

Reputation Concerns in Risky Experimentation*

CHIA-HUI CHEN
Kyoto University

JUNICHIRO ISHIDA
Osaka University

WING SUEN
University of Hong Kong

July 31, 2020

Abstract. High-ability agents are more likely to achieve early success in risky experimentation, but learn faster that their project is not promising. These counteracting effects give rise to a signaling model with double-crossing property. This property tends to induce homogenization of quitting times between types, which in turn leads to some pooling in equilibrium. Low-ability agents may hold out to continue their project for the prospect of pooling with the high type, despite having a negative instantaneous net payoff. A war-of-attrition mechanism causes low-ability agents to quit only gradually over time, and to stop quitting for a period immediately before all agents exit.

Keywords. double-crossing property; dynamic signaling; venture startups

JEL Classification. D82; D83; O31

*Part of this research was conducted when Suen visited the Institute of Social and Economic Research at Osaka University; he likes to thank the Institute for its hospitality. The first author acknowledges financial support from JSPS KAKENHI 17H13699. The second author acknowledges financial support from JSPS KAKENHI JP15H05728, JP18H03640, and JP19H01468 as well as the program of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

1. Introduction

Reputation concerns are important in many facets of our life, and they are even more so in the context of exploratory activities undertaken by experts, such as entrepreneurs, politicians, and scientists, whose reputation is an indispensable asset for advancing their careers. Reputation is established by both successes and failures, but in ways that are not as straightforward as they seem. Consider, for example, the decision to abandon a risky venture. Does a better entrepreneur persist longer with his project, because he is better able to implement a good idea? Or does he quit a risky venture earlier, because he is quicker to recognize its futility? Even success is not always an unblemished boost to one's reputation—success that arrives very late may be taken as a sign of mediocrity.

This paper aims to develop a unified framework to analyze the role of reputation concerns in risky experimentation. The setup is a standard bandit problem, where an agent engages in a project of unknown quality while at the same time attempts to signal his type to the market. If the project is good, the hazard rate of success (also called a “break-through” in this paper) is greater for the high type agent than for the low type. If the project is bad, neither type ever succeeds with the project. We identify two effects that are relevant for signaling in this context. The *ability effect* suggests that the high type has more incentive to persist longer with a project, because he is better at successfully implementing a good project. But a counteracting *learning effect* suggests the opposite: the high type becomes pessimistic about the project quality more quickly if success does not arrive, and therefore has an incentive to abandon it earlier.

The key to our analysis is not which effect outweighs the other. Rather, we show that the ability effect dominates in early stages of experimentation (when differences in beliefs of the two types are still small), while the learning effect dominates in later stages (when differences in conditional hazard rates are less relevant as the project is quite likely to be bad). As a result, the standard single-crossing property in signaling models does not hold. The high type initially benefits more from persisting with the project than does the low type, but eventually benefits less from persisting as he becomes increasingly pessimistic relative to the low type, generating a *double-crossing property* in signaling incentives. In this setting, therefore, perseverance is a sign of strength in early stages of experimentation but a sign of weakness in later stages.

1.1. Applications

Our model can be applied to a broad range of circumstances in which an agent embarks on an activity of exploratory nature to acquire new knowledge in the presence of reputation concerns. In each of the cases we discuss below, there is a tradeoff between quitting early and late, which stems from the aforementioned double-crossing property; as a consequence, it is hardly clear in which direction reputation concerns would influence the agent's quitting decisions. As we will argue in detail later, the insights we learn from standard models, which are deeply rooted in the single-crossing property, cannot be extended straightforwardly to these environments.

Venture startups. A leading example of our framework is venture startups. First, any business startup necessarily involves a process of experimentation in that success is uncertain and its timing is stochastic. In addition, there is ample evidence to suggest that reputation matters in entrepreneurship. Many recent studies report that serial entrepreneurs outperform first-time entrepreneurs (Chen, 2013; Rocha et al., 2015; Lafontaine and Shaw, 2016). Gompers et al. (2010) find that entrepreneurs with a track record of success are more likely to succeed than are first-time entrepreneurs or those who have failed previously. They also find that entrepreneurs exhibit performance persistence, and argue that entrepreneurship is a skill rather than luck. These findings point to the importance of the “success breeds success” mechanism, where the perception of performance persistence induces real performance persistence, suggesting that reputation is a crucial input for developing a new business.¹ In this type of environment, the entrepreneur needs to protect his reputation and hence has an incentive to strategically choose when to retreat from the business when it does not appear to take off anytime soon. At a glance, the problem faced by the entrepreneur looks deceptively similar to that of standard reputation models: he chooses an observable action, which is his quitting time, so as to look good in the eyes of the market. A closer look at it reveals, however, that there is no obvious way for the entrepreneur to establish (or salvage) his reputation, begging the type of questions we noted at the beginning.

Academia. Research and experimentation are almost synonym to each other, as researchers must take an unexplored path to discover new knowledge. There is little doubt that reputation is a crucial asset for researchers as their career opportunities, such as inter-

¹Gompers et al. (2010) argue that this is the case if suppliers and customers are more likely to commit resources to firms perceived to be more likely to succeed.

nal or external promotion, grants, and potential employment in industry, depend heavily on the perception of their expected future productivity. Consider an (early-career) economic theorist who tries to prove a conjecture which may or may not be true. A more capable theorist can prove it more quickly if the conjecture is indeed true, or else it is highly likely that the conjecture is false. With no sign of success, should the theorist, in need of signaling his analytical prowess to coauthors, colleagues and the entire academic community, quit earlier or persist longer with the project?

Political economy. Politicians often need to take actions whose benefits, if any, take time to materialize. Electoral accountability obviously feeds reputation concerns for office-seeking politicians when they conduct policy experimentation. For a wide range of issues for which politicians are held accountable, therefore, both experimentation and reputation formation play a crucial role. Consider, for instance, a politician who chooses to implement a new policy instead of following his predecessor's old policy. What if, after a while, things do not head in the right direction? The politician will face a difficult strategic decision of whether to persist with the current policy or to switch back to the old status quo. If he abandons the experimentation, what kind of message does it send to his constituents?

1.2. Main results

To the best of our knowledge, Matthews and Moore (1987) are the first to consider the double-crossing property in a multidimensional screening model.² We exploit this property to obtain a complete characterization of D1 equilibria in our setup. Equilibrium is unique, but it differs from the least-cost separating equilibrium. This result suggests that the way reputation concerns work in the context of risky experimentation is qualitatively different from that in other economic contexts characterized by single-crossing preferences.

A general lesson from our analysis is that the double-crossing property tends to induce homogenization of quitting times between types, compared to the "standard setup" characterized by the single-crossing property. This in turn leads to some pooling in equilibrium, making inferences based on quitting times less precise. Consider a situation in which separation requires the high type to quit the project inefficiently late (assuming the ability effect dominates in the relevant parameter region). Suppose reputation concerns become stronger in the sense that the difference in exit payoffs for the two perceived types becomes larger. Then separation would require the high type to quit later, but beyond a point the

²Daley and Green (2014) show that the double-crossing property may emerge in an education signaling model with grades as an additional source of information.

learning effect would give the high type more incentive to deviate by stopping earlier, making separation impossible under the D1 refinement. In effect the double-crossing property imposes an endogenous constraint on how late the high type can quit. But if the high type cannot quit too late, it becomes cheaper for the low type to mimic the high type, and pooling or semi-pooling is the result. In our model, the low type may “hold out” for the prospect of pooling with the high type, despite the fact that his instantaneous expected payoff from continuing with the project is strictly worse than the flow payoff from exiting, which in turn forces the high type to conform to avoid adverse inference. This is in contrast to the “standard setup,” where the low type never chooses a signaling action higher than the full-information optimal level and the high type can always go far enough to separate. It is also worth noting that the pooling equilibrium so determined is sensitive to the distribution of types, which allows us to discuss the role of prior reputation in risky experimentation.

Another general lesson is that signaling in risky experimentation is inherently a dynamic process, in which signaling incentives evolve over time. We distinguish between two cases in this paper, and discuss them separately in Sections 3 and 4. In the case of *exit signaling*, the exit payoff to an agent who abandons a project depends on his reputation while the reward to success does not. For example, a research scientist may be paid a fixed bonus for making a discovery but has to find alternative employment when he abandons his project. His exit payoff obviously will depend on the market’s perception of his ability. In the case of *breakthrough signaling*, reputation also matters for how large the reward to a breakthrough is. For a startup entrepreneur, the initial discovery or development of a prototype product from risky experimentation is only a first step toward ultimate success. How likely the initial discovery will turn into a profitable business is affected by whether the entrepreneur can attract financial capital and talents to work with him, which obviously will depend on his reputation. Exit signaling can be analyzed as if it were a static problem, in that the equilibrium quitting times are deterministic. On the other hand, breakthrough signaling is inherently dynamic, in that the equilibrium quitting times can be random, with separation of types occurring gradually through time.

To see this, suppose there is efficient separation in which each type stops at his full-information optimal timing (with the low type quitting first) if the reward to success does not depend on reputation. Introduce breakthrough signaling by making the reward to success much larger if the market perceives the agent to be a high type. Efficient separation can no longer be maintained in equilibrium, because the reward to success would

rise discontinuously after the low type is supposed to have quit, making it profitable for the low type to deviate by staying on longer. As low-type agents abandon the project, the reputation for those who remain and succeed becomes higher, therefore inducing these low-type agents not to abandon so soon. This war-of-attrition feature of the model implies that equilibrium will entail continuous randomization by low-type agents, which further exacerbates their incentive to hold out as it slows down the dynamic separation of types. If this effect is strong enough, the low type stops randomizing to “hold out for exit” as time draws close to the equilibrium quitting time of the high type. In this case equilibrium will exhibit a continuous randomization phase which gradually winnows out low-type agents, followed by a hold-out phase in which the remaining low-type agents get negative instantaneous payoffs, and finally leading to a mass exit by both types for the outside option.

After characterizing equilibria for the two cases, Sections 5.1 and 5.2 study comparative statics with respect to the magnitude of the reputation concern and to the prior reputation of the agent. Section 5.3 extends the model to allow for type-dependent project quality, and considers the implications of different types of ability (the ability to identify a good project versus the ability to successfully implement a good project) for organizational design. Section 5.4 extends the baseline model to the case with a continuum of types to show robustness of our findings. Finally, in Section 6, we provide a brief illustration of the implications of the analysis for venture startups—our leading example of risky experimentation—by offering potential ways to cope with the distortions caused by reputation concerns.

1.3. Related literature

Our model falls broadly into the growing literature on strategic experimentation where multiple parties are involved in the experimentation process. Since the seminal work of Keller et al. (2005), much of the literature builds on exponential bandits with constant hazard rates because of their simplicity and tractability. Our analysis suggests that in the signaling context this assumption can be substantially relaxed while still admitting a clear characterization of equilibria, thereby offering predictions that are robust to a range of specifications.³

Some recent works introduce reputation concerns into experimentation models as we

³Our model assumes that the conditional hazard rate is a general non-increasing function of time. See also Boyarchenko (2017) for a model of strategic experimentation in which the unconditional hazard is a humped-shaped function of time.

do. In Halac and Kremer (2020) agents privately observe good or bad news about their projects and want to maintain reputation for project quality. In a different vein, Bonatti and Hörner (2017) analyze a career concerns model with exponential learning which incorporates a moral hazard component. Bobtcheff and Levy (2017) and Thomas (2019) consider an exponential bandit model in which a decision maker chooses when to stop the experimentation, with the reputation payoff determined at the time of termination as in ours. An important distinction is that agents in those models are heterogeneous only in one particular dimension: they differ only in the ability to identify good projects (Thomas, 2019) or in the speed of learning (Bobtcheff and Levy, 2017). Our model encompasses and integrates these two distinct notions of ability and clarifies the differences they make.⁴

When agents differ in the speed of learning, the single-crossing property may break down due to the learning effect.⁵ In this sense, Bobtcheff and Levy (2017) share an important commonality with our model. In their model, however, the agent privately observes a conclusive signal which reveals that the project is bad, with the timing of investment serving as the sole signaling device. In contrast the timing of success in our model provides an inconclusive public signal (on top of the timing of project abandonment), so that the agent is subject to dynamic reputation concerns which further distort the quitting decision.

We can also relate our work to dynamic signaling models which incorporate additional sources of (noisy) information (Bar-Isaac, 2003; Daley and Green, 2012; Gul and Pesendorfer, 2012; Lee and Liu, 2013). Among them, our model is more closely related to Daley and Green (2012), who consider an environment where there are two types of seller, either high or low, and each type decides when to trade. Aside from technical differences (their information arrives via a diffusion process while ours arrives via a jump process), the key difference is the element of experimentation. In their model, each seller knows his own type (asset value), so that high-type sellers are always more confident to receive good news and hence more willing to wait than low-type sellers.⁶ In our model, the project quality is initially not known to anyone and needs to be uncovered via experimentation. While high-ability agents are more likely to achieve early success, they also learn more quickly that their project is not promising. Therefore, the marginal benefit of experimentation is

⁴Prendergast and Stole (1996) is one of the earliest works to explore this point, although in a very different context.

⁵Halac et al. (2016) also make this observation under the exponential bandit framework.

⁶In Daley and Green (2012), there is a lower belief threshold at which low-type sellers choose to separate from high-type sellers by accepting a low price. Meanwhile high-type sellers, who are more confident that the belief will recover in future, always wait until the belief becomes sufficiently favorable.

initially higher for high-ability agents, but this relationship flips as the game progresses, which generates the double-crossing property.

Finally, our model predicts pooling equilibria under a wide range of parameters, which can be interpreted as a form of conformity or herding. There is now a very diverse literature which explores this possibility in various ways. For instance, Bernheim (1994) considers a situation where agents care about status as well as intrinsic utility. He shows that agents with moderate preferences converge to a homogeneous standard of behavior to avoid an inference that they have undesirable extreme preferences. The logic of our model is different because agents attempt to signal competence, not preferences, and conform in order to avoid an inference that they are incompetent. In this sense, our model has a closer connection to Scharfstein and Stein (1990), who show that an agent mimics the behavior of predecessors in the presence of reputation concerns about his competence. Their focus is on informational conformity where agents take actions sequentially, while ours rests on stigma attached to off-path deviations.

2. The Model

2.1. Setup

An agent undertakes a risky project with an uncertain project quality. If the project quality is good, it will generate success, or a “breakthrough,” at some random time τ , provided the agent has not abandoned the project by that time. If the project quality is bad, it will never generate a breakthrough no matter how much time the agent spends working on it. The flow cost of working on the risky project for a small time interval of length dt is $c dt$.

The agent can be of two types, depending on his ability in implementing the risky project. We capture the difference between types by specifying different distributions of the stochastic time of success τ , given that the project quality is good. Specifically, let $f_H(\tau)$ be the density function of τ for the high type and $f_L(\tau)$ be the density function for the low type, conditional on good project quality (and let F_H and F_L represent the corresponding cumulative distribution functions). We assume that the conditional hazard rate given good project quality, $f_i(\cdot)/(1 - F_i(\cdot))$, is non-increasing for $i \in \{H, L\}$ (high type or low type). Moreover, the distributions for the two types are ordered by monotone likelihood ratio property, in the sense that $f_H(\cdot)/f_L(\cdot)$ is strictly decreasing. For example, in the exponential bandit model, a higher Poisson arrival rate of success implies that the likelihood ratio, $f_H(\tau)/f_L(\tau) = (\lambda_H/\lambda_L)e^{-(\lambda_H-\lambda_L)\tau}$, decreases in τ for $\lambda_H > \lambda_L$. While we

focus on the two-type case for most of our analysis, in Section 5.4, we will consider the case with a continuum of types to show robustness of our findings.

