Strategic investment and learning with private information[*]

Nicolas Klein[†] and Peter Wagner[‡]

July 5, 2018

Abstract

We study a two-player game of strategic experimentation in which agents choose the timing of investments which yield uncertain returns over time. Agents learn about future returns through privately observed signals, others' investment decisions and from public experimentation outcomes when returns are realized. We characterize symmetric equilibria, and we relate the extent of strategic delay of investments in equilibrium to the primitives of the information structure. Agents invest without delay when the most optimistic intermediate belief exceeds a threshold. Otherwise, delay in investments induces a negative learning feed-back which may either escalate or dampen beliefs and investment choices. We highlight how private information in strategic experimentation can increase ex ante welfare because of strategic uncertainty and due to an "encouragement effect of private information".

1 Introduction

Is transparency good for innovation? The question is at times controversially debated in the public square, often seeming to pit industry executives against regulators.[1] We propose a simple model of innovation in which, say, two pharmaceutical companies decide on making an irreversible investment, such as putting a certain medical substance on the market. The quality of the investment option is the same for both companies but initially unknown. Either company has some initial private information about its quality. As long as the substance is marketed, it generates some positive flow payoff. But, if the substance is of low quality, this will be discovered during its marketing phase after an exponentially distributed time, e.g, because a consumer is harmed. In this case, the drug is banned, the game ends, and the company responsible for causing the harm is fined so heavily that it regrets having marketed the drug.

---

[1]See e.g. the report *Ending Drug Trial Secrecy Is a Boon For R & D, Says EU Agency* (Reuters, 2013).

We show that, in our setting, keeping the companies' initial private information private may be beneficial for welfare.[2] Indeed, on account of the irreversibility of investments, agents wait too long before they invest in the hope that the other provides free information. Excessive delay *before investments* is the result of *leadership-aversion.* Indeed, the follower can get additional information concerning the quality of the investment by observing the leader's experience; moreover, conditionally on the project being bad, the leader is more likely to be subjected to the fine. *After the first agent has invested,* the other agent trades off the forgone returns from investing with the benefit of waiting for more information, without taking into account the positive externality her investment would exert on her partner, leading to a delay in the follower's investment through a classical *free-riding effect.*

As our main result shows, not knowing one's partner's type may mitigate both sources of inefficiency. Indeed, when agents have private information, delaying their investments has drawbacks. An agent can only benefit from waiting to invest if the other agent invests first. However, one agent's delay may signal bad information to the other, who, in response, may be less willing to invest, thereby reducing the first agent's incentive to delay. As a result, an agent who received a bad signal may have an incentive to invest without delay nonetheless, in order to encourage the other agent to invest earlier. A second drawback results from each agent's uncertainty about the other agent's behavior. For example, an agent who received a good signal may wait in the hope that the other agent invests first. The other agent, however, may have received a bad signal and therefore never invest. Or an agent may delay her investment despite having received a good signal, so as to learn whether or not the other agent invests. However, by delaying her investment, she discourages future investment by the other agent, even if the other has received a good signal. These negative effects of delay under private information can induce agents to invest earlier, and thereby reduce free-riding and leadership aversion. This reduction increases ex-ante social welfare as long as it does not result in inefficient over-investment. For example, an agent with a good signal may invest not knowing the other agent's signal, but regret investing if she learns that the other agent's signal is bad. In cases like these, it might be socially desirable to make agents' private information public.

Several papers have investigated the role of private information in games of informational externalities.[3] Rosenberg et al. (2013) investigate a game of strategic experimentation with exponential two-armed bandits, where players' action choices are public information, while the outcomes of these actions are private information. A player's switch to the safe arm is irreversible. They show that, in their setting, public information is unequivocally good for welfare. Their setup differs from ours inter alia by the fact that their players accrue private information over time, while ours are

---

[2]We do not include consumers in our model. Thus, if we wanted to include them in our welfare calculus, we should have to assume that the fine is set in such a way as to make the firm fully internalize the harm caused.

[3]The problem of strategic information acquisition in bandit games has been introduced by Bolton and Harris (1999) in a Brownian-motion environment. Keller et al. (2005) have extended the analysis to a Poisson setting, while Keller and Rady (2015) have introduced "bad news" Poisson events. Private information in this setting has also been analyzed in Rosenberg et al. (2007).

privately informed at the outset. Furthermore, there are no payoff externalities in Rosenberg et al. (2013), so that, once they have irreversibly exited the game, their players do not care about their partner's behavior any longer. Our players, by contrast, decide on when irreversibly to *enter* the market, and thus might have incentives to bias their partner's entry decision even conditional on their entering themselves.

Heidhues et al. (2015) investigate the problem with public reversible actions and private payoffs, while allowing for cheap-talk communication among players. They show that equilibria with public information can always be replicated under private information, so that private information is unequivocally good for welfare. This conclusion heavily depends on their assumption that players can communicate among each other. In contrast, Bonatti and Hörner (2011), analyze the case of unobservable and reversible actions and observable outcomes, and find that private information boosts welfare in their setting. The reason is that, with observable actions, a deviation by a player may render the other players more optimistic, and hence more willing to pick up the slack caused by the deviation. This in turn makes deviating more attractive under public information.

In Chamley and Gale (1994) and Murto and Välimäki (2011, 2013) , information is dispersed throughout society, and agents decide when to take an irreversible action. As in our setting, information is inefficiently aggregated because investors have an incentive to delay their decision so as to acquire more information by observing the behavior of others. The players do not observe the results of their partners' experimentation directly. Moreover, once an agent takes an action, she effectively exits the game, and is no longer affected by others decisions. In our setting, by contrast, players continue to be affected by others actions after taking their irreversible action, and thus have an incentive to influence their partners actions and beliefs when making their irreversible investment decision.

A closely related paper is Décamps and Mariotti (2004), who also study a two-player game of irreversible investments. Private information in their setting, however, pertains to the players' idiosyncratic investment costs, while all information concerning the common quality of the investment opportunity is public. Since, as in our setting, players get additional (public) information after the other player has invested, they prefer the role of follower and thus have an incentive to convince each other that their own costs of investment are high. Once they have made their irreversible investment decision, Décamps and Mariotti (2004)'s players do not care about their partner's actions any longer, as externalities are purely informational in their setting. In our setting, by contrast, there is a payoff externality, as only one player will incur the costs of an accident. Thus, even conditionally on having made the irreversible investment decision, a player prefers her partner to invest as soon as possible in our setting. Thus, in contrast to Décamps and Mariotti (2004), our players have incentives to render their partners as optimistic as possible concerning the common investment prospects.

Moscarini and Squintani (2010) consider a model of a winner-takes-all R&D competition in

which firms observe an initial private signal about the unknown type of a research project, which is drawn from a continuous distribution. Over time, firms learn about the project's type from their competitor's actions and the lack of success in the past, deciding when to exit irreversibly. They show that the aggregate duration of experimentation is longer under private information, when firms may also exit simultaneously, a non-generic outcome under public information.

The potential welfare improvement of private information is related to the "smoothing effect of uncertainty" (Morris and Shin, 2002). Teoh (1997) demonstrates this effect in a model of public-good provision, where the marginal return to agents' investments is determined by an uncertain state of the world. The author shows that non-disclosure of information may increase ex-ante welfare when the investment has marginally diminishing returns, because the loss resulting from a reduction in investment after the release of bad news outweighs the benefits from increased investment when the information is favorable. This is the same mechanism that drives the main result in our paper: when bad news is publicly disclosed, free-riding and leadership-aversion increase, leading to an over-proportional reduction in the expected value of investment.

2   Model

There are two agents, indexed $i = 1, 2$. Time $t \in \mathbb{R}_+$ is continuous, with an infinite horizon, and future payoffs are discounted at the common discount rate r. Each agent decides when to invest in a project which generates a payoff stream that depends on an unknown state of the world $\theta \in \{G, B\}$, which is either good ($\theta = G$) or bad ($\theta = B$). In either state, the project yields a certain flow return $y > 0$. In state B, an accident occurs at a random time corresponding to the first jumping time of a Poisson process with parameter $\gamma > 0$. An accident never occurs in state G. The game ends at the time of the first accident. The agent suffering the first accident incurs a lump-sum cost of $c > 0$. 0At the outset, each agent $i = 1, 2$ observes a private binary signal $s^i \in \{g, b\}$, which provides information about the realization of the state. The agents share a private initial belief that the state is good, which we denote by $p_0 \in (0, 1)$. Either agent's signal is correct (i.e., equals g [b] in state G [B]) with probability $\rho \in (1/2, 1)$. Conditionally on θ, the signal realizations are independent between the agents.

We model the continuous-time environment as a stopping game in two phases. At the beginning of the first phase, the agents play an attrition game in which each decides how long to wait before making the investment, conditional on the event that the other agent has not yet invested. The initial stage ends after the first agent (the *leader*) invests. In the second phase, the leader decides how long to wait before exiting if the other agent (the *follower*) doesn't invest, and the follower decides when to invest, conditional on the leader not having exited. To simplify the exposition, we assume that a leader who has exited cannot reinvest (thus, exiting is irreversible).

