather than slow signal acquisition. The general comparative static is that t^{**} is decreasing in $\overline{\lambda}/r$, is positive when this parameter ratio is sufficiently small, and is zero when it is sufficiently large.

Lemma 6. *Suppose $c \leq \overline{c}$. Then in any perfect Bayesian equilibrium,* $\min\{t_1^*, t_2^*\} \in \{t^{**}, \infty\}$ *and* $\max\{t_1^*, t_2^*\} = \infty$.

This result is derived in the following way. First, we establish that at most one t_i^* may be finite. For whenever the first firm stops investing, the second firm is in autarky and cannot benefit from waiting. Thus either neither firm ever waits, or a single firm eventually waits while the other firm always invests immediately.

So suppose $t_1^* < \infty$ while $t_2^* = \infty$. What remains is to pin down the values of t_1^* consistent with equilibrium. Once time t_1^* has been reached, firm 2 is in autarky and does not shirk or delay investment. (This conclusion requires an auxiliary result showing that firm 2's beliefs do not drop below π^A by t_1^*.) Thus since firm 1 finds it optimal to invest prior to t_1^* and to wait afterward, continuity of the value function combined with the definition of μ^* imply that $\mu^1(t_1^*) = \mu^*$. The lemma is then proven by showing that firm 2 must never shirk prior to t_1^*, in which case $\mu^1 = \mu^{\overline{\lambda}}$ and so $t_1^* = t^{**}$.

This final step requires a technical result proving, essentially, that any optimal prospecting strategy involves a single switch from working (i.e. prospecting at rate $\overline{\lambda}$) to shirking (prospecting at rate 0). It follows from a calculation involving the HJB equation establishing that once the value of shirking overtakes the value of working, the gap between the two values can only grow over time. This result relies crucially on the upper bound $\overline{c}$ on costs. With this result in hand, the fact that firm 2 eventually finds it optimal to work after time t_1^* implies that it must optimally work for all times.

3.3 The no-delay equilibrium

In this subsection we characterize all equilibria of the game satisfying $t_1^* = t_2^* = \infty$. Note that the only elements of an equilibrium strategy profile to be pinned down are prospecting rates following histories in which a firm has no signal and has seen no investment. All other elements of firm strategies are pinned down by Lemmas 2 and 3 and the assumption on the investment cutoff times.

It turns out that when $c \leq \bar{c}$, there exists exactly one equilibrium[10] with no delay in investment, which is characterized by a common cutoff time at which each firm switches from working to shirking. The cutoff turns out to be $t^A \equiv (\mu^{\overline{\lambda}})^{-1}(\pi^A)$, which is the time at which each firm's beliefs drop to π^A in equilibrium. This result is established formally in the following proposition, with equilibrium prospecting and investment strategies depicted schematically in Figure 1.

Proposition 1 (The no-delay equilibrium)**.** *Suppose $c \leq \bar{c}$. Then there exists an essentially unique perfect Bayesian equilibrium in which $t_1^* = t_2^* = \infty$. In this equilibrium, absent observing investment each firm prospects at rate $\overline{\lambda}$ until time t^A, after which it abandons prospecting forever.*

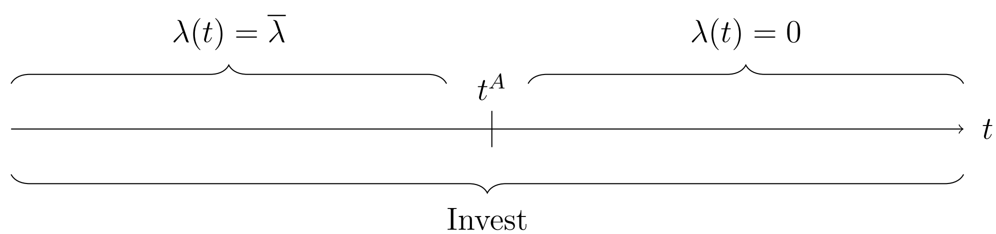

Figure 1: A timeline of prospecting and investment in the no-delay equilibrium

Intuitively, if neither firm ever finds it profitable to begin waiting to invest, then both firms must eventually abandon prospecting to limit the value of waiting. And this abandonment must happen precisely when beliefs drop to π^A. For otherwise either some firm would either prefer to continue on in autarky after abandonment at beliefs above π^A; or else some firm would prefer to stop buying a signal early when its beliefs drop below π^A.

In a bit more detail, recall from the discussion following Lemma 6 that in any equilibrium, each firm's prospecting policy must consist of a time threshold $t_i^\dagger$, before which it works and after which it shirks. Next, note that $t_i^\dagger$ must be finite for each firm, or else some firm's beliefs after receiving a High signal would eventually drop below $1/R$ and that firm would optimally wait for sufficiently large times. Suppose without loss that $t_1^\dagger \leq t_2^\dagger$. If $t_1^\dagger > t^A$, then some firm would be prospecting when its beliefs were below π^A, which cannot be optimal. On the other hand, if $t_1^\dagger < t^A$, the remaining firm would be in autarky with beliefs strictly above π^A, contradicting $t_2^\dagger < \infty$. So it must be that $t_1^\dagger = t^A$. Finally, if $t_2^\dagger > t^A$, then for

[10] Technically, there exist many equilibria, which may differ in their prospecting and investing policies on sets of times of measure zero. Our formal characterizations sidestep this issue by proving only essential uniqueness. In the text, we ignore this issue since the difference between uniqueness and essential uniqueness has no practical consequences.

times near t^A firm 1 would prefer to shirk and observe firm 2's actions rather than work to purchase a signal which barely recoups its prospecting costs. So both thresholds must be exactly t^A. The proof of the proposition establishes that these prospecting rules indeed hold together as an equilibrium.

Note that although neither firm delays investment, and although each firm abandons prospecting only when its value to the firm drops to zero, this outcome is nonethless socially inefficient. This is because each firm fails to internalize the value of its information-gathering to the other firm, and each could improve the other firm's welfare at no cost to itself by continuing to prospect.

3.4 The leader-follower equilibrium

We now characterize all equilibria of the game satisfying $\min\{t_1, t_2\} = t^{**}$ and $\max\{t_1, t_2\} = \infty$. Without loss suppose $t_2 = t^{**}$. (Since firms are symmetric, all equilibria with $t_1 = t^{**}$ can be found by swapping the roles of the two firms.) When costs are bounded above by $\bar{c}$, there turns out to be exactly one equilibrium satisfying these waiting thresholds. In this equilibrium, the two firms take on distinctly asymmetric roles. Firm 1, who we term the *leader*, never stops working (prospecting at rate $\overline{\lambda}$) prior to acquiring a signal or observing investment. By contrast firm 2, the *follower*, eventually shirks and passively observes the market to see what action the leader ultimately takes. The following proposition states the characterization formally. The follower's prospecting and investment strategies are depicted schematically in Figure 2.

Proposition 2 (The leader-follower equilibrium)**.** *Suppose $c \leq \bar{c}$. Then there exists a unique PBE in which $t_1^* = \infty$ and $t_2^* = t^{**}$. In this equilibrium, prior to observing investment each firm's prospecting policies are:*

- $\lambda^1(t) = \overline{\lambda}$ *for all t,*
- $\lambda^2(t) = \overline{\lambda}\,\mathbf{1}\{t < \bar{t}\}$ *for some $\bar{t} \geq 0$.*

Further, $\min\{\bar{t}, t^{**}\} < t^A$*, and if c is sufficiently small, $\bar{t} > t^{**}$.*

The proof of this proposition relies on two facts. First, as observed earlier, each firm's prospecting strategy must follow a threshold rule. Second, after time $\min\{\bar{t}, t^{**}\}$ the leader is in autarky, and it can be shown that its beliefs at this time are strictly above π^A. So the leader must eventually find it optimal to work, meaning it optimally works for all time. In this case the followers beliefs eventually fall below π^A, meaning it must eventually shirk.

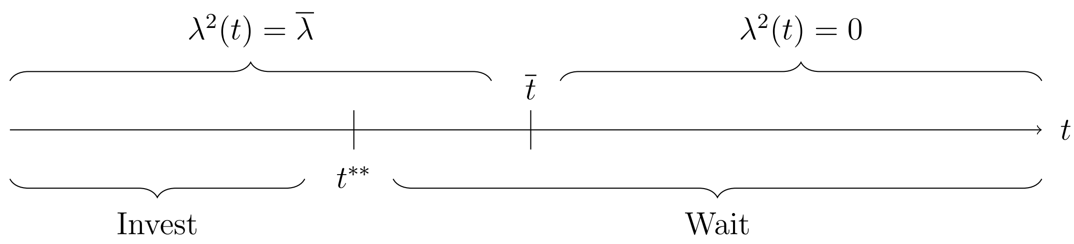

Figure 2: A timeline of prospecting and investing in the leader-follower equilibrium

The exact value of $\bar{t}$ is pinned down by a calculation of the option value of waiting for more information before prospecting. This option value rises as time progresses and the follower becomes more pessimistic about the quality of the project, and thus the likely information content of the signal.

A key feature of this equilibrium is that when prospecting costs are low, $\bar{t} > t^{**}$ and the follower continues prospecting for some time *after* it optimally begins waiting upon acquiring a signal. Hence with positive probability, on the equilibrium path the follower "waits to see" following acquisition of a positive signal about the project.

The force driving this result is the delay cost of acquiring a signal. Once the follower has seen the leader invest, acquiring an additional signal requires time and thus diminishes the value of investing in the project. By acquiring a signal ahead of time, the follower eliminates this delay cost and improves its payoff in the event the leader does invest. (This is true even if discounting is so large that the follower would not actually acquire a second signal in the event it had not by the time the leader invested.) On the other hand, acquiring a signal incurs prospecting costs, which are wasted in the event the leader does not ultimately invest. When these costs are low, the first force dominates for sufficiently high beliefs about the state, and the follower prefers to prospect for a time after t^{**}.

Finally, note that while the follower always eventually waits to invest following signal acquisition, there may be an initial period in which it both acquires a signal and invests immediately on the equilibrium path. This can occur only under substitutable signals, as under complementary signals $t^{**} = 0$. The robust prediction of this equilibrium is that waiting to see does occur with positive probability for sufficiently low prospecting costs, with the answer to the question of whether the follower *always* waits to see depending on the details of the signal structure.

