Discussion Paper No. 1060

REPUTATION CONCERNS IN RISKY EXPERIMENTATION

Chia-Hui Chen
Junichiro Ishida
Wing Suen

July 2019

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

Reputation Concerns in Risky Experimentation*

CHIA-HUI CHEN
Kyoto University

JUNICHIRO ISHIDA
Osaka University

WING SUEN
University of Hong Kong

July 1, 2019

Abstract. We develop a general model, with the exponential bandit as a special case, in which high-ability agents are more likely to achieve early success but also learn faster that their project is not promising. These counteracting effects give rise to a signaling model in which the single-crossing condition fails but a double-crossing property holds. We characterize the unique D1 equilibrium under double-crossing condition, and show that it tends to produce pooling. Ability to identify good projects and ability to execute a good project have different implications for the equilibrium allocation. Our model also incorporates public news, which generates dynamic distortions.

Keywords. double-crossing property; D1 equilibrium refinement; 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

The well-being of society depends ultimately on our ability to acquire new knowledge and expand the knowledge frontier. The process of knowledge acquisition, however, is inherently time-consuming and entails a great deal of uncertainty, as an experimenter must go through many alternative approaches and test them one by one without knowing if and when new discovery arrives. With this dynamic and random nature of experimentation, a critical decision an experimenter would eventually face, if things did not proceed as favorably as expected, is when to abandon the process altogether, which essentially determines the collective efficiency of knowledge acquisition. Our ability to acquire new knowledge is undermined when there exists a force which systematically distorts this decision one way or the other.

One potential impediment to this risky process of experimentation is reputation concerns. Since the ability to make new discovery is most likely heterogeneous, the outcome of risky experimentation inevitably reveals some information about the experimenter's "competence," either through the timing of making discovery in case of success or, more importantly, through the timing of abandoning the process in case of failure. If the experimenter benefits from holding a better reputation in some ways, this structure amounts to a dynamic signaling problem in which he strategically chooses the stopping time in order to enhance or protect his reputation. Although reputation concerns are important in many facets of our life, they are even more so in the context of knowledge acquisition, because most of our exploratory activities are undertaken by professionals and experts, such as entrepreneurs, politicians, lawyers, engineers, and scientists, for whom reputation is always an indispensable asset for advancing their careers.

This paper aims to develop a unified framework to analyze the role of reputation concerns in risky experimentation. The basic setup is a standard two-armed bandit problem. Specifically, we consider a dynamic game between an agent (experimenter) and a market (evaluator) in which the agent engages in a project of unknown quality while at the same time attempts to signal his "competence" to the market. The project is either good or bad, and produces a publicly observable breakthrough at some random time only if it is good. In this setup, therefore, the agent becomes pessimistic as failures accumulate ("no news is bad news") and may eventually be forced to retreat from the project without achieving a breakthrough. The game ends either when he achieves a breakthrough or when he abandons the project, and his continuation payoff is determined by his "reputation" at that time.

Our framework has three important features. First, we adopt a general specification of the underlying learning process which includes exponential bandits—the workhorse specification in the literature—as a special case. Second, our model encompasses two distinct notions of reputation that have been adopted in the literature. To be more precise, the agent in our model may differ in two dimensions, the ability to identify promising projects (i.e., the prior probability that the project is good) and the speed of learning (i.e., the rate of discovering the true project quality). Finally, our model also incorporates public news, in addition to the timing of project abandonment, stemming from the fact that the outcome of risky experimentation is publicly observable. Although this last feature is not a generalization *per se*,[1] it nonetheless plays an important role as many situations naturally involve publicly observable signals which (imperfectly) reveal the experimenter's hidden characteristics.

The main technical contribution of the paper lies in our equilibrium characterization, which can be applied to a broad class of learning processes and model specifications. The essential ingredient of our characterization is the double-crossing property of indifference curves.[2] We identify a general condition which gives rise to this property, and exploit it to obtain a complete characterization of D1 equilibria in our general setup. As it turns out, the equilibrium is always unique under D1 as in standard signaling models with single-crossing indifference curves, but the structure of equilibrium differs significantly in that the D1 criterion may not select the least-cost separating equilibrium. More specifically, the equilibrium selected by D1 entails some pooling when the double-crossing property holds in the relevant range. It is also worth noting that the equilibrium is highly sensitive to the prior belief about the agent's type, which allows us to discuss the role of prior reputation in risky experimentation.

The equilibrium allocation depends crucially on two countervailing forces that are inherent in risky experimentation: on one hand, high-ability agents are always more likely to find success than low-ability agents, conditional on the project being good; on the other hand, high-ability agents learn more quickly, and hence their posterior beliefs decline faster

[1] As we will note below, our approach can easily be extended to the case with private news.

[2] It is well known that the double-crossing property emerges in some signaling or screening models. To the best of our knowledge, Matthews and Moore (1987) are the first to consider the double-crossing property in a multidimensional screening model. Daley and Green (2014) show that the double-crossing may emerge in a signaling model with an additional source of information such as grades.

when there is no success. For clarity, we call the former the *ability effect* and the latter the *learning effect*. Note that the learning effect, which arises only when there is uncertainty in the project quality, plays an essential role in generating the double-crossing property. When the ability effect is more dominant, indifference curves become single-crossing in the relevant range, thereby inducing full separation between high and low types. As the learning effect gains more importance, on the other hand, indifference curves exhibit the double-crossing property, and the equilibrium may entail some pooling as a consequence.

This finding suggests that the equilibrium allocation depends crucially on what is embodied in the agent's "reputation." From a broader perspective, our analysis in general suggests that different notions of reputation lead to different equilibrium predictions, thereby necessitating different remedies to cope with different types of distortion. The ability effect tends to be stronger when agents differ more in the ability to identify promising projects. We argue that this is more likely in situations where: (i) an agent has discretion over which project to work on; or (ii) exploration of new ideas, rather than exploitation of old ideas, is more important. A prominent example which fits this description is academia, where researchers are conferred substantial discretion over what to do, and it is the novelty of ideas, rather than the efficiency of task implementation, which is indispensable for success. In this type of environment, separating equilibria are more likely to emerge, and the timing of project abandonment reflects the agent's private information. In contrast, when an agent simply undertakes tasks that are assigned to them or it is the efficiency of task implementation that is more tested, the learning effect becomes more dominant. In this case, pooling equilibria are more likely, and the timing of project abandonment conveys no useful information.

Whether equilibrium entails some separation or not gives some important insight for who should have the right to abandon projects. To address this issue, suppose that there is a principal (supervisor, investor) who decides whether to delegate this right to the agent or retain it to herself. It is clear in this context that delegation has no value if a pooling equilibrium is expected to prevail because the timing of project abandonment does not reflect the agent's private information; the principal can do no worse by retaining the right and simply quitting at her optimal timing. This is of course not the ideal strategy for the principal because the project is terminated once and for all independently of the agent's ability. When some separation is expected in equilibrium, therefore, there is scope for delegation to enable the principal to induce the agent's information. This argument then predicts that the right to abandon projects should be more delegated in situations where the ability effect is more important.

The presence of public information invites further complication when the reward upon success depends on the agent's reputation, in which case the agent must be concerned not only about the timing of project abandonment (which is his strategic choice) but also about the timing of success (which arrives stochastically over time and is beyond his control). This random nature of risky experimentation gives rise to dynamic reputation concerns that have not been discussed in existing reputation models of experimentation. We refer to this setup, where the reward upon success is contingent on the agent's reputation, as the dynamic case, in order to distinguish it from the static case where success yields a fixed reward. The equilibrium allocation in the dynamic case is substantially more complicated because the expected payoff now depends on the agent's evolving belief. As low-ability agents abandon the project, the reputation for those who remain and succeed becomes higher, therefore inducing these low-ability agents not to abandon so soon. This war-of-attrition feature of the model implies that equilibrium will entail continuous randomization by low-ability agents. Nevertheless the main insight derived from this is relatively clear: high-ability agents quit prematurely while low-ability agents persist excessively. The presence of public news thus raises additional efficiency and policy implications. We later draw on venture startups, which constitute a leading example of our framework, and discuss some potential remedies that are particularly relevant for this industry.