The state of nature in our model is therefore two-dimensional, defined over the project quality (good or bad) and the ability type (high or low). Neither the agent nor anyone else knows the quality of the project. The common prior belief that the project quality is good is p_0 . We will later allow the prior to depend on agent type in Section 5.3; for now assume that it is type-independent. The agent knows his own ability type, but this private information is unavailable to the market. The market's prior belief that the agent is a high type is q_0 . Generally, we use q_t to represent the updated belief at time t that the agent is a high type. We sometimes refer to this belief as the agent's "reputation."

Time is continuous. At each point in time, the agent decides whether to continue working on the risky project or to abandon it. For $i \in \{H, L\}$, let the agent's strategy be represented by $\sigma_i : [0, \infty) \rightarrow [0, 1]$, with $\sigma_i(t)$ being the probability that the agent is still working on the risky project at time t conditional on no breakthrough having occurred. Abandoning the project is irreversible; i.e., $\sigma_i(t)$ is non-increasing in t . If the agent adopts a pure strategy, there is a unique stopping time s_i such that $\sigma_i(t) = 1$ for all $t < s_i$ and $\sigma_i(t) = 0$ thereafter. We sometimes abuse notation by saying that the agent's strategy is s_i . (The agent will stop before this time if he has already achieved success in the risky project.) The agent's decision at each point in time is publicly observable. The arrival of success is also publicly observable, representing an additional source of information.

2.2. Payoffs

The game ends either when the risky project generates a breakthrough, or when the agent abandons the project. We assume that there is a competitive market that pays an agent commensurate with his expected output whenever the game ends. Let W_i denote the productivity (in present value) of type i after he achieves a breakthrough and w_i the productivity after he abandons the project. Therefore, if the project generates a breakthrough at time τ , the agent's payoff is $W(\tau) = q_\tau W_H + (1 - q_\tau)W_L$. If the agent's reputation when he abandons the project at time t is q_t , his outside opportunity from quitting is $w(t) = q_t w_H + (1 - q_t)w_L$. As explained below, we take a reduced-form approach and regard $W(\tau)$ and $w(t)$ as representing the continuation payoffs of success and failure, respectively. Although we assume that the continuation payoffs depend only on the agent's reputation with no type-specific payoffs, adding a type-specific component does not change

the substance of our analysis.⁷

We make several assumptions regarding the productivities. First, we generally assume $w_H > w_L$, so that the agent's reputation always matters in case of failure. In case of success, we assume $W_H \geq W_L$ and consider two cases. If $W_H = W_L$, then all reward to success is non-contingent on the agent's reputation at the time of success. If $W_H > W_L$, part of the reward will depend on the agent's reputation at the time of success. As we will see, the difference between the two cases proves to be critical, and we analyze these cases separately. Finally, we also assume that achieving a breakthrough is always better than the outside opportunity of working in the labor market. Specifically,

$$W_L > w_H. \quad (1)$$

Assumption (1) implies that $W(t) > w(t)$ at any time t and for any market belief, ensuring that the net value of success is always positive.⁸

The agent discounts future payoffs at rate ρ . For $i \in \{H, L\}$, the expected payoff to an agent of type i if he plans to abandon the project at time s is

$$U_i = \int_0^s e^{-\rho\tau} p_0 f_i(\tau) [W(\tau) - C(\tau)] d\tau + e^{-\rho s} (1 - p_0 F_i(s)) [-C(s) + w(s)], \quad (2)$$

where $C(t) = c(e^{\rho t} - 1)/\rho$ is the accumulated cost of working with the risky project for a period of length t .

2.3. Interpretation of the model

To provide some suitable interpretation of our theoretical framework, consider the case of a venture startup. A breakthrough in this context refers to reaching a milestone over the course of the project, such as successfully developing a prototype product, establishing distribution channels, gaining market recognition, and so on. If the entrepreneur achieves a breakthrough, he may build on this success and take the next step, the expected payoff of which is summarized by $W(\tau)$. If he abandons the project, he will be forced to take an exit option, such as starting a new business all over again or switching to the labor market as a worker. In either case, the expected payoff of this contingency depends on his reputation and is summarized by $w(t)$.

⁷See the Online Appendix for the analysis of this case.

⁸The role of this assumption is purely technical, as the same economic intuition applies even if it does not hold.

As noted above, whether the reward to success is contingent on the agent’s reputation or not is crucial in our model, and we make a clear distinction between $W_H = W_L$ (which we label as the case of *exit signaling*) and $W_H > W_L$ (the case of *breakthrough signaling*). Exit signaling is primarily motivated by the difference in exit payoffs (w_H versus w_L) for agents who fail the experimentation, while the difference in reward to success (W_H versus W_L) also matters for the case of breakthrough signaling. The difference between the two cases reflects, among other things, where in the experimentation process the agent stands. For instance, in the case of a venture startup, successfully developing a prototype product is only a step toward a bigger goal. Reputation still plays an essential role as it facilitates access to complementary inputs and makes the entrepreneur even more likely to succeed in the future, the extent of which is captured by $W_H - W_L$. On the other hand, gaining global market recognition is valuable in and of itself. Once this level of success is achieved, reputation matters much less, and $W_H - W_L$ is regarded as relatively small.

2.4. Equilibrium selection

A pair of strategies (\hat{s}_H, \hat{s}_L) and the beliefs $\{q_t\}$ constitute an equilibrium if \hat{s}_i maximizes U_i given $\{q_t\}$ for $i \in \{H, L\}$ and if the beliefs $\{q_t\}$ are consistent with Bayes’ rule and the strategies whenever applicable. As is typical in signaling models, these requirements do not pin down a unique equilibrium, because we can construct an arbitrary equilibrium by manipulating off-equilibrium beliefs. Consider, for instance, a pooling equilibrium in which both type quits at some t_p . For a range of values of t_p , we can construct such an equilibrium by assigning off-equilibrium belief that an agent is a low type if he quits at any $t \neq t_p$. To obtain sharper and more reasonable predictions, therefore, we adopt the D1 criterion (Cho and Kreps, 1987) for equilibrium refinement. While there are many specific cases we need to examine, we first note a result that generally holds in our setup (the proof will be provided when we discuss each specific case).

Proposition 1. *There always exists a unique equilibrium which survives the D1 criterion.*

A crucial feature of our model, which distinguishes it from a standard signaling model, is that preferences do not satisfy the single-crossing property. Instead, indifference curves for the two types cross twice in the relevant space. We call this a *double-crossing property*. It turns out that the D1 criterion does not always select the least-cost separating equilibrium in this environment.⁹ Since the problem becomes substantially more complicated when

⁹It is interesting to note that D1 and the Intuitive Criterion offer slightly different predictions in our

$W_H > W_L$, we first consider the case of exit signaling with $W_H = W_L$ in Section 3 to focus only on the issue arising from the double-crossing property. Then, in Section 4, we let $W_H > W_L$ to deal with the inherently dynamic issues arising from the case of breakthrough signaling.

3. Exit Signaling

3.1. Reputation concerns

Throughout Section 3 we assume that $W_H = W_L = W > w_H$, so that the reward to success in the risky project is simply W . In this case, the agent's reputation matters only when he abandons the project and takes up his exit option, the timing of which can be controlled perfectly by him. This feature simplifies the problem to a great extent, and allows us to clarify the key mechanism of the double-crossing property.

The updated belief about an agent's type at the time of project abandonment depends on (i) inferences based on the agent's choice and its consistency with the equilibrium strategies of the two types; and (ii) observation about the timing of success τ . We use

$$\hat{q} = \Pr[\text{high type} \mid \sigma_L, \sigma_H, \text{ stops at } t]$$

to denote the *interim belief* based on equilibrium inference alone, and use

$$\begin{aligned} r(t; \hat{q}) &= \Pr[\text{high type} \mid \sigma_L, \sigma_H, \text{ stops at } t, \tau > t] \\ &= \frac{\hat{q}(1 - p_0 F_H(t))}{\hat{q}(1 - p_0 F_H(t)) + (1 - \hat{q})(1 - p_0 F_L(t))} \end{aligned}$$

to denote the belief based on both (i) and (ii). If the interim belief for an agent who stops at time t is \hat{q} , the final reputation $r(t; \hat{q})$ of this agent incorporates the information from the event that he stops without a breakthrough, i.e., $\tau > t$. Monotone likelihood ratio property implies that $F_H(t) \geq F_L(t)$, and hence $r(t; \hat{q}) \leq \hat{q}$. Failure to achieve a breakthrough is bad news for the agent's ability.

It is straightforward to verify that $r(t; \hat{q})$ decreases in t if and only if $g_H(t) > g_L(t)$, where

$$g_i(t) = \frac{p_0 f_i(t)}{1 - p_0 F_i(t)} = \left(\frac{f_i(t)}{1 - F_i(t)} \right) \left(\frac{p_0(1 - F_i(t))}{p_0(1 - F_i(t)) + 1 - p_0} \right) \quad (3)$$

model, even though our model has only two types. All equilibria we identify survive the Intuitive Criterion, but equilibrium is no longer unique under it. The Intuitive Criterion cannot rule out the least-cost separating equilibrium when there is a D1 pooling equilibrium.

is the *unconditional* hazard rate of success in the risky project for type $i \in \{H, L\}$. The first term is the conditional hazard function for type i given that the project quality is good. The second term is the posterior belief that the project quality is good given that an agent of type i fails to obtain success by time t . Monotone likelihood ratio property implies that the conditional hazard rate for the high type is always greater than that for the low type. We call this the *ability effect*. However, the same property also implies that the posterior belief about project quality upon failure to obtain success is smaller for the high type than that for the low type. We call this the *learning effect*, because more able agents learn more quickly that they are likely to be working on a bad project than do less able ones if success has not been already observed. Note that the learning effect disappears when $p_0 = 1$, pointing to the essential role of experimentation in our setup.

The following result is crucial for our subsequent analysis.

Lemma 1. *The hazard rate $g_i(\cdot)$ is strictly decreasing for $i \in \{H, L\}$. There exists a unique \hat{t} such that $g_H(t) > g_L(t)$ if and only if $t < \hat{t}$, and $g_H(t) < g_L(t)$ if and only if $t > \hat{t}$.*

Proof. The first term of equation (3) is non-increasing by assumption, and the second term is strictly decreasing. Hence $g_i(\cdot)$ is strictly decreasing. At $t = 0$, $g_H(0)/g_L(0) = f_H(0)/f_L(0) > 1$. As t approaches infinity, $g_H(t)/g_L(t)$ approaches $\lim_{t \rightarrow \infty} f_H(t)/f_L(t)$, which is less than 1. Therefore, there exists \hat{t} such that $g_H(\hat{t})/g_L(\hat{t}) - 1 = 0$. The derivative of $g_H(t)/g_L(t)$ at $t = \hat{t}$ has the same sign as the derivative of $f_H(t)/f_L(t)$ at $t = \hat{t}$, which is negative. This shows that $g_H(t)/g_L(t) - 1$ is single-crossing from above. ■

The posterior belief about project quality is decreasing in t . Thus, differences in ability to implement a good project becomes less relevant over time, and Lemma 1 shows that the learning effect dominates the ability effect beyond time \hat{t} . If a high-ability agent has been working on a project for a long time without achieving a breakthrough, he should have recognized that the project is not promising and quit. As noted in the introduction, in this environment, perseverance is a sign of strength in early stages of experimentation, but a sign of weakness in later stages.

3.2. Double-crossing property

The objective function (2) for type $i \in \{H, L\}$ in the case $W_H = W_L = W$ reduces to

$$U_i(s, \hat{q}) = \int_0^s e^{-\rho\tau} p_0 f_i(\tau) [W - C(\tau)] d\tau + e^{-\rho s} (1 - p_0 F_i(s)) [-C(s) + w_L + r(s; \hat{q})(w_H - w_L)]. \quad (4)$$

The marginal rate of substitution between stopping time s and interim belief \hat{q} , denoted $MRS_i(s, \hat{q})$, is given by

$$\frac{g_i(s)[W - w_L - r(s; \hat{q})(w_H - w_L)] - \rho(w_L + r(s; \hat{q})(w_H - w_L)) - c + (\partial r / \partial s)(w_H - w_L)}{(\partial r / \partial \hat{q})(w_H - w_L)}.$$

Observe that the marginal rate of substitution depends on agent type only through the hazard rate $g_i(s)$. By assumption (1), the term in square brackets is positive. Lemma 1 then implies that $MRS_H(s, \hat{q}) > MRS_L(s, \hat{q})$ if and only if $s < \hat{t}$, and $MRS_H(s, \hat{q}) < MRS_L(s, \hat{q})$ if and only if $s > \hat{t}$. Thus, indifference curves of the two types cross twice: before time \hat{t} the opportunity cost of quitting is larger for the high type than for the low type; the opposite is true beyond \hat{t} .

Because $g_H(t) - g_L(t)$ is single-crossing from above, the slope of the indifference curve changes more rapidly with respect to the quitting time s for the high type than for the low type. At \hat{t} , the slopes of the two indifference curves are equal; moreover the indifference curve of the high type is “more convex” than that of the low type. See Figure 1 for illustration. An implication of this greater convexity is that, for any interim belief \hat{q} , whenever the high type is indifferent between some outcome (s', \hat{q}') and the outcome (\hat{t}, \hat{q}) , the low type strictly prefers the former to the latter. Formally,

$$U_H(s', \hat{q}') = U_H(\hat{t}, \hat{q}) \implies U_L(s', \hat{q}') > U_L(\hat{t}, \hat{q}). \quad (5)$$

If there is an equilibrium in which both types quit at \hat{t} , the set of interim beliefs that would support deviation to $s' \neq \hat{t}$ by the high type is strictly contained in the corresponding set for the low type. By the D1 criterion, the market should assign off-equilibrium belief that an agent is a low type if he quits at s' .

3.3. Equilibrium with exit signaling

To set the stage for equilibrium analysis, we first consider the full-information outcome. Let s_i^* represent the stopping time chosen by type $i \in \{H, L\}$ if his type is known. The optimal




Figure 1. The indifference curves U_L and U_H are tangent to each other at \hat{t} , with U_H “more convex” than U_L at that point. The point $(s_H^*, 1)$ is the “bliss point” for the high type. In panel (a), indifference curves are single-crossing to the left of \hat{t} . The only equilibrium is a separating equilibrium: the low type quits at s_L^* and the high type quits at \bar{s} . In panel (b), the indifference curve U_L' crosses U_H twice—first from below to the left of \hat{t} , then from above to the right of \hat{t} . The only equilibrium is a pooling equilibrium with both types quitting at \hat{t} .

stopping rule can be obtained by value-matching and smooth-pasting (or, equivalently, by the first-order condition for maximizing the objective function (4) for the relevant type), which gives

$$g_i(s_i^*)(W - w_i) - (\rho w_i + c) = 0. \quad (6)$$

The first term on the left-hand-side is the expected capital gain from extending the risky project for a small interval of time. The second term is the opportunity cost of doing so. Note that s_L^* can be greater than or less than s_H^* .