4 Welfare comparison

We have seen that there are exactly three equilibria of the investing game. One involves no delay in investment and no shirking by either firm, but does involve inefficiently early abandonment of the project by each firm. The other two involve one firm which investigates the project efficiently, but another who free-rides off the leader. How do these equilibria compare in terms of individual and aggregate firm welfare, and to what extent do they improve on autarky?

Let V^A be the autarky value, V^{ND} be the expected payoff of each firm in the no-delay equilibrium, and V^L and V^F be the expected payoffs the leader and follower, respectively, in the leader-follower equilibrium. The following propositions summarize the rankings of each of these payoffs, depending on whether signals are substitutes or complements.

Proposition 3. *Suppose signals are complements. Then:*

- $V^{ND} = V^A$,
- $V^F > V^L = V^A$.

This proposition yields the surprising result that social learning provides no gains over autarky in the no-delay equilibrium. At first glance, each firm should be able to improve on their autarky payoff by learning from whether the other firm invests. However, this learning occurs only when it does not impact the firm's prospecting or investment decisions. At no point does the firm become pessimistic enough from lack of investment by the other firm to optimally cease prospecting or refrain from investing after obtaining a High signal. Of course, each firm does eventually stop prospecting in equilibrium, but this is only *weakly* optimal; and since beliefs remain fixed at π^A after t^A, another optimal strategy would involve prospecting forever. And on the other hand, under complements observing investment by the other firm does not lead a firm to stop prospecting for its own signal. So each firm might as well ignore the other firm's actions when formulating its own prospecting and investment plan, thus gains nothing in equilibrium from social learning.

Meanwhile, in the leader-follower equilibrium exactly one firm gains from social learning. It's easy to see that the leader does not benefit from social learning, because under complements $t^{**} = 0$ and so the leader does not learn from the follower at all. On the other hand, because the leader never stops prospecting, the follower's beliefs eventually drop below π^A. Thus the follower *does* optimally condition its strategy on the actions of the leader, halting prospecting following prolonged lack of investment and boosting its payoffs over autarky.

Note that in the case of complementary signals, the leader-follower equilibrium actually Pareto-dominates the no-delay equilibrium. This provides a natural rationale for firms to coordinate on such a structure in order to capture the benefits of social learning.

Proposition 4. *Suppose signals are substitutes and $c \leq \bar{c}$. Then:*

- $V^{ND} > V^A$,
- $V^F > V^L \geq V^A$, *and* $V^L > V^A$ *iff* $\min\{\bar{t}, t^{**}\} > 0$,
- $V^F > V^{ND} > V^L$.

Under substitutable signals, firms do benefit from social learning in the no-delay equilibrium. The reason is that now observing investment induces a firm to switch from prospecting to investing, reducing prospecting and delay costs and boosting overall payoffs. This result highlights a key difference between equilibrium behavior under complementary and substitutable signals. Under complements only bad news is payoff-relevant for each firm, and the no-delay equilibrium never delivers enough bad news to be pivotal. But under substitutes, both good and bad news are payoff-relevant, and the no-delay equilibrium does deliver good news with positive probability.

The relative ranking of leader and follower payoffs is the same as in the complements case, and for the same reasons. The main difference is that now the leader also potentially benefits over autarky, due to a conjunction of two factors. First, under substitutes the leader benefits from observing good news from the follower investing. Second, it becomes possible that the follower invests before the leader on the equilibrium path. In particular, if $\Delta(\pi_0) > 0$ and c is sufficiently small, this behavior arises. So the leader gains under autarky whenever model parameters are such that it observes good news from the follower with positive probability on the equilibrium path.

Finally, the two equilibria are no longer Pareto-ranked. In particular, the follower prefers its payoff to the no-delay payoff, while the leader has the opposite preferences. The follower's preferences are straightforward, because it observes pivotal bad news in the leader-follower equilibrium but not the no-delay equilibrium. As for the leader, the key to this result is the fact that $\min\{\bar{t}, t^{**}\} < t^A$. This implies that social learning from the follower stops strictly earlier in the leader-follower equilibrium than the no-delay equilibrium, meaning the leader's payoffs diminish.

While the two equilibria are no longer Pareto-ranked, continuity of equilibrium payoffs in model parameters implies that if signals are not *too* substitutable, then it will continue to be true that $V^L + V^F > 2V^{ND}$. In this case aggregate welfare is higher in the leader-follower equilibrium. Firms therefore continue to have an incentive to coordinate on this outcome if they have some means of sharing the total surplus, either through transfers or through rotation of the leadership role across multiple investments.

5 Conclusion

We study a model of strategic investment timing with costly dynamic information acquisition. We find that in equilibrium, firms may both free-ride on information acquisition and delay investment even after acquiring good news about the project. While each of these effects individually constitutes an inefficiency, in our model the two effects interact such that delay in investment can mitigate the welfare costs of free-riding.

While our model is stylized, it is consistent with features of several important economic settings, such as the market for venture capital. It also make predictions that are consistent with evidence on behavior in these markets. In particular, it captures the phenomenon of lead investors, who act quickly, and followers, who wait and see, in venture capital markets.

It would be interesting to extend our model to capture potential payoff externalities often present in investment environments. For instance, moving first may yield some advantage relative to following, or conversely projects may perform better when they are better-funded by a larger group of backers. Another extension would be to incorporate more firms, in particular to investigate whether equilibria in large markets might involve only a few active investigators and a large fringe of passive firms who don't prospect at all.

References

Aghamolla, Cyrus and Tadashi Hashimoto (2018). "Information arrival, delay, and clustering in dynamic freeriding". Working paper.

Ali, S Nageeb (2018). "Herding with costly information". In: *Journal of Economic Theory* 175, pp. 713–729.

Bolton, Patrick and Christopher Harris (1999). "Strategic Experimentation". In: *Econometrica* 67.2, pp. 349–374.

Bonatti, Alessandro and Johannes Hörner (2011). "Collaborating". In: *American Economic Review* 101.2, pp. 632–63.

— (2017). "Career concerns with exponential learning". In: *Theoretical Economics* 12.1, pp. 425–475.

Chamley, Christophe and Douglas Gale (1994). "Information Revelation and Strategic Delay in a Model of Investment". In: *Econometrica* 62.5, pp. 1065–1085.

Chari, V.V. and Patrick J. Kehoe (2004). "Financial crises as herds: overturning the critiques". In: *Journal of Economic Theory* 119.1, pp. 128–150.

Feld, B. and J. Mendelson (2016). *Venture Deals: Be Smarter Than Your Lawyer and Venture Capitalist*. John Wiley & Sons, Incorporated.

Frick, Mira and Yuhta Ishii (2016). "Innovation Adoption by Forward-Looking Social Learners". In: Working paper.

Horner, Johannes and Andy Skrzypacz (2017). "Learning, Experimentation and Information Design". In: *Advances in Economics and Econometrics: Volume 1: Eleventh World Congress*. Ed. by Bo Honoré et al. Cambridge University Press, pp. 63–98.

Keller, Godfrey and Sven Rady (2010). "Strategic experimentation with Poisson bandits". In: *Theoretical Economics* 5.2, pp. 275–311.

— (2015). "Breakdowns". In: *Theoretical Economics* 10.1, pp. 175–202.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005). "Strategic experimentation with exponential bandits". In: *Econometrica* 73.1, pp. 39–68.

Murto, Pauli and Juuso Välimäki (2011). "Learning and information aggregation in an exit game". In: *The Review of Economic Studies* 78.4, pp. 1426–1461.

— (2013). "Delay and information aggregation in stopping games with private information". In: *Journal of Economic Theory* 148.6, pp. 2404–2435.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille (2007). "Social Learning in One-Arm Bandit Problems". In: *Econometrica* 75.6, pp. 1591–1611.

Appendices

A An alternate form for the belief updating equation

It will be useful to express a firm's prospecting strategy in terms of the evolution of the opposing firm's beliefs. Solving the Bayes' rule relationship between $\mu^i(t)$ and $\Omega^{-i}(t)$ given in Section 2.2 for the latter variable and performing some simplifying manipulations yields

$$\Omega^{-i}(t) = \frac{l(\pi_0)}{h(\pi_0)} \frac{\mu^i(t) - \pi_-}{\pi_+ - \mu^i(t)}.$$

Note that Ω^{-i} is not in general differentiable everywhere, as λ^i is not continuous. However, it is certainly absolutely continuous, and its derivative is defined and equal to $-\lambda^i(t)\Omega(t)$ a.e. So differentiate the definition of Ω^{-i} to obtain $\dot{\Omega}^{-i}(t) = -\lambda^{-i}(t)\Omega^{-i}(t)$. Inserting this result into the previous expression for Ω^{-i} one obtains

$$\lambda^{-i}(t) = -\left(\frac{\pi_+ - \pi_-}{\mu^i(t) - \pi_-}\right)\left(\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}\right).$$

(As Ω^i is only absolutely continous, so is $\dot{\mu}$, and so this expression holds only a.e.) These results reflect the fact that specifying cumulative effort is equivalent to specifying the opponent's current beliefs, and specifying instantaneous effort is equivalent to specifying the rate of change of the opponent's beliefs (conditional on their current beliefs).