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 (with jump processes) builds on exponential/Poisson bandits because of its simplicity and tractability. It is perhaps safe to say that this assumption, which implies that success arrives at a constant rate irrespective of how much one has worked in the past, is often too restrictive from a practical point of view. It is thus important to see which predictions are robust to alterations in the underlying learning process. Our analysis suggests that signaling environments provide one context in which 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.

Some recent works introduce reputation concerns into experimentation models as we do (Bobtcheff and Levy, 2017; Thomas, 2019; Halac and Kremer, 2019).[3] Aside from the

[3] Halac and Kremer (2019) are different from the other two works as well as from ours in that the agent receives a flow payoff (wage) that is contingent on his reputation. In a different vein, Bonatti and Hörner

fact that they all build on exponential/Poisson bandits, an important distinction is that agents in those models are heterogeneous only in one particular dimension: using our terminology, they differ only in the ability to identify projects (Thomas, 2019; Halac and Kremer, 2019) or in the speed of learning (Bobtcheff and Levy, 2017). In contrast, our model encompasses and integrates these two distinct notions of reputation and clarify the differences it makes.[4]

When agents differ in the speed of learning, the learning effect kicks in and the single-crossing property may break down as a consequence.[5] In this sense, Bobtcheff and Levy (2017) shares an important commonality with the static version of our model. They consider an environment in which a liquidity-constrained agent learns the potential value of a project at a privately known speed and decides when to make investment. In their setup, as in ours, the single-crossing property may fail to hold because of the learning effect, although they follow a different characterization approach.[6] An important difference is that their model is a private-news model where 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.[7] This draws clear contrast to our setting where the timing of success 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 timing of project abandonment.

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).[8] Among them, our model is more closely related to Daley

(2017) also analyze a career concerns model with exponential learning which incorporates a moral hazard component.

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

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

[6] Bobtcheff and Levy (2017) restrict their attention to the case where the difference in the expected payoffs stems only from the difference in the option values of experimentation, and exploit this feature to characterize equilibria.

[7] To be more precise, whether news is public or private makes no difference in their setup because the agent never invests if he observes a bad signal. On the other hand, Thomas (2019) incorporates both public and private news, but the reward to success (or failure) in her model is independent of the agent's reputation (as in the static case of our model).

[8] Daley and Green (2014) consider a static signaling model with "grades" which serve as an additional information source.

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. A crucial assumption in their model is that there is an exogenous public information process, called "news," which stochastically reveals the seller's type via a Brownian diffusion process. As in our model, therefore, their model has two sources of information: the timing of trade and news about the seller's asset.[9] 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, which plays an essential role in generating the double-crossing property in our analysis.

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, and an agent's status depends on public perceptions of his predispositions. It is shown 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.[10]

2. The Model

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

[9] It should be noted, however, that only one of the two sources is available in our model while both of them are always available in theirs.

[10] In addition, their model assumes symmetric information, so that it is not a signaling model. Less related are models of informational herding such as Bikhchandani et al. (1992) and Banerjee (1992). They are driven purely by information externality and social learning with no reputation concerns.

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. This assumption implies that, if the project quality is good, the probability that success arrives before any given time t is greater for the high type than for the low type. For example, in the exponential bandit model (Keller et al., 2005), a higher Poisson arrival rate of success would imply that the likelihood ratio, $f_H(\tau)/f_L(\tau) = (\lambda_H/\lambda_L)e^{-(\lambda_H-\lambda_L)\tau}$, decreases in τ for $\lambda_H > \lambda_L$.

Neither the agent nor anyone else knows the quality of the project that he is undertaking. The common prior belief that the project quality is good is p_0. We will later allow the prior to depend on agent type, but for now assume that it is type-independent. The agent knows his own type, but this private information is unavailable to the labor 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 also 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 $\sigma_i : [0, \infty) \to [0, 1]$, with $\sigma_i(t)$ being the probability that the agent is still working on the risky project at time t. 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.) We assume that the agent's decision at each point in time is publicly observable. The arrival of success is publicly observable, representing an additional source of information.

The game ends either when the risky project generates success, or when the agent abandons the project without success. 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 present value paid to an agent of type i after he achieves success. For example, the agent may become an entrepreneur or manager after obtaining success with his risky project, and W_H and W_L (with $W_H \geq W_L$) are, respectively, the expected lifetime earnings of a high-type and a low-type entrepreneur or manager. We also allow that, in addition to the post-project income from the competitive market, the risky project can produce some reward to the agent when it generates success, which is included in the definition of W_i. If the agent abandons the project without success, he switches to work in the labor market and is paid a sum that reflects his expected productivity as a worker. We let w_H and w_L (with $w_H > w_L$) represent the present value of productivity of a high-type and a low-type in the labor market. Therefore, if the project generates success at time τ, the agent's payoff is $W(\tau) = q_\tau W_H + (1 - q_\tau) W_L$. If the agent's reputation at time t is q_t, his outside opportunity from working in the labor market is $w(t) = q_t w_H + (1 - q_t) w_L$.

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 shortly, the difference between the two cases proves to be critical, and we analyze these cases separately. Finally, we also assume that achieving success is always better than the outside opportunity of working in the labor market. Specifically,

$$W_L > w_H. \tag{1}$$

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:

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

where $C(t) \equiv c(e^{\rho t} - 1)/\rho$ is the accumulated cost of working with the risky project for a period of length t. 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 $(\hat{s}_H, \hat{s}_L)$ whenever applicable. As is typical in signaling models, these requirements do not pin down a unique equilibrium. We adopt the D1 criterion (Banks and Sobel, 1987; 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.*

There are two features of our model that distinguish it from a standard signaling model. First, we will show 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. Instead, given the double-crossing property, D1 tends to produce equilibrium outcomes that exhibit pooling. Second, the decision to engage in risky experimentation is not a static decision made once and for all. Instead the agent's decision to continue working with the project or not is determined partly by the market belief about his type, which evolves over time. This dynamic consideration is relevant when the reward to success depends on the agent's reputation (i.e., when $W_H > W_L$), but is not relevant when the reward to success is fixed (i.e., when $W_H = W_L$). To separate the second feature from the first, we first consider the model 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 issues arising from dynamic considerations.

3. Double-Crossing Property and Static Signaling Equilibrium

3.1. Reputation concern

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, since the reward to success does not depend on the market's evolving belief about the agent's type, the stopping decision of the agent can be analyzed as if it were a static decision. We call the case of $W_H = W_L$ the "static" case.

The updated belief about an agent's ability type at the time of project abandonment depends on (1) inferences based on the agent's choice (i.e., to continue working on the risky project or not) and its consistency with the equilibrium strategies of the two types; and (2) observation about the timing of success τ. We use $\hat{q}$ to denote the *interim belief* based on equilibrium inference alone, and consider how inference based on the timing of success further modifies the interim belief.