In any signaling equilibrium, the low type cannot do worse than choosing s_L^* and revealing himself to be a low type, which gives him a utility of $U_L(s_L^*, 0)$. It is useful to define two thresholds, $\underline{s} < \bar{s}$, such that

$$U_L(\underline{s}, 1) = U_L(s_L^*, 0) = U_L(\bar{s}, 1).$$

The low type never wants to stop before \underline{s} or after \bar{s} , even if by doing so he could successfully mimic the high type.¹⁰ Therefore, if s_H^* and s_L^* are sufficiently far apart, then neither type

¹⁰It is possible that $U_L(0, 1) > U_L(s_L^*, 0)$, in which case \underline{s} is defined to be equal to 0. For ease of exposition, we assume that W is sufficiently large that \underline{s} is positive.

wants to deviate from his optimal stopping time, and the full-information outcome is the unique equilibrium outcome in the model. To focus on the interesting case, we assume in the following that the incentive compatibility constraint for the low type is binding, in the sense that $U_L(s_L^*, 0) < U_L(s_H^*, 1)$, so that the full-information outcome is not achievable. This is equivalent to $s_H^* \in (\underline{s}, \bar{s})$.

Consider first the case of separating equilibrium.¹¹

Proposition 2. *In the case of exit signaling, suppose the full-information outcome is not feasible.*

- (a) *If $\hat{t} \leq \underline{s}$, the equilibrium is separating, with the high type quitting at \underline{s} and the low type quitting at s_L^* .*
- (b) *If $\hat{t} \geq \bar{s}$, the equilibrium is separating, with the high type quitting at \bar{s} and the low type quitting at s_L^* .*

Proof. When $\hat{t} \leq \underline{s}$, we have $MRS_H(s, \hat{q}) \leq MRS_L(s, \hat{q})$ for all $s \in [\underline{s}, \bar{s}]$. Because the single-crossing property is satisfied, with the high type having a greater incentive to quit earlier than the low type does. It follows from a standard refinement argument (Cho and Kreps, 1987) that the least-cost separating equilibrium (corresponding to the stopping times (\underline{s}, s_L^*) for high type and low type, respectively) is the only equilibrium that satisfies the D1 criterion. When $\hat{t} \geq \bar{s}$, we have $MRS_H(s; \hat{q}) \geq MRS_L(s; \hat{q})$ for all $s \in [\underline{s}, \bar{s}]$. The single-crossing property is again satisfied, but with the high type having less incentive to quit earlier than the low type does. The least-cost separating equilibrium in this case is for the high type to quit at \bar{s} , and for the low type to quit at s_L^* , and this is the only equilibrium that satisfies D1. ■

The value of \hat{t} is determined entirely by the statistical properties of the hazard functions and is independent of the payoff parameters. If \hat{t} is very low, the learning effect dominates throughout the relevant region. Because the high type has more incentive to quit early in this case, the equilibrium outcome is determined by the condition that the high type quits prematurely at \underline{s} to just deter the low type from mimicking. If \hat{t} is very high, the ability effect dominates throughout the relevant region. Because the high type has more incentive to stay longer with the project, the equilibrium outcome is determined by the condition that the high type stays inefficiently long until \bar{s} to just deter the low type from mimicking.

¹¹Throughout the analysis, when we say separating equilibrium, it means full separation between the two types where the quitting time fully reveals the agent's type.

Although equilibrium is separating in both cases of Proposition 2, the direction of how the high type separates from the low type differs. In this regard, our result is different from the standard setup. Panel (a) of Figure 1 illustrate the least-cost separating equilibrium (with the high type quitting later) corresponding to case (b) of Proposition 2.

The values of \underline{s} and \bar{s} needed to deter the low type from mimicking the high type become more extreme as the signaling incentive (measured by $w_H - w_L$) becomes stronger, while the value of \hat{t} remains unchanged. When the signaling incentive is strong enough, we will have $\hat{t} \in (\underline{s}, \bar{s})$. In this case the ability effect and the learning effect dominate in different stages of experimentation, and our model predicts pooling (or semi-pooling) between the different types.

Proposition 3. *In the case of exit signaling, suppose the full-information outcome is not feasible, and suppose $\hat{t} \in (\underline{s}, \bar{s})$.*

- (a) *If $U_L(\hat{t}, q_0) \geq U_L(s_L^*, 0)$, the equilibrium is full pooling, with both types quitting at \hat{t} .*
- (b) *If $U_L(\hat{t}, q_0) < U_L(s_L^*, 0)$, the equilibrium is semi-pooling, with the high type quitting at \hat{t} and the low type randomizing between quitting at \hat{t} and s_L^* .*

Proof. We first show that there cannot be a separating equilibrium. Suppose otherwise, and let the high type quit at some time t in this equilibrium. If $t \in (\underline{s}, \bar{s})$, the low type could profitably deviate by stopping at t . If $t \leq \underline{s}$, the high type could profitably deviate by stopping later at $t + \epsilon$ for some small positive ϵ , because according to the D1 criterion the off-equilibrium belief associated with such a deviation is that it comes from a high type. Similarly, if $t \geq \bar{s}$, the high type could profitably deviate by stopping a bit earlier.

Next, if the two types pool (or partially pool) by both stopping at the same time t with positive probability, then we must have $t = \hat{t}$. Otherwise, by stopping a little later (if $t < \hat{t}$) or a little earlier (if $t > \hat{t}$), an agent could obtain a discrete improvement in the market's belief of his type from some $\hat{q} < 1$ to 1.

Finally, condition (5) shows that if there is a semi-pooling equilibrium in which the high type randomizes between quitting at \hat{t} and quitting at some other time t' , the low type strictly prefers to quit at t' . This contradicts our earlier conclusion that the two types cannot partially pool by both stopping at $t' \neq \hat{t}$. Hence, in equilibrium, the high type quits at \hat{t} with probability 1. Given that the high type quits only at \hat{t} , by D1 the market assigns interim belief $\hat{q} = 0$ to an agent who quits at $t \neq \hat{t}$. For such interim belief, if a semi-pooling equilibrium exists, the low type must quit at s_L^* . This leaves us with two possible

types of equilibrium: (i) full pooling in which both types quit at \hat{t} ; or (ii) semi-pooling in which the low type quits at both \hat{t} and s_L^* with positive probability, and the high type quits at \hat{t} .

If $U_L(\hat{t}, q_0) \geq U_L(s_L^*, 0)$, semi-pooling is not an equilibrium. In a semi-pooling equilibrium, because the high type quits with probability 1 while the low type quits with probability strictly less than 1 at \hat{t} , the interim belief $\hat{q}_{\hat{t}}$ based on equilibrium inference is strictly higher than the initial belief q_0 . Thus, $U_L(\hat{t}, q_0) \geq U_L(s_L^*, 0)$ implies $U_L(\hat{t}, \hat{q}_{\hat{t}}) > U_L(s_L^*, 0)$, meaning that the low type strictly prefers quitting at \hat{t} to quitting at s_L^* , a contradiction. Therefore, it is a unique equilibrium for both types to quit at \hat{t} , and the market assigns an interim belief q_0 upon observing an agent quitting at \hat{t} . Neither type could profitably deviate because quitting at another time would be interpreted as deviation by a low type.

If $U_L(\hat{t}, q_0) < U_L(s_L^*, 0)$, full pooling is not an equilibrium, because the low type could profitably deviate by quitting at s_L^* instead. Given that there exists a unique $q \in (q_0, 1)$ such that $U_L(\hat{t}, q) = U_L(s_L^*, 0)$, the semi-pooling equilibrium is unique, with the high type quitting at \hat{t} with probability 1, and the low type doing the same with some positive probability so that the interim belief about an agent who quits at \hat{t} is exactly q . The remaining low types quit at s_L^* . By construction, the low type is indifferent between quitting at \hat{t} and s_L^* . The high type strictly prefers quitting at \hat{t} to quitting at another time, because such deviation would be interpreted as made by a low type. ■

When indifference curves of the two types exhibit the double-crossing property, Proposition 3 shows that the D1 refinement does not yield the least-cost separating equilibrium as the unique equilibrium outcome. Instead, equilibrium entails pooling at \hat{t} (i.e., the point where the indifference curves of the two types are tangent to one another), supported by the belief that an agent who abandons the project at any time other than \hat{t} is a low type. See panel (b) of Figure 1 for an illustration of a full pooling equilibrium corresponding to case (a) of Proposition 3. It is interesting that the equilibrium time \hat{t} for both types to quit depends only on the distributions of the timing of success (i.e., on p_0 , F_L and F_H), but not on the costs and benefits of risky experimentation.

4. Breakthrough Signaling

Beginning from this section, we drop the assumption that $W_H = W_L$ and assume $W_H > W_L$ instead. Because the reward to success, $W(\tau)$, is a function of the reputation of the agent who is staying to work with the risky project at the time of success τ , this introduces a

dynamic element into the signaling model that is absent in the case discussed in Section 3: the agent’s reputation now matters even when he achieves a breakthrough, the timing of which arrives stochastically (as opposed to the timing of quitting, which is deterministic) and is beyond his control. Since the timing of achieving a breakthrough now has signaling value,¹² we refer to this case as “breakthrough signaling.”

Given the equilibrium strategies σ_L and σ_H of the two types, we use

$$\tilde{q} = \Pr[\text{high type} \mid \sigma_L, \sigma_H, \text{has not stopped by } t]$$

to represent the interim belief about an agent who has not abandoned the project by time t based on equilibrium inference alone, and let

$$\begin{aligned} R(t; \tilde{q}) &= \Pr[\text{high type} \mid \sigma_L, \sigma_H, \text{has not stopped by } t, \tau = t] \\ &= \frac{\tilde{q}f_H(t)}{\tilde{q}f_H(t) + (1 - \tilde{q})f_L(t)} \end{aligned}$$

represent the final belief that incorporates the information from the event that a breakthrough occurs at time t (i.e., $\tau = t$). Because $f_H(t)/f_L(t)$ is decreasing, $R(t; \tilde{q})$ is decreasing in t .

It is worth noting that the reputation upon success may be higher or lower than the reputation upon failure to obtain success. In particular, it is straightforward to show that, for any interim belief q , $R(t; q) \geq r(t; q)$ if and only if $t \leq \hat{t}$, and $R(t; q) \leq r(t; q)$ if and only if $t \geq \hat{t}$. In other words, success that comes too late may be worse for an agent’s reputation than no success. This is because late success is evidence for low-ability type, while no success can be partly attributable to bad project quality.

For the case of breakthrough signaling, we continue to have the following result. The logic of this result is the same as that for the case of exit signaling; and the proof is relegated to the Appendix.

Lemma 2. *If both types of agent abandon the risky project at some time t with positive probability in equilibrium, then $t = \hat{t}$.*

Lemma 2 implies that a full pooling equilibrium must have both types quit at \hat{t} . Nevertheless, the possibility of breakthrough signaling makes a non-trivial difference to our

¹²By this we mean that the timing of achieving a breakthrough reveals payoff-relevant information. It also reveals some information about the agent’s type in the case of exit signaling, but this information is not payoff-relevant because the reward to success is fixed and independent of the agent’s reputation.

analysis whenever equilibrium entails some separation between the two types. To see this, note that the instantaneous benefit from working with the risky project for a small time interval of length dt for an agent of type $i \in \{H, L\}$ is

$$g_i(t)[W_L + R(t; \tilde{q})(W_H - W_L)] dt,$$

which depends on the value of \tilde{q} when $W_H > W_L$. But the interim belief \tilde{q} about an agent who stays evolves over time as different types of agents quit at different times to separate from one another. If the high type quits before the low type, then \tilde{q} falls as the high type quits, and this can reduce the incentive of the remaining agents to continue working with the risky project. The opposite is true if the low type quits before the high type.

For $i \in \{H, L\}$, let $s_i^*(\tilde{q})$ represent the solution to the following equation:

$$g_i(s_i)[W_L + R(s_i; \tilde{q})(W_H - W_L) - w_i] - (\rho w_i + c) = 0. \quad (7)$$

Note that s_L^* defined in equation (6) is the same as $s_L^*(0)$; and s_H^* is the same as $s_H^*(1)$. When $W_H > W_L$, equation (7) gives the optimal stopping rule for an agent of type i , if by continuing the market belief about his type would be \tilde{q} and by quitting he would reveal his true type. Because the left-hand-side of (7) decreases in s_i and increases in \tilde{q} , $s_i^*(\tilde{q})$ is increasing in \tilde{q} . A higher interim reputation for stayers raises the reward to success and tends to delay quitting.

The analysis of breakthrough signaling depends crucially on which type quits first. It turns out that if $s_L^*(q_0) > \hat{t}$, the high type quits first. The incentive for the low type to stay falls when the high type quits because the former can no longer pool with the latter. Thus, the low type may quit *before* $s_L^*(q_0)$. If $s_L^*(q_0) < \hat{t}$, the low type quits first. But when the low type quits, the interim belief about stayers improves, which makes quitting by the low type self-defeating. Equilibrium in this case generally involves continuous randomization by the low type, who ends up quitting *after* $s_L^*(q_0)$. We discuss these two cases in turn.

4.1. High type quits first

In this subsection, we consider the case where $s_L^*(q_0) \geq \hat{t}$. This case obtains, for example, when the outside option for the low type w_L or the flow cost of working with the project c is small. It can be shown that $s_L^*(q_0) \geq \hat{t}$ implies $s_L^*(q_0) > s_H^*(q_0)$,¹³ so the high type tends to quit before the low type.

¹³From equation (7), $w_L < w_H$ implies $g_L(s_L^*(q_0)) < g_H(s_H^*(q_0))$. If $s_L^*(q_0) \geq \hat{t}$, this implies $g_H(s_L^*(q_0)) < g_H(s_H^*(q_0))$ and therefore $s_L^*(q_0) > s_H^*(q_0)$.

For $t \leq s_L^*(0)$, define the following function:

$$v(t; s_L^*(0)) = \int_t^{s_L^*(0)} e^{-\rho(\tau-t)} \frac{p_0 f_L(\tau)}{1 - p_0 F_L(\tau)} [W_L - C(\tau - t)] d\tau \\ + e^{-\rho(s_L^*(0)-t)} \frac{1 - p_0 F_L(s_L^*(0))}{1 - p_0 F_L(t)} [-C(s_L^*(0) - t) + w_L].$$

Conditional on the project not having achieved a breakthrough by time t , the function $v(t; s)$ gives the payoff (from the perspective of time t) from continuing with it until time s , given that the agent is known to be a low type. Recall that $s_L^*(0)$ is the optimal stopping time when an agent is known to be a low type, so $v(t; s) \leq v(t; s_L^*(0))$ for any $s \geq t$. In the following, we simply write $v(t)$ for $v(t; s_L^*(0))$ when there is no confusion. It is straightforward to verify that $v(t)$ strictly decreases in t for $t < s_L^*(0)$, with $v'(s_L^*(0)) = 0$ and $v(s_L^*(0)) = w_L$.

Define \underline{t} such that $v(\underline{t}) = w_H$. (If no such \underline{t} exists, we let $\underline{t} = 0$.) When $s_H^*(q_0) \leq \underline{t}$, we have $v(s_H^*(q_0)) \geq w_H$. The incentive compatibility constraint for the low type is not binding, because the low type prefers staying with the risky project until time $s_L^*(0)$ to quitting at $s_H^*(q_0)$ to obtain an outside market wage of w_H . In this case, it is an equilibrium for the high type to quit at $s_H^*(q_0)$ and the low type to quit at $s_L^*(0)$.

The following proposition characterizes the equilibrium when the incentive compatibility constraint is binding. The proof is relegated to the Appendix.