B The HJB equation under immediate investing

In this subsection we characterize the HJB equation when each firm invests immediately upon obtaining a High signal. Let $V^i(t)$ be firm i's equilibrium continuation value function conditional on receiving no signal and seeing no investment up to time t. Let

$$\overline{V} \equiv \max\left\{\pi_+ R - 1, \frac{\overline{\lambda}}{\overline{\lambda} + r}(h(\pi_+)(\pi_{++}R - 1) - c)\right\}$$

be i's continuation value upon seeing an investment by $-i$. Then assuming the value function is differentiable, V^i satisfies the HJB equation

$$rV^i(t) = \max_{\lambda \in [0,\overline{\lambda}]} \left\{ \lambda (h(\mu^i(t))(\mu^i_+(t)R - 1)_+ - c - V^i(t)) \right\} + \frac{\Omega^{-i}(t)}{\Omega^{-i}(t) + (1 - \Omega^{-i}(t))l(\pi_0)} \lambda^{-i}(t) h(\pi_0)(\overline{V} - V^i(t)) + \dot{V}^i(t).$$

If V^i satisfies this equation for all time along with $0 \le V^i \le \overline{V}$, then it must be firm i's equilibrium value function. Note that whenever $\mu^i_+(t) < 1/R$, non-negativity of V^i implies that the optimal effort choice is $\lambda = 0$. Thus

$$\begin{aligned} & \max_{\lambda \in [0,\overline{\lambda}]} \left\{ \lambda (h(\mu^i(t))(\mu^i_+(t)R - 1)_+ - c - V^i(t)) \right\} \\ = & \max_{\lambda \in [0,\overline{\lambda}]} \left\{ \lambda (h(\mu^i(t))(\mu^i_+(t)R - 1) - c - V^i(t)) \right\} \\ = & \overline{\lambda} \left(h(\mu^i(t))(\mu^i_+(t)R - 1) - c - V^i(t) \right)_+ . \end{aligned}$$

Recall that $h(\mu)(\mu_+ R - 1) - c = K(\mu - \pi^A)$. The HJB equation may then be written more simply as

$$rV^i(t) = \overline{\lambda} \left(K(\mu^i(t) - \pi^A) - V^i(t) \right)_+ + \frac{\Omega^{-i}(t)}{\Omega^{-i}(t) + (1 - \Omega^{-i}(t))l(\pi_0)} \lambda^{-i}(t) h(\pi_0)(\overline{V} - V^i(t)) + \dot{V}^i(t),$$

with associated optimal control

$$\lambda^{i*}_t = \begin{cases} \overline{\lambda}, & V^i(t) < K(\mu^i_t - \pi^A), \\ [0, \overline{\lambda}], & V^i(t) = K(\mu^i_t - \pi^A), \\ 0, & V^i(t) > K(\mu^i_t - \pi^A). \end{cases}$$

At this point the HJB equation includes both the firm's beliefs μ^i about the state, as well as the (conjectured) instantaneous and cumulative efforts λ^{-i} and Ω^{-i} of the other firm. It will be useful to re-express the latter two in terms of μ^i and its time derivative $\dot{\mu}^i$. Recall that by Bayes' rule

$$\mu^i_t = \frac{\Omega^{-i}(t)\pi_0 + (1 - \Omega^{-i}(t))l(\pi_0)\pi_-}{\Omega^{-i}(t) + (1 - \Omega^{-i}(t))l(\pi_0)}.$$

Solving this expression for $\Omega^{-i}(t)$ and performing some simplifying manipulations yields

$$\Omega^{-i}(t) = \frac{l(\pi_0)}{h(\pi_0)} \frac{\mu^i(t) - \pi_-}{\pi_+ - \mu^i(t)}.$$

Now, differentiating the definition of Ω^{-i} yields $\dot{\Omega}^{-i}(t) = -\lambda^{-i}(t)\Omega^{-i}(t)$, so by inserting the previous expression for Ω^{-i} one obtains

$$\lambda^{-i}(t) = -\left(\frac{\pi_+ - \pi_-}{\mu^i(t) - \pi_-}\right)\left(\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}\right).$$

This relationship tells us that from firm i's perspective, $-i$'s choice of effort is equivalent to controlling the rate at which i's beliefs about the state deteriorate. We therefore treat $\dot{\mu}^i$ as $-i$'s control variable and drop λ^{-i} from firm i's optimization problem. The resulting reduced HJB equation is

$$rV^i(t) = \overline{\lambda}\left(K(\mu^i(t) - \pi^A) - V^i(t)\right)_+ - \frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}(\overline{V} - V^i(t)) + \dot{V}^i(t).$$

The bounds $\lambda^{-i}(t) \in [0, \overline{\lambda}]$ correspond to $\dot{\mu}^i(t) \in \left[-\overline{\lambda}(\pi_+ - \mu^i(t))\frac{\pi_+ - \pi_-}{\mu^i(t) - \pi_-}, 0\right]$. (At this point one can probably prove that for any μ^i induced by firm $-i$, this HJB equation admits a solution bounded between 0 and $\overline{V}$, which ensures that the HJB equation is necessary as well as sufficient.)

C Proofs

C.1 Proof of Lemma 2

First consider a firm who has received a Low signal. Then regardless of his beliefs about the content of any signal received by the other firm, his posterior belief that the state is Good cannot be higher than π_0. As $\pi_0 R - 1 < 0$ by assumption, investing is never optimal at any point in the future.

To prove the second part of the lemma, we first show that in equilibrium firms are always able to use Bayes' rule to update their beliefs about the state no matter the history of the game. Fix an equilibrium σ, and suppose by way of contradiction that at some time t^* and following some history, 1) neither firm has invested by time t^*; 2) Bayes' rule applies for both firms at all $t < t^*$, but at t^* firm i cannot use Bayes' rule to form a posterior probability about θ. These two conditions imply that according to σ, firm $-i$ should have invested with probability 1 prior to t^* absent any investment by firm i. (Otherwise Bayes' rule would still be applicable at time t^*.) We already know that in any PBE, firm $-i$ will never invest if in possession of a Low signal. Therefore if firm $-i$ had a Low signal with some probability by time t^*, Bayes' rule would still apply for firm i. Hence firm $-i$ cannot have obtained a signal, hence cannot have prospected prior to t^*. In this case any updating to firm $-i$'s beliefs prior to t^* must come solely from social learning due to the absence of firm i's investment.

Bayes' rule applies for firm $-i$ at all $t < t^*$, meaning that with some probability under σ^i, firm i did not invest prior to t^*. And we already know that in any PBE, firm i never invests when in possession of a Low signal. Thus no matter how frequently it would invest when in possession of no signal or a High signal, the inference $-i$ must make about the state from lack of investment is weakly negative. Hence firm $-i$'s beliefs at all times prior to t^* must be no higher than π_0, meaning the payoff from investing is no higher than $R\pi_0 - 1 < 0$. This contradicts the assumption that σ is an equilibrium.

So under any equilibrium, if neither firm has invested by time t^*, and Bayes' rule applies for both firms at all prior times, then each firm must be able to use Bayes' rule to form beliefs at time t^*. Since this reasoning applies for every t^* and every history, it must be that Bayes' rule applies for both firms at all times following all histories in which neither firm has invested. It then follows from the argument of the previous two paragraphs that a firm in possession of no signal and seeing no investment must at all times have beliefs no higher than π_0, and hence must never find immediate investment profitable.

C.2 Proof of Lemma 3

Upon observing firm $-i$ invest, Lemma 2 implies that firm i's beliefs that the state is Good are π_+ if $s^i_t = \emptyset$. Suppose first that $\pi_{+-} \geq 1/R$. Then regardless of what signal firm i eventually obtains, investment upon obtaining that signal is optimal. Hence costly acquisition of a signal cannot be profitable, and so firm i must optimally invest immediately. So assume that $\pi_{+-} < 1/R$.

The continuation payoff from investing immediately is $\pi_+ R - 1$. Meanwhile receiving a signal has expected continuation payoff $h(\pi_+)(\pi_{++}R - 1)$, as the firm invests only if the signal is High. Let V be firm i's continuation value from an optimal policy. Then V satisfies the HJB equation

$$rV = \max\left\{ r(\pi_+ R - 1), \max_{\lambda \in [0, \overline{\lambda}]} \{\lambda(h(\pi_+)(\pi_{++}R - 1) - c - V)\} \right\}.$$

Suppose first that $\pi_+ R - 1 \geq \frac{\overline{\lambda}}{\overline{\lambda}+r}(h(\pi_+)(\pi_{++}R - 1) - c)$. Then as $V \geq \pi_+ R - 1$, it must be that

$$h(\pi_+)(\pi_{++}R - 1) - c - V \leq \frac{r}{\overline{\lambda}}(\pi_+ R - 1),$$

so the second term on the rhs of the HJB equation is weakly smaller than $r(\pi_+ R - 1)$, and is strictly smaller except in the knife-edge case. Hence $V = \pi_+ R - 1$ is the unique solution to the HJB equation, with immediate investment a corresponding optimal policy. Except in the knife-edge case this is the unique optimal policy; in the knife-edge case the two arguments of the max are equal, hence any stopping rule for investing prior to receiving a signal is optimal, so long as prospecting is undertaken at rate $\overline{\lambda}$.

On the other hand, suppose that $\pi_+ R - 1 < \frac{\overline{\lambda}}{\overline{\lambda}+r}(h(\pi_+)(\pi_{++}R - 1) - c)$. If $V = \pi_+ R - 1$ then

$$h(\pi_+)(\pi_{++}R - 1) - c - V > \frac{r}{\overline{\lambda}}(\pi_+ R - 1),$$

so that the second term on the rhs of the HJB equation is strictly greater than $r(\pi_+ R - 1)$, a contradiction. Hence

$$rV = \max_{\lambda \in [0,\overline{\lambda}]} \{\lambda(h(\pi_+)(\pi_{++}R - 1) - c - V)\}$$

and the rhs is strictly greater than $r(\pi_+ R - 1)$, meaning $\lambda = \overline{\lambda}$. Solving for V yields

$$V = \frac{\overline{\lambda}}{\overline{\lambda} + r}(h(\pi_+)(\pi_{++}R - 1) - c),$$

which by assumption is strictly greater than $r(\pi_+ R - 1)$. Thus this value function is the unique solution to the HJB equation, and the associated optimal policy is to prospect at rate $\overline{\lambda}$ until receiving a signal.

C.3 Proof of Lemma 4

We show first that, for any time t such that $\mu_+^i(t) > 1/R$, any best reply for i either invests immediately or waits until $-i$ invests. Suppose by way of contradiction that firm i had a best reply such that upon receiving a High signal at time t, i invests at a random time $\tau \in [t, \infty) \cup \{\infty\}$ conditional on no investment by firm $-i$, with $\Pr(\tau \in (t, \infty)) > 0$. Then there must exist another best reply such that firm i waits until some time $t' \in (t, \infty)$ and then invests w.p. 1 at time t' conditional on no investment by firm $-i$. In particular, it must be that $\mu_+^i(t')R - 1 \geq 0$.

Let $\tau' \in \mathbb{R}_+ \cup \{\infty\}$ be the random time at which $-i$ invests. Then whenever $\tau' > t'$, i's ex post continuation payoff as of time t is $e^{-rt'}(\mu^i_+(t')R - 1) \leq \mu^i_+(t')R - 1$. In particular, if $\Pr(\tau' > t' \mid \tau' > t, s_i = H) = 1$, then $\mu^i_+(t') = \mu^i_+(t)$ and the previous inequality is strict. And whenever $\tau' \in (t, t']$, i's continuation payoff is $e^{-r\tau'}(\pi_{++}R - 1) < \pi_{++}R - 1$. Then i's continuation payoff from this best reply is strictly less than

$$U' = \Pr(\tau' \leq t' \mid \tau' > t, s_i = H)(\pi_{++}R - 1) + \Pr(\tau' > t' \mid \tau' > t, s_i = H)(\mu^i_+(t')R - 1).$$

As

$$\mu^i_+(t) = \Pr(\tau' \leq t' \mid \tau' > t, s_i = H)\pi_{++} + \Pr(\tau' > t' \mid \tau' > t, s_i = H)\mu^i_+(t'),$$

$U' = \mu^i_+(t)R - 1$, which is exactly i's payoff from investing immediately at time t. Thus waiting until t' and then investing cannot be cannot be a best reply, yielding the desired contradiction.