Let $q_t = r(t; \hat{q})$ represent the belief about an agent's type when he abandons the project without success at time t (i.e., $\tau > t$) and when the interim belief about such a quitter's type is $\hat{q}$. We have

$$r(t;\hat{q}) = \frac{\hat{q}(1-p_0F_H(t))}{\hat{q}(1-p_0F_H(t))+(1-\hat{q})(1-p_0F_L(t))}.$$

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 success 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

$$g_i(t) = \frac{p_0 f_i(t)}{1-p_0F_i(t)}$$

is the hazard rate of success in the risky project for type $i \in \{H, L\}$. When there is no uncertainty about project quality (i.e., when $p_0 = 1$), monotone likelihood ratio property implies that $g_H(t)$ is greater than $g_L(t)$, so that the agent's reputation is lower if he abandons the project at a later time, because high-type agents are expected to have achieved success early. However, this is no longer always true when the quality of the risky project is uncertain. To see why, we can decompose the hazard function into two parts:

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

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 $i \in \{H, L\}$. Moreover 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 $\lim_{t\to\infty} f_H(t)/f_L(t)$, 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 $g_H(t)/g_L(t)$ with respect to t, when evaluated 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, the learning effect becomes stronger if the agent fails to achieve success after trying for a long time. Lemma 1 shows that the learning effect dominates the ability effect whenever the agent abandons the project beyond time $\hat{t}$. The unconditional hazard rate for achieving success is smaller for the high type than for the low type beyond $\hat{t}$ because the project quality is likely to be bad if a high-type agent has been working on this project for very long without success. It is this interaction between the two effects which makes it ambiguous which type of agent has an incentive to quit earlier. In this environment, it is not clear whether perseverance is a sign of strength or weakness.

3.2. Double-crossing indifference curves

Because of the counteracting ability effect and learning effect, the single-crossing property used in standard signaling models does not hold in our setup. Thus it is not immediately obvious whether a high-type agent can signal his type by abandoning a risky project early or by abandoning it late. On one hand, quitting early can potentially signal high type because a high-type agent learns quickly through his failure to achieve success that his risky project is not very promising. On the other hand, quitting late can also potentially signal high type because a high-type agent is more confident of his ability to achieve success from a good project. We analyze the implications of these incentives for equilibrium strategies in this section.

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

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

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

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

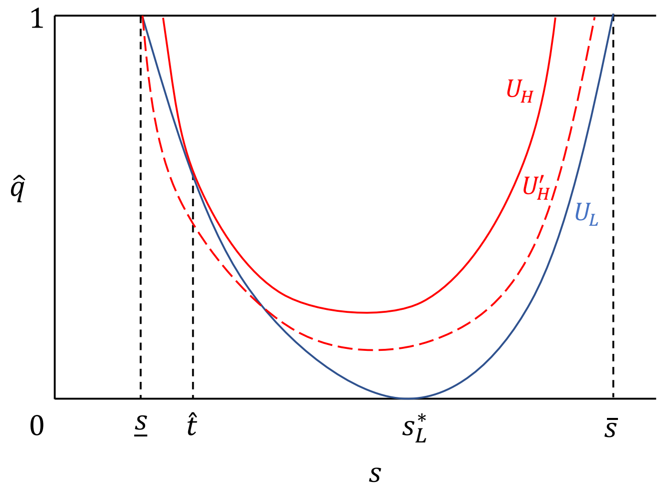

Figure 1. *The indifference curve U_H' for the high type crosses the indifference curve U_L twice, once from above to the left of $\hat{t}$ and from once below to the right of $\hat{t}$. The indifference curves of the two types are tangent to one another at $s = \hat{t}$.*

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. Therefore, Lemma 1 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 this model does not exhibit the single-crossing property used in most signaling models. See Figure 1 for illustration.

Because $g_H(t) - g_L(t)$ is single-crossing *from above* at $t = \hat{t}$, this implies that the marginal rate of substitution $MRS_H(s,\hat{q})$ for the high type is more sensitive to changes in the quitting time s at $s = \hat{t}$ than is the marginal rate of substitution for the low type. In Figure 1, the indifference curve U_H is "more convex" than the indifference curve U_L at the tangency point at $\hat{t}$. Formally, we have the following result.

Lemma 2. *For any interim beliefs $\hat{q}$ and $\hat{q}'$ and any stopping time $s' \neq \hat{t}$,*

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

Proof. For $i \in \{H, L\}$, let $\phi_i(s)$ represent the solution to the differential equation, $\phi_i'(s) = -MRS_i(s,\phi_i(s))$, with initial condition $\phi_i(\hat{t}) = \hat{q}$. Lemma 1 implies that $\phi_L'(\hat{t}) - \phi_H'(\hat{t}) = 0$ and $\phi_L''(\hat{t}) - \phi_H''(\hat{t}) < 0$. This means that $\phi_L(s) - \phi_H(s)$ reaches a local maximum of 0 at

$s = \hat{t}$. Therefore, there exists $\epsilon > 0$ such that $\phi_L(s') < \phi_H(s')$ for all $s' \in [\hat{t} - \epsilon, \hat{t} + \epsilon]$ and $s' \neq \hat{t}$. Suppose now $s' > \hat{t} + \epsilon$. By Lemma 1, $-MRS_L(s,q) < -MRS_H(s,q)$ for any q and any $s \in [\hat{t} + \epsilon, s']$. Since $\phi_L(\hat{t} + \epsilon) < \phi_H(\hat{t} + \epsilon)$, the comparison theorem for differential equations (e.g., McNabb, 1986) implies that $\phi_L(s') < \phi_H(s')$. Similarly, if $s' < \hat{t} - \epsilon$, then $-MRS_L(s,q) > -MRS_H(s,q)$ for any q and any $s \in [s', \hat{t} - \epsilon]$. Since $\phi_L(\hat{t} - \epsilon) < \phi_H(\hat{t} + \epsilon)$, the same comparison theorem again implies that $\phi_L(s') < \phi_H(s')$.

Note that $U_H(s', \hat{q}') = U_H(\hat{t}, \hat{q})$ if and only if $\hat{q}' = \phi_H(s')$. Because $\phi_H(s') > \phi_L(s')$, we have

$$U_L(s', \hat{q}') = U_L(s', \phi_H(s')) > U_L(s', \phi_L(s')) = U_L(\hat{t}, \hat{q}). \qquad \blacksquare$$

Lemma 2 has an important implication for off-equilibrium inference under the D1 criterion. 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'.

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

$$g_i(s_i^*)[W - w_i] - \rho w_i - c = 0. \tag{5}$$

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 and the third terms are the opportunity cost and direct cost of doing so.

Lemma 3. *If* $\max\{s_H^*, s_L^*\} > \hat{t}$, *then* $s_L^* > s_H^*$.

Proof. Because $w_H > w_L$, we have $g_H(s_H^*) > g_L(s_L^*)$. If $s_H^* > \hat{t}$, Lemma 1 implies that $g_L(s_H^*) > g_H(s_H^*) > g_L(s_L^*)$, which implies $s_L^* > s_H^*$ because $g_L(\cdot)$ is strictly decreasing. To show that $s_L^* > \hat{t}$ implies $s_L^* > s_H^*$, suppose $s_H^* \geq s_L^*$ instead. But we have already shown that $s_H^* \geq s_L^* > \hat{t}$ implies $s_L^* > s_H^*$, a contradiction. $\blacksquare$

The value of $\hat{t}$ is determined entirely by the statistical properties of the arrival time of success, while s_i^* depends also on the reward and cost of achieving success. Other things equal, raising the reward W from the risky project will raise both s_H^* and s_L^*. For sufficiently large W, we will have $\max\{s_H^*, s_L^*\} > \hat{t}$. Lemma 3 therefore implies that the low type would stop later than the high type if the reward to the risky project is sufficiently large.

3.3. Signaling equilibrium in the static case

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 $\bar{s}$.[11] To focus on the interesting case, we assume that the incentive compatibility constraint for the low type is binding, in the sense that $U_L(s_L^*, 0) < U_L(s_H^*, 1)$. This is equivalent to $s_H^* \in (\underline{s}, \bar{s})$.[12]

Consider first the case of fully separating equilibrium.