Formally, define an investment history at any time t to be a profile $h = (\tau^1, \tau^2)$, where τ^i represents the increasing sequence of "switching times" at which agent i changed his investment

decision. Specifically, we have $\tau^i = 0$ if agent i did not invest in the past, $\tau^i = (0, t_0)$ if agent i invested at time $t_0 \geq 0$ but has not exited, and $\tau^i = (0, t_0, t_1)$ if agent i invested at time t_0 and exited at time $t_1 \geq t_0$. For any history, the most recent switching time $T(h) = \max\{h\}$ denotes the beginning of the current game phase. Each agent's strategy specifies a distribution $F_i(t|s_i, h)$ over switching times on $[T(h), \infty)$ for any investment history $h \in H$. Thus, the strategy of agent i is characterized by a family of cumulative distribution function $\{F_i(t|s_i, h)\}_{h \in H}$, where $F_i(t|s_i, h)$ represents the probability that agent i with signal s_i takes action (invests or exits) before time t after investment history h, conditional on (i) the other agent $-i$ not taking action before t and (ii) agent i has not exited in the past. Our solution concept is symmetric perfect Bayesian equilibrium. In any equilibrium, agents will indefinitely delay or immediately exit prior investment upon the arrival of an accident, which we take as given in the subsequent analysis.

As an restriction on beliefs, we impose that after one agents' deviation from the equilbrium path, the other agent's belief about the state does not change beyond what he could learn about the deviator's private information.[4] We further assume that any deviation resulting in an unexpectedly early investment of one agent conveys good news to the competitor, and unexpected delay conveys bad news. This restriction is innocuous, and it's sole purpose is to rule out a class of artificial equilibria in which agents are deterred from deviating because of the other agent's off-equilibrium belief after their deviation. For example, it is straightforward to construct an equilibrium in which the good type of each agent does not invest because the other agent would then be convinced that the investing agent's type is bad, and would therefore delay her own investment excessively, so that it is no longer profitable for good types to invest.

Throughout, we denote by p_t the (ex ante random) equilibrium public belief that $\theta = G$ at time t in the first phase of the game. Similarly, we denote by q_{it} the posterior belief assigned to agent i's type being g in the first phase of the game (we omit the index i whenever the belief is the same for each agent). Furthmore, we write $p_t(s)$ and $q_t(s)$ for the posterior probabilities conditional on a private signal $s \in \{g, b\}$, and we denote by $p_t(s, s')$ and $q_{it}(s, s')$ the posterior probabilities conditional on any pair of private signals $(s, s') \in \{g, b\}^2$. Note that since signals are i.i.d. and symmetric, we have $p_t = p_t(g, b) = p_t(b, g)$ and $q_{it} = q_{it}(g, b) = q_{it}(b, g)$.

3 Equilibrium analysis

3.1 Public information

We first consider the case in which both signals are publicly observable. The agents then share a common initial belief about the state, which we denote by $\check{p}_0 := p_0(s^1, s^2)$. Once both agents have made the investment, there are no additional moves to make, and the game continues indefinitely or until an accident occurs.

---

[4]This is the standard "no signalling what you don't" condition from **?**

Each agent benefits from the other agent's project, because of the possibility that they learn the true state from another agent's accident without incurring a loss themselves. The expected net present value for the leader at belief p is

$$v_l(p, \tau) = py + (1 - p)\left((1 - e^{-(r+\gamma)\tau})\lambda_1 + e^{-(r+\gamma)\tau}\lambda_2\right)(y - \gamma c), \tag{1}$$

when the other agent delays her investment by τ, where $\lambda_k = r/(r + k\gamma)$ for each $k = 1, 2$. Similarly, the expected present value for the follower is

$$v_f(p, \tau) = e^{-r\tau}py + e^{-(r+\gamma)\tau}(1 - p)\lambda_2(y - c\gamma). \tag{2}$$

To state the following lemma, which reports basic properties of the functions v_l and v_f, we define the log-likelihood ratio $\phi(p) := \ln\left(\frac{p}{1-p}\right)$.

Lemma 1. *The function $v_l(p, \tau)$ is linearly increasing in p, convex and decreasing in τ for every $p \in (0, 1)$ and supermodular in (p, τ). The function $v_f(p, \tau)$ is linearly increasing in p and has a single peak in τ at*

$$\tau^*(p) = \begin{cases} \left(\phi(p_f^*) - \phi(p)\right)/\gamma & \text{if } p < p_f^* \\ 0 & \text{if } p \geq p_f^* \end{cases} \tag{3}$$

for every $p \in (0, 1)$, where

$$p_f^* = \frac{\lambda_2(\gamma c - y)}{\lambda_1 y + \lambda_2(\gamma c - y)}. \tag{4}$$

All proofs are found in the appendix. We write $v_f^*(p) = v_f(p, \tau^*(p))$ and $v_l^*(p) = v_l(p, \tau^*(p))$ for the values of the leader and the follower, respectively, given the follower uses the optimal delay. Since τ^* is weakly decreasing in p, and v_l and v_f are strictly increasing in p as well as decreasing in τ, it follows that v_l^* and v_f^* are strictly increasing functions in p. Moreover, v_l^* is continuous, positive if $p = 1$, and negative if $p = 0$. Hence, it has a unique root on $(0, 1)$, which we denote by p_l^*. Further, denote by $\underline{p}$ the threshold above which the value of both agents investing immediately is positive:

$$\underline{p} = \frac{\lambda_2(\gamma c - y)}{y + \lambda_2(\gamma c - y)}.$$

Lemma 2. $0 < \underline{p} < p_l^* < p_f^* < 1$.

It follows from Lemma 1, that, in any equilibrium, both agents invest immediately if $\check{p}_0 \geq p_f^*$. Indeed, if the belief exceeds this threshold, then it is optimal for the follower not to delay investing, so that it is optimal for each agent to invest immediately. If $\check{p}_0 \leq p_l^*$, neither agent invests in any equilibrium, by definition of p_l^*. If $p_l^* < \check{p}_0 < p_f^*$, there can be no symmetric equilibrium in

6

pure strategies, as the best response to the other agent's investment would be not to invest and vice-versa. The same argument rules out atoms in mixed-strategy equilibrium. Thus, a symmetric equilibrium must be in atomless mixed strategies, with each agent investing at a rate that renders the other agent indifferent between investing immediately and delaying her investment by any length of time. As long as neither agent has made the investment, no new information becomes available, so that the agents' equilibrium flow rate of investment β is constant over time. We can immediately calculate the equilibrium investment rate, using the fact that each agent must be indifferent between making the investment immediately and never making the investment, i.e.,

$$v_l^*(\check{p}_0) = \int_0^\infty \beta e^{-(\beta+r)t} v_f^*(\check{p}_0) dt.$$

Solving the equation for β gives the equilibrium investment rate as a function of the common belief $\check{p}_0$

$$\beta^*(\check{p}_0) = \max \left\{ \frac{r v_l^*(\check{p}_0)}{v_f^*(\check{p}_0) - v_l^*(\check{p}_0)}, 0 \right\}. \tag{5}$$

Observe that the rate of investment is positive whenever the value of becoming the leader is greater than 0, because it follows from $\check{p}_0 < p_f^*$ and Lemma 1 that the denominator of β^* is always strictly positive. Moreover, the difference between the value of the follower and the value of the leader converges to 0 as $\check{p}_0$ approaches the threshold p_f^*, so that the equilibrium rate of investment approaches infinity. The following theorem characterizes the unique symmetric equilibrium under public information.

Theorem 1 (Symmetric equilibrium with public information). *There is a unique symmetric equilibrium. If $\check{p}_0 \geq p_f^*$, both agents invest immediately in* every *equilibrium. If $p_l^* < \check{p}_0 < p_f^*$, each agent invests at constant rate $\beta^*(\check{p}_0)$ given by Equation* (5) *in the first phase, while the follower starts the project with delay $\tau^*(\check{p}_0)$. If $\check{p}_0 \leq p_l^*$, neither agent invests in* any *equilibrium.*

The equilibrium dynamics are similar to those of standard attrition games. The agents never invest if their are very pessimistic about the state of the world, and both invest immediately if they are very optimistic. For intermediate beliefs, each agent prefers the other to invest first, to enjoy the opportunity to gather additional information about the state before making a final decision. The informational externality thus generates a second-mover advantage. In equilibrium, the agents delay their investments strategically, and the rates of investment are such that each agent is indifferent between investing and waiting at each instant of time.

3.2 Private information

We now turn to symmetric equilibria in the case in which agents' signals are private. We here focus on the parameter values that are relevant for our main result showing that welfare can be higher under asymmetric information.[5] Thus, we shall focus on the case $p_0 > p_l^*$ and a relatively uninformative signal, i.e., low ρ. The following theorem summarizes our findings for this case.