We now treat the $\bar{t} < \infty$ case. For any t such that $\mu^i_+(t) < 1/R$, trivially the unique best response is for i to wait forever, since investing at any time after t yields a strictly negative payoff. It is also true that whenever $\mu^i_+(t) = 1/R$, waiting forever is i's unique best response. For investing at any time before $-i$ invests yields a non-positive continuation payoff, whereas given $\bar{t} < \infty$ there is a positive probability that $-i$ invests in the future, so that waiting for $-i$ to invest yields a strictly positive payoff.

Next fix any time t for which investing immediately is a best response for i. In this case $\mu^i_+(t) > 1/R$. We claim that for every $t' < t$, investing immediately is i's unique best response. Suppose by way of contradiction that for some $t' < t$, there existed a best response which does not invest immediately. Then by earlier results, waiting until $-i$ invests must be a best response. But by assumption upon reaching time t with no investment by $-i$, investing immediately is a best response. Hence waiting until time t and then investing must also be a best response. This is a contradiction of earlier results, as desired.

Hence when $\bar{t} < \infty$, the set of best responses by i must have a simple structure - there exists a cutoff time $t^* \leq \underline{t}$ such that any best response by i invests immediately for $t < t^*$ and waits for any $t > t^*$.

The argument for the case $\bar{t} = \infty$ is a minor modification of the steps above. In this case it is not true that $\mu^i(t) = 1/R$ implies that i waits forever as a unique best reply. Indeed at this point all continuation strategies yield the same payoff, as $-i$ never acts again. However, whenever $\mu^i_+(t) > 1/R$ and investing immediately at t is a best reply, it is still true that investing immediately is a unique best reply for all $t' < t$. Hence there still exists a cutoff time t^* such that i invests immediately prior to t^*, and such that for $(t^*, \underline{t})$ firm i does not invest. Therefore all best replies by firm i entail waiting until at least time $\underline{t}$ to invest upon receiving a signal after t^*.

C.4 Proof of Lemma 5

It can be shown that

$$\Delta(\mu) = \frac{\mu_+ - \pi_{+-}}{\pi_{++} - \pi_{+-}} \frac{\overline{\lambda}}{\overline{\lambda} + r} (\pi_{++}R - 1) - (\mu_+ R - 1),$$

which is continuous in μ. Differentiating this expression yields

$$\Delta'(\mu) = \frac{1}{\pi_{++} - \pi_{+-}} \frac{\overline{\lambda}}{\overline{\lambda} + r} (\pi_{++}R - 1) - R.$$

By assumption $\pi_{+-} < 1/R < \pi_{++}$, so

$$\Delta'(\mu) < -\frac{r}{\overline{\lambda} + r} R < 0.$$

Further, $\Delta(\pi_-) = -(\pi_{+-}R - 1) > 0$. Finally,

$$\Delta(\pi_0) = \frac{\pi_+ - \pi_{+-}}{\pi_{++} - \pi_{+-}} \frac{\overline{\lambda}}{\overline{\lambda} + r} (\pi_{++}R - 1) - (\pi_+ R - 1),$$

and by Bayes' rule

$$\begin{aligned}
\frac{\pi_+ - \pi_{+-}}{\pi_{++} - \pi_{+-}} &= 1 - \frac{\pi_{++} - \pi_+}{\pi_{++} - \pi_{+-}} \\
&= 1 - \frac{\frac{q^H \pi_+}{h(\pi_+)} - \pi_+}{\frac{q^H \pi_+}{h(\pi_+)} - \frac{(1-q^H)\pi_+}{l(\pi_+)}} \\
&= 1 - l(\pi_+) \frac{q^H - h(\pi_+)}{q^H l(\pi_+) - (1 - q^H) h(\pi_+)} \\
&= 1 - l(\pi_+) = h(\pi_+).
\end{aligned}$$

So

$$\Delta(\pi_0) = \frac{\overline{\lambda}}{\overline{\lambda} + r} h(\pi_+)(\pi_{++}R - 1) - (\pi_+ R - 1).$$

When signals are complements

$$\frac{\overline{\lambda}}{\overline{\lambda}+r}h(\pi_+)(\pi_{++}R-1) \geq \pi_+R-1+\frac{\overline{\lambda}}{\overline{\lambda}+r}c > \pi_+R-1,$$

so $\Delta(\pi_0) > 0$.

C.5 Proof of Lemma 6

We begin by proving a series of lemma, which build gradually toward the result.

Lemma 7. *Suppose $c \leq \bar{c}$. Fix any perfect Bayesian equilibrium. For any firm i, define $f_i(t) \equiv V^i(t) - K(\mu^i(t) - \pi^A)$ and $\hat{t}^i \equiv \inf\{t \;:\; \mu^i(t) \leq \pi^A\}$. Then for almost every $t \in [0, \min\{t_i^*, t_{-i}^*, \hat{t}_i\}]$, either $f_i(t) < 0$ or $f_i'(t) > 0$.*

Proof. First note that as V^i is C^1 and μ^i is absolutely continuous, f_i is absolutely continuous as well and f_i' is defined a.e. Suppose that for a positive measure set of times $T^* \subset [0, \min\{t_i^*, t_{-i}^*, \hat{t}_i\}]$, $f_i(t) \geq 0$ and $f_i'(t) \leq 0$. Whenever $f_i(t) \geq 0$ and $f_i'(t)$ exists, the HJB equation characterizing V^i for $t \leq \min\{t_i^*, t_{-i}^*\}$ becomes

$$rV^i(t) = -\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}(\overline{V} - V^i(t)) + \dot{V}^i(t).$$

The rhs may be rewritten in terms of f_i as

$$rV^i(t) = -\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}(\overline{V} - K(\pi_+ - \pi^A) - f_i(t)) + f_i'(t).$$

Now, the bound $c \leq \bar{c}$ implies $\overline{V} \leq h(\pi_+)(\pi_{++}R-1) - c = K(\pi_+ - \pi^A)$, so when $f_i(t) \geq 0$ and $f_i'(t) \leq 0$ the rhs is bounded above by 0. Finally, whenever $\mu^i(t) > \pi^A$ it must be that $V^i(t) > 0$, since a firm's value function is bounded below by their autarky payoff. Hence the HJB equation must be violated on T^*, a contradiction. □

Lemma 8. *Suppose $c \leq \bar{c}$. In any perfect Bayesian equilibrium such that $t_1^* \leq t_2^*$, if $t_1^* < \infty$ then $t_2^* = \infty$, $\mu^2(t_1^*) \geq \pi^A$, and $\mu^1(t_1^*) \leq \mu^*$.*

Proof. Suppose that $\mu^2(t_1^*) < \pi^A$. In this case clearly $t_1^* > 0$ given that $\mu^2(0) = \pi_0 > \pi^A$. Let $\check{t}^2 \equiv \sup\{t \;:\; \mu^2(t) \geq \pi^A\}$. By continuity $\check{t}^2 < t_1^*$. Then as $V^2 \geq 0$, for $t \in (\check{t}^2, t_1^*]$ it must be that $V^2(t) > K(\mu^2(t) - \pi^A)$, so that by the HJB equation 2's essentially unique optimal prospecting policy is $\lambda^2(t) = 0$ for $t \in (\check{t}^2, t_1^*]$. And after t_1^* firm 2 is in autarky with beliefs below the autarky threshold, so it must also be that $\lambda^2(t) = 0$ for almost every $t \geq t_1^*$.

Then on the equilibrium path firm 2 never invests first after $\check{t}^2$, meaning firm 1 is in autarky with constant beliefs $\mu^1(t) = \mu^1(\check{t}^2)$ for all $t > \check{t}^2$. But then $\mu^1_+(\check{t}^2) = 1/R$, otherwise it could not be optimal for firm 1 to invest immediately prior to t_1^* and wait forever after t_1^*. In this case $\mu^1(\check{t}^2) < \pi^A$, so $\lambda^1(t) = 0$ for all $t \geq \check{t}^2$. But then on the equilibrium path firm 1 does not invest first on $[\check{t}^2, t_1^*]$, implying $\mu^2(\check{t}^2) = \mu^2(t_1^*)$. This contradicts our assumption on $\mu^2(t_1^*)$ and the definition of $\check{t}^2$. So it must be that $\mu^2(t_1^*) \geq \pi^A$.

As firm 1 does not invest first on the equilibrium path after t_1^*, firm 2 is in autarky with constant beliefs $\mu^2(t) \geq \pi^A$ for all $t \geq t_1^*$. Hence $\mu^2_+(t) > 1/R$ and so immediate investing is strictly superior to waiting forever for every $t \geq t_1^*$. In other words, $t_2^* = \infty$.

Now suppose by way of contradiction that $\mu^1(t_1^*) > \mu^*$. Then by continuity, $\mu^1(t) > \mu^*$ for t larger than but sufficiently close to t_1^*. But by definition, after t_1^* firm 1's optimal investment strategy upon receiving a High signal is to wait forever for firm 2 to invest. The continuation payoff to this strategy is bounded above by

$$\frac{\mu^1_+(t) - \pi_{+-}}{\pi_{++} - \pi_{+-}} \frac{\overline{\lambda}}{\overline{\lambda} + r}(\pi_{++}R - 1),$$

which is its continuation payoff from the strategy supposing firm 2 prospects at the maximum possible rate and never delays investment. Meanwhile the continuation payoff from investing immediately is $\mu^1_+(t)R - 1$. By definition of μ^*, the second quantity is larger than the first whenever $\mu^1(t) > \mu^*$, contradicting the optimality of waiting forever subsequent to time t_1^*. So $\mu^1(t_1^*) \leq \mu^*$ in any case. □

In principle a multiplicity of equilibria might be supported by leaving some firm indifferent between different prospecting strategies, for instance by keeping a firm at beliefs π^A forever after some threshold time. The following lemma rules out certain strategies of this sort when $t_1^* < \infty$.