Proposition 4. *In the case of breakthrough signaling, suppose $s_L^*(q_0) \geq \hat{t}$ and $s_H^*(q_0) > \underline{t}$.*

- (a) *If $\hat{t} \leq \underline{t}$, the equilibrium is separating. The high type quits at \underline{t} and the low type quits at $s_L^*(0)$.*
- (b) *If $\hat{t} \in (\underline{t}, s_L^*(0))$ and $v(\hat{t}) > w_L + r(\hat{t}; q_0)(w_H - w_L)$, the equilibrium is semi-pooling. The high type quits at \hat{t} and the low type randomizes between quitting at \hat{t} and $s_L^*(0)$.*
- (c) *If $\hat{t} \in (\underline{t}, s_L^*(0))$ and $v(\hat{t}) \leq w_L + r(\hat{t}; q_0)(w_H - w_L)$, or if $\hat{t} \in [s_L^*(0), s_L^*(q_0)]$, the equilibrium is full pooling. Both types quit at \hat{t} .*

The logic of Proposition 4 is very similar to that described in the case of exit signaling. Case (a) is a least-cost separating equilibrium in which the high type quits first, because the learning effect dominates in the relevant region when both $s_L(q_0)$ and $s_H^*(q_0)$ exceed \hat{t} . In cases (b) and (c), the double-crossing property induces pooling in equilibrium under the D1 criterion. In case (b), there is an interim belief $\hat{q} > q_0$ such that if a fraction of low

types quit at \hat{t} , then the remaining low types are indifferent between quitting at that time and continuing with the risky project until $s_L^*(0)$. So the equilibrium is semi-pooling. In case (c), the low type prefers to quit at \hat{t} and pool with the high type in the labor market than to continue with the project until $s_L^*(0)$. So there is a full pooling equilibrium.

4.2. Low type quits first

We now consider the case where $s_L^*(q_0) < \hat{t}$. This case obtains when w_L or c is large relative to the reward to success. Note also that $s_L^*(q_0) < \hat{t}$ implies $s_H^*(q_0) < \hat{t}$.¹⁴ Since both $s_L^*(q_0)$ and $s_H^*(q_0)$ are in the region where the ability effect dominates, the low type has an incentive to quit before the high type. Nevertheless we will show that any separation in this case must occur gradually over time.

There cannot be an equilibrium in which the high type separates from the low type by quitting before $s_L^*(q_0)$; otherwise the low type would profitably mimic the high type.¹⁵ Lemma 2 also establishes that if the two types pool, then it must occur at $\hat{t} > s_L^*(q_0)$. These observations imply that, by the time the game reaches time $s_L^*(q_0)$, the high type has not abandoned the project yet. At this time, the low type would prefer to stop if the reputation of stayers were fixed at q_0 . However, if the low type stops with positive probability, the interim belief \tilde{q} about those who stays at time $s_L^*(q_0)$ would jump up, which means that a low type could profitably deviate by staying a little bit longer instead of quitting at $s_L^*(q_0)$. The only way to eliminate this deviation incentive is to have the low type exit continuously at some atomless rate (i.e., $\dot{\sigma}_L(t) < 0$) when $t \geq s_L^*(q_0)$.

The low-type's payoff is pinned down by the outside option w_L when $\dot{\sigma}_L(t) < 0$. If $\tilde{q}(t)$ is the interim belief about an agent who is still staying at time t , the low type must be indifferent between staying and quitting whenever $\dot{\sigma}_L(t) < 0$. This condition can be written as

$$g_L(t)[W_L + R(t; \tilde{q}(t))(W_H - W_L) - w_L] - (\rho w_L + c) = 0. \quad (8)$$

The interim belief $\tilde{q}(t)$ satisfies

$$\tilde{q}(t) = \frac{q_0}{q_0 + (1 - q_0)\sigma_L(t)}. \quad (9)$$

As t increases from $s_L^*(q_0)$ to $s_L^*(1)$, $\tilde{q}(t)$ must rise continuously from q_0 to 1 to maintain the indifference condition (8), and $\sigma_L(t)$ must fall continuously from 1 to 0 according to equation (9).

¹⁴If $s_L^*(q_0) < \hat{t}$ and $s_H^*(q_0) \geq \hat{t}$, we would have $g_L(s_L^*(q_0)) > g_H(s_H^*(q_0))$, which contradicts equation (7).

¹⁵See Fact 1 in the proof of Proposition 5 for a detailed argument behind this claim.

For $t \in [s_L^*(q_0), \hat{t}]$, we define the following function:

$$V(t; \hat{t}) = \int_t^{\hat{t}} e^{-\rho(\tau-t)} \frac{p_0 f_L(\tau)}{1 - p_0 F_L(\tau)} [W_L + R(\tau; \tilde{q}(t))(W_H - W_L) - C(\tau - t)] d\tau \\ + e^{-\rho(\hat{t}-t)} \frac{1 - p_0 F_L(\hat{t})}{1 - p_0 F_L(t)} [-C(\hat{t} - t) + w_L + r(\hat{t}; \tilde{q}(t))(w_H - w_L)],$$

where $\tilde{q}(t)$ follows equation (8) for $t \leq s_L^*(1)$ and is equal to 1 for $t > s_L^*(1)$. The function $V(t; \hat{t})$ gives the payoff to a low type if he stays with the risky project between time t and \hat{t} , and then pools with the high type to quit at \hat{t} . Whether there will be some separation or not will depend on whether $V(s_L^*(q_0); \hat{t})$ is less than or greater than w_L . If $V(s_L^*(q_0); \hat{t}) \geq w_L$, the low type prefers to wait and pool with high types at time \hat{t} than to quit at $s_L^*(q_0)$. There will be full pooling. If $V(s_L^*(q_0); \hat{t}) < w_L$, some separation will occur gradually in equilibrium.

Proposition 5. *In the case of breakthrough signaling, suppose $s_L^*(q_0) < \hat{t}$.*

- (a) *If $V(s_L^*(q_0); \hat{t}) \geq w_L$, the equilibrium is full pooling. Both types quit at \hat{t} .*
- (b) *If $V(s_L^*(q_0); \hat{t}) < w_L$, there exists a unique $t_0 \in (s_L^*(q_0), \hat{t})$ such that $V(t_0; \hat{t}) = w_L$.*
 - (i) *If $t_0 < s_L^*(1)$, the equilibrium is semi-pooling. The high type quits at \hat{t} with probability 1. The low type's behavior has four phases: 1) stays (i.e., $\sigma_L(t) = 1$) for $t < s_L^*(q_0)$; 2) randomizes between quitting and staying (with $\sigma_L(t)$ determined by equations (8) and (9) and decreasing continuously) for $t \in [s_L^*(q_0), t_0]$; 3) stays (with $\sigma_L(t) = \sigma_L(t_0) > 0$) for $t \in (t_0, \hat{t})$; and 4) quits with probability 1 (i.e., $\sigma_L(t) = 0$) at $t = \hat{t}$.*
 - (ii) *If $t_0 \geq s_L^*(1)$, the equilibrium is separating. There exists a unique $t_1 \in (s_L^*(1), \hat{t})$ such that $V(s_L^*(1); t_1) = w_L$ and the high type quits at $\max\{t_1, s_H^*(1)\}$. The low type's behavior has three phases: 1) stays for $t < s_L^*(q_0)$; 2) randomizes between quitting and staying for $t \in [s_L^*(q_0), s_L^*(1)]$; and 3) quits with probability 1 at $t = s_L^*(1)$.*

The strategy of the low type described in Proposition 5 is qualitatively different from that described in the earlier propositions. In particular the low type always quits at or before the optimal time $s_L^*(q_0)$ in other cases, but Proposition 5 shows that it is possible for the low type to stay beyond $s_L^*(q_0)$. We leave the details of the proof of Proposition 5 to the Appendix. Intuitively, if $V(s_L^*(q_0); \hat{t}) < w_L$, the low type would prefer to quit than to stay at $t = s_L^*(q_0)$. Because the reward to success would rise as the low type quits, continuous

randomization is needed to sustain equilibrium. During the phase of randomization, the payoff to the low type is pinned down by the exit option w_L . As time t gets closer to \hat{t} , the prospect of pooling with the high type by quitting at \hat{t} becomes increasingly attractive; so the randomization phase switches to a “hold out” phase in which the low type stays with probability 1 for $t \in (t_0, \hat{t})$. Because $g(\cdot)$ and $R(\cdot; \tilde{q}(t_0))$ are strictly decreasing; the indifference condition (8) at time $t = t_0$ implies that, for $t \in (t_0, \hat{t})$,

$$g_L(t)[W_L + R(t; \tilde{q}(t_0))(W_H - W_L) - w_L] - c < \rho w_L.$$

In other words the instantaneous expected payoff from continuing with the project is strictly worse than the flow payoff from the exit option. Nevertheless the low type strictly prefers to stay because he is “holding out” in order to pool with the high type to obtain a better exit payoff at a later date.

If \hat{t} is large, the t_0 that satisfies $V(t_0; \hat{t}) = w_L$ will also be large. For $t_0 \geq s_L^*(1)$, the rate of exit determined by equations (8) and (9) implies that all low types would have quit by time $s_L^*(1)$. Moreover, since $V(s_L^*(1); \cdot)$ is decreasing on $[s_L^*(1), \hat{t}]$ with $V(s_L^*(1); t_1) = w_L$, the low type cannot profitably deviate to any $t' \geq t_1$. Because the high type can always benefit from moving towards $\max\{t_1, s_H^*(1)\}$, the only separating equilibrium that can survive the D1 criterion is the one in which the high type quits at $\max\{t_1, s_H^*(1)\}$.

4.3. Dynamic distortions

In our model, the dynamic separation of types is incomplete because we have a public information source which imperfectly reveals the agent’s type—the arrival of a breakthrough is a noisy signal that can come from either type. The presence of public news gives rise to dynamic reputation concerns which further distort the timing of project abandonment. Although the equilibrium characterization with breakthrough signaling is more complicated, the basic insight is relatively clear: the high type quits too early, and the low type quits too late (relative to the full-information benchmark).

If the high type quits before the low type, Proposition 4 states that the high type quits at $\max\{t, \hat{t}\}$, which is earlier than $s_H^*(q_0)$. Since $s_H^*(q_0) < s_H^*(1)$, the high type quits prematurely compared to the full-information benchmark. This is because the “reputational value of success” is necessarily lower when there are more low-type agents around, which reduces the continuation payoff of risky experimentation and forces the high type to abandon the project too early.

The timing of project abandonment is even more distorted when the low type quits

before the high type. In the case of exit signaling the low type quits once and for all at $s_L^*(0)$. In breakthrough signaling, the low type has an incentive to inefficiently wait until $s_L^*(q_0)$ because the “reputational value of success” is higher with more high-type agents around. Moreover, when a low-type agent quits, it raises the interim belief about the ability type of stayers, and hence the continuation payoff. The game thus resembles a war of attrition in that each low-type agent is waiting for others to drop out. Because of this, the low type must randomize over time, causing the separation of types to occur only gradually, and even later than $s_L^*(q_0)$.

It is worth emphasizing that as a consequence of these forces, there exists no efficient (full-information) equilibrium under breakthrough signaling, i.e., the one in which the low type quits at $s_L^*(0)$ and the high type at $s_H^*(1)$, no matter how far apart they are from each other. This is a stark difference from the case under exit signaling, where efficient separation is feasible as long as s_L^* and s_H^* are sufficiently far apart from each other.

5. Discussion

5.1. The role of reputation concerns

One of the most important predictions of our model is the presence of pooling equilibria in which high-type agents quit at the same time as (some of the) low-type agents. But our analysis also predicts *homogenization of quitting times* across types, in comparison to the standard setup. Below, we will illustrate this point and its efficiency implications.

The extent of reputation concerns in our model can be measured by $w_H - w_L$, where the agent possesses stronger reputation concerns when the productivity difference between the two types is larger. Here, we examine how equilibrium varies with changes in the extent of reputation concerns by increasing w_H for a given w_L (while holding $W_H = W_L$).¹⁶ There are two cases, depending on whether s_L^* is larger than \hat{t} or not. See Figure 2 for illustration. Note that in both cases, the space of w_H is divided into four regions.¹⁷

¹⁶The effect of an increase in W_H for a given W_L also has a similar effect and induces more pooling, but its effect is more complicated and less clear. In this section, therefore, we focus on the impact of $w_H - w_L$ while holding $W_H = W_L$ fixed.

¹⁷It is easy to see that: (1) s_L^* is independent of w_H while s_H^* decreases in w_H (higher exit payoff induces the high type to quit earlier); (2) \underline{s} decreases and \bar{s} increases in w_H ; (3) both s_H^* and \underline{s} decrease and tend to 0 as w_H increases. We can also show that \underline{s} crosses \hat{t} before s_H^* crosses it in panel (a), while s_H^* and \underline{s} never cross each other in panel (b).




Figure 2. Equilibrium as a function of w_H for a given w_L with $W_H = W_L$. Panel (a) depicts the case of $s_L^* > \hat{t}$. Panel (b) depicts the case of $s_L^* < \hat{t}$. In both cases, in the range where the incentive compatibility is binding ($s_H^* \in (\underline{s}, \bar{s})$), the equilibrium changes from separating to semi-pooling and to full pooling as w_H increases.

The key objects to look at in this figure are \underline{s} and \bar{s} , which indicate the quitting times needed to prevent the low type from mimicking the high type to obtain the exit payoff w_H instead of w_L . In the standard setup, these represent the expected points of separation: the high type either quits excessively early at \underline{s} or excessively late at \bar{s} to separate from the low type while the low type always quits at the full-information optimal point s_L^* . This is in fact what happens in equilibrium of our model when $w_H - w_L$ is relatively small and hence reputation concerns are relatively weak.

As $w_H - w_L$ becomes large, the gap between \underline{s} and \bar{s} widens to contain \hat{t} , at which point full separation is no longer feasible. The high type cannot quit too early or too late for the sake of separation, because any deviation towards \hat{t} is attributed to the high type under D1. There is hence a gravitating force towards \hat{t} which tends to homogenize or compress quitting times between types and consequently breaks down any separation that spans over \hat{t} . Since an increase in w_H makes pooling a more attractive option, the low type gradually switches from s_L^* to \hat{t} as w_H increases, and full pooling is realized when w_H becomes high enough.

This exercise raises an important observation on the role of reputation concerns in risky experimentation: the difference in quitting times is compressed and bounded from above. This homogenizing force yields both positive and negative efficiency effects. To see this, note that in the standard setup, the inefficiency in quitting time arises solely from

the high type as the low type always quits at the efficient time. For illustration, suppose $\hat{t} > s_L^*$ (the case of panel (b)). In this case, under single-crossing preferences, the high type quits later and later to separate from the low type as w_H increases. As a consequence, the equilibrium quitting time diverts away from s_H^* without bounds, which leads to a serious loss of efficiency. In the current setting which generates the double-crossing property, on the other hand, there is compression of quitting times towards \hat{t} , which constrains how late the high type can quit. Note, however, that this comes with a cost: knowing that the high type cannot persist for too long, the low type now has an incentive to inefficiently hold out to pool with the high type.

Although the above observation is more about immediate or short-run consequences of reputation concerns, there are also long-run consequences. In the context of risky experimentation, separation is harder to achieve due to the homogenizing force of the double-crossing property, and less information about agent type will be revealed as a result. Our exercise in particular suggests that we cannot learn much from observing that an agent abandons his risky project. This means that separation of types is less complete and take more time in environments where success is rare, as is often the case in venture startups.