Theorem 2 (Symmetric equilibrium with private information)**.**

(1.) *Suppose $p_0 > 0$ and $\rho \in (1/2, 1)$ are such that $p_0(g, g) \geq p_f^*$. Then there exists a symmetric equilibrium in which type g invests immediately and type b of each agent invests immediately with some probability $\eta^* \in [0, 1]$, and with probability $1 - \eta^*$, she invests at a random time arriving at constant rate $\beta^*(p_0(b, b))$.*

(2.) *Suppose $p_0 > 0$ and $\rho \in (1/2, 1)$ are such that $p_0(g, g) < p_f^*$. Then there exists a symmetric equilibrium in which type g of each agent invests at rate $\mu_t^* \geq 0$ given by* (15). *Type b delays investment until (possibly infinite) $t^* > 0$, given by* (16), *and invests at constant rate $\beta^*(p_0(b, b))$ thereafter.*

The preceding theorem reveals that the delay of investments in equilibrium depends on the most optimistic intermediate belief $p_0(g, g)$, which is the belief conditioned on the event that each agent observed the good signal g. If this belief exceeds the follower threshold p_f^*, then the equilibrium exhibits "aggressive investment" by optimistic agents. The equilibrium may be pooling, partially separating, or fully separating, depending on the prior belief and the informativeness of signals. A high prior and weak signals typically lead to pooling equilibria. Indeed, if the intermediate belief of pessimistic agents, conditional on their own signal, is high enough then it is optimal to invest if both types of the other agent invest, and thus pooling is an equilibrium. On the other hand, for any prior belief and highly informative signals, the equilibrium is partially or fully separating. Intuitively, an informative good signal provides a strong incentive to an agent to invest, and an informative bad signal makes it costly for a pessimistic agent to imitate an optimistic agent.

A pessimistic agent then faces a choice: either to invest immediately or to reveal his pessimism to his competitor by waiting. For for a pessimistic agent, the downside of delaying the investment is that immediate investment communicates good news that makes the other more willing to invest. This additional investment generates value through the information from more experimentation. In other words, agents who privately hold bad news choose to invest to avoid discouraging others from experimenting. This effect on incentives is reminiscent of the encouragement effect known from models of strategic experimentation (Bolton and Harris, 1999; Keller and Rady, 2010). The encouragement effect describes the incentive of an agent to experiment, driven by the possibility of generating good news in the near future to induce more experimentation of others in the long-run.

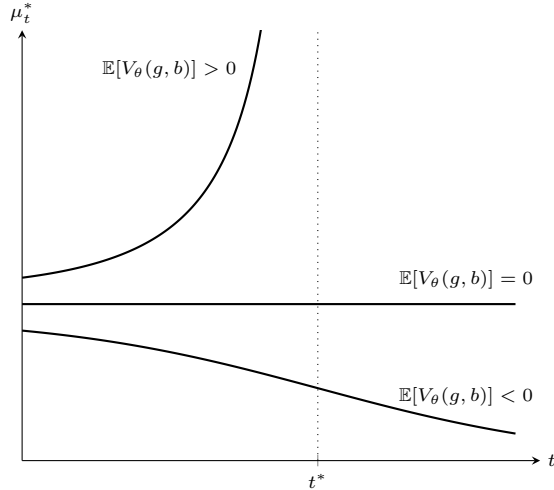[5]Results for the other parameter values are available from the authors upon request.

Figure 1: Three branches of equilibrium investment rates.

Here, the mechanism is slightly different. Agents who hold private information have an incentive to experiment to conceal bad news from others to encourage them to experiment more. Accordingly, we label this effect as encouragement through signaling.

When the most optimistic belief $p_0(g, g)$ lies below the follower threshold, then agents invest cautiously. Optimistic agents still invest, but with delay, because they benefit from the possibility, that the other invests first and then provides free information. Pessimistic agents strictly prefer to wait for the other agent to invest. Because different types invest with different probabilities, they separate gradually over time. While both agents wait for the first investment, they continuously update their beliefs about each other's type. The longer an agent has waited, the more likely it is that she observed a bad signal.

The dynamics of investments are determined by the evolution of the agents' belief about the state of the world θ. Each agent obtains information about the state from the investment behavior of the other. The longer the game proceeds without investment, the more pessemistic agents become both about type of the other agent and the state of the world. This change in beliefs in return affects their incentives to invest. The result of this mutual interaction is a learning feed-back loop.

The effects of the feed-back loop can be seen in the dynamics of investment rates, illustrated in Figure ??. The upper branch represents the case in which the leader value is positive, even if the other agent's signal was bad: $\mathbb{E}[V_\theta(g, b)] > 0$. If this is the case, then one agent's inactivity has a motivating effect on the other. As agents become more convinced that the other agent's type is bad, their value from waiting diminishes, and thus they are more willing to invest. This change in incentives is reflected by the increasing slope of investment rates. As investment rates increase, learning accelerates in a self-reinforcing manner and eventually explodes at a finite time t^*.

The lower branch corresponds to the case in which the worst-case continuation payoff $\mathbb{E}[V_\theta(g, b)]$

9

from immediate investment is negative. In this case, inactivity has a dampening effects on investments. As agents become more convinced that the other's signal is bad after a period of inactivity, it becomes more likely that their investment will result in losses. They are thus less willing to invest, and also reveal less information about their type. The investment rates eventually converge to zero, and type are never fully revealed in the long run.

That the investment rates are constant when $\mathbb{E}[V_\theta(g, b)]$ appears natural in the diagram, but is not obvious a priori. Inactivity still makes agents more pessimistic about the state of the world. At the same time, the opportunity cost of waiting decreases. In equilibrium, these two effects balance each other out perfectly.

4   Welfare and transparency

The following theorem characterizes the cooperative benchmark, which corresponds to a setting in which agents pool their private information and commit to a strategy at the outset of the game, seeking to maximize the sum of their payoffs. We shall see that, when signals are publicly observable, all equilibria are ex-post efficient when the prior belief that the state is good is either very high or very low.

Theorem 3 (Cooperative benchmark)**.** *In the cooperative benchmark, both agents invest immediately if $\check{p}_0 > \underline{p}$ and never invest if $\check{p}_0 < \underline{p}$. If $\check{p}_0 = \underline{p}$, either solution is optimal.*

Thus, as a comparison with the equilibrium in Theorem 2 shows, equilibrium under private information may be inefficient for intermediate prior beliefs. Inefficiencies arise because of leadership aversion: agents prefer being the follower and therefore wait if there is a sufficiently high probability that the other agent invests first. Free-riding aggravates its effect, because it lowers the value of being the leader, and thus reduces the rate at which the agents are willing to make the first investment.[6] It is therefore natural to ask if there is a social gain from increasing transparency. Naïve logic may suggest that more transparency should unambiguously lead to better outcomes, as it allows the agents to make better-informed decisions. This view, however, disregards the fact that agents wait inefficiently long for others to provide information for free when they are uncertain about the state, which negatively affects the efficiency of learning. Asymmetric information, as it turns out, can have positive side-effects. On the one hand, uncertainty introduces strategic risk about the other agent's future investment decisions, which lowers the incentives to delay initial investments. On the other hand, when signals are private, an early investment can signal good news, giving rise to an *encouragement-though-signalling effect,* by which an agent invests in order to motivate her partner to invest as well. Thus, there is a tradeoff between the gains from information aggregation and limiting public access to information in order to reduce the agents' incentives to free-ride.

---

[6]There is also a third type of inefficiency as agents may not invest in equilibrium although investment is optimal. With private signals, agents may also invest although investing is not ex-post socially optimal.

In the following, we state our main result, which formalizes this insight. It shows that private information is socially preferable if the prior belief that the state is good is high enough, and the signals are not too informative. To this end, denote by $W_0(s^1, s^2)$ the expected social surplus generated in the unique symmetric equilibrium under public information, when $p_0(s_1, s_2)$ is the common initial belief that the state is good. Let $\tilde{W}$ denote the ex-ante expected social surplus generated in a symmetric equilibrium in which each agent's signal is private information. The following theorem shows that, for some parameters, $\tilde{W} \geq \mathbb{E}[W_0(s_1, s_2)]$.

Theorem 4 (Welfare improvement through private information). *Let $p_0 > p_l^*$. There exists a $\rho^* > 1/2$ such that, for all signal precisions $\rho \in (1/2, \rho^*)$, we have $\tilde{W} \geq \mathbb{E}[W_0(s_1, s_2)]$. For $p_l^* < p_0 < p_f^*$, the inequality is strict.*

The condition on the signal precision ensures that it is socially optimal for both agents to invest immediately. Thus uncertainty only affects strategic interaction. If $p_0 > p_f^*$, there always exists a range of signal precisions guaranteeing that both types invest immediately under private information, as in the efficient benchmark. If $p_0 < p_f^*$, then there signal precisions such that the symmetric equilibrium has delayed entry. In this case, the welfare improvement results from strategic uncertainty and a spiraling effect, in which increasing pessimism about each other's types and increasing eagerness to invest result in ex-ante earlier investment.

5   Conclusion

We have investigated how private information affects agents' incentives to make irreversible investments. Agents benefit from their partner's experimentation because it provides them free information. However, delaying one's investment can signal bad news, thus discouraging experimentation by one's partner. We have shown that private information may increase social welfare in our setting by deterring delays in investments, as they may be interpreted as signaling bad news.