Lemma 9. *Suppose $c \leq \overline{c}$. In any perfect Bayesian equilibrium such that $t_1^* \leq t_2^*$, if $t_1^* < \infty$ then $\mu^2(t_1^*) > \pi^A$, $\mu^1(t_1^*) = \mu^*$, and $\lambda^2(t) = \overline{\lambda}$ for almost every $t \geq t_1^*$.*

Proof. For each i, let $\hat{t}^i \equiv \inf\{t : \mu^i(t) \leq \pi^A\}$ be the first time firm i's beliefs hit the autarky threshold. By continuity $\mu^i(\hat{t}^i) = \pi^A$ for each i. By Lemma 8 we already know that $\mu^2(t_1^*) \geq \pi^A$ and $t_2^* = \infty$. Assume by way of contradiction that $\mu^2(t_1^*) = \pi^A$. Clearly $t_1^* > 0$ in this case given $\mu^2(0) = \pi_0 > \pi^A$, and further $\hat{t}^2 \leq t_1^*$ given the definition of $\hat{t}^2$.

Next observe that in equilibrium, firm 1 never invests first after $\hat{t}^2$. This is automatically true for $t \geq t_1^*$, so the remaining thing to show is that $\lambda^1(t) = 0$ on $[\hat{t}^2, t_1^*)$ in the case that this interval is non-empty. But $\mu^2(t)$ is constant on this interval, so given that firm 1 invests immediately upon obtaining a High signal its prospecting rate must be zero. Thus $V^2(\hat{t}^2) = 0$, since at time $\hat{t}^2$ firm 2 is in autarky with beliefs π^A. Also $V^2(t) > 0$ for all $t < \hat{t}^2$. For at any such time $\mu^2(t) > \pi^A$, so firm 2's autarky payoff as of time t is strictly positive, and this payoff is a lower bound on $V^2(t)$.

Define $f_2(t) \equiv V^2(t) - K(\mu^2(t) - \pi^A)$. By Lemma 7, for almost every $t \in [0, \hat{t}^2]$ either $f_2(t) < 0$ or $f_2'(t) > 0$. Note also that $V^2(\hat{t}^2) = 0$ and $\mu^2(\hat{t}^2) = \pi^A$ implies $f_2(\hat{t}^2) = 0$. We next establish that $\lambda^2(t) = \overline{\lambda}$ a.e. on $[0, \hat{t}^2]$.

Suppose first that $f_2(t') > 0$ for some $t' \in [0, \hat{t}^2)$. Then as $f_2(\hat{t}^2) = 0$ and f is absolutely continuous, there must exist a positive-measure subset of $[t', \hat{t}^2]$ on which $f_2(t) > 0$ and $f_2'(t) < 0$, a contradiction. So certainly $f_2(t) \leq 0$ on $[0, \hat{t}^2]$. If $f_2(t) = 0$ on a positive-measure subset of $[0, \hat{t}^2]$, then a.e. on this set $f_2'(t) > 0$. But whenever $f_2(t) = 0$ and $f_2'(t) > 0$, the definition of $f_2'(t)$ implies that $f_2(t + \varepsilon) > f_2(t) = 0$ for sufficiently small $\varepsilon > 0$, a contradiction. So $f_2(t) < 0$ for almost every $t \in [0, \hat{t}^2]$. Hence from the HJB equation $\lambda^2(t) = \overline{\lambda}$ for almost every $t \in [0, \hat{t}^2]$.

This means in particular that $\mu^1(t) \leq \mu^2(t)$ for $t \in [0, \hat{t}^2]$ and therefore $\hat{t}^1 \leq \hat{t}^2$, no matter the prospecting policy firm 1 follows. If $\hat{t}^1 < \hat{t}^2$, then the fact that firm 2 prospects with positive intensity and invests immediately on $[\hat{t}^1, \hat{t}^2]$ means $V^1(\hat{t}^1) > 0$. Then as $\mu^1(t) \leq \pi^A$ for $t \geq \hat{t}^1$, the HJB equation requires that $\lambda^1(t) = 0$ for almost every $t \in [\hat{t}^1, \hat{t}^2]$. But then μ^2 is constant on this interval, a contradiction of the fact that $\hat{t}^2$ is the first time μ^2 hits π^A. So $\hat{t}^1 = \hat{t}^2 = \hat{t}$, which can only hold if $\lambda^1(t) = \overline{\lambda}$ for almost every $t \in [0, \hat{t}]$.

If $V^1(\hat{t}) > 0$, then given continuity of V^1 and μ^1, for sufficiently large $t < \hat{t}$ it would be the case that $V^1(t) > K(\mu^1(t) - \pi^A)$. But then $\lambda^1(t) = 0$ by the HJB equation, a contradiction. So $V^1(\hat{t}) = 0$. But as $t_2^* = \infty$ by Lemma 8, this can only be true if $\lambda^2(t) = 0$ for a.e. $t > \hat{t}$. As $\mu^1(\hat{t}^1) = \pi^A$ by definition of $\hat{t}^1$, it must be that $\mu^1(t) = \pi^A$ for all $t > \hat{t}$, so $\mu_+^1(t) > 1/R$ on this time range. As firm 2 never invests along the equilibrium path after $\hat{t}$, this means that investing immediately upon receiving a High signal must yield a strictly higher continuation payoff for firm 1 than waiting forever, contradicting $t_1^* < \infty$. This is the desired contradiction ruling out $\mu^2(t_1^*) = \pi^A$.

Now note that for $t > t_1^*$ firm 2 is in autarky with constant beliefs $\mu^2(t) = \mu^2(t_1^*) > \pi^A$, meaning 2's essentially unique optimal prospecting policy subsequent to t_1^* is $\lambda^2(t) = \overline{\lambda}$.

Finally, suppose $\mu^1(t^*) < \mu^*$. Then by definition of μ^* and given firm 2's strategy subsequent to t_1^*, for $t \geq t_1^*$ firm 1's continuation value from waiting forever upon obtaining a High signal, say $\widetilde{V}^1(t)$, is strictly larger than $\mu_+^1(t)R - 1$. Also, $t_1^* > 0$ given $\mu^1(0) = \pi_0 > \mu^*$. Then as $\widetilde{V}^1$ is continuous in t regardless of 2's prospecting strategy, $\widetilde{V}^1(t) > \mu_+^1(t)R - 1$ for sufficiently large $t < t_1^*$, contradicting the optimality of immediate investing prior to t_1^*. So

$\mu^1(t_1^*) \geq \mu^*$, meaning $\mu^1(t_1^*) = \mu^*$ in light of Lemma 8. □

We are now ready to prove Lemma 6. Without loss suppose that $t_1^* < \infty$ and $t_1^* \leq t_2^*$. We will make frequent use of the results, established in Lemmas 8 and 9, that $t_2^* = \infty$, $\mu^2(t_1^*) > \pi^A$, $\mu^1(t_1^*) = \mu^*$, and $\lambda^2(t) = \overline{\lambda}$ for almost every $t \geq t_1^*$.

Consider first the case $\Delta(\pi_0) \geq 0$, i.e. $\mu^* = \pi_0$. In this case $\mu^*(t_1^*) = \mu^*$ immediately implies $t_1^* = 0 = t^{**}$, which is the desired result.

Now consider the case when $\Delta(\pi_0) < 0$, i.e. $\mu^* < \pi_0$. Note that $\mu^1(t_1^*) = \mu^*$ necessarily implies $t_1^* > 0$ in this case. We establish that $V^\dagger(t) \equiv K(\mu_2(t) - \pi^A)$ is a supersolution to firm 2's HJB equation on $[0, t_1^*]$, no matter firm 1's prospecting policy. On this time range 2's HJB equation is

$$rV^2(t) = \overline{\lambda}\left(K(\mu^2(t) - \pi^A) - V^2(t)\right)_+ - \frac{\dot{\mu}^2(t)}{\pi_+ - \mu^2(t)}\left(\overline{V} - V^2(t)\right) + \dot{V}^2(t).$$

Inserting $V^\dagger$ into the rhs yields $-\frac{\dot{\mu}^2(t)}{\pi_+ - \mu^2(t)}(\overline{V} - K(\pi_+ - \pi^A))$, which is bounded above by 0 when $c \leq \overline{c}$. Meanwhile inserting $V^\dagger$ into the lhs yields a positive quantity given $\mu^2(t_1^*) > \pi^A$. So indeed $V^\dagger$ is a supersolution to the HJB equation.

Now note that as firm 2 is in autarky beginning at time t_1^* and $\mu_2(t_1^*) > \pi^A$, it follows that

$$V^2(t_1^*) = \frac{\overline{\lambda}}{\overline{\lambda} + r}K(\mu_2(t^*) - \pi^A) < \overline{V}(t^*).$$

This result combined with the fact that $V^\dagger$ is a supersolution to the HJB equation on $[0, t_1^*]$ implies that $V^\dagger(t) > V^2(t)$ for all $t \leq t_1^*$. Thus from the HJB equation firm 2's essentially unique optimal prospecting strategy is $\lambda^2(t) = \overline{\lambda}$ for all time. Given $t_2^* = \infty$, it follows that $\mu^1 = \mu^{\overline{\lambda}}$. Then $\mu^1(t_1^*) = \mu^*$ entails $t_1^* = t^{**}$.

C.6 Proof of Proposition 1

This proof proceeds in two steps. First, we prove in the following lemma that the prospecting strategies in the proposition statement are *necessary* for equilibrium. Afterward, we prove that they are also sufficient.

Lemma 10. *Fix any perfect Bayesian equilibrium in which* $t_1^* = t_2^* = \infty$*. Then* $\lambda^1(t) = \lambda^2(t) = \overline{\lambda}$ *for almost every* $t < (\mu^{\overline{\lambda}})^{-1}(\pi^A)$*, while* $\lambda^1(t) = \lambda^2(t) = 0$ *for almost every* $t > (\mu^{\overline{\lambda}})^{-1}(\pi^A)$*.*

Proof. For each i, let $\hat{t}^i \equiv \inf\{t \,:\, \mu^i(t) \leq \pi^A\}$ be the first time firm i's beliefs hit the autarky threshold. Also let $f_i(t) \equiv V^i(t) - K(\mu^i(t) - \pi^A)$ be the term in firm i's HJB equation whose sign determines i's optimal prospecting rate. As V^i is C^1 and μ^i is absolutely continuous, f_i is absolutely continuous as well and f_i' is defined a.e.