Proposition 2. *Suppose $W_L = W_H$ and $s_H^* \in (\underline{s}, \bar{s})$.*

(a) If $\hat{t} \leq \underline{s}$, the equilibrium is fully 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 fully 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 in the relevant region, for any $\hat{q}$, the high type prefers 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 in the relevant region, but with the high type having a stronger incentive to quit later. 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. ∎

[11] It 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.

[12] When the incentive constraint is not binding, the full-information outcome is the unique equilibrium outcome in this model.

If $\hat{t}$ is very high or very low, our model produces full separation between types. When low types have incentive to mimic high types, the latter separate themselves by abandoning the project earlier than the optimal quitting time s_H^* if $\hat{t}$ is very low, or by abandoning the project later than s_H^* if $\hat{t}$ is very high. Although equilibrium is fully separating in both two 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 a standard signaling model in which the single-crossing property holds.

For intermediate values of $\hat{t}$, our model produces pooling (or semi-pooling) between the different types.

Proposition 3. *Suppose $W_L = W_H$, $s_H^* \in (\underline{s}, \bar{s})$, and $\hat{t} \in (\underline{s}, \bar{s})$.*

(a) If $U_L(\hat{t}, q_0) \geq U_L(s_L^*, 0)$, the equilibrium is complete 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, note that there cannot be a semi-pooling equilibrium in which the high type randomizes between quitting at $\hat{t}$ and at some other time t'. This follows from Lemma 2, which establishes that whenever the high type is indifferent between quitting at $\hat{t}$ and t', the low type strictly prefers quitting at t' to quitting at $\hat{t}$. 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 $\hat{t}$, by D1, the market assigns interim belief $\hat{q} = 0$ to an agent who quits at $t \neq \hat{t}$. Given such an interim belief, if a semi-pooling equilibrium ever exists, the low type must quit at

s_L^*. This leaves us with two possible types of equilibrium: (1) complete pooling in which both types quit at $\hat{t}$; or (2) semi-pooling in which the low type quits at both $\hat{t}$ and s_L^* with positive probability, and the high type quits once and for all at $\hat{t}$.

If $U_L(\hat{t}, q_0) \geq U_L(s_L^*, 0)$, the semi-pooling equilibrium cannot exist, because $\hat{q}_{\hat{t}} > q_0$, and thus $U_L(\hat{t}, \hat{q}_{\hat{t}}) > U_L(s_L^*, 0)$, meaning that the low type strictly prefers quitting at $\hat{t}$ to quitting at s_L^*. 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.

Similarly, if $U_L(\hat{t}, q_0) < U_L(s_L^*, 0)$, the complete pooling equilibrium cannot exist, because in such an equilibrium the low type could profitably deviate by quitting at s_L^* instead. Therefore, given that there exists a unique $q \in (q_0, 1)$ such that $U_L(\hat{t}, q) = U_L(s_L^*, 0)$, the equilibrium must be unique. In equilibrium, the high type quits at $\hat{t}$ with probability 1, and the low type does so with some positive probability, so that the interim belief about an agent who quits at $\hat{t}$ is 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 satisfy the double-crossing property rather than the single-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. 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. Dynamic 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 static case discussed in Section 3. We label the case of $W_H > W_L$ the "dynamic" case.

Let $q_t = R(t; \tilde{q})$ represent the belief about an agent's type when he stays with the risky project up to time t and gets a success at that time (i.e., $\tau = t$) and when the interim belief about such a stayer's type is $\tilde{q}$.[13] We have

$$R(t; \tilde{q}) = \frac{\tilde{q} p_0 f_H(t)}{\tilde{q} p_0 f_H(t) + (1-\tilde{q}) p_0 f_L(t)}.$$

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 than or lower than the reputation upon failure to obtain success. In particular, it is straightforward to show that, for any interim belief q, $R(t; q) > r(t; q)$ if and only if $t < \hat{t}$, and $R(t; q) < r(t; q)$ if and only if $t > \hat{t}$. In other words, success that comes too late may be worse for an agent's reputation than a decision to abandon the project. This is due to the learning effect dominating the ability effect, as the probability that a high-type agent is still working on a good quality project after a protracted period of time is quite low.

For the case $W_H > W_L$, the objective function for type $i \in \{H, L\}$ given by equation (2) can be written as $U_i(s, \hat{q}, \tilde{q})$, where the two interim beliefs enters into the agent's payoff because $w(s) = w_L + r(s; \hat{q})(w_H - w_L)$ and $W(\tau) = W_L + R(\tau; \tilde{q})(W_H - W_L)$. For any fixed $\tilde{q}$, the indifference curves of the two types in the $(s, \hat{q})$-space cross twice, with $MRS_H(s, \hat{q}; \tilde{q}) > MRS_L(s, \hat{q}; \tilde{q})$ if and only if $s < \hat{t}$ and $MRS_H(s, \hat{q}; \tilde{q}) < MRS_L(s, \hat{q}; \tilde{q})$ if and only if $s > \hat{t}$. This double-crossing property leads to the following result.

Lemma 4. *If both types of agent abandon the risky project at some time t with positive probability in equilibrium, then $t = \hat{t}$.*

The logic of this result is the same as that for Proposition 3 in the static case, and the proof is relegated to the Appendix. Lemma 4 implies that a complete pooling equilibrium in the dynamic case must have both types quit at $\hat{t}$, just like in the static case. Nevertheless, the dynamic case is different from the static case whenever equilibrium entails some separation between the two types.

To see why dynamic considerations matter, note that the instantaneous benefit from continuing to work with the risky project for a small time interval of length $\mathrm{d}t$ for an agent of type $i \in \{H, L\}$ is

$$g_i(t)\left[W_L + R(t; \tilde{q})(W_H - W_L)\right] \mathrm{d}t,$$

[13] At a given time, $\tilde{q}$ can be equal to or different from $\hat{q}$. The former is the interim belief about an agent who stays with working on the risky project up to a certain time t and continues to stay at time t, while the latter is the interim belief about an agent who stays up to time t and then quits at that time.

which depends on the value of $\tilde{q}$ when $W_H > W_L$.[14] 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 earlier than the low type does, 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. On the other hand, if the low type quits earlier than the high type does, then $\tilde{q}$ rises as the low type quits, and this can raise the incentive of the remaining agents to continue working with the risky project.

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

$$g_i(s_i)\,[W_L + R(s_i;\tilde{q})(W_H - W_L) - w_i] - \rho w_i - c = 0. \tag{6}$$

Note that the s_L^* defined in equation (5) for the static case is the same as $s_L^*(0)$; and similarly, s_H^* in the static case the same as $s_H^*(1)$. For the dynamic case, equation (6) 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 (6) decreases in s_i and increases in $\tilde{q}$, $s_i^*(\tilde{q})$ is increasing in $\tilde{q}$. A higher interim reputation for stayers tends to delay the quitting time if quitting would reveal an agent's type.

The analysis of signaling equilibrium in the dynamic case 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 $s_L^*(q_0) \geq \hat{t}$. By the same logic behind Lemma 3 for the static case, one can verify that $\max\{s_H^*(q), s_L^*(q)\} > \hat{t}$ implies $s_L^*(q) > s_H^*(q)$ for any q in the dynamic case. Hence, $s_L^*(q_0) \geq s_H^*(q_0)$ for the case considered here.

[14] If the outcome of experimentation can only be observed privately, but the reward is type-contingent, then the instantaneous benefit is given by $g_i(t)W_i\,\mathrm{d}t$. In this case, $\hat{t}$ can no longer be pinned down entirely from the properties of g_i but also depends on W_i. Aside from this, since there are no dynamic concerns, the analysis of the private-news case is quite similar to that of the static case.