One might wonder about the impact of allowing communication between agents in our setting. We conjecture that allowing for cheap talk between agents would not change the outcome. Indeed, if it did, either agent would always send the most positive signal to induce the other one to invest early. We should furthermore anticipate communication with verifiable messages to lead to unraveling, as agents would prefer to communicate their signal only if it is good, so that each agent can infer the other agent's signal perfectly. A deeper analysis of the role of communication in our setting would thus likely require that one abandon either our binary state space or the binary signal structure. We commend such an investigation to future research.

6 Proofs

Proof of Lemma 1. (i) That v_l is linear in p is obvious from its definition in Equation (1) and v_l is increasing in p because it follows from $\gamma c > y > 0$ that the second term in Equation (1) is negative. To see that v_l is decreasing in τ, note that $\lambda_2 < \lambda_1$, and hence

$$\frac{d}{d\tau}v_l(p,\tau) = -(r+\gamma)(1-p)(\lambda_1 - \lambda_2)e^{-(r+\gamma)\tau}(\gamma c - y) < 0$$

for all $p \in (0,1)$ and $\tau \geq 0$. That v_l is convex in τ for all $p \in (0,1)$ follows from

$$\frac{d^2}{d^2\tau}v_l(p,\tau) = (r+\gamma)^2(1-p)(\lambda_1 - \lambda_2)e^{-(r+\gamma)\tau}(\gamma c - y) > 0.$$

Finally, supermodularity holds because

$$\frac{d^2}{dpd\tau}v_l(p,\tau) = (r+\gamma)(\lambda_1 - \lambda_2)e^{-(r+\gamma)\tau}(\gamma c - y) > 0.$$

(ii) Linearity of v_f in p is obvious from its definition in (2) and it is increasing in p because $\gamma c > y > 0$ implies that the first term in Equation (2) is positive and the second term is negative. For fixed $p \in (0,1)$, the derivative of v_f with respect to τ is

$$\frac{d}{d\tau}v_f(p,\tau) = -e^{-r\tau}\left[rpy - (r+\gamma)e^{-\gamma\tau}(1-p)\lambda_2(c\gamma - y)\right].$$

Let $\hat{\tau}(p)$ be the (finite) solution to the first order condition $dv_f(p,\tau)/d\tau = 0$. The second term in brackets is positive, so that $dv_f(p,\tau)/d\tau > 0$ if $\tau < \hat{\tau}(p)$ and $dv_f(p,\tau)/d\tau < 0$ if $\tau > \hat{\tau}(p)$. Hence, v_f attains a global maximum at $\hat{\tau}(p)$. If $\hat{\tau}(p) \geq 0$, then solving

$$rpy + (r+\gamma)e^{-\gamma\hat{\tau}^*(p_0)}(1-p)(\lambda_2(y - c\gamma)) = 0$$

for $\hat{\tau}(p)$ shows that $\hat{\tau}(p) = \tau^*(p)$. If $\hat{\tau}(p) < 0$, then $v_f(p,\cdot)$ is strictly decreasing on $[0,\infty)$, and therefore assumes its maximum at 0.

(iii) If $p < p_f^*$ then $\tau^*(p) > 0$. Since v_l is strictly decreasing in τ and because $\tau^*(p)$ maximizes v_f, we have $v_f(p,\tau^*(p)) > v_f(p,0) = v_l(p,0) > v_l(p,\tau^*(p))$. If $p \geq p_f^*$, then $\tau^*(p) = 0$, and therefore $v_f(p,\tau^*(p)) = v_f(p,0) = v_l(p,0) = v_l(p,\tau^*(p))$. $\square$

Proof of Lemma 2. From the definition of v_f, it follows that $v_f(p,0) > 0$ if and only if $p > \underline{p}$. Inspection reveals that $p_f^* > \underline{p}$ which implies $v_f(p_f^*,0) > 0$. Moreover, $v_l(p,0) = v_f(p,0)$ for all p, and therefore $v_l^*(p_f^*) = v_l(p_f^*,0) > 0$. Because v_l is continuous and strictly increasing, it follows from $v_l^*(p_l^*) = 0$ that $p_l^* < p_f^*$. Moreover, we have $0 = v_l^*(p_l^*) < v_l(p_l^*,0) = v_f(p_l^*,0)$. By Lemma 1, v_f is linearly increasing in its first argument. Since $v_f(\underline{p},0) = 0$, it follows from $v_f(p_l^*,0) > 0$ that

$\underline{p} < p_l^*$. □

Proof of Theorem 1. (i) Let $\check{p}_0 \geq p_f^*$. By definition of p_f^*, the follower's unique best response is to invest as soon as the other player has invested. It remains to be shown that there exists no equilibrium in which both players either never invest or invest (simultaneously) at some time $t > 0$. If a player deviates to investing at time 0, she would get a payoff of $w^* = \check{p}_0 y + (1 - \check{p}_0) \lambda_2 (y - \gamma c) > 0$, where the inequality holds because $\check{p}_0 \geq p_f^* > \frac{\lambda_2(\gamma c - y)}{y + \lambda_2(\gamma c - y)}$. In the conjectured equilibrium, by contrast, she obtains $e^{-rt} w^* < w^*$ ($t \in (0, \infty) \cup \{\infty\}$). Thus, she has an incentive to deviate and the only equilibrium is one in which both agents invest immediately.

(ii) Let $p_l^* < \check{p}_0 < p_f^*$. By definition of p_l^*, an agent would want to invest immediately if the other agent never invested. If the other agent invested, however, an agent would prefer not to invest, as this would give her $v_f^*(\check{p}_0)$ instead of $v_f(\check{p}_0, 0) < v_f^*(\check{p}_0)$, where the inequality holds because of $\check{p}_0 < p_f^*$ by definition of p_f^*. As shown in the main text, there is a unique investment rate $\beta^*(\check{p}_0)$ that renders the other agent indifferent between investing and not investing, which establishes the claim.

(iii) For $\check{p}_0 \leq p_l^*$, the claim follows immediately from the definition of p_l^*.

□

Lemma 3. $v_f^*(p) - v_l^*(p)$ *is a weakly decreasing function.*

Proof. Define $D(p, \tau) = v_f(p, \tau) - v_l(p, \tau)$. If $p \geq p_f^*$, then $\tau^*(p) = 0$ and therefore $D(p, \tau^*(p)) = 0$. Suppose $p < p_f^*$. Taking the derivative gives

$$\frac{dD(p, \tau^*(p))}{dp} = \frac{\partial D(p, \tau^*(p))}{\partial p} + \frac{\partial D(p, \tau^*(p))}{\partial \tau} \frac{d\tau^*(p)}{dp}. \tag{6}$$

Lemma 1 implies that the leader value is decreasing in τ, i.e., $dv_l(p, \tau^*(p))/d\tau < 0$, the optimal delay is decreasing, $d\tau^*(p)/dp < 0$, and that $\tau^*(p)$ maximizes the follower value, i.e., $dv_f(p, \tau^*(p))/d\tau = 0$. Therefore, $D(p)$ is strictly decreasing on $(0, p_f^*)$ if $\partial D(p, \tau^*(p))/\partial p < 0$. Using the definition of v_f in Equation (2), the derivative of the follower value is

$$\frac{\partial v_f(p, \tau^*(p))}{\partial p} = e^{-r\tau^*(p)} y + e^{-(r+\gamma)\tau^*(p)} \lambda_2 (\gamma c - y).$$

From Equation (1), the derivative of the leader's value is

$$\frac{\partial v_l(p, \tau^*(p))}{\partial p} = y + (\lambda_1 + (\lambda_2 - \lambda_1) e^{-(r+\gamma)\tau^*(p)})(\gamma c - y).$$

The partial derivative of D with respect to p is

$$\frac{\partial D(p, \tau^*(p))}{\partial p} = -(1 - e^{-r\tau^*(p)}) y - \lambda_1 (1 - e^{-(r+\gamma)\tau^*(p)})(\gamma c - y) < 0.$$

13

$\square$

Proof of Theorem 2.

Part (i). Suppose $p_0 > 0$ and $\rho > 1/2$ such that $p_0(g, g) \geq p_f^*$. We proceed to verify that the following strategies and beliefs are part of an equilibrium. In the first phase, an agent with signal g invests immediately. An agent with signal b invests immediately with probability $\eta \in [0, 1]$. With probability $1 - \eta$ she invests at a random time drawn from an exponential distribution with parameter $\beta^*(p_0(b, b)) \vee 0$. In the second phase of the game, a follower with posterior belief p delays investment by $\tau^*(p)$, and the beliefs at t are updated via Bayes' rule whenever possible. A leader with signal g reverses her investment immediately at $t = 0$ if and only if her posterior belief about θ is lower than p_l^*. A leader with signal b reverses her investment immediately at $t = 0$ with some probability $\nu \in [0, 1)$. Either agent exits in the second phase after the occurrence of a failure. In either phase, each agent i's belief with signal s assigns probability $p_0(b, s)$ to state $\theta = H$ after an off-equilibrium deviation of agent $-i$.