5.2. The role of prior reputation

In the standard setup, the D1 criterion always selects the least-cost separating equilibrium, or the Riley outcome, in which the low type chooses his (full-information) optimal investment level and the high type invests just enough to separate from the low type. This prediction is somewhat disturbing because the equilibrium allocation in the least-cost separating equilibrium is independent of the prior belief, implying that the agent's prior reputation has no real consequences for signaling. In contrast, in our model, the prior belief q_0 plays a crucial role in shaping the equilibrium outcome.

There are basically two ways in which the prior belief affects the equilibrium allocation in our setup. First, an increase in q_0 directly raises the value of pooling for the low type, $U_L(\hat{t}, q_0)$, relative to the value of quitting at $s_L^*(0)$, and hence favors a pooling equilibrium. This in turn forces the high type to also quit at \hat{t} in order to avoid adverse inference. This effect is present even in the pure exit signaling case.

With breakthrough signaling, there arises an additional effect of the prior belief when a separating equilibrium prevails. This effect can work either positively or negatively, depending on which type quits first. In a separating equilibrium in which the high type

quits first, the high type quits at $s_H^*(q_0)$, which is earlier than the full-information optimal time $s_H^*(1)$. The extent of inefficiency thus diminishes as q_0 increases. In a separating equilibrium in which the low type quits first, on the other hand, the low type starts quitting at $s_L^*(q_0)$, which is later than $s_L^*(0)$. An increase in q_0 thus raises the extent of inefficiency.

To sum up, the equilibrium selected by D1 is sensitive to the agent's prior reputation. The effect of an increase in q_0 is hardly straightforward, and can often be negative as it provides a stronger incentive for the low type to stay and mimic the high type. This points to the difficulty in predicting the outcome of risky experimentation from publicly observable traits when reputation matters, because the timing of project abandonment could be related in some complicated and non-monotonic way to the experimenter's prior reputation. On top of the inherent randomness of risky experimentation, this fact may explain why it is so hard to predict success of startup businesses (Kerr et al., 2014).

5.3. Implementation ability versus identification ability

We now consider an extension of our baseline model to allow for type-dependent project quality. Suppose that there is a project-selection stage before the game begins, where the agent chooses which project to work on among possible alternatives. Let p_0^i denote the prior probability that a project handled by type i is of good quality. Here, we assume that $1 > p_0^H > p_0^L > 0$; i.e., the high type is better at discovering ideas or identifying promising projects.¹⁸ Also, throughout this subsection, we focus on the exponential bandit specification where $f_i(\tau) = \lambda_i e^{-\lambda_i \tau}$. It is easy to verify that the double-crossing property still holds, and we can follow the same procedure to characterize equilibria.

With this modification, high-type and low-type agents are different along two dimensions: the ability to implement a project (λ_i), and the ability to identify a good project (p_0^i). Which one is more important depends on the context. One obvious factor is who has the right to choose projects: if the agent has discretion over which project to work on, the prior quality of the project p_0^i most likely will depend on the agent's type; if the agent has no such discretion and simply works on the project assigned, the prior should not differ much between the two types. We also argue that identification ability matters more in areas where exploration of new ideas is required, and implementation ability matters more in areas where exploitation of existing ideas is sufficient. For any $p_0^H \geq p_0^L$, the

¹⁸From the technical point of view, it is straightforward to extend the analysis to the case where $p_0^L > p_0^H$ (while $\lambda_H > \lambda_L$). The double-crossing property is preserved and our analysis carries through as long as $p_0^H \lambda_H > p_0^L \lambda_L$. If this condition fails, the model then reduces to the standard setup.

dimension concerning the implementation ability becomes relatively more important as λ_H becomes farther apart from λ_L . Below, we conduct comparative statics with respect to λ_H , to show that the equilibrium allocation depends crucially on what is captured by the agent’s “reputation.”

The equilibrium outcome of our model is determined largely by which type quits first, or alternatively whether \hat{t} is larger or smaller than $s_L^*(q_0)$. Since $s_L^*(q_0)$ is independent of λ_H , we only need to look at how \hat{t} varies with λ_H . For clarity, define $\hat{t}(\lambda_H)$ explicitly as a function of λ_H , which solves

$$\frac{\lambda_H p_0^H e^{-\lambda_H \hat{t}}}{1 - p_0^H + p_0^H e^{-\lambda_H \hat{t}}} = \frac{\lambda_L p_0^L e^{-\lambda_L \hat{t}}}{1 - p_0^L + p_0^L e^{-\lambda_L \hat{t}}}.$$

In the proof of the proposition below (provided in the Appendix), we show that $\hat{t}(\lambda_H)$ strictly decreases in λ_H , with $\lim_{\lambda_H \rightarrow \lambda_L} \hat{t}(\lambda_H) = \infty$ and $\lim_{\lambda_H \rightarrow \infty} \hat{t}(\lambda_H) = 0$ for any $p_0^H > p_0^L$. Thus \hat{t} is relatively small when the reputation reflects the implementation ability more (a high λ_H), and increases as the identification ability gains more importance (a low λ_H). Since $s_L^*(q_0)$ is independent of λ_H , this result, along with Propositions 4 and 5, immediately leads to the following statement.

Proposition 6. *For any λ_L , there exist $\bar{\lambda}$ and $\underline{\lambda}$ (with $\bar{\lambda} > \underline{\lambda} > \lambda_L$) such that:*

- (a) *if $\lambda_H < \underline{\lambda}$, the equilibrium entails some separation, and the low type quits first, starting from $s_L^*(q_0)$;*
- (b) *if $\lambda_H \in [\underline{\lambda}, \bar{\lambda}]$, the equilibrium entails full pooling, and both types quit at \hat{t} ;*
- (c) *if $\lambda_H > \bar{\lambda}$, the equilibrium entails some separation, and the high type quits first.*

Moreover, if $w_H - w_L$ is sufficiently large, $\bar{\lambda} \rightarrow \infty$ so that only cases (a) and (b) are relevant.

In existing reputation models of experimentation, agents are assumed to differ only in one dimension, either in implementation ability (Bobtcheff and Levy, 2017) or in identification ability (Thomas, 2019). Our analysis suggests that this difference could crucially affect the outcome of experimentation. There must be some separation if λ_H is sufficiently close to λ_L (identification ability is relatively more important). However, if $w_H - w_L$ is large, the incentive for the low type to pool with the high type is strong. A larger gap $w_H - w_L$ expands the range of parameters which support a pooling equilibrium. For sufficiently large $w_H - w_L$, Proposition 6 provides a simple prediction: there exists some $\underline{\lambda}$

such that the equilibrium is full pooling if and only if $\lambda_H > \underline{\lambda}$. In this case pooling occurs when the difference in implementation abilities of the two types is large ($\lambda_H > \underline{\lambda}$), while some separation occurs when the difference in implementation abilities is small ($\lambda_H < \underline{\lambda}$).

Whether equilibrium entails some separation or not offers important implications for who should retain the right to abandon the project. To put this idea in context, consider a principal who decides whether to delegate or not at the outset of the game. Suppose that the principal earns a payoff of W_p when the agent succeeds and incurs a flow cost c_p as long as the project is implemented. The payoff in case of failure is normalized to 0. The expected payoff to the principal if the agent is of type i and the project is abandoned at time s is

$$U_i^P(s) = \int_0^s e^{-\rho\tau} p_0^i f_i(\tau) [W_p - C_p(\tau)] d\tau - e^{-\rho s} (1 - p_0^i F_i(s)) C_p(s),$$

where $C_p(t) = c_p(e^{\rho t} - 1)/\rho$. When the principal retains the authority, she chooses s to maximize $q_0 U_H^P(s) + (1 - q_0) U_L^P(s)$. Let s_p^* denote the optimal stopping time for the principal. The costs and benefits of centralization are clear in this setup: by centralization, the principal must abandon the project without any information about the agent's type, but can choose to stop at her optimal timing s_p^* . This suggests that delegation is of value only if the agent uses his private information to decide the timing of project abandonment. Alternatively, for a range of parameters under which the equilibrium is full pooling, the principal cannot be worse off by retaining the authority (centralization), even if she acquires no additional information of her own along the way.¹⁹ Proposition 6 clarifies when centralization is more effective relative to delegation. Given that $w_H - w_L$ is relatively large, centralization is unambiguously the better choice when implementation ability is more important. As identification ability gains more importance, however, the equilibrium entails separation, which tends to raise the value of delegation.

5.4. Model with a continuum of types

As another extension, we may consider a case with more than two types, which would allow us to check robustness of our main findings. Of particular interest from the analytical point of view is the case with a continuum of types. Consider an agent who is characterized by his type $\theta \in [\underline{\theta}, \bar{\theta}]$, where we assume that the distribution of types is continuous with

¹⁹In many cases, the principal often has means to evaluate the agent's productivity over time. When the principal has access to an additional information source, centralization performs even better.

full support. Suppose further that the distribution of arrival times conditional on a good quality project is exponential with parameter θ .

As in the baseline model, the marginal rate of substitution depends on agent type only through the conditional hazard rate, which is given by $g(t; \theta) = p_0 \theta e^{-\theta t} / (1 - p_0 e^{-\theta t})$. For any two types, θ' and θ'' , $g(t; \theta') - g(t; \theta'')$ has the same sign as

$$e^{-(\theta' - \theta'')t} \frac{1 - p_0 e^{-\theta''t}}{1 - p_0 e^{-\theta't}} - \frac{\theta''}{\theta'},$$

which is single-crossing from above in t for any $\theta' > \theta''$. Therefore, there exists $\hat{t}(\theta', \theta'')$ such that $g(t; \theta') > g(t; \theta'')$ for $t < \hat{t}(\theta', \theta'')$ and $g(t; \theta') < g(t; \theta'')$ for $t > \hat{t}(\theta', \theta'')$. Moreover,

$$\frac{\partial g(t; \theta)}{\partial \theta} = \frac{p_0 e^{-\theta t} (1 - p_0 e^{-\theta t} - \theta t)}{(1 - p_0 e^{-\theta t})^2}$$

is single-crossing from above in θ , meaning that the marginal rate of substitution is non-monotone in θ . These properties ensure that the double-crossing property holds globally: between any two types $\theta' > \theta''$, indifference curves of the two types cross twice, with that of type θ' being more convex.

In a companion paper (Chen et al., 2020), we provide a general analysis of signaling under double-crossing preferences with a continuum of types which includes this extended model as a special case. We show that equilibrium under double-crossing preferences exhibits a particular form of pooling, which can be viewed as a generalized version of LSHP (Low types Separate High types Pool) equilibrium introduced by Kartik (2009): separation is feasible only at the lower end of types, whereas some forms of pooling emerge at the higher end. The equilibrium quitting time, as a function of type θ , jumps at most once and is continuous elsewhere. As such, there is a threshold type below which types are fully separated and above which they are clustered together in some ways.²⁰

The general analysis raises two important points. First, it shows that pooling is a robust prediction of signaling under double-crossing preferences. Second, it is also important to recall that equilibrium in the current setting is characterized by the low type's indifference condition, where some of low-type agents choose to separate and some others choose to pool with high-type agents in a semi-pooling equilibrium. This structure has a precise

²⁰Unlike in Kartik (2009), the model with double-crossing preferences admits various forms of pooling above the threshold, though with some regularities. See Chen et al. (2020) for a more precise characterization.

counterpart in the continuous-type model: equilibrium in this case is also characterized by the indifference condition of the threshold type, where types below the threshold choose to fully separate and those above choose to pool with higher types. Types are thus divided into two segments in a manner similar to the two-type case. As a consequence, the continuous-type model exhibits similar equilibrium properties, i.e., homogenization of quitting times and some pooling in equilibrium, implying that our predictions are robust in a qualitative sense and carry over to more general settings.

6. Policy implications

Our model illustrates how reputation concerns distort the timing of project abandonment. There are many potential remedies for those distortions, with centralization of decision-making rights being one of them. The problem is that centralization is a second-best solution in that the equilibrium quitting time does not reflect the agent’s private information. The principal can do better if she can commit to and enforce a more sophisticated incentive scheme. In this section, we explore this possibility and discuss some policy implications when the principal is equipped with more tools to manipulate the underlying payoff structure. Since the case of venture startups provides a leading example of our model, we focus on two remedies that are particularly relevant for this industry.

The valley of death. Consider a situation where the principal can raise the cost of continuing the project at some predetermined point in time. Specifically, suppose she can set a deadline and charge a “fee” when the agent continues the project past the deadline. We argue that this type of intervention can be quite effective in regulating the dynamic distortion when $s_H^*(1) > s_L^*(0)$, or alternatively when identification ability matters more. In this case, the low type persists for too long in holding out to partially pool with the high type. If the principal can credibly enforce this scheme, the optimal solution is conceptually straightforward: set a deadline at $s_L^*(0)$ and charge a fee high enough to make the continuation payoff for the low type nonpositive. Under this scheme, it is indeed optimal for the low type to quit at $s_L^*(0)$, which in turn allows the high type to separate and quit at the optimal timing $s_H^*(1)$. The optimal scheme can thus restore efficient separation and achieve the full-information outcome.²¹

²¹The effect of the valley of death is most pronounced when there are only two types as we assume here, as it can achieve efficient separation. When we have more types, it is in general not feasible to achieve efficient separation by a simple scheme like the valley of death, but the basic ideas remain valid. In particular, when low-type agents tend to over-experiment, it is always beneficial to let them exit earlier. Moreover, by doing

This type of midterm screening is often observed in venture financing, where many venture startups face difficulty raising follow-on (“series A”) funding after initial (seed) funding. In many cases, it is relatively easy to obtain seed funding from public or angel sources; past this stage, many startups struggle to raise follow-on funding. This funding gap, which inevitably raises the operational cost of running a business, is often referred to as the “valley of death,” as most startups cannot survive past this phase. However, the reason why such funding gap exists is not immediately obvious, especially given the fact that seed funding typically comes from angel and public sources that are not driven by profit maximization. The presence of the valley of death thus indicates that the expected return of a startup business, conditional on its survival, follows a U-shaped path and stagnates in the middle, so much so that even those funding sources are reluctant to continue.

Our analysis provides a mechanism through which the valley of death naturally emerges, and offers an important efficiency rationale of this funding gap.²² In the presence of reputation concerns, relatively less productive entrepreneurs tend to “hold out” for too long even though their projects are in hopeless shape, which could create an interval where the expected return of a startup business stagnates and dips below the efficient level. Given this, the valley of death can be efficiency-enhancing as it can work as a screening device to differentiate entrepreneurs with different degrees of vision and confidence. It prevents less efficient entrepreneurs from over-experimenting out of reputation concerns. It also raises the expected return for the remaining, more efficient, entrepreneurs, which is important when the participation constraint is binding for some of them. Although there is now a heated debate over how to bridge this gap, with some calling for active public interventions (Murphy and Edwards, 2003; Butler, 2008), our analysis suggests a positive role of the valley of death as a screening device, which is particularly effective in areas where exploration of new ideas is crucial for success.

Startup subsidies. When implementation ability matters more, we have $s_L^*(0) > s_H^*(1)$, in which case the high type quits too early because the reputational value of success is not sufficient. In contrast to our previous discussion about the “valley of death,” the issue here is to induce high-type agents to persist longer to fully explore the true worth of their projects. In this case, startup subsidies which lower the operational cost of continuing a

so, it raises the reputational value of success for the remaining types, which induces those more efficient types to enter the game.