We show first that each $\hat{t}^i < \infty$. To establish this, maintain for the time being the opposite assumption, that some $\hat{t}^i = \infty$, say $i = 1$.

By Lemma 7, for almost every t either $f_1(t) < 0$ or $f_1'(t) > 0$. If $f_1(t) < 0$ a.e., then the HJB equation would imply that $\lambda^1(t) = \overline{\lambda}$ a.e. But in this case eventually $\mu_+^2(t) < 1/R$, meaning $t_2^* < \infty$, a contradiction of $t_1^* = \infty$. So on some positive-measure set of times $T^\dagger$, $f_1(t) \geq 0$ and $f_1'(t) > 0$. But whenever $f_1(t) \geq 0$ and $f_1'(t) > 0$, the definition of $f_1'(t)$ implies that $f_1(t + \varepsilon) > 0$ for sufficiently small $\varepsilon > 0$. Therefore $t^0 \equiv \inf\{t \ : \ f_1(t) > 0\} < \infty$.

Suppose $f_1(t') \leq 0$ for some $t' > t^0$. Then by definition of t^0 there exists a $t'' < t'$ such that $f_1(t'') > 0$, and so absolute continuity of f_1 implies that $f_1(t) > 0$ and $f_1'(t) < 0$ on some positive-measure subset of $[t'', t']$. This contradicts our earlier finding, so $f_1(t) > 0$ for all $t > t^0$. Hence the HJB equation implies that $\lambda^1(t) = 0$ for a.e. $t > t^0$.

Further, by definition of t^0 it must be that $f_1(t) \leq 0$ for all $t < t^0$. If $f_1(t) = 0$ on some positive-measure subset of $[0, t^0)$, then $f_1'(t) > 0$ a.e. on this set and so there would exist a $t' < t^0$ such that $f_1(t') > 0$ a contradiction. Hence $f_1(t) < 0$ for almost every $t \in [0, t^0]$, implying by the HJB equation that $\lambda^1(t) = \overline{\lambda}$ almost everywhere on $[0, t^0]$.

The fact that firm 1 does not prospect subsequent to t^0 means that firm 2 is in autarky after t^0. If $\mu^2(t^0) < \pi^A$, then its optimal prospecting rate is 0. But the assumption of $\hat{t}^1 = \infty$ means $\mu^1(t^0) > \pi^A$, so firm 1 would also also be autarky after t^0 with beliefs above the autarky threshold, contradicting the optimality of $\lambda^1(t) = 0$. On the other hand, if $\mu^2(t^0) > \pi^A$, then firm 2's optimal prospecting rate after t^0 is $\overline{\lambda}$ forever. But in this case eventually $\mu_+^1(t) < 1/R$, meaning $t_1^* < \infty$, a contradiction of our assumption. So it must be that $\mu^2(t^0) = \pi^A$. This implies in particular that $V^2(t^0) = 0$ and $f_2(t^0) = 0$ given that firm 2 is in autarky after t^0.

Meanwhile $\mu^2(t) > \pi^A$ for all $t < t^0$ given that firm 1 prospects at a strictly positive rate and invests immediately until t^0. So $t^0 = \hat{t}^2$, and Lemma 7 tells us that for almost every $t \in [0, t^0]$, either $f_2(t) < 0$ or $f_2'(t) > 0$. If ever $f_2(t') > 0$ for some $t' \in [0, t^0)$, then $f_2(t^0) = 0$ and absolute continuity of f_2 would imply $f_2(t) > 0$ and $f_2'(t) < 0$ on a positive-measure subset of $[t', t^0]$, a contradiction. So $f_2(t) \leq 0$ on $[0, t^0]$, and further $f_2(t) < 0$ for almost every $t \in [0, t^0]$. Thus by the HJB equation $\lambda^2(t) = \overline{\lambda}$ a.e. on $[0, t^0]$, implying $\mu^1(t^0) = \pi^A$ given that $\mu^2(t^0) = \pi^A$ and both firms employ the same prospecting and investing strategy prior to t^0. This yields the desired contradiction of our hypothesis that $\hat{t}^1 = \infty$.

So it must be that each $\hat{t}^i < \infty$. Wlog we assume that $\hat{t}^1 \leq \hat{t}^2$ going forward. We next prove that $V^1(\hat{t}^1) = 0$. For the time being, suppose instead that $V^1(\hat{t}^1) > 0$.

Given the hypothesis on $V^1(\hat{t}^1)$, the definition of $\hat{t}^1$, and the continuity of V^1 and μ^1, for sufficiently large $t < \hat{t}^1$ it must be that $f_1(t) > 0$. Hence $\lambda^1(t) = 0$ from the HJB equation, meaning $\mu^2(t)$ is constant and therefore $\hat{t}^2 > \hat{t}^1$. If there existed a $t' \in (\hat{t}^1, \hat{t}^2)$ such that $\mu^1(t) < \pi^A$, then $\lambda^1(t) = 0$ a.e. on $[t', \hat{t}^2]$, meaning $\mu^2(t)$ would be constant on that interval, a contradiction of the definition of $\hat{t}^2$. So $\mu^1(t) = \pi^A$ on $[\hat{t}^1, \hat{t}^2]$, reducing the HJB equation for V^1 to $rV^1(t) = \dot{V}^1(t)$, with solution $V^1(t) = e^{r(t-\hat{t}^1)}V^1(\hat{t}^1)$. So $V^1(t) > 0$ on $[\hat{t}^1, \hat{t}^2]$, meaning that $\lambda^1(t) = 0$ a.e. on the interval, yielding a constant μ^2 and another contradiction. So it must be that $V^1(\hat{t}^1) = 0$ and hence also $f_1(\hat{t}^1) = 0$.

Given $f_1(\hat{t}^1) = 0$, a nearly identical argument to that applied to f_2 prior to t^0 implies that $\lambda^1(t) = \overline{\lambda}$ a.e. on $[0, \hat{t}^1]$. Then surely $\hat{t}^2 \leq \hat{t}^1$ no matter what prospecting policy firm 2 chooses, meaning $\hat{t}^1 = \hat{t}^2 = \hat{t}$ and $\lambda^2(t) = \overline{\lambda}$ a.e. on $[0, \hat{t}]$. Then by definition of $\hat{t}$ it must be that $\hat{t} = (\mu^{\overline{\lambda}})^{-1}(\pi^A)$.

Finally, suppose that for some i and some positive-measure subset of $[\hat{t}, \infty)$, that $\lambda^i(t) > 0$. Then given $t_i^* = \infty$, firm $-i$'s value at $\hat{t}$ would be strictly positive. But then for sufficiently large $t < \hat{t}$ we would have $f_{-i}(t) > 0$, contradicting $\lambda^{-i} = \overline{\lambda}$. So it must be that $\lambda^1(t) = \lambda^2(t) = 0$ a.e. on $[\hat{t}, \infty)$. □

The lemma just proven shows that the prospecting strategies in the proposition statement are the only ones possibly consistent with any PBE satisfying $t_1^* = t_2^* = \infty$. It remains only to establish that the prospecting and investment strategies together actually constitute an equilibrium.

Fix a firm i, and consider any continuation game in which it has already obtained a High signal. Because $\mu^i \geq \pi^A$, therefore $\mu_+^i(t) > 1/R$ for all time. So investment is always profitable at each future time, regardless of whether the other firm has invested or not. Therefore the payoff of any investment strategy which occasionally never invests is dominated by the payoff of an investment strategy which always eventually invests, and due to time discounting all strategies involving delay in investment yield a strictly lower payoff than a strategy which invests immediately. So investing immediately is an optimal continuation strategy in all such continuation games, implying optimality of $t_i^* = \infty$.

Now consider firm i's optimal prospecting problem prior to obtaining a signal. Subsequent to the cutoff time t^A its beliefs are exactly π^A, so no prospecting is trivially an optimal strategy at this point. So consider times prior to t^A. We first show that $V^\dagger(t) = K(\mu^i(t) - \pi^A)$ is a supersolution to firm i's HJB equation on $[0, t^\dagger]$. Recall that the HJB equation for firm

i in this regime is

$$rV^i(t) = \overline{\lambda}\left(K(\mu^i(t) - \pi^A) - V^i(t)\right)_+ - \frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}\left(\overline{V} - V^i(t)\right) + \dot{V}^i(t).$$

Inserting $V^\dagger$ into the rhs yields $-\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)}(\overline{V} - K(\pi_+ - \pi^A))$, which is bounded above by 0 when $c \leq \bar{c}$. Meanwhile $\mu^i(t) > \pi^A$ for times before t^A, so inserting $V^\dagger$ into the lhs yields a strictly positive quantity. So indeed $V^\dagger$ is a supersolution to the HJB equation. And $V^\dagger(t^A) = 0$ by definition of t^A, while also $V^i(t^A) = 0$ given that firm i is in autarky with beliefs π^A subsequent to t^A. Therefore $V^\dagger(t) \geq V^i(t)$ for all $t \in [0, t^A]$. The HJB equation then implies that prospecting at the maximum rate prior to t^A is an optimal strategy.

As both the prospecting and investment strategy of each firm under the specified strategy profile are best responses to the other firm's strategy, the strategy profile constitutes a perfect Bayesian equilibrium.

C.7 Proof of Proposition 2

In this proof we will assume that $t_1^* < t_2^*$, i.e. that firm 1 takes on the role of the follower. The proposition can be obtained by swapping the roles of the two firms.

This proof first develops a series of lemma characterizing $\bar{t}$ and showing both necessity of sufficiency of the structure in the strategy profile in the proposition statement for equilibrium, under the assumption that $\mu^2(t_1^*) > \pi^A$, as required in any equilibrium such that $t_1^* < t_2^*$ by Lemma 9. We conclude the proof by showing that the strategy profile in the proposition statement in fact satisfies this condition.

Our first lemma shows that whenever $t_1^* < \infty$ and $\mu^2(t_1^*) > \pi^A$, then firm 2 optimally takes on the role of the leader.