We show that there exist μ, ν such that above strategies and associated beliefs characterize a perfect Bayesian equilibrium. Consider first the second phase, taken as given that each type g invests at $t = 0$ with probability 1, and each type b invests at $t = 0$ with probability η and waits with probability $1 - \eta$. As shown in Lemma 1, the function $\tau^*(p)$ is the optimal delay of the follower with posterior p, and thus given the leader stays in the game, a follower cannot gain from deviating in the second phase of the game. If the leader exits immediately in the second phase, the follower cannot influence that decision. In the second phase, suppose agent i with signal s invests at time $t = 0$ in the first phase, and the other agent $-i$ does not, so that at $t = 0$ in the second phase, agent i is the leader and $-i$ the follower. If both expect that the other agent follows the strategy described in the previous paragraph, then the posterior belief of agent i with signal s is $p_0(b, s)$, and the posterior belief of type s of agent $-i$ is, by Bayes' rule,

$$\tilde{p}(\eta, s) = \frac{p_0(s)(\rho + \eta(1 - \rho))}{p_0(s)(\rho + \eta(1 - \rho)) + (1 - p_0(s))(1 - (1 - \eta)\rho)}. \tag{7}$$

If $v_l(p_0(b, b), \tau^*(p(\eta, b))) > 0$, then the continuation payoffs for either type of agent i is positive (since $v_l(p_0, \tau^*(\tilde{p}(\eta\nu, b))) > v_l(p_0(b, b), \tau^*(\tilde{p}(\eta\nu, b))))$, and thus it is optimal for both types of agent i to stay in. If, on the other hand, $v_l(p_0(b, b), \tau^*(\tilde{p}(\eta, b))) < 0 < v_l(p_0(b, b), \tau^*(p_0))$, then type b of agent i remains in the game with probability $\nu^* \in (0, 1)$ solving $v_l(p_0(b, b), \tau^*(\tilde{p}(\eta\nu^*, b))) = 0$.[7] If $v_l(p_0(b, b), \tau^*(p_0)) < 0$, then type b of agent i exits for sure, and type g remains if $v_l(p_0, \tau^*(p_0)) = v_l^*(p_0) \geq 0$. These actions are optimal for each agent by construction.

It remains to be shown that there exists a value for η such that neither agent can gain by deviating from the specified strategies in the first phase. Denote by $V_\theta(\eta)$ the value of investing

---

[7]Note that $\tilde{p}(\eta\nu^*, b)$ is the posterior belief of type b of agent $-i$ that the state is H, conditional on the joint event that agent i invested in the first phase and remains in the second phase, where $\eta\nu^*$ is the probability that type b of agent i does this.

at $t = 0$ for an agent in state θ, and let $W_\theta(\eta, \tau)$ be the value of waiting at $t = 0$ when the agent delays investment by τ as follower. Note here that τ refers to the delay of the investment, given that the other agent invests immediately at $t = 0$. When neither agent invests, then each agent is convinced that the other agents' type is bad, and thus there is no longer any uncertainty about others' private information. Note here that when the state and strategies are given, the payoff is independent of private signals. Note also that for a follower in the second phase, the optimal delay for each type b is $\tau^*(\tilde{p}(\eta, b))$ when type b of the other agent invests with probability η. We write

$$W_\theta^*(\eta) := \max_\tau W_\theta(\eta, \tau) = W_\theta(\eta, \tau^*(\tilde{p}(\eta, b))).$$

for the payoff from waiting when using the optimal delay.

There is a pooling equilibrium if and only if $p_0(b) \geq p_f^*$, since in this case each type of each agent is willing to invest immediately, if the other agent invests for sure, and thus reveals no information. Thus, consider the case $p_0(b) < p_f^*$. Note that in this case, we have $\mathbb{E}[W_\theta^*(1)|b] \geq \mathbb{E}[V_\theta(1)|b]$, since bad types always have weak incentives to wait when the other agent invests with probability one. There are two cases to consider, $\mathbb{E}[W_\theta^*(0)|b] < \mathbb{E}[V_\theta(0)|b]$ and $\mathbb{E}[W_b^*(0)|b] \geq \mathbb{E}[V_\theta(0)|b]$.

1.) Suppose $\mathbb{E}[W_\theta^*(0)|b] < \mathbb{E}[V_\theta(0)|b]$, i.e., an agent with a bad signal prefers to invest immediately if the other agent invests with zero probability after a bad signal. gIn this case, there exists a partial (or full) pooling equilibrium in which type g always invests while type b invests with probability $\eta^* \in [0, 1]$. Note that the functions $\mathbb{E}[W_\theta^*(\eta)|b]$ and $\mathbb{E}[V_\theta(\eta)|b]$ are convex combinations of continuous functions and hence continuous. Thus, there exists an $\eta^* \in (0, 1]$ such that $\mathbb{E}[V_\theta(\eta^*) - W_\theta^*(\eta^*)|b]$, so that an agent with type b is indifferent between investing and not investing, given the other agent invests with probability η^* after observing signal b. We shall now verify an agent of type g's incentives to invest. Note that we have the following inequality:

$$
\begin{aligned}
0 = \mathbb{E}[V_\theta(\eta^*) - W_\theta^*(\eta^*)|b] &\leq \mathbb{E}[V_\theta(\eta^*) - W_\theta(\eta^*, \tau^*(\tilde{p}(\eta^*, g)))|b] \\
&= p_0(b)\big(V_H(\eta^*) - W_H(\eta^*, \tau^*(\tilde{p}(\eta^*, g)))\big) + (1 - p_0(b))\big(V_L(\eta^*) - W_L(\eta^*, \tau^*(\tilde{p}(\eta^*, g)))\big) \\
&\leq p_0(g)\big(V_H(\eta^*) - W_H(\eta^*, \tau^*(\tilde{p}(\eta^*, g)))\big) + (1 - p_0(g))\big(V_L(\eta^*) - W_L(\eta^*, \tau^*(\tilde{p}(\eta^*, g)))\big) \\
&= \mathbb{E}[V_\theta(\eta^*) - W_\theta^*(\eta^*)|g],
\end{aligned}
$$

where the first inequality follows from the fact that $\mathbb{E}[W_\theta^*(\eta^*)|b] \geq \mathbb{E}[W_\theta(\eta^*, \tau)|b]$ for all $\tau \geq 0$ by definition, and the second inequality from $p_0(g) > p_0(b)$, and from the fact that investing immediately is strictly better than waiting if and only if the state is H (since there is no gain from delay in state H, and no gain from investing in state L).

2.) Now, suppose $\mathbb{E}[W_b^*(0)|b] \geq \mathbb{E}[V_\theta(0)|b]$, so that agents with signal b prefer to wait if the other

agent invests only if his signal is g. For agents with signal g in this case we have

$$\mathbb{E}[V_\theta(0) - W_\theta(0)|g] = q_0(g)\big(v_l(p_0(g,g),0) - v_f^*(p_0(g,g))\big)$$
$$+ (1 - q_0(g))\big(0 \vee v_l^*(p_0) - 0 \vee v_l(p_0, \tau^*(p_0(b,b)))\big).$$

Since $p_0(g,g) \geq p_f^*$ by assumption, we have $v_l(p_0(g,g),0) - v_f^*(p_0(g,g)) = 0$, and thus $\mathbb{E}[V_\theta(0) - W_\theta(0)|g] \geq 0$. Thus, in this case, we have a fully separating equilibrium in which g-types invest at $t = 0$, whereas b-types do not.

Part (2.). Suppose $p_0 > 0$ and $\rho \in (1/2, 1)$ are such that $p_0(g,g) < p_f^*$. Consider again the second phase first. The strategies outlines in the theorem constitute a separating equilibrium in the first phase, so that an agent who invests in the first phase and becomes the leader reveals to be of type g. Suppose that after type g of agent i invests at time t in the first phase, the agents use the following continuation strategy:

1. If $v_l^*(p_{t+\Delta}) \geq 0$, where $\Delta = \tau^*(p_0(g,g))$, and

$$p_{t+\Delta} = \frac{p_0}{p_0 + (1 - p_0)e^{-\gamma\Delta}},$$

   then type g of agent i remains in the game indefinitely, and each type s of agent $-i$ invests with delay $\tau^*(p_0(s,g))$

2. If $v_l^*(p_{t+\Delta}) < 0 < v_l(p_{t+\Delta},0)$, then type g of agent i remains indefinitely if agent $-i$ invests with delay $\tau^*(p_0(g,g))$ and exits otherwise. Each type of agent $-i$ invests with delay Δ.[8]

3. If $v_l(p_{t+\Delta},0) < 0$, then type g of agent i remains in until $t + \Delta$ if agent $-i$ invests with delay Δ, and agent i exits otherwise. Type g of agent $-i$ invests with delay Δ, and type b of agent $-i$ invests with delay $\tau^*(p_0)$.