²²In a different framework, Chen and Ishida (2018) also discuss a positive role of the valley of death to screen out less confident entrepreneurs. In their model, there are no reputation concerns in that the payoff from success or failure is fixed independently of the market belief.

startup business can be productive as they allow high-type agents to continue up to their optimal timing.

More precisely, consider a startup subsidy which lowers the flow cost of experimentation from c to \tilde{c} , possibly for some duration of time. Assume $\hat{t} < \underline{t}$ and let $v(t; \tilde{c})$ to denote its dependence on \tilde{c} . If there exists a \tilde{c} such that $v(s_H^*(1); \tilde{c}) = w_H$, we can restore efficiency by providing subsidies for $t \in [0, s_L^*(0)]$, with the cost reverting back to c after $s_L^*(0)$. Under this scheme, the low type is indifferent between quitting at $s_H^*(1)$ and $s_L^*(0)$, and the high type has no incentive to quit before $s_H^*(1)$. We can thus achieve efficient separation in which the high type quits at $s_H^*(1)$ and the low type quits at $s_L^*(0)$.²³

Startup subsidies or grants are ubiquitous in both developed and developing economies. In the United States, for instance, there are several federally funded programs aiming at getting small startup businesses off the ground. In addition to those public funding sources, angel investors also play an essential role in early stages of venture financing, accounting for a substantial fraction of seed money supplied to the venture market. The primary rationale for startup subsidies is to relax credit constraints which small startup firms may face due to market imperfections. Empirical support for this channel is not strong, however, as recent evidence suggests that credit constraints may not play as important a role as was previously believed (Kerr and Nanda, 2011). Our analysis provides an alternative rationale for startup subsidies, which is to allow more promising projects to persist longer and live up to their full potential. Note that our argument differs from the conventional one in an important way, as it stems from dynamic reputation concerns and holds irrespective of whether there are credit market imperfections. This implies that our argument can be applied to a range of situations, outside of venture financing, where credit constraints are not a crucial factor.

7. Conclusion

This paper provides a framework to analyze the role of reputation concerns in risky experimentation. We develop a general approach which encompasses a broad class of learning processes and model specifications and obtain a complete characterization of unique D1 equilibrium. Our analysis suggests that signaling incentives in the context of risky exper-

²³While we focus on the simplest scheme here for the sake of illustration, there are many other schemes that can equally achieve efficient separation. In particular, if there is a positive transfer cost and the government would like to minimize the amount of subsidies, the scheme must be time-contingent: letting \tilde{c}_t denote the flow cost at time t , the cost-minimizing scheme must satisfy $v(t; \tilde{c}_t) = w_H$ for all $t \in [s_H^*(1), s_L^*(0)]$.

imentation are qualitatively different from those in other more standard contexts characterized by the single-crossing property, with reputation concerns inducing homogenization of quitting times between types. We study the implications of these inefficient termination decisions for organizational design, and offer potential remedies to correct the distortions. The implications obtained here are far-reaching and can be applied broadly to many situations of interest, such as an entrepreneur experimenting with a business startup, a politician with a policy reform, an engineer with a new product design, and a researcher with a scientific hypothesis.

As a final remark, we would like to offer a broader perspective for the scope of our analysis. The insights obtained from our analysis rely crucially on the double-crossing property of indifference curves. This property naturally arises in the context of risky experimentation as long as the time horizon is long enough. We work with an infinite-horizon model in our analysis, but there is sometimes an exogenously imposed deadline for experimentation. The political-economy example is a case in point, where politicians are often subject to a time frame defined by election cycle. In those instances, the speed of learning becomes an additional factor to determine the form of equilibrium. For some policy issues, it takes a long time, sometimes decades, to observe outcomes of policy experimentation; many environmental or foreign policies supposedly fall into this category. In such a case, the deadline (or the election date) binds, and the single-crossing property dictates equilibrium behavior as in the standard setup.²⁴ For some other issues, on the other hand, we can often expect to learn quickly which policy measures work and which do not, so that the time horizon is less likely to be an issue; an example may be domestic economic policy for which there are many readily available indicators. The insights from our analysis are more applicable to this type of environment.

²⁴If the deadline arrives before \hat{t} , preferences are single-crossing for all feasible choices of t . The model then becomes a standard signaling model with a constrained action space.

References

- Bar-Isaac, Heski**, “Reputation and Survival: Learning in a Dynamic Signalling Model,” *Review of Economic Studies*, 2003, 70 (2), 231–251.
- Bernheim, B. Douglas**, “A Theory of Conformity,” *Journal of Political Economy*, 1994, 102 (5), 841–877.
- Bobtcheff, Catherine and Raphael Levy**, “More Haste, Less Speed? Signaling through Investment Timing,” *American Economic Journal: Microeconomics*, 2017, 9 (3), 148–186.
- Bonatti, Alessandro and Johannes Hörner**, “Career Concerns with Exponential Learning,” *Theoretical Economics*, 2017, 12 (1), 425–475.
- Boyarchenko, Svetlana**, “Strategic Experimentation with Erlang Bandits,” 2017. Working Paper, University of Texas, Austin.
- Butler, Declan**, “Translational Research: Crossing the Valley of Death,” *Nature*, 2008, 453 (7197), 840–842.
- Chen, Chia-Hui and Junichiro Ishida**, “Hierarchical Experimentation,” *Journal of Economic Theory*, 2018, 177, 365–404.
- , —, and **Wing Suen**, “Signaling under Double-Crossing Preferences,” 2020. mimeo.
- Chen, Jing**, “Selection and Serial Entrepreneurs,” *Journal of Economics and Management Strategy*, 2013, 22 (2), 281–311.
- Cho, In-Koo and David M. Kreps**, “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 1987, 102 (2), 179–221.
- Daley, Brendan and Brett Green**, “Waiting for News in the Market for Lemons,” *Econometrica*, 2012, 80 (4), 1433–1504.
- and —, “Market Signaling with Grades,” *Journal of Economic Theory*, 2014, 151, 114 – 145.
- Gompers, Paul, Anna Kovner, Josh Lerner, and David Scharfstein**, “Performance Persistence in Entrepreneurship,” *Journal of Financial Economics*, 2010, 96 (1), 18–32.

- Gul, Faruk and Wolfgang Pesendorfer**, “The War of Information,” *Review of Economic Studies*, 2012, 79 (2), 707–734.
- Halac, Marina and Ilan Kremer**, “Experimenting with Career Concerns,” *American Economic Journal: Microeconomics*, 2020, 12 (1), 260–288.
- , **Navin Kartik, and Qingmin Liu**, “Optimal Contracts for Experimentation,” *Review of Economic Studies*, 2016, 83 (3), 1040–1091.
- Kartik, Navin**, “Strategic Communication with Lying Costs,” *Review of Economic Studies*, 2009, 76 (4), 1359–1395.
- Keller, Godfrey, Sven Rady, and Martin Cripps**, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, 73 (1), 39–68.
- Kerr, William R. and Ramana Nanda**, “Financial Constraint and Entrepreneurship,” in D. Audretsch, O. Falck, and S. Heblich, eds., *Handbook of Research on Innovation and Entrepreneurship*, Cheltenham: Edward Elgar, 2011, pp. 88–103.
- , – , and **Matthew Rhodes-Kropf**, “Entrepreneurship as Experimentation,” *Journal of Economic Perspectives*, 2014, 28 (3), 25–48.
- Lafontaine, Francine and Kathryn Shaw**, “Serial Entrepreneurship: Learning by Doing?,” *Journal of Labor Economics*, 2016, 34 (S2), S217–S254.
- Lee, Jihong and Qingmin Liu**, “Gambling Reputation: Repeated Bargaining with Outside Options,” *Econometrica*, 2013, 81 (4), 1601–1672.
- Matthews, Steven and John Moore**, “Monopoly Provision of Quality and Warranties: An Exploration in the Theory of Multidimensional Screening,” *Econometrica*, 1987, 55 (2), 441–467.
- Murphy, Lonna M. and Peter M. Edwards**, “Bridging the Valley of Death: Transitioning from Public to Private Sector Financing,” 2003. Research Report MP-720-34036, National Renewable Energy Laboratory.
- Prendergast, Canice and Lars Stole**, “Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning,” *Journal of Political Economy*, 1996, 104 (6), 1105–1134.

Rocha, Vera, Anabela Carneiro, and Celeste Amorim Varum, “Serial Entrepreneurship, Learning by Doing and Self-Selection,” *International Journal of Industrial Organization*, 2015, 40 (1), 91–106.

Scharfstein, David S. and Jeremy C. Stein, “Herd Behavior and Investment,” *American Economic Review*, 1990, 80 (3), 465–479.

Thomas, Caroline, “Experimentation with Reputation Concerns—Dynamic Signalling with Changing Types,” *Journal of Economic Theory*, 2019, 179, 366–415.

Appendix

Proof of Lemma 2. Since both types quit with positive probability at time t , the interim belief assigned to an agent who quits at that time satisfies $\hat{q} \in (0, 1)$.

Suppose $t < \hat{t}$. Pick a small $\epsilon > 0$. There are two cases: (a) $\sigma_L(t')$ is constant on $t' \in [t, t + \epsilon)$. By D1, the market assigns interim belief 1 to an agent who quits at $t' \in (t, t + \epsilon)$. Because the gain in interim belief is discrete while the change in payoff from delaying to quit is infinitesimal, such a deviation would be profitable. (b) $\sigma_L(t')$ is strictly decreasing on $t' \in [t, t + \epsilon)$. Since $t < \hat{t}$, if the low type is indifferent between continuing and quitting, the high type strictly prefers to continue. This implies that $\sigma_H(t')$ is constant on $[t, t + \epsilon)$. The interim belief assigned to one who quits at $t' \in (t, t + \epsilon)$ must be 0. But then this cannot be optimal for the low type to quit at t' because he would gain by deviating to quit at t and obtain an interim belief of \hat{q} instead of 0.

Suppose $t > \hat{t}$. There are two cases: (a) $\sigma_L(t')$ is constant on $t' \in (t - \epsilon, t)$. By D1, the market assigns interim belief 1 to an agent who quits at $t' \in (t - \epsilon, t)$. It would pay for an agent to deviate by quitting slightly earlier at t' instead of t . (b) $\sigma_L(t')$ is strictly decreasing on $t' \in (t - \epsilon, t)$. Since $t > \hat{t}$, if the low type is indifferent between quitting at t' and quitting at t , the high type strictly prefers quitting at t' . Hence, the interim belief assigned to one who quits at $t' \in (t - \epsilon, t)$ must be 0. But then the low type could gain by deviating to quit at t instead of t' and obtain an interim belief of \hat{q} instead of 0. ■

Proof of Proposition 4. We show that the posited strategy profiles constitute an equilibrium that satisfies D1 in each case. The proof that equilibrium is unique follows the same logic as that for Propositions 2 and 3 and is omitted for brevity.

(a) *Separating equilibrium.* The equilibrium strategy profile is that the high type quits at \underline{t} and the low type quits at $s_L^*(0)$. At time \underline{t} , the low type has no strict incentive to deviate by quitting immediately to get w_H , because by construction $v(\underline{t}; s_L^*(0)) = w_H$. Moreover, when the low type is indifferent between quitting at \underline{t} and quitting at $s_L^*(0)$, the high type strictly prefers quitting at the earlier time \underline{t} , because both \underline{t} and $s_L^*(0)$ are greater than \hat{t} . Now, consider deviating to $s' < \underline{t}$. Because \underline{t} is lower than $s_H^*(q_0)$, quitting at such s' is equilibrium dominated by quitting at \underline{t} for the high type. This means we can assign an off-equilibrium belief of 0 to such deviation. Given this off-equilibrium belief, no type has an incentive to deviate to $s' < \underline{t}$. If $s' > \underline{t}$, then in this case we have $s' > \hat{t}$, and D1 requires that deviations to s' be assigned an off-equilibrium belief of 0. Thus, a low type

that deviates to $s' > \underline{t}$ will get a payoff of $v(\underline{t}; s')$ from the perspective of time \underline{t} . Because it is optimal to quit at $s_L^*(0)$ when the belief is 0, we have $v(\underline{t}; s') < v(\underline{t}; s_L^*(0))$, meaning that deviating to s' is unprofitable for the low type. Further, since $g_H(s') < g_L(s')$ for any $s' > \underline{t} \geq \hat{t}$, the payoff from deviating to $s' > \underline{t}$ for the high type is strictly less than $v(\underline{t}; s')$, which in turn is strictly less than $v(\underline{t}; s_L^*(0)) = w_H$. This means that it is not profitable for the high type to deviate to s' either.

(b) *Semi-pooling equilibrium.* The strategy profile is that the high type quits at \hat{t} and the low type randomizes between \hat{t} and $s_L^*(0)$. The belief associated with quitting at \hat{t} is some $\hat{q} > q_0$ such that $v(\hat{t}; s_L^*(0)) = w_L + r(\hat{t}; \hat{q})(w_H - w_L)$, and the belief associated with quitting at $s_L^*(0)$ is 0. By construction, \hat{q} is chosen such that the low type is indifferent between the two equilibrium quitting times. Further, because $s_L^*(0) > \hat{t}$, indifference of the low type implies that the high type strictly prefers to quit at the earlier time \hat{t} . Now, we consider off-equilibrium deviations. Because both types quit at \hat{t} with positive probability, condition (5) and D1 require that any deviation s' be assigned an off-equilibrium belief of 0. Given such belief, for any quitting time after \hat{t} , the one that is optimal for the low type is by definition $s_L^*(0)$. Thus any deviation $s' > \hat{t}$ is unprofitable for the low type. By the same reasoning as in part (a), this implies that such deviation is unprofitable for the high type as well. Similarly, both types strictly prefer quitting at \hat{t} to deviating to quit at $s' < \hat{t}$.

(c) *Full pooling equilibrium.* The strategy profile is that both types quit at \hat{t} . The belief associated with quitting at \hat{t} is q_0 , and the off-equilibrium belief associated with quitting at any other time is 0. The low type does not want to deviate to $s_L^*(0)$, because the deviation payoff $v(\hat{t}; s_L^*(0))$ would be weakly lower than the equilibrium payoff from quitting immediately. If the low type does not want to deviate to $s_L^*(0)$, he does not want to deviate to any other $s' \neq \hat{t}$. The same argument as before implies that the high type does not want to deviate to any $s' \neq \hat{t}$ either. ■

Proof of Proposition 5. We begin by showing that $V(\cdot; \hat{t}) - w_L$ is single-crossing from below. To see this, note that the derivative of $V(t; \hat{t})$ with respect to t is

$$\begin{aligned} \frac{\partial V(t; \hat{t})}{\partial t} &= (\rho V(t; \hat{t}) + c - g_L(t) [W_L + R(t; \tilde{q}(t))(W_H - W_L) - V(t; \hat{t})]) \\ &\quad + \left(\int_t^{\hat{t}} e^{-\rho(\tau-t)} \frac{p_0 f_L(\tau)}{1 - p_0 F_L(t)} (W_H - W_L) \frac{\partial R(\tau; \tilde{q}(t))}{\partial \tilde{q}} d\tau \right. \\ &\quad \left. + e^{-\rho(\hat{t}-t)} \frac{1 - p_0 F_L(\hat{t})}{1 - p_0 F_L(t)} (w_H - w_L) \frac{\partial r(\hat{t}; \tilde{q}(t))}{\partial \tilde{q}} \right) \frac{d\tilde{q}(t)}{dt}. \end{aligned}$$

If $V(t; \hat{t}) = w_L$ at $t < s_L^*(1)$, the first term is 0 and the second term is positive. If $V(t; \hat{t}) = w_L$ at $t > s_L^*(1)$, the first term is positive and the second term is 0. This shows that $V(\cdot; \hat{t}) - w_L$ is single-crossing from below. Hence, $V(s_L^*(q_0); \hat{t}) < w_L$ and $V(\hat{t}; \hat{t}) > w_L$ imply that $t_0 \in (s_L^*(q_0), \hat{t})$ exists and is unique, with $V(t; \hat{t}) > w_L$ for all $t \in (t_0, \hat{t}]$.