Lemma 11. *Assume $c \leq \bar{c}$. Suppose firm 1 employs a strategy such that $t_1^* < \infty$ and a prospecting rule inducing $\mu^2(t_1^*) > \pi^A$. Then firm 2's essentially unique best response sets $\lambda^2(t) = \overline{\lambda}$ for all time and $t_2^* = \infty$.*

Proof. First note that $\mu^2(t_1^*) > \pi^A$ implies $\mu_+^2(t_1^*) > 1/R$. Hence $\mu_+^2(t) > 1/R$ for all time given that μ^2 is decreasing and is constant subsequent to t_1^*. So consider any continuation game in which firm 2 has already obtained a High signal. Given $\mu^2 > 1/R$, investment is always profitable at each future time, regardless of whether firm 1 has invested or not. Therefore the payoff of any investment strategy which occasionally never invests is dominated by the payoff of an investment strategy which always eventually invests, and due to time discounting all strategies involving delay in investment yield a strictly lower payoff than a strategy which invests immediately. So investing immediately is an optimal continuation strategy in all such continuation games, implying optimality of $t_2^* = \infty$.

Now consider firm 2's optimal prospecting problem prior to obtaining a signal. As 2 is in autarky subsequent to t_1^* with beliefs strictly above the autarky threshold, prospecting at rate $\overline{\lambda}$ is the essentially unique optimal strategy for $t \geq t_1^*$. So consider times $t < t_1^*$. We will show that $V^\dagger(t) = K(\mu^2(t) - \pi^A)$ is a supersolution to firm 2's HJB equation on $[0, t_1^*]$. Recall that the HJB equation for firm i in this regime is

$$rV^2(t) = \overline{\lambda}\left(K(\mu^2(t) - \pi^A) - V^2(t)\right)_+ - \frac{\dot{\mu}^2(t)}{\pi_+ - \mu^2(t)}\left(\overline{V} - V^2(t)\right) + \dot{V}^2(t),$$

and is well-defined for almost every $t \in [0, t_1^*]$. Inserting $V^\dagger$ into the rhs yields $-\frac{\dot{\mu}^2(t)}{\pi_+ - \mu^2(t)}(\overline{V} - K(\pi_+ - \pi^A))$, which is bounded above by 0 when $c \leq \overline{c}$. Meanwhile $\mu^2 > \pi^A$, so inserting $V^\dagger$ into the lhs yields a strictly positive quantity. So indeed $V^\dagger$ is a supersolution to the HJB equation.

Now note that as firm 2 is in autarky at time t_1^*, its value function at this point is $V^2(t_1^*) = \frac{\overline{\lambda}}{\overline{\lambda}+r}K(\mu^2(t_1^*) - \pi^A) < V^\dagger(t_1^*)$. This fact combined with the fact that $V^\dagger$ is a supersolution to the HJB equation implies $V^\dagger(t) > V^i(t)$ for all $t \in [0, \overline{t}]$. The HJB equation then implies that prospecting at the maximum rate prior to t_1^* is the essentially unique optimal strategy for firm 2. □

Now we characterize the follower's best reply to the leader behavior characterized in the previous lemma. We first characterize a threshold $\overline{t}$, which we then prove is the cutoff at which the follower optimally begins shirking.

Define a function $\check{V}(\mu)$ for $\mu \in [\pi_-, \pi_0]$ by

$$\check{V}(\mu) \equiv \max\left\{h(\mu)(\mu_+ R - 1), \frac{\mu - \pi_-}{\pi_+ - \pi_-}\frac{\overline{\lambda}}{\overline{\lambda} + r}h(\pi_+)(\pi_{++}R - 1)\right\}.$$

This function is the continuation value of a firm which has just received a signal when its current beliefs are μ and it has not seen the other firm invest, supposing the other firm prospects at rate $\overline{\lambda}$ and invests immediately forever after. The maximum captures the firm's choice of either investing immediately or waiting for the other firm to invest supposing its signal turns out to be High.

Note that $\check{V}(\mu)$ does *not* condition on the signal received having been High. The following lemma shows that when when the probability of receiving a High signal is factored out, $\check{V}$

can be written in terms of Δ.

Lemma 12. $\check{V}(\mu) = h(\mu)(\mu_+ R - 1 + \max\{\Delta(\mu), 0\})$.

Proof. Algebra. $\square$

By the previous lemma and the definition of μ^*, $\check{V}(\mu) = h(\mu)(\mu_+ R - 1)$ whenever $\mu \in (\mu^*, \pi_0]$, and otherwise $\check{V}(\mu) = \frac{\mu - \pi_-}{\pi_+ - \pi_-} \frac{\overline{\lambda}}{\overline{\lambda} + r} h(\pi_+)(\pi_{++} R - 1)$.

Now define a function $\widetilde{\Delta}$ for $\mu \in [\pi_-, \pi_0]$ by

$$\widetilde{\Delta}(\mu) \equiv \frac{\mu - \pi_-}{\pi_+ - \pi_-} \frac{\overline{\lambda}}{\overline{\lambda} + r} \overline{V} - \check{V}(\mu) + c.$$

Lemma 13. $\widetilde{\Delta}$ *is a strictly decreasing function and* $\widetilde{\Delta}(\pi_-) > 0$.

Proof. Let

$$\hat{\Delta}(\mu) \equiv \frac{\mu - \pi_-}{\pi_+ - \pi_-} \frac{\overline{\lambda}}{\overline{\lambda} + r} (\overline{V} - h(\pi_+)(\pi_{++} R - 1)) + c.$$

Differentiate $\hat{\Delta}$ to obtain

$$\hat{\Delta}'(\mu) = \frac{1}{\pi_+ - \pi_-} \frac{\overline{\lambda}}{\overline{\lambda} + r} (\overline{V} - h(\pi_+)(\pi_{++} R - 1)).$$

Given the assumption that $D \leq 0$, the final term is bounded above by $-c$, so $\hat{\Delta}'(\mu) < 0$ for all μ.

Now, $\widetilde{\Delta}(\mu) = \hat{\Delta}(\mu)$ for $\mu \leq \mu^*$, while $\widetilde{\Delta}(\mu) \leq \hat{\Delta}(\mu)$ for $\mu > \mu^*$. Clearly $\widetilde{\Delta}'(\mu) < 0$ for $\mu < \mu^*$. Meanwhile as $\widetilde{\Delta}$ is continuous at μ^* and an affine function of μ on $[\mu^*, \pi_0]$, to ensure $\widetilde{\Delta} \leq \hat{\Delta}$ it must be that $\widetilde{\Delta}'(\mu) = \widetilde{\Delta}'(\mu^*+) \leq \hat{\Delta}'(\mu^*) < 0$ for $\mu \in (\mu^*, \pi_0]$. Hence $\widetilde{\Delta}$ is a strictly decreasing function. Finally, note that $\widetilde{\Delta}(\pi_-) = c > 0$. $\square$

In light of the previous lemma, define a belief threshold $\overline{\mu} \in (\pi_-, \pi_0]$ by

$$\overline{\mu} \equiv \begin{cases} \pi_0, & \widetilde{\Delta}(\pi_0) \geq 0, \\ \Delta^{-1}(0), & \widetilde{\Delta}(\pi_0) < 0. \end{cases}$$

Lemma 14. $\overline{\mu} < \mu^*$ *when* c *is sufficiently small.*

Proof. Note that c appears nowhere in the definition of Δ, and so μ^* is independent of c. First suppose that $\mu^* = \pi_0$. Note that $\widetilde{\Delta}(\pi_0)$ may be written

$$\widetilde{\Delta}(\pi_0) = h(\pi_0) \left(\frac{\overline{\lambda}}{\overline{\lambda} + r} \overline{V} - \max\left\{ \pi_+ R - 1, \frac{\overline{\lambda}}{\overline{\lambda} + r} h(\pi_+)(\pi_{++} R - 1) \right\} \right) + c.$$

When $c \downarrow 0$, $\overline{V} \to \max\left\{\pi_+ R - 1, \frac{\overline{\lambda}}{\overline{\lambda}+r} h(\pi_+)(\pi_{++}R - 1)\right\}$. Thus the first term approaches a strictly negative value in this limit, while the second term approaches zero. This means $\widetilde{\Delta}(\pi_0) < 0$ for small c, i.e. $\overline{\mu} < \pi_0 = \mu^*$.

Next suppose $\mu^* < \pi_0$. In this case $\Delta(\mu^*) = 0$ and hence $\check{V}(\mu^*) = \frac{\mu^* - \pi_-}{\pi_+ - \pi_-} \frac{\overline{\lambda}}{\overline{\lambda}+r} h(\pi_+)(\pi_{++}R - 1)$. So $\widetilde{\Delta}(\mu^*)$ may be written

$$\widetilde{\Delta}(\mu^*) = \frac{\mu^* - \pi_-}{\pi_+ - \pi_-} \frac{\overline{\lambda}}{\overline{\lambda} + r} \left(\overline{V} - h(\pi_+)(\pi_{++}R - 1)\right) + c.$$

$\overline{V}$ is decreasing in c, but due to time discounting $\overline{V} < h(\pi_+)(\pi_{++}R - 1)$ even in the limit as $c \downarrow 0$. So the first term is negative and bounded away from 0 for all c, meaning that for sufficiently small c, it must be that $\widetilde{\Delta}(\mu^*) < 0$. Hence $\overline{\mu} < \mu^*$ in this case as well. □

Define a corresponding time threshold $\bar{t} \equiv (\mu^{\overline{\lambda}})^{-1}(\overline{\mu})$. Note that the previous lemma implies $\bar{t} > t^{**}$ for sufficiently small c. THe following lemma shows that the follower's best reply to the leader's behavior is to work until time $\bar{t}$, then shirk forever after.

Lemma 15. *Assume $c \le \bar{c}$. Suppose firm 2 chooses a strategy such that $\lambda^2(t) = \overline{\lambda}$ for all time and $t_2^* = \infty$. Then firm 1's essentially unique best response sets $t_1^* = t^{**}$ and $\lambda(t) = \overline{\lambda}\mathbf{1}\{t < \bar{t}\}$.*

Proof. Given firm 2's strategy, μ_+^2 must eventually drop below $1/R$. Then by Lemma 4, we know that in any subgame in which firm 2 has received a High signal but seen no investment, its optimal continuation strategy is either to invest immediately or to wait forever for investment by firm 2. Given this, the definition of t^{**} implies that $t_1^* = t^{**}$ is firm 1's optimal investment cutoff.

Now consider firm 1's optimal prospecting strategy. Note that 1's beliefs evolve as $\mu^1 = \mu^{\overline{\lambda}}$. Let $\widetilde{V}^1(t)$ be firm 1's continuation value at time t conditional on having received a High signal but seen no investment. Then

$$\widetilde{V}^1(t) = \max\left\{\mu_+^{\overline{\lambda}}(t)R - 1, \frac{\mu_+^{\overline{\lambda}}(t) - \pi_{+-}}{\pi_{++} - \pi_{+-}} \frac{\overline{\lambda}}{\overline{\lambda} + r}(\pi_{++}R - 1)\right\}.$$

Lemma 12 establishes that $h(\mu^{\overline{\lambda}}(t))\widetilde{V}^1(t) = \check{V}(\mu^{\overline{\lambda}}(t))$.