Henceforth, for each of the three cases above, denote by $V_\theta(s_i, s_{-i})$ the value of agent i conditional on (i) state θ, (ii) agent i following the strategy for agents with signal of becoming the leader conditional on the other agent being of type s_{-i}, and let $W_\theta(s_i, s_{-i})$ be the value of becoming the follower. Note that these payoffs are independent of time on the first phase, since the signal pair (s_i, s_{-i}) encapsulates all information that is exchanged in the first phase, so that $p_t(s_i, s_{-i}) = p_0(s_i, s_{-i})$.

---

[8]This case uses as threat that the leader cannot reinvest after dropping out. There also is a slightly more involved equilibrium that does not require this restriction. In this equilibrium, the agents enter another war of attrition, in which the follower delays investment in such a way that at each point in time, the leader is just willing to stay in, up until the optimal delay of the follower is reached. This would make the follower better off, and result in a welfare loss.

Thus, each agent's expected value of becoming the leader is given by

$$U(q_t(g)) = q_t(g)\mathbb{E}[V_\theta(g,g)|s_i = g, s_{-i} = g] + (1 - q_t(g))\mathbb{E}[V_\theta(g,b))|s_i = g, s_{-i} = b].$$

Type g of each agent is willing to randomize if she is indifferent between investing immediately and waiting for another instant. Hence, the value function for type g of the agent must satisfy the indifference condition

$$U(q_t(g)) = \mu_t q_t(g)\mathbb{E}[W_\theta(g,g)|s_i = g, s_{-i} = g]dt + (1 - rdt - \mu_t q_t(g)dt)U(q_{t+dt}(g)). \qquad (8)$$

By Ito's Lemma, the indifference condition (8) can be written as

$$U(q_{t+dt}(g)) = U(q_t(g)) + dU(q_t(g))dq_t(g), \qquad (9)$$

where by definition of U, we have $dU(q_t(g))/dq_t(g) = \mathbb{E}[V_\theta(g,g)] - \mathbb{E}[V_\theta(g,b)]$. Bayes' rule implies that the posterior belief at $t + dt$ is

$$q_{t+dt}(g) = \frac{q_t(g)(1 - \mu_t dt)}{1 - q_t(g)\mu_t dt}.$$

The differential change in belief is therefore

$$\frac{dq_t(g)}{dt} \equiv \lim_{dt \to 0} \frac{q_{t+dt}(g) - q_t(g)}{dt} = -\mu_t q_t(g)(1 - q_t(g)). \qquad (10)$$

If we now substitute equations (9) and (10) in the indifference condition (8) and ignore higher order terms, we obtain the expression

$$rU(q_t(g)) = \mu_t q_t(g)\mathbb{E}[W_\theta(g,g) - V_\theta(g,g)|s_i = g, s_{-i} = g]. \qquad (11)$$

Since $p_0(g,g) < p_f^*$, the right-hand side of this equation is strictly positive. Simplifying and solving the equation for μ_t yields

$$\mu_t = \frac{rU(q_t(g))}{q_t(g)\mathbb{E}[W_\theta(g,g) - V_\theta(g,g)|s_i = g, s_{-i} = g]}. \qquad (12)$$

Here, μ_t is the rate of investment for type g of each agent in the symmetric equilibrium at a given belief $q_t(g)$. Substituting this last expression into Equation (10), we obtain the evolution of the posterior $q_t(g)$ in equilibrium:

$$dq_t(g) = -(1 - q_t(g))\frac{rU(q_t(g))}{\mathbb{E}[W_\theta(g,g) - V_\theta(g,g)|s_i = g, s_{-i} = g]}dt. \qquad (13)$$

We obtain the equilibrium belief and equilibrium investment rate at each time t by solving Equation (13) with given initial belief q_0. The initial value problem (13) has the unique solution

$$q_t(g) = \frac{e^{-t\hat{\beta}}U(q_0(g)) + (1 - q_0(g))\mathbb{E}[V_\theta(g, b)]}{e^{-t\hat{\beta}}U(q_0(g)) - (1 - q_0(g))\mathbb{E}[V_\theta(g, g) - V_\theta(g, b)|s_i = g]}. \tag{14}$$

We now substitute $q_t(g)^*$ into Equation (12) and simplify to obtain the equilibrium rate of investment

$$\mu_t^* = \frac{e^{-t\hat{\beta}}U(q_0(g))}{e^{-t\hat{\beta}}U(q_0(g)) + (1 - q_0)\mathbb{E}[V_\theta(g, b)|s_i = g]}\hat{\beta}. \tag{15}$$

If $r\mathbb{E}[V_\theta(g, b)] > 0$ then the investment rate μ_t^* diverges to $+\infty$ as $t \to t^*$, where

$$t^* = \log\left(1 + \frac{p_0(g, g)}{p_0}\frac{\mathbb{E}[V_\theta(g, g)]}{\mathbb{E}[V_\theta(g, b)]}\right)^{\beta^*(p_0(g,g))}. \tag{16}$$

If, on the other hand, $\mathbb{E}[V_\theta(g, b)] < 0$, then Equation (15) converges to 0 as $t \to \infty$. Thus, $t^* = \infty$.

It remains to be shown that it is optimal for type b of each agent to delay investment until t^*. For agents with signal b, the incremental opportunity cost from waiting is $\mathbb{E}[V_\theta(b, s)|s_i = b]dt$. The expected incremental gain from waiting for this type is $\mu_t^* q_t(b)\mathbb{E}[W_\theta(b, g) - V_\theta(b, g)|s_i = b]dt$. We show that when agents with signal g invest at rate μ_t^*, then agents with signal b prefer to wait:

$$r\mathbb{E}[V_\theta(b, s)|s_i = b] \leq \mu_t^* q_t(b)\mathbb{E}[W_\theta(b, g) - V_\theta(b, g)|s_i = b, s_{-i} = g]. \tag{17}$$

Because flow values are positive in state H and negative in state L, i.e., $y \geq 0 \geq y - \gamma c$, we have $V_H(s_i, s_{-i}) \geq 0 \geq V_L(s_i, s_{-i})$. Therefore:

$$\begin{aligned}
\mathbb{E}_t[V_\theta(s_i, s_{-i})|s_i = b] &= p_t(b)\mathbb{E}_t[V_H(s_i, s_{-i})|s_i = b] + (1 - p_t(b))\mathbb{E}_t[V_L(s_i, s_{-i})|s_i = b] \\
&\leq p_t(g)\mathbb{E}_t[V_H(s_i, s_{-i})|s_i = b] + (1 - p_t(g))\mathbb{E}_t[V_L(s_i, s_{-i})|s_i = b] \\
&\leq p_t(g)\mathbb{E}_t[V_H(b, s_{-i})|s_i = g] + (1 - p_t(g))\mathbb{E}_t[V_L(b, s_{-i})|s_i = g] \\
&\leq p_t(g)\mathbb{E}_t[V_H(g, s_{-i})|s_i = g] + (1 - p_t(g))\mathbb{E}_t[V_L(g, s_{-i})|s_i = g] \\
&= \mathbb{E}[V_\theta(s_i, s_{-i})|s_i = g].
\end{aligned}$$

The first inequality follows because $p_t(g) > p_t(b)$. Note that according to our prescribed strategies, agents with signal g invest earlier than agents with signal b, and exit later. Therefore, we have $V_\theta(s_i, g) \geq V_\theta(s_i, b)$. Since $q_t(g) > q_t(b)$, we thus have $\mathbb{E}[V_\theta(b, s_{-i})|s_i = g] \geq \mathbb{E}[V_\theta(b, s_{-i})|s_i = b]$ for each θ, which explains the second inequality. The last inequality follows because the strategy of

type g is constructed to maximize the continuation payoff after investing. Note that

$$\frac{1}{q_t(s_i)}\mathbb{E}[V_\theta(s_i, s_{-i})|s_i] = \frac{p_t(s_i)}{q_t(s_i)}\mathbb{E}[V_H(s_i, s_{-i})|s_i] + \frac{1 - p_t(s_i)}{q_t(s_i)}\mathbb{E}[V_L(s_i, s_{-i})|s_i]$$

Because $\rho > 1/2$ and $p_t(g) > p_t(b)$, it follows that

$$\frac{p_t(b)}{q_t(b)} = \frac{p_t(b)}{\rho p_t(b) + (1 - \rho)(1 - p_t(b))} < \frac{p_t(g)}{\rho p_t(g) + (1 - \rho)(1 - p_t(g))} = \frac{p_t(g)}{q_t(g)}$$

and similarly,

$$\frac{1 - p_t(b)}{q_t(b)} > \frac{1 - p_t(g)}{q_t(g)}$$

Combining the previous results, we find that

$$\begin{aligned}
\frac{1}{q_t(b)}\mathbb{E}[V_\theta(s_i, s_{-i})|s_i = b] &= \frac{p_t(b)}{q_t(b)}\mathbb{E}[V_H(s_i, s_{-i})|s_i = b] + \frac{1 - p_t(b)}{q_t(b)}\mathbb{E}[V_L(s_i, s_{-i})|s_i = b] \\
&\leq \frac{p_t(g)}{q_t(g)}\mathbb{E}[V_H(s_i, s_{-i})|s_i = b] + \frac{1 - p_t(g)}{q_t(g)}\mathbb{E}[V_L(s_i, s_{-i})|s_i = b] \\
&\leq \frac{p_t(g)}{q_t(g)}\mathbb{E}[V_H(s_i, s_{-i})|s_i = g] + \frac{1 - p_t(g)}{q_t(g)}\mathbb{E}[V_L(s_i, s_{-i})|s_i = g] \\
&= \frac{1}{q_t(g)}\mathbb{E}[V_\theta(s_i, s_{-i})|s_i = g]
\end{aligned}$$