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

$$v(t) = \int_t^{s_L^*(0)} e^{-\rho(\tau-t)} \frac{p_0 f_L(\tau)}{1-p_0F_L(t)} [W_L - C(\tau-t)]\, \mathrm{d}\tau$$
$$+ e^{-\rho(s_L^*(0)-t)} \frac{1-p_0F_L(s_L^*(0))}{1-p_0F_L(t)} \left[-C(s_L^*(0)-t) + w_L\right].$$

Recall that $s_L^*(0)$ is the optimal stopping time when an agent is known to be a low type. Conditional on the project not having achieved success by time t, the function $v(t)$ gives the maximum payoff (from the perspective of time t) from continuing with the project given that the agent is known to be a low type. We note that $v(t)$ is the solution to the differential equation:

$$v'(t) = \rho v(t) + c - g_L(t)[W_L - v(t)],$$

with terminal condition $v(s_L^*(0)) = w_L$. It is straightforward to verify that $v'(t)$ is single-crossing from below, and that the definition of $s_L^*(0)$ implies that $v'(t) = 0$ when $t = s_L^*(0)$. This implies that $v(t)$ strictly decreases in t for $t < s_L^*(0)$.

Define $\underline{t}$ such that $v(\underline{t}) = w_H$. (If no such $\underline{t}$ exist, we let $\underline{t} = 0$.) When $s_H^*(q_0) \leq \underline{t}$, we have $v(s_H^*(q_0)) > w_H$. The incentive compatibility constraint for the low type is not binding, because the low type prefers to continue working with the risky project until time $s_L^*(0)$ to mimicking the high type by 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.

Proposition 4. *Suppose $W_H > W_L$, $s_L^*(q_0) \geq \hat{t}$, and $s_H^*(q_0) > \underline{t}$.*

(a) If $\hat{t} \leq \underline{t}$, then the equilibrium is fully separating, with the high type quitting at $\underline{t}$ and the low type quitting at $s_L^*(0)$.

(b) If $\hat{t} \in (\underline{t}, s_L^*(0))$ and $v(\hat{t}) > w_L + r(\hat{t}; q_0)(w_H - w_L)$, then the equilibrium is semi-pooling, with the high type quitting at $\hat{t}$ and the low type randomizing between quitting at $\hat{t}$ and $s_L^*(0)$.

(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 $\hat{t} \in [s_L^*(0), s_L^*(q_0)]$, then the equilibrium is complete pooling, with both types quitting at $\hat{t}$.

Case (a) of Proposition 4 describes a least-cost separating equilibrium, as the indifference curves satisfy the single-crossing property in the relevant region when $\hat{t} \leq \underline{t}$. In cases (b) and (c), double-crossing indifference curves induce 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 complete pooling equilibrium. Because the logic of Proposition 4 follows the same reasoning as that for Propositions 2 and 3 in the static case, we omit the proof of Proposition 4 for brevity.

4.2. Low type quits first

We now consider the case where $s_L^*(q_0) < \hat{t}$ in this subsection. Note that Lemma 3 implies $s_H^*(q_0) \leq \hat{t}$ for this case.

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. Lemma 4 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 $t \geq s_L^*(q_0)$.

To determine the rate of exit by the low type, note that 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:

$$g_L(t)\left[W_L + R(t;\tilde{q}(t))(W_H - W_L) - w_L\right] - \rho w_L - c = 0. \tag{7}$$

Equation (7) holds at $t = s_L^*(q_0)$ and $\tilde{q}(t) = q_0$. The left-hand-side of (7) is decreasing in t and increasing in $\tilde{q}(t)$. Thus, as t increases beyond $s_L^*(q_0)$, $\tilde{q}(t)$ must rise to maintain the indifference condition. The interim belief $\tilde{q}(t)$ satisfies:

$$\tilde{q}(t) = \frac{q_0}{q_0 + (1 - q_0)\sigma_L(t)}. \tag{8}$$

As t increases from $s_L^*(q_0)$ to $s_L^*(1)$, $\tilde{q}(t)$ increases continuously from q_0 to 1 according to equation (7), and $\sigma_L(t)$ decreases continuously from 1 to 0 according to equation (8).

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

$$V(t;\hat{t}) = \int_t^{\hat{t}} e^{-\rho(\tau-t)} \frac{p_0 f_L(\tau)}{1-p_0F_L(t)} [W_L + R(\tau;\tilde{q}(t))(W_H - W_L) - C(\tau - t)]\, \mathrm{d}\tau$$
$$+ e^{-\rho(\hat{t}-t)} \frac{1-p_0F_L(\hat{t})}{1-p_0F_L(t)} \left[-C(\hat{t}-t) + w_L + r(\hat{t};\tilde{q}(t))(w_H - w_L)\right],$$

where $\tilde{q}(t)$ follows equation (7) for $t \le s_L^*(1)$ and is equal to 1 for $t > s_L^*(1)$. The function $V(t;\hat{t})$ gives the payoff to a low type if all low types stop quitting at time t and continue working on the risky project until time $\hat{t}$, at which point agents of all types quit. Whether a low type agent will quit at some rate beginning at time $s_L^*(q_0)$ or not depends on whether $V(s_L^*(q_0);\hat{t})$ is less than or greater than w_L. If $V(s_L^*(q_0);\hat{t}) \ge w_L$, then all low types prefer to wait and pool with high types at time $\hat{t}$ than to quit at $s_L^*(q_0)$. There will be a complete pooling equilibrium. If $V(s_L^*(q_0);\hat{t}) < w_L$, then some separation will occur in equilibrium.

Proposition 5. *Suppose $W_H > W_L$ and $s_L^*(q_0) < \hat{t}$.*

(a) If $V(s_L^*(q_0);\hat{t}) \ge w_L$, then the equilibrium is complete pooling with both types quitting at $\hat{t}$.

(b) If $V(s_L^*(q_0);\hat{t}) < w_L$, then 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 with the high type quitting at $\hat{t}$ with probability 1*. For the low type, if $t < s_L^*(q_0)$, then $\sigma_L(t) = 1$; if $t \in [s_L^*(q_0), t_0]$, then $\sigma_L(t)$ is determined by equations (7) and (8); if $t \in (t_0, \hat{t})$, then the low type staya and $\sigma_L(t) = \sigma_L(t_0)$; at $t = \hat{t}$ the low type quits with probability* 1 *and $\sigma_L(t) = 0$ for $t \ge \hat{t}$.*

(ii) If $t_0 \ge s_L^*(1)$, there exists a unique $t_1 \in (s_L^*(1), \hat{t})$ such that $V(s_L^*(1); t_1) = w_L$. The equilibrium is fully separating with the high type quitting at $\max\{t_1, s_H^*(1)\}$*. For the low type, if $t < s_L^*(q_0)$, then $\sigma_L(t) = 1$; if $t \in [s_L^*(q_0), s_L^*(1)]$, then $\sigma_L(t)$ is determined by equations (7) and (8); and if $t > s_L^*(1)$, then $\sigma_L(t) = \sigma_L(s_L^*(1)) =$* 0.