We next show that $V(s_L^*(1); \cdot)$ is decreasing on $[s_L^*(1), \hat{t}]$. The derivative of $V(s_L^*(1); s)$ with respect to s is

$$\begin{aligned} \frac{\partial V(s_L^*(1); s)}{\partial s} &= e^{-\rho(s-s_L^*(1))} \frac{1-p_0 F_L(s)}{1-p_0 F_L(s_L^*(1))} (g_L(s)[W_H - w_H] - \rho w_H - c) \\ &< e^{-\rho(s-s_L^*(1))} \frac{1-p_0 F_L(s)}{1-p_0 F_L(s_L^*(1))} (g_L(s)[W_H - w_L] - \rho w_L - c), \end{aligned}$$

which is non-positive for $s \in [s_L^*(1), \hat{t}]$. Because $V(s_L^*(1); s_L^*(1)) > w_L$, and $t_0 > s_L^*(1)$ implies $V(s_L^*(1); \hat{t}) < w_L$, we can conclude that $t_1 \in (s_L^*(1), \hat{t})$ exists and is unique, with $V(s_L^*(1); t') \leq w_L$ for any $t' \geq t_1$.

Equilibrium. We first show that the strategies described in cases (a), (b)(i), and (b)(ii) of the proposition constitute an equilibrium of the corresponding cases.

(a) Let $J_i(t)$ be the value function for type $i \in \{H, L\}$ corresponding to the strategy profile of full pooling (i.e., both types do not quit until time \hat{t}). By the Bellman equation,

$$J_i(t) = -c dt + g_i(t) dt [W_L + R(t; q_0)(W_H - W_L)] + (1 - g_i(t) dt) e^{-\rho dt} J_i(t + dt).$$

From this, we obtain the differential equation:

$$J_i'(t) = \rho J_i(t) + c - g_i(t) [W_L + R(t; q_0)(W_H - W_L) - J_i(t)], \quad (\text{A1})$$

with terminal condition $J_i(\hat{t}) = w_L + r(\hat{t}; q_0)(w_H - w_L)$. By construction, $J_L(s_L^*(q_0)) = V(s_L^*(q_0); \hat{t}) \geq w_L$ for the low type. Moreover, we can write

$$J_L'(t) = (\rho + g_L(t))(J_L(t) - w_L) + (\rho w_L + c - g_L(t) [W_L + R(t; q_0)(W_H - W_L) - w_L]).$$

From the definition of $s_L^*(q_0)$, the second term is equal to zero at $t = s_L^*(q_0)$, and it is positive when $t > s_L^*(q_0)$ and negative when $t < s_L^*(q_0)$. For $t \in (s_L^*(q_0), \hat{t}]$, the second term is positive, and so $J_L(t) - w_L$ is single-crossing from below. But $J_L(t) - w_L$ is non-negative at $t = s_L^*(q_0)$ and is positive at $t = \hat{t}$. We therefore must have $J_L(t) > w_L$ for all $t \in (s_L^*(q_0), \hat{t}]$. For $t \in [0, s_L^*(q_0))$, the second term is negative, and so $J_L(t) - w_L$ is single-crossing from above. But since $J_L(t) - w_L$ is non-negative at $t = s_L^*(q_0)$, we must

have $J_L(t) > w_L$ for all $t \in [0, s_L^*(q_0))$. Since $J_L(t) \geq w_L$ for all $t \leq \hat{t}$, it is indeed optimal for the low type not to quit until \hat{t} . Since $g_H(t) > g_L(t)$ for all $t < \hat{t}$, we also have $J_H(t) > J_L(t) \geq w_L$. Thus, the high type also has no incentive to quit until time \hat{t} .

(b)(i) For $i \in \{H, L\}$, let $J_i(t)$ be the solution to the differential equation (A1), with q_0 in (A1) and in the terminal condition replaced by $\tilde{q}(t_0)$. Since the high type never quits until time \hat{t} , $J_H(t)$ is the value function for the high type.

For the low type, let $\tilde{J}_L(t)$ be the solution to the differential equation (A1), with terminal condition $\tilde{J}_L(s_L^*(q_0)) = w_L$. Then, the value function for the low type is given by

$$J_L^*(t) = \begin{cases} \tilde{J}_L(t) & \text{if } t \in [0, s_L^*(q_0)), \\ w_L & \text{if } t \in [s_L^*(q_0), t_0], \\ J_L(t) & \text{if } t \in (t_0, \hat{t}]. \end{cases}$$

By the same argument as in part (a), we have $J_L^*(t) > w_L$ for $t \in [0, s_L^*(q_0))$ and for $t \in (t_0, \hat{t}]$. Thus, it is optimal for the low type not to quit for such t . For $t \in [s_L^*(q_0), t_0]$, equation (8) ensures that the low type is indifferent between quitting and staying. Thus, the strategy of the low type is indeed a best response. Furthermore, $J_H(t) > J_L^*(t) \geq w_L$ for all $t < \hat{t}$. Thus, the high type has no incentive to quit until time \hat{t} .

(b)(ii) Fix any $t' \in [\max\{t_1, s_H^*(1)\}, \hat{t}]$. We have already shown that $t_1 < \hat{t}$. Moreover, $s_L^*(1) \leq t_0 < \hat{t}$ implies $s_H^*(1) < \hat{t}$. Thus, the interval is non-empty.

For $t \leq t'$, the value function for the high type is given by the solution to the differential equation:

$$J_H'(t) = \rho J_H(t) + c - g_H(t)[W_H - J_H(t)],$$

with terminal condition $J_H(t') = w_H$. For the low type, the value function is given by

$$J_L^*(t) = \begin{cases} \tilde{J}_L(t) & \text{if } t \in [0, s_L^*(q_0)), \\ w_L & \text{if } t \in [s_L^*(q_0), t'), \\ w_H & \text{if } t = t'; \end{cases}$$

where $\tilde{J}_L(t)$ is as defined in part (b)(i). Note that no agent quits at time $t \in (s_L^*(1), t')$. We assign off-equilibrium belief $\hat{q}(t) = 0$ for an agent who quits at such time, which is consistent with the D1 criterion because $t < \hat{t}$. Because $J_L^*(t) > w_L$ for $t < s_L^*(q_0)$, the low type strictly prefers continuing with the risky project than quitting. For $t \in [s_L^*(q_0), s_L^*(1)]$, we have $J_L^*(t) = w_L$. Therefore the strategy $\sigma_L(t)$ that satisfies equations (8) and (9) is

indeed a best response. At $t = s_L^*(1)$, the gain for a low type from deviating to quit at t' instead is $V(s_L^*(1), t') - w_L \leq V(s_L^*(1), t_1) - w_L = 0$. Therefore, the low type cannot gain from deviating to wait until t' to quit.

Uniqueness. We now show that the equilibria described in cases (a), (b)(i), and (b)(ii) of the proposition are the only candidates for equilibrium for the corresponding cases. To this end, let \underline{t}_i and \bar{t}_i be the earliest and latest possible time, respectively, for type i to quit on the equilibrium path. We first establish the following facts.

Fact 1. $\underline{t}_H \geq s_L^*(q_0)$.

Proof. Suppose $\underline{t}_H < s_L^*(q_0)$. From the proof of Lemma 2, we know that there cannot be any pooling before time \hat{t} . Given $s_L^*(q_0) < \hat{t}$, this means that we cannot have $\underline{t}_H = \underline{t}_L < \hat{t}$. Therefore, $\underline{t}_H \neq \underline{t}_L$. Moreover, in this case, it is not possible to have $\underline{t}_L < \underline{t}_H$, for otherwise $\hat{q}_{\underline{t}_L} = 0 < \hat{q}_{\underline{t}_H}$, $\tilde{q}_{\underline{t}_L} = q_0 < \tilde{q}_{\underline{t}_H}$, and the low type could profitably deviate by quitting at \underline{t}_H . This means that we only need to look at the case where $\underline{t}_H < \underline{t}_L$.

With some abuse of notation, let $V_i(\underline{t}_H; \underline{t}_L)$ denote the expected payoff if a type i agent continues to work with the risky project from \underline{t}_H to \underline{t}_L on the equilibrium path. We obtain

$$\begin{aligned} V_i(\underline{t}_H; \underline{t}_L) &= e^{-\rho(\underline{t}_L - \underline{t}_H)} \frac{1 - p_0 F_i(\underline{t}_L)}{1 - p_0 F_i(\underline{t}_H)} [J_i(\underline{t}_L) - C(\underline{t}_L - \underline{t}_H)] \\ &\quad + \int_{\underline{t}_H}^{\underline{t}_L} e^{-\rho(\tau - \underline{t}_H)} \frac{p_0 f_i(\tau)}{1 - p_0 F_i(\underline{t}_H)} [W_L + R(\tau; \tilde{q}_\tau)(W_H - W_L) - C(\tau - \underline{t}_H)] d\tau, \end{aligned}$$

where $J_i(\underline{t}_L)$ is the equilibrium payoff to type i agent at time \underline{t}_L . Since an agent earns w_H by quitting at \underline{t}_H , the high type quitting at \underline{t}_H and the low type quitting at \underline{t}_L imply $V_L(\underline{t}_H; \underline{t}_L) \geq w_H \geq V_H(\underline{t}_H; \underline{t}_L)$.

We now show that this condition cannot be satisfied. To this end, observe first that $\underline{t}_L \leq \hat{t}$. Suppose otherwise. Then, since \tilde{q}_t is weakly decreasing on $[\hat{t}, \underline{t}_L)$ and $s_H^*(\tilde{q}_{\hat{t}}) < \hat{t}$, there cannot exist t' and t'' in $[\hat{t}, \underline{t}_L)$ such that the high type is indifferent between quitting at t' and t'' . Therefore, there must exist some interval $(\underline{t}_L - \epsilon, \underline{t}_L)$ such that no type quits, and as $\underline{t}_L > s_L^*(q_0)$, the low type can profitably deviate by quitting at $t \in (\underline{t}_L - \epsilon, \underline{t}_L)$.

We note that, for $t \leq \underline{t}_L$, $V_i(t; \underline{t}_L)$ solves the differential equation,

$$V_i'(t) = \rho V_i(t) + c - g_i(t)[W_L + R(t; \tilde{q}_t)(W_H - W_L) - V_i(t)],$$

with terminal condition $V_i(\underline{t}_L) = J_i(\underline{t}_L)$. Note also that $\underline{t}_L \leq \hat{t}$ implies that $g_L(t) < g_H(t)$ for $t \in [\underline{t}_H, \underline{t}_L)$. Hence, for any $t \in [\underline{t}_H, \underline{t}_L)$, the right-hand-side of the above is strictly higher

for $i = L$ than for $i = H$. Moreover, because it is feasible for the high type to quit at \underline{t}_L , we have $J_H(\underline{t}_L) \geq w_H + r(\underline{t}_L; \hat{q}_{\underline{t}_L})(w_H - w_L) = J_L(\underline{t}_L)$. Therefore, by the comparison theorem for differential equations, we have $V_L(\underline{t}_H; \underline{t}_L) < V_H(\underline{t}_H; \underline{t}_L)$. This contradicts $V_L(\underline{t}_H; \underline{t}_L) \geq w_H \geq V_H(\underline{t}_H; \underline{t}_L)$. \square

Fact 2. Either $\underline{t}_L = \underline{t}_H = \hat{t}$, or $\underline{t}_L = s_L^*(q_0) < \underline{t}_H$.

Proof. Suppose $\underline{t}_H < \underline{t}_L$. Then, $\hat{q}_{\underline{t}_L} < \hat{q}_{\underline{t}_H} = 1$ and $\tilde{q}_\tau < q_0$ for $\tau \in (\underline{t}_H, \underline{t}_L)$. Moreover, we know that $\underline{t}_H \geq s_L^*(q_0)$ from Fact 1. Given that $\underline{t}_L > \underline{t}_H \geq s_L^*(q_0) > s_L^*(\tilde{q}_{\underline{t}_L})$, and also that \tilde{q}_t is weakly decreasing (as only the high type may quit before \underline{t}_L), the low type could receive a higher payoff by quitting at \underline{t}_H than by quitting at \underline{t}_L , a contradiction. Therefore, $\underline{t}_L \leq \underline{t}_H$. If $\underline{t}_L = \underline{t}_H$, we must have $\underline{t}_L = \underline{t}_H = \hat{t}$ by Lemma 2. If $\underline{t}_L < \underline{t}_H$, then it is optimal for the low type to start quitting at $s_L^*(q_0)$. \square

Fact 3. Either $\bar{t}_L = \bar{t}_H = \hat{t}$, or $\bar{t}_L < \bar{t}_H$.

Proof. If $\bar{t}_H < \bar{t}_L$, we have $\tilde{q}_{\bar{t}_L} = 0$. Then, since $\bar{t}_H \geq \underline{t}_H > s_L^*(q_0)$, the low type could profitably deviate by quitting at $t \in (\bar{t}_H, \bar{t}_L)$. Therefore, $\bar{t}_L \leq \bar{t}_H$. Lemma 2 implies that if $\bar{t}_L = \bar{t}_H$, then both are equal to \hat{t} . \square

Fact 4. $\underline{t}_H = \bar{t}_H$.

Proof. Suppose $\underline{t}_H < \bar{t}_H$. Suppose further that both types quit at \hat{t} with positive probability. This means that the high type is indifferent among quitting at \underline{t}_H , $\hat{t} \in [\underline{t}_H, \bar{t}_H]$, and \bar{t}_H , but condition (5) then implies that the low type must strictly prefer quitting either at \underline{t}_H or \bar{t}_H to quitting at \hat{t} , a contradiction. This rules out pooling or partial pooling with $\underline{t}_H < \bar{t}_H$.

Now suppose that $\underline{t}_H < \bar{t}_H$ and the equilibrium is separating, in which case we have either (1) $\underline{t}_L = s_L^*(q_0) < \bar{t}_L < \underline{t}_H < \bar{t}_H$; or (2) $\underline{t}_L = s_L^*(q_0) < \underline{t}_H < \bar{t}_L < \bar{t}_H$.

In case (1), we have $\bar{t}_L = s_L^*(1)$, $\hat{q}_{\bar{t}_L} = 0$, $\hat{q}_{\underline{t}_H} = 1$, $\tilde{q}_t = 1$ for $t \in (\bar{t}_L, \bar{t}_H)$, and the low type weakly prefers quitting at \bar{t}_L to quitting at \underline{t}_H . This implies that for $t \in (\underline{t}_H, \bar{t}_H)$, no belief \hat{q}_t can give the low type a higher payoff than the equilibrium payoff. Moreover, since the high type is indifferent between quitting at \underline{t}_H and \bar{t}_H , the expected payoff must go up first and then go down for $t \in (\underline{t}_H, \bar{t}_H)$ (with $\hat{q}_t = \tilde{q}_t = 1$ over this interval), suggesting that the set of beliefs \hat{q}_t that give the high type a higher payoff than the equilibrium payoff is not empty for this interval. Therefore, $\hat{q}_t = 1$ for $t \in (\underline{t}_H, \bar{t}_H)$ by D1, and the high type could profitably deviate by quitting at any t in this interval.