The previous inequalities, together with (11), imply

$$\frac{1}{q_t(b)}r\mathbb{E}[V_\theta(s_i, s_{-i})|s_i = b] \leq \frac{1}{q_t(g)}r\mathbb{E}[V_\theta(s_i, s_{-i})|s_i = g] \tag{18}$$

$$= \mu_t^* \mathbb{E}[W_\theta(g, g) - V_\theta(g, g)|s_i = g, s_{-i} = g]. \tag{19}$$

On the other hand, again by $y > 0 > y - \gamma c$, waiting has negative value in state H, and positive value in state L, so that

$$W_H(s_i, g) - V_H(s_i, g) \leq 0 \leq W_L(s_i, g) - V_L(s_i, g). \tag{20}$$

The continuation payoff $\mathbb{E}[W_\theta(b, g)|s_i = b]$ for agent i with signal b in the event that agent $-i$ invests must maximize i's payoff, and therefore,

$$\mathbb{E}[W_\theta(g, g)|s_i = b] \leq \mathbb{E}[W_\theta(b, g)|s_i = b]. \tag{21}$$

Since the continuation strategies of a leader are constructed to be optimal in the continuation game,

we have

$$\mathbb{E}[V_\theta(s_i, s_{-i})|s_i = b, s_{-i} = g] \geq \mathbb{E}[V_\theta(g, s_{-i})|s_i = b, s_{-i} = g]. \tag{22}$$

It then follows that

$$\frac{1}{q_t(b)} r\mathbb{E}[V_\theta(s_i, s_{-i})|s_i = b] < \mu_t^* \mathbb{E}[W_\theta(s_i, s_{-i}) - V_\theta(s_i, s_{-i})|s_i = g, s_{-i} = g]$$

$$\leq \mu_t^* \mathbb{E}[W_\theta(g, s_{-i}) - V_\theta(g, s_{-i})|s_i = b, s_{-i} = g]$$

$$\leq \mu_t^* \mathbb{E}[W_\theta(s_i, s_{-i}) - V_\theta(s_i, s_{-i})|s_i = b, s_{-i} = g]$$

The first inequality follows from (18). The second inequality follows from (20). The last inequality follows from (21) and (22). The chain of inequalities implies (17) shich was to be shown.

Part (3.). As before, the function $\tau^*(p)$ is the optimal delay of the follower with posterior p, and thus given the leader stays in the game, a follower cannot gain from deviating in the second phase of the game. Assuming that the follower always uses delay $\tau^*(p)$ at posterior p, the value of becoming the leader for an agent with signal g is $\mathbb{E}[v_l^*(p_0(b, s))|g]$. If this value is negative, then no agent is willing to become a leader, so no investment takes places at any $t \geq 0$. $\qquad\square$

Proof of Theorem 3. The team's objective is to maximize

$$w(p, \tau) = \max\{0, v_l(p, \tau) + v_f(p, \tau)\},$$

where τ denotes the delay with which the follower invests. Vsing Equations (1) and (2), we can write explicitly:

$$w(p, \tau) = (1 + e^{-r\tau})py + (1 - p)\left[\lambda_1 + (2\lambda_2 - \lambda_1)e^{-(r+\gamma)\tau}\right](y - \gamma c).$$

It follows from the definitions of λ_1 and λ_2 that $2\lambda_2 - \lambda_1 = \lambda_1\lambda_2$. Therefore, the marginal value of delaying the second investment is

$$\frac{\partial w(p, \tau)}{\partial \tau} = -re^{-r\tau}[py + (1 - p)\lambda_2(y - \gamma c)e^{-\gamma\tau}]. \tag{23}$$

The expression in brackets is strictly increasing in τ. This implies that $\partial w(p, \tau)/\partial \tau < 0$ for all $\tau \geq 0$ whenever $p > \underline{p}$, where

$$\underline{p} := \frac{\lambda_2(\gamma c - y)}{y + \lambda_2(\gamma c - y)}, \tag{24}$$

in which case it is socially optimal to make the second investment immediately. If $p \leq \underline{p}$, then the socially optimal delay solves the first-order condition $dw(p, \tau)/d\tau = 0$. Thus, the optimal delay is

given by

$$\tau^s(p) = \begin{cases} [\phi(\underline{p}) - \phi(p)]/\gamma & \text{if } p < \underline{p}, \\ 0 & \text{if } p \geq \underline{p}. \end{cases}$$

It remains to be shown that, if $p \leq \underline{p}$, $w(p, \tau^s(p)) = 0$. Indeed,

$$w(p, \tau^s(p)) = \max\left\{ py + (1-p)\lambda_1(y - \gamma c) + py\left(1 - \frac{\lambda_1}{1-p}\right)\left(\frac{p(1-\underline{p})}{\underline{p}(1-p)}\right)^{\frac{r}{\gamma}}, 0 \right\}.$$

Since $p \leq \underline{p}$, it is sufficient to show that $py + (1-p)\lambda_1(y - \gamma c) + py\left(1 - \lambda_1/(1-p)\right) \leq 0$ and $py + (1-p)\lambda_1(y - \gamma c) \leq 0$. As both inequalities are satisfied for $p \leq \underline{p}$, this concludes the proof. □

Proof of Theorem 4. Let P be the distribution over signals for given parameters p_0 and ρ. We write $P(s)$ for the probability that a given agent's signal is s and $P(s_1, s_2)$ for the probability that the pair of signals is (s_1, s_2).

(1.) Suppose $p_0 > p_f^*$. Then, choose $\rho^* > 1/2$ such that $p_0(b) > p_f^*$ and $p_0(b, b) > \underline{p}$. By Theorem 2, it follows that for all $\rho < \rho^*$, there exists a pooling equilibrium in which each type of each agent invests immediately. By Theorem 3, this equilibrium is efficient. Hence $\tilde{W} \geq E[W_\theta(s_1, s_2)]$.

(2.) Let $p_f^* > p_0 > p_l^*$. Choose $\rho^* > 1/2$ such that $p_f^* > p_0(g, g)$ and $p_0(b, b) \geq \underline{p}$. By Theorem 2, there exists an equilibrium with delayed entry. It follows from arguments in the proof of Theorem 2 that the expected equilibrium value of the good type of each agent is $\mathbb{E}[V_\theta(g, s_{-i})] \geq \mathbb{E}[v_l^*|g]$. The inequality follows from the fact that leaders have the option to exit. By Equation (17), bad types strictly prefer to delay investment at each $t < t^*$. The expected payoff for an agent of type b who deviates by investing before time t^* is given by

$$q_t(b)v_l(p_0, \tau^*(p_0(g, g))) + (1 - q_t(b))v_l^*(p_0(b, b), \tau^*(p_0)) > \mathbb{E}[v_l^*|b].$$

Therefore, the expected social surplus for each agent is

$$\tilde{W} > P(g)\mathbb{E}[v_l^*|g] + P(b)\mathbb{E}[v_l^*|b] = \mathbb{E}[W_\theta(s^1, s^2)].$$

(3.) Let $p_0 = p_f^*$. Since $p_0(g, g) > p_f^*$, each agent with signal g invests immediately. For ρ close to $1/2$, agents with type g never exit in equilibrium. We show that there exists $\rho^* > 1/2$ such that $\tilde{W} > E[W_\theta(s_1, s_2)]$ for all $\rho \in (1/2, \rho^*)$. The equilibrium is with immediate investment, since $p_0(g) > p_f^*$. There cannot be a pooling eqilibrium, since $p_0(b) < p_0 = p_f^*$. If it was fully separating, then an agent with signal b would strictly prefer to invest immediately to induce the other agent to invest. Hence it must be partially separating. By indifference for type b,

we can write the social welfare per agent as

$$\tilde{W} = P(g)\Big[q_0(g)v_l(p_0(g,g),0) + (1 - q_0(b))(\eta v_l(p_0,0) + (1 - q_0(g))(1 - \eta)v_l(p_0,\tau^*(\tilde{p}_\eta(b)))\Big]$$

$$+ P(b)\Big[q_0(b)v_l(p_0,0) + (1 - q_0(b))(\eta v_l(p_0(b,b),0) + (1 - \eta)(v_l(p_0(b,b),\tau^*(p_\eta(b))))\Big].$$