Given the strategy of firm 2, V^1 satisfies the HJB equation

$$rV^1(t) = \overline{\lambda}\left(\check{V}(\mu^{\overline{\lambda}}(t)) - c - V^1(t)\right)_+ - \frac{\dot{\mu}^{\overline{\lambda}}(t)}{\pi_+ - \mu^{\overline{\lambda}}(t)}(\overline{V} - V^1(t)) + \dot{V}^1(t)$$

for all time.

We now show that $V^\dagger(t) \equiv \frac{\mu^{\overline{\lambda}}(t)-\pi_-}{\pi_+-\pi_-}\frac{\overline{\lambda}}{\overline{\lambda}+r}\overline{V}$ satisfies the HJB equation for $t \geq \overline{t}$. By definition of $\overline{t}$, $\check{V}(\mu^{\overline{\lambda}}(t)) - c - V^\dagger(t) = -\widetilde{\Delta}(\mu^{\overline{\lambda}}(t)) \leq 0$ for $t \geq \overline{t}$. So the first term on the rhs of the HJB equation vanishes when $V^\dagger$ is inserted. Then after eliminating $\dot{\mu}^{\overline{\lambda}}(t)$ using the identity

$$-\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)} = \lambda^{-i}(t)\frac{\mu^i(t) - \pi_-}{\pi_+ - \pi_-},$$

a little algebra shows that $V^\dagger$ satisfies the HJB equation.

As $V^\dagger$ is a bounded C^1 function, it follows by a standard verification argument that $V^1(t) = V^\dagger(t)$ for $t \geq \overline{t}$. Further, $\check{V}(\mu^{\overline{\lambda}}(t)) - c - V^1(t) = -\widetilde{\Delta}(\mu^{\overline{\lambda}}(t)) < 0$ for $t > \overline{t}$, hence the essentially unique optimal prospecting strategy for $t \geq \overline{t}$ is $\lambda^1(t) = 0$.

If $\overline{t} = 0$, then this finishes the proof of the lemma. So assume $\overline{t} > 0$ going forward. We next show that $V^\ddagger(t) \equiv \check{V}(\mu^{\overline{\lambda}}(t)) - c$ is a strict supersolution to the HJB equation for $t \leq \overline{t}$. First note that $t \leq \overline{t}$ implies $\widetilde{\Delta}(\mu^{\overline{\lambda}}(t)) \geq 0$ and therefore

$$V^\ddagger(t) = \check{V}(\mu^{\overline{\lambda}}(t)) - c \geq \frac{\mu^{\overline{\lambda}}(t) - \pi_-}{\pi_+ - \pi_-}\frac{\overline{\lambda}}{\overline{\lambda} + r}\overline{V} > 0,$$

so the lhs of the HJB equation is always strictly positive when $V^\ddagger$ is inserted.

Now consider times $t \leq t^{**}$. On this time range $\check{V}(\mu^{\overline{\lambda}}(t)) - c = K(\mu^{\overline{\lambda}}(t) - \pi^A)$, and so the rhs of the HJB equation reduces to $-\frac{\dot{\mu}^{\overline{\lambda}}(t)}{\pi_+-\mu^{\overline{\lambda}}(t)}(\overline{V} - K(\pi_+ - \pi^A))$, which is bounded above by 0 when $c \leq \overline{c}$. Thus $V^\ddagger$ is indeed a strict supersolution on $[0, \min\{t^{**}, \overline{t}\}]$. If $\overline{t} \leq t^{**}$ then this is all we need to show, so suppose instead that $t^{**} < \overline{t}$. Recall that for $t \in (t^{**}, \overline{t}]$, $\check{V}(\mu^{\overline{\lambda}}(t)) = \frac{\mu^{\overline{\lambda}}(t)-\pi_-}{\pi_+-\pi_-}\frac{\overline{\lambda}}{\overline{\lambda}+r}h(\pi_+)(\pi_{++}R - 1)$. Now insert $V^\ddagger$ into the HJB equation and use the identity

$$-\frac{\dot{\mu}^i(t)}{\pi_+ - \mu^i(t)} = \lambda^{-i}(t)\frac{\mu^i(t) - \pi_-}{\pi_+ - \pi_-}$$

to eliminate $\dot{\mu}^{\overline{\lambda}}(t)$. Some algebra then shows that the lhs minus the rhs of the HJB equation is equal to the quantity

$$F(t) \equiv -\overline{\lambda}\frac{\mu^{\overline{\lambda}}(t) - \pi_-}{\pi_+ - \pi_-}(\overline{V} - K(\pi_+ - \pi^A)) - rc.$$

Now, for $t \in (t^{**}, \overline{t}]$ it must be that $\mu^{\overline{\lambda}}(t) \in [\overline{\mu}, \mu^*)$, and therefore

$$\widetilde{\Delta}(\mu^{\overline{\lambda}}(t)) = \frac{\mu^{\overline{\lambda}}(t) - \pi_-}{\pi_+ - \pi_-}\frac{\overline{\lambda}}{\overline{\lambda} + r}(\overline{V} - K(\pi_+ - \pi^A)) + c \leq 0,$$

or equivalently

$$-\overline{\lambda}\frac{\mu^{\overline{\lambda}}(t)-\pi_-}{\pi_+-\pi_-}(\overline{V}-K(\pi_+-\pi^A))\geq(\overline{\lambda}+r)c.$$

This bounds $F(t)$ from below as $F(t)\geq\overline{\lambda}c>0$. So indeed $V^{\ddagger}$ is a strict supersolution for all $t\leq\overline{t}$.

Finally, note that $V^{\ddagger}(\overline{t})=\frac{\mu^{\overline{\lambda}}(t)-\pi_-}{\pi_+-\pi_-}\frac{\overline{\lambda}}{\overline{\lambda}+r}\overline{V}=V^1(\overline{t})$ from the definition of $\overline{t}$ and the earlier characterization of $V^1(t)$ for $t\geq\overline{t}$. This fact combined with the result that $V^{\ddagger}$ is a strict supersolution for $t\leq\overline{t}$ implies that $V^{\ddagger}(t)>V^1(\overline{t})$ for all $t<\overline{t}$. The HJB equation then implies that the essentially unique optimal prospecting policy is $\lambda^1(t)=\overline{\lambda}$ for $t\leq\overline{t}$, completing the proof. □

By Lemma 6, any equilibrium such that $t_1^*<\infty$ must satisfy $t_1^*=t^{**}$ and $t_2^*=\infty$. Further, by Lemma 9 firm 1's strategy must induce $\mu^2(t_1^*)>\pi^A$. Under these conditions Lemma 11 uniquely pins down firm 2's optimal strategy, and then Lemma 15 uniquely pins down firm 1's best reply.

The only step that must be taken to confirm that this strategy profile is actually an equilibrium is to verify that firm 1's strategy profile in fact induces $\mu^2(t_1^*)>\pi^A$. Suppose first that $\overline{t}\leq t^{**}$. In this case firm 2's beliefs remain fixed after $\overline{t}$, and so $\mu^2(t_1^*)=\mu^2(\overline{t})=\overline{\mu}$. If $\overline{\mu}=\pi_0$ then the result is automatic, so assume $\overline{\mu}<\pi_0$. Recall that in this case $\overline{\mu}$ is pinned down by the condition $\widetilde{\Delta}(\overline{\mu})=0$, and given $\overline{t}\leq t^{**}$ and therefore $\overline{\mu}\geq\mu^*$ this condition may be written

$$0=\frac{\overline{\mu}-\pi_-}{\pi_+-\pi_-}\frac{\overline{\lambda}}{\overline{\lambda}+r}\overline{V}-K(\overline{\mu}-\pi^A).$$

As $\overline{\mu}>\pi_-$, the first term on the rhs is strictly positive, meaning so must be the second. Hence $\overline{\mu}>\pi^A$ in this case.

Now suppose instead that $t^{**}<\overline{t}$, in which case $\mu^2(t_1^*)=\mu^*$. Now $\mu^*>\overline{\mu}$, meaning $\widetilde{\Delta}(\mu^*)<0$. And this condition is just

$$\frac{\mu^*-\pi_-}{\pi_+-\pi_-}\frac{\overline{\lambda}}{\overline{\lambda}+r}\overline{V}-K(\mu^*-\pi^A)<0,$$

implying $\mu^*>\pi^A$ given $\mu^*>\pi_-$. Hence $\mu^2(t_1^*)=\mu^*>\pi^A$. So in all cases $\mu^2(t_1^*)>\pi^A$, proving the proposition.

C.8 Proof of Proposition 3

Consider first the symmetric equilibrium. In this equilibrium, one best response is the equilibrium strategy of prospecting at the maximum rate until time t^A and then shutting down forever. However, for $t \geq t^A$ each firm's beliefs are constant at $\mu(t) = \pi^A$, with no investment by the other firm. Hence another best response after time t^A is to continue prospecting at the maximum rate forever and to invest if a High signal is received. Hence there exists a best response of the following sort - as long as there is no investment by the other firm, prospect forever until receiving a signal; if the other firm invests, continue prospecting forever until receiving a signal; and when a signal is received, invest immediately iff it is High, otherwise never invest. But this strategy is exactly the autarky strategy of ignoring the other firm's behavior, prospecting forever, and investing immediately iff a High signal is received. And as the presence of the other firm brings infformational externalities but no payoff externalities, this strategy must yield the autarky payoff. Hence each firm's equilibrium payoff must be the same as the autarky payoff.

In the leader-follower equilibrium, the follower never invests before the leader. Thus the leader is effectively in autarky and receives his autarky payoff. Meanwhile, by Lemma 15 and the fact that $t^{**} = 0$ under complements, it is an essentially unique best response for the follower to wait forever upon receiving a High signal at any time. Note that this is true regardless of the follower's prospecting strategy. So consider modifying the follower's autarky strategy to preserve the prospecting rule but to wait forever rather than investing immediately after obtaining a High signal. This must strictly improve on the autarky strategy. And the follower's equilibrium payoff must be at least as high as this modified autarky strategy, so in equilibrium the follower earns strictly more than under autarky.