In case (2), the proof of Fact 1 shows that if the two types quit separately before \hat{t} , the

low type must quit before the high type. By a similar argument, we can also show that if the two types quit separately after \hat{t} , the high type must quit before the low type. These facts imply that given $\underline{t}_L < \underline{t}_H < \bar{t}_L$, we must have $\underline{t}_L < \hat{t}$ and $\bar{t}_L > \hat{t}$. If $\bar{t}_L > \hat{t}$, however, we cannot have $\bar{t}_L < \bar{t}_H$, which implies that $\underline{t}_H < \bar{t}_H$ cannot occur in equilibrium. \square

Now define $t_H = \underline{t}_H = \bar{t}_H$. The above facts show that there can only be three types of equilibria: (1) full pooling in which $\underline{t}_L = t_H = \hat{t}$; (2) semi-pooling in which $\underline{t}_L = s_L^*(q_0) < \bar{t}_L = t_H = \hat{t}$; or (3) separating in which $\underline{t}_L = s_L^*(q_0) < \bar{t}_L = s_L^*(1) < t_H$. Drawing on this fact, we show that the equilibrium is unique in each of the cases.

(a) We start with the case where $V(s_L^*(q_0); \hat{t}) \geq w_L$. In this case, since $V(\cdot, \hat{t}) - w_L$ is single-crossing from below, $V(s_L^*(q_0); \hat{t}) \geq w_L$ implies that $V(t; \hat{t}) > w_L$ for all $t \in (s_L^*(q_0), \hat{t}]$. Suppose first that $\underline{t}_L = s_L^*(q_0) < \bar{t}_L = t_H = \hat{t}$, so that the equilibrium is semi-pooling. Then, since $V(t; \hat{t}) > w_L$ for all $t \in (s_L^*(q_0), \hat{t}]$, the low type strictly prefers quitting at \hat{t} to quitting at $\underline{t}_L = s_L^*(q_0)$. Next suppose that $\bar{t}_L = s_L^*(1) < t_H$, so that the equilibrium is separating. Since $s_L^*(1) \leq \hat{t}$ implies $s_H^*(1) < \hat{t}$ and $s_L^*(1) > \hat{t}$ implies $s_H^*(1) < s_L^*(1)$, we have $s_H^*(1) < \max\{\hat{t}, s_L^*(1)\}$. Moreover, $t_H > \max\{\hat{t}, s_L^*(1)\}$, because $V(s_L^*(1); \hat{t}) > w_L$ if $s_L^*(1) < \hat{t}$, and thus t_H must be greater than \hat{t} in order to prevent the low type from deviating. By D1, $\hat{q}_t = 1$ for $t \in (\max\{\hat{t}, s_L^*(1)\}, t_H)$, but then the high type could profitably deviate by quitting at $t \in (\max\{\hat{t}, s_L^*(1)\}, t_H)$. This shows that the only possible equilibrium in this case is full pooling.

(b) If $V(s_L^*(q_0); \hat{t}) < w_L$, the full pooling equilibrium cannot exist, because the low type could profitably deviate by quitting at $s_L^*(q_0)$. We then need to consider two cases, depending on whether t_0 is larger or smaller than $s_L^*(1)$.

(b)(i) Suppose $t_0 < s_L^*(1)$. Suppose further that a separating equilibrium exists, i.e., $\bar{t}_L = s_L^*(1)$. The high type must then quit at $t_H > \hat{t}$ such that $V(s_L^*(1), t_H) = w_L$. Recall that $s_H^*(1) < \max\{\hat{t}, s_L^*(1)\}$, and thus the equilibrium cannot survive the D1 criterion. This shows that the only possible equilibrium in this case is semi-pooling.

(b)(ii) If $t_0 \geq s_L^*(1)$, the semi-pooling equilibrium cannot exist, since the low type must prefer quitting at $s_L^*(q_0)$ to quitting at \hat{t} . Therefore, only the separating equilibrium is feasible in this case, although there is still a continuum of separating equilibria that are feasible. To further reduce the set of equilibria, note that since $\bar{t}_L = s_L^*(1)$, t_H must be weakly greater than t_1 so that the low type would not deviate. Moreover, since $V(s_L^*(1); \cdot)$ is decreasing on $[s_L^*(1), \hat{t}]$, the low type strictly prefers quitting at \bar{t}_L to quitting at any $t \in (t_1, \hat{t}]$ even if $\hat{q}_t = 1$. If $t_H \in (t_1, \max\{t_1, s_H^*(1)\})$ (which implies that $s_H^*(1) > t_1$), by

D1, we assign off-equilibrium belief $\hat{q}_{t'} = 1$ to an agent who quits at any $t' \in (t_H, s_H^*(1))$, since compared to the equilibrium payoff, only the high type could possibly benefit from such a deviation. Then the high type could profitably deviate by quitting at t' . If $t_H \in (\max\{t_1, s_H^*(1)\}, \hat{t})$, by D1, we also assign off-equilibrium belief $\hat{q}_{t''} = 1$ to an agent who quits at any $t'' \in (\max\{t_1, s_H^*(1)\}, t_H)$, and thus the high type could profitably deviate by quitting at t'' . Therefore, we must have $t_H = \max\{t_1, s_H^*(1)\}$, which uniquely pins down the equilibrium. ■

Proof of Proposition 6. We first show $\hat{t}(\lambda_H)$ strictly decreases in λ_H , with $\lim_{\lambda_H \rightarrow \lambda_L} \hat{t}(\lambda_H) = \infty$ and $\lim_{\lambda_H \rightarrow \infty} \hat{t}(\lambda_H) = 0$.

To show $\lim_{\lambda_H \rightarrow \lambda_L} \hat{t}(\lambda_H) = \infty$, suppose that \hat{t} is bounded from above when λ_H approaches λ_L . Then, in the limit, we must have

$$\frac{p_0^H}{1 - p_0^H + p_0^H e^{-\lambda_L \hat{t}}} = \frac{p_0^L}{1 - p_0^L + p_0^L e^{-\lambda_L \hat{t}}}.$$

For any $p_0^H > p_0^L$, this is a contradiction because the equality cannot be satisfied for any \hat{t} . Similarly, to show that $\lim_{\lambda_H \rightarrow \infty} \hat{t}(\lambda_H) = 0$, suppose that there exists some $\epsilon > 0$ such that $\lim_{\lambda_H \rightarrow \infty} \hat{t}(\lambda_H) > \epsilon$. For any such ϵ , however, we have $\lim_{\lambda_H \rightarrow \infty} \lambda_H e^{-\lambda_H \epsilon} = 0$ by l'Hopital's rule. This implies that

$$\lim_{\lambda_H \rightarrow \infty} \frac{\lambda_H p_0^H e^{-\lambda_H \epsilon}}{1 - p_0^H + p_0^H e^{-\lambda_H \epsilon}} = 0 < \lim_{\lambda_H \rightarrow \infty} \frac{\lambda_L p_0^L e^{-\lambda_L \epsilon}}{1 - p_0^L + p_0^L e^{-\lambda_L \epsilon}},$$

which is a contradiction.

We next show that $\hat{t}(\lambda_H)$ is decreasing. Observe that $g_H(\hat{t}(\lambda_H)) = g_L(\hat{t}(\lambda_H))$. Taking derivative of both sides with respect to λ_H , we obtain

$$\left(\frac{\partial g_L}{\partial t} - \frac{\partial g_H}{\partial t} \right) \frac{d\hat{t}}{d\lambda_H} = \frac{\partial g_H}{\partial \lambda_H}.$$

We know that, evaluated at $t = \hat{t}(\lambda_H)$, $\partial g_L / \partial t > \partial g_H / \partial t$. This means that $d\hat{t}/d\lambda_H$ has the same sign as $\partial g_H / \partial \lambda_H$ (evaluated at $t = \hat{t}(\lambda_H)$). Therefore, $d\hat{t}/d\lambda_H$ has the same sign as

$$(1 - p_0^H + p_0^H e^{-\lambda_H \hat{t}(\lambda_H)}) - \lambda_H \hat{t}(\lambda_H) (1 - p_0^H).$$

This shows that $d\hat{t}/d\lambda_H$ is single-crossing from above as λ_H increases, because the above expression is decreasing in λ_H if $d\hat{t}/d\lambda_H = 0$. As λ_H approaches λ_L , we have shown that

$\hat{t}(\lambda_H)$ approaches infinity, and hence the sign of $d\hat{t}/d\lambda_H$ is negative. Together with the fact that $d\hat{t}(\lambda_H)/d\lambda_H$ is single-crossing from above, the fact that $\lim_{\lambda_H \rightarrow \lambda_L} d\hat{t}(\lambda_H)/d\lambda_H < 0$ implies $d\hat{t}(\lambda_H)/d\lambda_H < 0$ for all $\lambda_H > \lambda_L$.

We consider how equilibrium varies as the value of \hat{t} increases. Suppose first that \hat{t} is sufficiently small and $\hat{t} \leq s_L^*(q_0)$. Moreover, suppose $s_L^*(0) > 0$ and $\underline{t} > 0$ (both of which are independent of λ_H). By Proposition 4, the equilibrium is full pooling if \hat{t} is sufficiently close to $s_L^*(0)$.

Now suppose $\hat{t} > s_L^*(q_0)$. First, it is clear that the equilibrium is full pooling if \hat{t} is sufficiently close to $s_L^*(0)$, because

$$\lim_{\hat{t} \rightarrow s_L^*(q_0)} V(s_L^*(q_0); \hat{t}) = w_L + r(s_L^*(0); q_0)(w_H - w_L) > w_L.$$

Note also that the derivative of $V(s_L^*(q_0); \hat{t})$ with respect to \hat{t} has the same sign as

$$g_L(s) [W_L + R(\hat{t}; q_0)(W_H - W_L) - w_L - r(\hat{t}; q_0)(w_H - w_L)] - \rho [w_L + r(\hat{t}; q_0)(w_H - w_L)] - c,$$

which is negative for any $\hat{t} > s_L^*(q_0)$ by definition. Moreover, since

$$\lim_{\hat{t} \rightarrow \infty} V(s_L^*(q_0); \hat{t}) = \int_{s_L^*(q_0)}^{\infty} e^{-\rho(\tau - s_L^*(q_0))} \frac{p_0 f_L(\tau)}{1 - p_0 F_L(s_L^*(q_0))} [W_L + R(\tau; q_0)(W_H - W_L) - C(\tau - s_L^*(q_0))] d\tau < w_L,$$

there exists a unique \hat{t}_0 such that $V(s_L^*(q_0); \hat{t}_0) = w_L$. At this point, the equilibrium must be semi-pooling. Finally, to satisfy $V(t_0; \hat{t}) = w_L$, the benefit of pooling evaluated at t_0 must be strictly positive, implying that $e^{-\rho(\hat{t} - t_0)}$ must be bounded away from 0. This means that $t_0 \rightarrow \infty$ as $\hat{t} \rightarrow \infty$. It then follows that the equilibrium must be separating when \hat{t} is sufficiently large.

We have shown that, as \hat{t} increases, the equilibrium changes from that described in case (c) to that described in case (b) and then to case (a). Since \hat{t} decreases in λ_H , therefore the proposition follows.

To show that $\bar{\lambda}$ can be arbitrarily large, observe that, for any fixed w_L , $\underline{t} = 0$ if w_H is sufficiently large. Moreover, for w_H sufficiently large, $v(\hat{t}) \leq w_L + r(\hat{t}; q_0)(w_H - w_L)$. Thus, case (c) of Proposition 4 applies and the equilibrium must be full pooling. ■

Online Appendix to “Reputation Concerns in Risky Experimentation”

Suppose that the agent’s payoff depends partly on his actual type. To illustrate how it changes our analysis clearly, we focus on the case of exit signaling ($W_H = W_L = W$).

For $i \in \{H, L\}$, let $w_i(s; \hat{q})$ denote the agent’s payoff when he quits at time s and when the interim belief is \hat{q} . This is now type-contingent and is given by

$$w_i(s; \hat{t}) = \alpha w(s; \hat{q}) + (1 - \alpha)w_i = \alpha[w_L + r(s; \hat{q})(w_H - w_L)] + (1 - \alpha)w_i,$$

where the parameter $\alpha \in (0, 1)$ measures the importance of reputation. The objective function becomes

$$U_i(s, \hat{q}) = \int_0^s e^{-\rho\tau} p_0 f_i(\tau) [W - C(\tau)] d\tau + e^{-\rho s} (1 - p_0 F_i(s)) [-C(s) + w_i(s; \hat{q})].$$

From this, the marginal rate of substitution between stopping time s and interim belief \hat{q} is obtained as

$$\frac{g_i(s)[W - w_i(s; \hat{q})] - \rho w_i(s; \hat{q}) - c + \alpha(\partial r / \partial s)(w_H - w_L)}{\alpha(\partial r / \partial \hat{q})(w_H - w_L)}.$$

The difference in the marginal rates of substitution between the high type and the low type has the same sign as

$$D(s; \hat{q}) = \frac{\alpha}{1 - \alpha} [g_H(s) - g_L(s)][W - w(s; \hat{q})] + g_H(s)(W - w_H) - g_L(s)(W - w_L) - \rho(w_H - w_L).$$

Observe that $D(s; \hat{q})$ may no longer be single-crossing from above when it crosses 0. Define $\tilde{t}(\hat{q})$, if it exists, such that

$$MRS_H(\tilde{t}(\hat{q}), \hat{q}) = MRS_L(\tilde{t}(\hat{q}), \hat{q}),$$

for a given α . Observe that $\tilde{t}(\hat{q})$ depends not only on the statistical properties of the underlying experimentation process but also on other primitives such as W , w_i and ρ .

Recall that $g_H(s) - g_L(s) = 0$ at $s = \hat{t}$. Therefore, we have $D(\hat{t}; \hat{q}) < 0$, and hence $\tilde{t}(\hat{q}) > \hat{t}$ for any α . Clearly, as $\alpha \rightarrow 1$, $\tilde{t}(\hat{q}) \rightarrow \hat{t}$, suggesting that all of our results hold in a qualitative sense for α close to 1. If α is small or ρ is large, on the other hand, the marginal rates may never cross. In this case, the indifference curves are single-crossing, and there can only exist a separating equilibrium (when $\alpha = 0$, the agent does not care about his reputation and simply stops at the full-information optimal point).

When α is in an intermediate range, it is possible to have multiple tangency points along a single indifference curve. Let \tilde{t}_j , $j = 1, 2, \dots, k$, be the tangency point along the equilibrium indifference curve for the high type, where $\tilde{t}_j < \tilde{t}_{j+1}$. If j is even, the indifference curve for the high type stays below that for the low type, and hence it does not constitute a D1 equilibrium. If j is odd, the indifference curve stays above, so it survives D1 for small deviations. Among them, however, only the one that gives the low type the lowest payoff would survive D1 globally. Let this point be denoted by \tilde{t}_{j^*} . Now consider j which is odd and $j \neq j^*$. This does not survive D1 because the indifference curves that go through this point cross somewhere between \hat{t}_j and \hat{t}_{j^*} , giving the high type an incentive to deviate. As such, there is generically a unique pooling equilibrium even if α is in an intermediate range.