Note that $P(g)q_0(g) = P(g,g)$, $P(g)(1 - q_0(g)) = P(b)q_0(b) = P(b,g)$ and $P(b)(1 - q_0(b)) = P(b,b)$. Using $q_0(b)p_0 + (1 - q_0(b))p_0(b,b) = p_0(b)$ together with the linearity of $v_l(p,0)$, we can write

$$\tilde{W} = P(g,g)v_l(p_0(g,g),0) + P(b,g)v_l(p_0,0)$$

$$+ (P(g,b) + P(b,b))\Big[\eta v_l(p_0(b),0) + (1 - \eta)v_l(p_0(b),\tau^*(\tilde{p}(\eta,b)))\Big].$$

When signals are public, then, after each realized pair of signals resulting in the posterior belief $\check{p}_0$, each agent's equilibrium payoff is $v_l^*(\check{p}_0)$. Thus, the expected welfare under public information can be written as

$$E[W_\theta(s_1, s_2)] = P(g,g)v_l^*(p_0(g,g)) + P(b,g)v_l^*(p_0)$$

$$+ (P(g,b) + P(b,b))\Big[q_0(b)v_l^*(p_0) + (1 - q_0(b))v_l^*(p_0(b,b))\Big]. \quad (25)$$

Using the definition of τ^* in [Lemma 1](#), we have

$$v_l^*(p) = v_l(p,\tau^*(p)) = py + (1 - p)\lambda_1(y - \gamma c) + (1 - p)e^{-(r+\gamma)\tau^*(p)}(\lambda_2 - \lambda_1)(y - \gamma c).$$

Since $p_0 = p_f^*$, we have $\tau^*(p_0) = \tau^*(p_0(g,g)) = 0$. We have that $\tilde{W} > E[W_\theta(s_1, s_2)]$ iff

$$\eta v_l(p_0(b),0) + (1 - \eta)v_l(p_0(b),\tau^*(\tilde{p}(\eta,b))) > q_0(b)v_l(p_0,0) + (1 - q_0(b))v_l^*(p_0(b,b))). \quad (26)$$

We define $\psi(p) := \frac{p}{1-p} = e^{\phi(p)}$ and $\alpha := \frac{r+\gamma}{\gamma} > 1$. Then, the left-hand side of Inequality [(26)](#) can be written as

$$\eta v_l(p_0(b),0) + (1 - \eta)v_l(p_0(b),\tau^*(\tilde{p}(\eta,b)) = p_0(b)y + (1 - p_0(b))\lambda_1(y - \gamma c)$$

$$+ (1 - p_0(b))\Big[\eta + (1 - \eta)\psi(\tilde{p}(\eta,b))^\alpha\psi(p_f^*)^{-\alpha}\Big](\lambda_2 - \lambda_1)(y - \gamma c) \quad (27)$$

From Bayes' rule and the definition of ψ, we have

$$\psi(p_0(b)) = \frac{p_0}{1 - p_0}\frac{1 - \rho}{\rho} = \psi(p_0)/\psi(\rho), \quad \psi(\tilde{p}(\eta,b)) = \psi(p_0(b))\left(\frac{\rho + (1 - \rho)\eta}{1 - \rho + \rho\eta}\right).$$

If we now use the previous equalities to factor out $\psi(p_0(b))^\alpha \psi(p_f^*)^{-\alpha}$ from the square brackets in (27), we obtain

$$\eta v_l(p_0(b), 0) + (1 - \eta)v_l(p_0(b), \tau^*(\tilde{p}(\eta, b))) = p_0(b)y + (1 - p_0(b))\lambda_1(y - \gamma c)$$

$$+ (1 - p_0(b))\psi(p_0(b))^\alpha \psi(p_f^*)^{-\alpha}\left[\eta\psi(\rho)^\alpha + (1 - \eta)\left(\frac{\rho + (1 - \rho)\eta}{1 - \rho + \rho\eta}\right)^\alpha\right](\lambda_2 - \lambda_1)(y - \gamma c). \quad (28)$$

The right-hand side of Inequality (26) is given by

$$q_0(b)v_l(p_0, 0) + (1 - q_0(b))v_l^*(p_0(b, b)) = p_0(b)y + (1 - p_0(b))\lambda_1(y - \gamma c)$$

$$+ \left[q_0(b)(1 - p_0) + (1 - q_0(b))(1 - p_0(b, b))\psi(p_0(b, b))^\alpha \psi(p_f^*)^{-\alpha}\right](\lambda_2 - \lambda_1)(y - \gamma c). \quad (29)$$

From Bayes' rule it follows that

$$\frac{q_0(b)}{1 - p_0(b)} = \frac{p_0(b)\rho + (1 - p_0(b))(1 - \rho)}{1 - p_0(b)} = \left(\frac{p_0}{1 - p_0}\frac{1 - \rho}{\rho}\right)\rho + (1 - \rho) = \frac{1 - \rho}{1 - p_0},$$

$$\frac{1 - q_0(b)}{1 - p_0(b)} = \frac{p_0(b)(1 - \rho) + (1 - p_0(b))\rho}{1 - p_0(b)} = \left(\frac{p_0}{1 - p_0}\frac{1 - \rho}{\rho}\right)(1 - \rho) + \rho = \frac{\rho}{1 - p_0(b, b)}.$$

Using these equalities together with the identity

$$\psi(p_0(b, b)) = \frac{p_0(b)}{1 - p_0(b)}\frac{1 - \rho}{\rho} = \psi(p_0(b)\psi(1 - \rho)$$

to factor out $(1 - p_0(b))\psi(p_0(b))^\alpha \psi(p_f^*)^{-\alpha}$ from the square brackets in (29), we obtain

$$q_0(b)v_l(p_0, 0) + (1 - q_0(b))v_l^*(p_0(b, b)) = p_0(b)y + (1 - p_0(b))\lambda_1(y - \gamma c)$$

$$+ (1 - p_0(b))\psi(p_0(b))^\alpha \psi(p_f^*)^{-\alpha}\left[(1 - \rho)\psi(\rho)^\alpha + \rho\psi(\rho)^{-\alpha}\right](\lambda_2 - \lambda_1)(y - \gamma c).$$

Define the functions,

$$h(\eta, \rho) := \eta\psi(\rho)^\alpha + (1 - \eta)\left(\frac{\rho + (1 - \rho)\eta}{1 - \rho + \rho\eta}\right)^\alpha, \quad g(\rho) := (1 - \rho)\psi(\rho)^\alpha + \rho\psi(\rho)^{-\alpha}.$$

Condition (26) is thus equivalent to $\inf_\eta h(\eta, \rho) > g(\rho)$. One calculates that the partial derivative of h at $\rho = 1/2$ is $\lim_{\rho \to 1/2} \partial_\rho h(\eta, \rho) = 4\alpha\left(2\eta^2 - \eta + 1\right)/(\eta + 1)$. The function $\lim_{\rho \to 1/2} \partial_\rho h(\eta, \rho)$ has its minimum in η at $\sqrt{2} - 1$ and is thus larger than $4\left(4\sqrt{2} - 5\right) > 0$. On the other hand $g'(1/2) = 0$. Thus, there exists a $\rho^* > 1/2$ such that for all $\rho \in (1/2, \rho^*)$, we have $\tilde{W} > E[W_\theta(s_1, s_2)]$.

□

References

Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.

Bonatti, A. and Hörner, J. (2011). Collaborating. *American Economic Review*, 101(2):632–63.

Chamley, C. and Gale, D. (1994). Information Revelation and Strategic Delay in a Model of Investment. *Econometrica*, 62(5):1065–85.

Décamps, J.-P. and Mariotti, T. (2004). Investment timing and learning externalities. *Journal of Economic Theory*, 118(1):80–102.

Heidhues, P., Rady, S., and Strack, P. (2015). Strategic experimentation with private payoffs. *Journal of Economic Theory*, 159:531–551.

Keller, G. and Rady, S. (2010). Strategic experimentation with Poisson bandits. *Theoretical Economics*, 5(2):275–311.

Keller, G. and Rady, S. (2015). Breakdowns. *Theoretical Economics*, 10(1):175–202.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic Experimentation with Exponential Bandits. *Econometrica*, 73(1):39–68.

Morris, S. and Shin, H. H. S. (2002). Social Value of Public Information. *The American Economic Review*, 92(5):1521–1534.

Moscarini, G. and Squintani, F. (2010). Competitive experimentation with private information: The survivor's curse. *Journal of Economic Theory*, 145(2):639–660.

Murto, P. and Välimäki, J. (2011). Learning and Information Aggregation in an Exit Game. *The Review of Economic Studies*, 78(4):1426–1461.

Murto, P. and Välimäki, J. (2013). Delay and information aggregation in stopping games with private information. *Journal of Economic Theory*, 148(6):2404–2435.

Rosenberg, D., Salomon, A., and Vieille, N. (2013). On games of strategic experimentation. *Games and Economic Behavior*, 82:31–51.

Rosenberg, D., Solan, E., and Vieille, N. (2007). Social Learning in One-Arm Bandit Problems. *Econometrica*, 75(6):1591–1611.

Teoh, S. (1997). Information disclosure and voluntary contributions to public goods. *RAND Journal of Economics*, 28(3):385–406.