The strategy of the low type described in Proposition 5 is qualitatively different from that described in Proposition 4 of the earlier subsection. In the case of $s_L^*(q_0) \ge \hat{t}$, Proposition 4 shows that the low type always quits at $s_L^*(q_0)$ or earlier. In the case of $s_L^*(q_0) < \hat{t}$ in

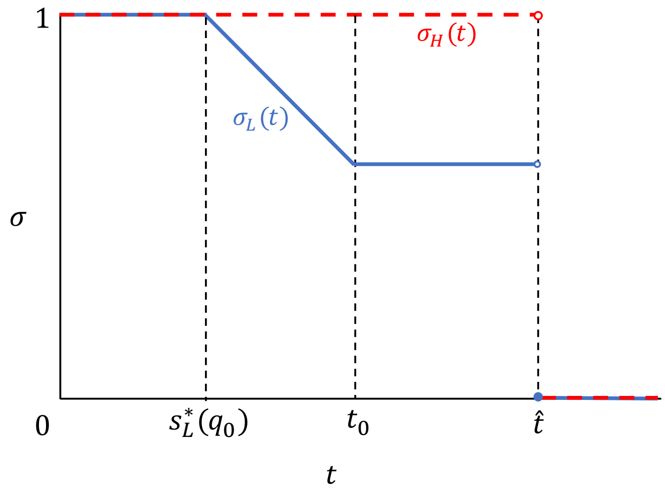

Figure 2. *In the semi-pooling equilibrium described in case (b)(i) of Proposition 5, low types are quitting continuously at a positive rate between $s_L^*(q_0)$ and t_0. The remaining low types strictly prefer to stay between t_0 and $\hat{t}$. At time $\hat{t}$, a positive mass of low types quit to pool with the high types.*

this subsection, Proposition 5 shows that the low type quits after $s_L^*(q_0)$. Moreover, when low types separate from high types by quitting first, as described in part (b) of Proposition 5, they can only quit at a positive flow rate but not with positive probability. See Figure 2 for an illustration. This is because the payoff from continuing with the risky project increases as the low types quit, so continuous randomization is needed to sustain equilibrium.

We leave the details of the proof of Proposition 5 to the Appendix. Intuitively, if $V(s_L^*(0);\hat{t}) < w_L$, the low type would prefer to quit than to stay at $t = s_L^*(q_0)$. But our discussion suggests that the low type can only quit gradually and continuously at some atomless rate. As this happens the interim belief $\tilde{q}(t)$ about stayers rises, and the payoff $V(t;\hat{t})$ from staying until $\hat{t}$ improves over time. For time t beyond t_0 specified in part (b) of Proposition 5, we have $V(t;\hat{t}) > w_L$. For $t \in (t_0, \hat{t})$, the low type strictly prefers to stay because the prospective gain from staying until $\hat{t}$ to pool with the high type exceeds the immediate losses from continuing with the risky project.

If $\hat{t}$ is very large, the t_0 that satisfies $V(t_0;\hat{t}) = w_L$ will also be very large. Whenever $t_0 \geq s_L^*(1)$, the rate of exit determined by equations (7) and (8) implies that all low types would have quit by time $s_L^*(1)$ (i.e., $\sigma_L(s_L^*(1)) = 0$). 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 $t' \geq t_1$ even if $\hat{q}_{t'} = 1$ is assigned to that deviation. Since the high type can always benefit from moving towards $\max\{t_1, s_H^*(1)\}$, the only fully 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—success 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 for this dynamic case 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). The precise way the dynamic inefficiency works is determined by which type is willing to persist longer, as we highlight below.

Suppose that the high type quits before the low type. In this case, by Proposition 4, the high type quits at $\max\{\underline{t}, \hat{t}\}$, 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. Note that when there is full separation in the static case, the low type quits once and for all at $s_L^*(0)$. This is no longer true in the dynamic case, because 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 types around. Moreover, when a low-type agent quits, it raises the interim belief for 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 generically exists no efficient (full-information) equilibrium in the dynamic case, 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 static case, where efficient separation is feasible as long as s_L^* and s_H^* are sufficiently far apart from each other.

4.4. Equilibrium payoffs

Let U_i^* denote the equilibrium payoff for type $i \in \{H, L\}$. Also, let U_i^{**} denote the full-information payoff, i.e., the payoff when the low type quits at $s_L^*(0)$ and the high type quits at $s_H^*(1)$. The argument in the previous subsection implies that the full-information payoffs cannot be achieved in equilibrium when $W_H > W_L$. A closer look at this reveals who gains and who loses from dynamic reputation concerns.

Proposition 6. *If $W_H = W_L$, then $U_H^* \leq U_H^{**}$ and $U_L^* \geq U_L^{**}$. If $W_H > W_L$, then $U_H^* < U_H^{**}$ and $U_L^* > U_L^{**}$.*

Proof. In general, a low-type agent can always quit at $s_L^*(0)$ and reveals his true type. If $W_H = W_L$, this would give the low type exactly the full-information payoff, and hence $U_L^* \geq U_L^{**}$ as there would be a profitable deviation otherwise. That the high type is weakly worse off directly follows from this, because otherwise there would be an allocation which Pareto dominates the full-information efficient benchmark, a contradiction.

We can essentially apply the same argument to the dynamic case, except that the payoff for the low type of quitting at $s_L^*(0)$ is strictly larger than U_L^{**} because of a higher reputational value of success up to that point. This assures $U_L^* > U_L^{**}$ and by the same argument $U_H^* < U_H^{**}$. ■

Proposition 6 suggests that the high type unambiguously loses as the low type can benefit at the expense of the high type, and this tendency is even stronger in the dynamic case where the high type can never realize the full-information payoff. This can potentially be a matter of concern from the efficiency point of view, if the participation constraint is binding for some high-type agents. Given that high-type agents have a better outside option, they may choose not to enter the market if U_H^* is not high enough, possibly to the extent that the market collapses entirely.[15] This possibility thus calls for some active intervention to restore efficiency. We will discuss this issue and ways to alleviate this problem in Section 5.3.

[15] This policy implication must be interpreted with some caution, however, as the effect of having more high-type agents is not necessarily positive in this setup, as we will note in the next section.

5. Discussion

5.1. The role of prior reputation

In the canonical signaling model, the D1 criterion always selects the least-cost separating equilibrium, or the Riley outcome, in which the low type selects 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, which allows us to derive some efficiency implications.

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. We call this the static effect of the prior belief because this effect is present even when $W_H = W_L$.

In the dynamic version of our model, there arises an additional effect of the prior belief when a separating equilibrium prevails. This dynamic effect can work either positively or negatively depending on which type quits first. Consider a fully separating equilibrium in which the high type quits first. In this equilibrium, the high type quits at $s_H^*(q_0)$, which is smaller than the full-information optimal level $s_H^*(1)$, while the low type quits at $s_L^*(0)$. The extent of inefficiency thus diminishes as q_0 increases. In a fully 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 larger than $s_L^*(0)$. An increase in q_0 thus further delays the low type's exit and makes the dynamic separation of types less complete, which may result in inefficient pooling if the low type's incentive to stay is excessively strong.

To sum up, in our model, the equilibrium selected by D1 is sensitive to the agent's prior reputation. The effect of an increase in q_0 (having "better reputation") 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, possibly to the extent that only pooling equilibria can be sustained. 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. Our model in particular suggests that high types tend to quit prematurely while low types tend to over-experiment, which makes the correlation between ability and success weaker. 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).

5.2. Implementation ability versus identification ability

We now consider an extension of our baseline model to allow for type-dependent project quality. Specifically, let p_0^i denote the prior probability that a project handled by type i is of good quality. Here, we assume that $1 > p_0^H \geq p_0^L > 0$; i.e., the high type is possibly better at discovering ideas or identifying promising projects. 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 under this specification, Lemma 1 continues to hold for any pair (p_0^L, p_0^H) such that $1 > p_0^H \geq p_0^L > 0$. This means that the double-crossing property still holds, and we can essentially 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 carries more importance depends on the underlying 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. Alternatively, we also argue that identification ability matters more in areas where exploration of new ideas is required whereas implementation ability matters more in areas where exploitation of existing ideas is sufficient. Note that for any $p_0^H \geq p_0^L$, the 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, by letting λ_H increase from λ_L to infinity, to show that the equilibrium allocation depends crucially on what is captured by the agent's "reputation."

As we have seen, 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

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

under the current specification. The following result is useful for our subsequent analysis; the proof is in the Appendix.

Lemma 5. *For any λ_L and any $p_0^H \geq p_0^L$, $\hat{t}(\lambda_H)$ strict decreases in λ_H, with $\lim_{\lambda_H \to \lambda_L} \hat{t}(\lambda_H) = \infty$ and $\lim_{\lambda_H \to \infty} \hat{t}(\lambda_H) = 0$.*

Lemma 5 suggests that $\hat{t}$ 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, which clarifies when the equilibrium entails some separation. We leave the proof to the Appendix.

Proposition 7. *For any λ_L, q_0, and any $p_0^H \geq p_0^L$, there exist $\overline{\lambda}$ and $\underline{\lambda}$ (with $\overline{\lambda} > \underline{\lambda} > \lambda_L$) 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 $\lambda_H \in [\underline{\lambda}, \overline{\lambda}]$, the equilibrium entails complete pooling, and both types quit at $\hat{t}$;*
- *(c) if $\lambda_H > \overline{\lambda}$, the equilibrium entails some separation, and the high type quits first.*

Moreover, $\overline{\lambda} \to \infty$ if $s_L^*(0)$ is sufficiently close to 0.

Whether equilibrium entails some separation or not offers crucial implications for who should retain the right to abandon the project. To explore this possibility, suppose that there is an additional player, called the principal, who can either delegate the decision-making right regarding project termination to the agent (delegation) or retain it to herself (centralization) at the outset. For example, the principal may be an employer or a supervisor if the agent works for a firm or an organization, or the principal may be an investor if the agent is an entrepreneur attempting to start a venture business. In this type of setting, it is clear that delegation is of value only if the agent uses his private information to decide the timing of project abandonment. Alternatively, this argument suggests that for a range of parameters under which the equilibrium is complete 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.[16]

Proposition 7 is useful in this regard, as it clarifies when delegation is more valuable. The proposition indicates that there must be some separation if λ_H is sufficiently close to λ_L (the identification ability is relatively more important). As λ_H increases (the implementation ability becomes more important), and the equilibrium becomes either pooling or separating. However, if $s_L^*(0)$ is sufficiently small, then the range of parameters which support the pooling equilibrium expands, thereby favoring centralization over delegation. In particular, $s_L^*(0) = 0$ if

$$p_0^L \lambda_L (W_L - w_L) < \rho w_L + c, \tag{9}$$

which is more likely to be satisfied if the low type is relatively unproductive in that either $p_0^L \lambda_L$ or W_L is low.[17] When condition (9) holds, Proposition 7 provides an even simpler prediction: there exists some $\underline{\lambda}$ such that the equilibrium is complete pooling if and only if $\lambda_H > \underline{\lambda}$.

This result implies that delegation is more valuable in environments where the identification ability is more important. One prominent example which fits this description is academia, where it is the novelty of ideas, rather than the efficiency of task implementation, which is indispensable for success. Intuitively, as the identification ability gains more importance, the ability effect becomes more dominant relative to the learning effect, thereby making the indifference curves single-crossing in the relevant region. In this type of environment, therefore, the value of centralization becomes more ambiguous, with a tradeoff between the loss of control and the loss of information, and it makes more sense to let the agent decide when to stop. In contrast, in environments where the implementation ability is more important, the learning effect plays a larger role and gives rise to the double-crossing property. In this case, there is less gain from delegation as a pooling equilibrium is more likely to emerge.

In addition, our analysis also sheds new light on the complementary relationship between the right to pick a project and the right to abandon it. As discussed above, the

[16] The presumption here is that the principal gains no information about the agent over the course of the project. In many cases, this is not the case as the principal often has means to evaluate the agent's productivity over time. When the principal has access to an additional information source, centralization then performs even better.

[17] Observe that this is a sufficient condition. Even if this condition fails to hold, we have $\overline{\lambda} \to \infty$ as long as $p_0^L \lambda_L$ or W_L is sufficiently small.

prior quality of the project p_0^i should depend more on the ability of agent i if the project is selected by him. This implies that if it is the agent who chooses what to do, it should also be the agent who chooses when to stop. Again, an example may be academia, where researchers are typically conferred substantial discretion over which projects to work on. On the other hand, in most firm organizations, workers are often assigned to tasks in a top-down manner and have little discretion over what to do. In this type of environment, only the implementation ability is relevant, and the decision-making right should be more centralized as the double-crossing property is more likely to be satisfied.

5.3. An application to venture startups

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 allocation 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 subsection, 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, we suppose that the principal can set a deadline and charge a "fee" when the agent is to continue 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 the identification ability matters more. In this case, the low type persists too long to partially pool with the high type.

If the principal can credibly enforce this scheme, the optimal solution is conceptually straightforward: the principal should simply 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 realize the full-information outcome.

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, however, 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 there exists such a funding gap is not immediately clear, 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.[18] In the presence of reputation concerns, relatively less productive entrepreneurs tend to persist longer than optimal 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. There are at least two virtues of the valley of death. First, it prevents less efficient entrepreneurs from over-experimenting out of reputation concerns. Second, this also raises the expected return for 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 the 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

[18] In 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 (project termination) is fixed independently of the belief.

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

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 necessarily 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.

6. Conclusion

This paper provides a unified 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 D1 equilibria based on the double-crossing property of indifference curves. Our analysis illustrates how reputation concerns on the part of experimenters constrain our ability to acquire new knowledge, in both static and dynamic ways, and offers potential remedies to correct the distortions. The implications obtained here are far-reaching as our framework is flexible enough to 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.

References

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Banks, Jeffrey S. and Joel Sobel, "Equilibrium Selection in Signaling Games," *Econometrica*, 1987, *55* (3), 647–661.

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.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascade," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Bobtcheff, Catherine and Raphael Levy, "More Haste, Less Speed? Signaling through Investment Timing," *American Econbomic Journal: Microeconomics*, 2017, *9* (3), 148–186.

Bonatti, Alessandro and Johannes Hörner, "Career Concerns with Exponential Learning," *Theoretical Economics*, 2017, *12* (1), 425–475.

Butler, Declan, "Translational Research: Crossing the Valley of Death," *Nature*, 2008, *453*, 840–842.

Chen, Chia-Hui and Junichiro Ishida, "Hierarchical Experimentation," *Journal of Economic Theory*, 2018, *177*, 365–404.

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 Lemmons," *Econometrica*, 2012, *80* (4), 1433–1504.

_ **and** _ , "Market Signaling with Grades," *Journal of Economic Theory*, 2014, *151*, 114 – 145.

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*, 2019, *forthcoming.*

_ , Navin Kartik, and Qingmin Liu, "Optimal Contracts for Experimentation," *Review of Economic Studies*, 2016, *83* (3), 1040–1091.

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 "Handbook of Research on Innovation and Entrepreneurship," Cheltenham, UK: Edward Elgar Publishing, 2011.

_ , _ , and Matthew Rhodes-Kropf, "Entrepreneurship as Experimentation," *Journal of Economic Perspectives*, 2014, *28* (3), 25–48.

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.

McNabb, Alex, "Comparison Theorems for Differential Equations," *Journal of Mathematical Analysis and Applications*, 1986, *119* (1–2), 417–428.

Murphy, Lonna M. and Peter M. Edwards, "Bridging the Valley of Death: Transitioning from Public to Private Sector Financing," 2003. NREL/MP-700-NREL/MP-34036.

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.

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 4. 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 $t' \in [t, t+\epsilon)$. Then, by the D1 criterion, the market must assign 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 the D1 criterion, the market must assign 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 weakly prefers quitting to continuing at t, the high type strictly prefers to quit at t. This implies that $\sigma_H(t) = 0$. Hence, the interim belief assigned to one who quits at $t' \in (t, t+\epsilon)$ 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 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:

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

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

$$\begin{aligned}\frac{\partial V(s_L^*(1);s)}{\partial s} &= e^{-\rho(s-s_L^*(1))}\frac{1-p_0F_L(s)}{1-p_0F_L(s_L^*(1))}\left(g_L(s)\left[W_H-w_H\right]-\rho w_H-c\right)\\ &< e^{-\rho(s-s_L^*(1))}\frac{1-p_0F_L(s)}{1-p_0F_L(s_L^*(1))}\left(g_L(s)\left[W_H-w_L\right]-\rho w_L-c\right),\end{aligned}$$

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 $V(s_L^*(1);t') \le w_L$ for any $t' \ge t_1$.

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 complete pooling (i.e., both types do not quit until time $\hat{t}$). By the principle of optimality,

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

From this, we obtain the differential equation:

$$J_i'(t) = \rho J_i(t) + c - g_i(t)\left[W_L + R(t;q_0)(W_H-W_L) - J_i(t)\right],$$

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

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

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 $J_L(t) \ge w_L$ for all $t \le \hat{t}$, 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 $J_H(t) > J_L(t) \ge w_L$. 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,

$$J_i'(t) = \rho J_i(t) + c - g_i(t)\left[W_L + R(t; \tilde{q}(t_0))(W_H - W_L) - J_i(t)\right],$$

with terminal condition $J_i(\hat{t}) = w_L + r(\hat{t}; \tilde{q}(t_0))(w_H - w_L)$. 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:

$$\tilde{J}_L'(t) = \rho \tilde{J}_L(t) + c - g_L(t)\left[W_L + R(t; q_0)(W_H - W_L) - \tilde{J}_L(t)\right],$$

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 (7) 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 $s_H^*(1) < \hat{t}$ by Lemma 3. 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:

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

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 $\tilde{J}_H(t)$ is the same solution to the differential equation specified in part (b). 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 (7) and (8) 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 $\overline{t}_i$ 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 $\underline{t}_H < s_L^*(q_0)$, where $\underline{t}_H \neq \underline{t}_L$ by Lemma 4. 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

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

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,

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

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 (McNabb, 1986), 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 4. 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 4 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 Lemma 2 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 fully 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 $\underline{t}_L < \underline{t}_H < \overline{t}_L$, we must have $\underline{t}_L < \hat{t}$ and $\overline{t}_L > \hat{t}$. If $\overline{t}_L > \hat{t}$, however, we cannot have $\overline{t}_L < \overline{t}_H$, which implies that $\underline{t}_H < \overline{t}_H$ cannot occur in equilibrium. $\square$

Now let $t_H \equiv \underline{t}_H = \overline{t}_H$. The above facts show that there can only be three types of equilibria: (1) complete pooling in which $\underline{t}_L = t_H = \hat{t}$; (2) semi-pooling in which $\underline{t}_L = s_L^*(q_0) < \overline{t}_L = t_H = \hat{t}$; or (3) fully separating in which $\underline{t}_L = s_L^*(q_0) < \overline{t}_L = s_L^*(1) < t_H$. 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 $\underline{t}_L = s_L^*(q_0) < \overline{t}_L = t_H = \hat{t}$, 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 $\overline{t}_L = s_L^*(1) < t_H$, so that the equilibrium is fully separating. Then, Lemma 3 implies that $s_H^*(1) \leq \hat{t}$ if $s_L^*(1) \leq \hat{t}$ and $s_H^*(1) < s_L^*(1)$ if $s_L^*(1) > \hat{t}$. Therefore, $s_H^*(1) \leq \max\{\hat{t}, s_L^*(1)\}$. 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 complete pooling.

(b) If $V(s_L^*(q_0); \hat{t}) < w_L$, the complete 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 fully separating equilibrium exists, i.e., $\overline{t}_L = s_L^*(1)$. The high type must then quit at $t_H > \hat{t}$ such that $V(s_L^*(1), t_H) = w_L$. Following the same argument as in (a), $s_H^*(1) \leq \max\{\hat{t}, s_L^*(1)\}$ by Lemma 3, 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 fully separating equilibrium is feasible in this case, although there is still a continuum of fully separating equilibria that are feasible. To further reduce the set of equilibria, note that since $\overline{t}_L = s_L^*(1)$, 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 $\overline{t}_L$ 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 Lemma 5. To show $\lim_{\lambda_H \to \lambda_L} \hat{t}(\lambda_H) = \infty$, 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 $\lim_{\lambda_H \to \infty} \hat{t}(\lambda_H) = 0$, suppose that $\hat{t}$ is bounded from below while λ_H tends to infinity. We can find λ_H that is large enough (yet finite) to satisfy

$$\frac{\lambda_H p_0^H e^{-\lambda_H \epsilon}}{1 - p_0^H + p_0^H e^{-\lambda_H \epsilon}} < \frac{\lambda_L p_0^L e^{-\lambda_L \epsilon}}{1 - p_0^L + p_0^L e^{-\lambda_L \epsilon}},$$

for any $\epsilon > 0$, 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

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

We know that, evaluated at $t = \hat{t}(\lambda_H)$, $\partial g_L / \partial t > \partial g_H / \partial t$. This means that $\mathrm{d}\hat{t}/\mathrm{d}\lambda_H$ has the same sign as $\partial g_H / \partial \lambda_H$ (evaluated at $t = \hat{t}(\lambda_H)$). Therefore, $\mathrm{d}\hat{t}/\mathrm{d}\lambda_H$ 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 $\mathrm{d}\hat{t}/\mathrm{d}\lambda_H$ is single-crossing from above as λ_H increases, because the above expression is decreasing in λ_H if $\mathrm{d}\hat{t}/\mathrm{d}\lambda_H = 0$. As λ_H approaches λ_L, we have shown that $\hat{t}(\lambda_H)$ approaches infinity, and hence the sign of $\mathrm{d}\hat{t}/\mathrm{d}\lambda_H$ is negative. Together with the fact that $\mathrm{d}\hat{t}(\lambda_H)/\mathrm{d}\lambda_H$ is single-crossing from above, the fact that $\lim_{\lambda_H \to \lambda_L} \mathrm{d}\hat{t}(\lambda_H)/\mathrm{d}\lambda_H < 0$ implies $\mathrm{d}\hat{t}(\lambda_H)/\mathrm{d}\lambda_H < 0$ for all $\lambda_H > \lambda_L$. ■

Proof of Proposition 7. We consider how the 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 fully separating if $\hat{t}$ sufficiently close to 0. As $\hat{t}$ increases and exceeds $\underline{t}$, the equilibrium becomes semi-pooling for any $q_0 \in (0,1)$. Finally, the equilibrium is complete 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 complete pooling if $\hat{t}$ is sufficiently close to $s_L^*(0)$, because

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

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

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

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

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

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 $t_0 \to \infty$ as $\hat{t} \to \infty$. It then follows that the equilibrium must be fully 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). Lemma 5 establishes $\hat{t}$ decreases in λ_H, and therefore the proposition follows. To show that $\overline{\lambda}$ can be arbitrarily large, observe that $\underline{t} = 0$ if $s_L^*(0)$ is sufficiently small. Moreover, since $\lim_{s_L^*(0) \to 0} v(0) = w_L < w_L + r(0; q_0)(w_H - w_L)$, the equilibrium must be complete pooling if $s_L^*(0)$ is sufficiently small. ■