ISSN (Print) 0473-453X
Discussion Paper No. 1270
ISSN (Online) 2435-0982

SIGNALING VISION: KNOWING WHEN TO QUIT

Junichiro Ishida
Wing Suen

December 2024

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

Signaling Vision: Knowing When to Quit

Junichiro Ishida
Osaka University

Wing Suen
University of Hong Kong

December 22, 2024

Abstract: We study a signaling game where agents signal their type by choosing when to quit pursuing an uncertain project. High types observe news about project quality and quit when bad news arrives. This creates opportunities for low types who do not observe any news to mimic high types by quitting strategically. In equilibrium, there is a mimicking phase of time when low types quit continuously. The reputation dynamics may exhibit non-monotonicity, with agents who quit either very early or very late carrying a higher reputation than do agents who quit near the optimal time for low types. Our analysis offers a unifying explanation for how and when both early and late quitting can enhance reputation and suggests novel welfare and policy implications.

Keywords: perseverance, behavioral diversity, mimicking, non-monotone reputation dynamics, bandit experimentation

JEL Classification: D82, O31

1. Introduction

Deciding when to quit is one of the most critical decisions in any exploratory endeavor, but its decision process becomes complicated in the presence of reputation concerns. Consider, for example, an entrepreneur working on a business startup that shows little promise. Should the entrepreneur be decisive and quit immediately or be resilient and persist longer? What message does this decision send to the market? The answers to these questions are not straightforward, as there are at least two schools of thought on when it is best to quit. On one hand, as an old Chinese proverb goes, "A wise man changes his mind; a fool never will." A contemporary version of this idea is further popularized by Godin (2007), where he stresses the importance of reversing previous decisions that turn out to be wrong and emphasizes that "winners quit fast and often." On the other hand, there are numerous discourses highlighting the importance of perseverance. Management texts tout Thomas Edison as an example of an entrepreneur who "failed" his way to success through tenacity and perspiration (Merriman, 2017), where hundreds of different filament materials were tried before his signature invention, the incandescent light bulb, became a commercial success. Duckworth (2016) highlights the concept of grit—a special blend of passion and perseverance—as the essential driver of success. According to this latter view, quitting early can be detrimental to the entrepreneur's reputation as it can be perceived as a lack of perseverance.

Given these seemingly conflicting views, it is not immediately clear how reputation concerns affect or distort the process of knowledge exploration, which has become increasingly more important in modern-day economies. In this paper, we propose a framework that can reconcile these views and illuminate the intricate nature of experimentation with reputation concerns. Specifically, we consider an environment where an agent (e.g., an entrepreneur, a politician, or a scientist) engages in a project of unknown quality while simultaneously attempting to signal his "vision"—defined as the ability to synthesize complex information and correctly foresee the future course of action—to the market. The project is either good or bad and succeeds at some random time only if it is good. We augment this otherwise standard bandit problem with the possibility that the agent may also observe some news that reveals the project's true worth at some random time. The agent's ability to observe news depends on his vision: it can be detected only by visionary agents but goes unnoticed by non-visionary ones.[1] The experimentation process continues

[1]Damiano et al. (2020) consider a model in which an agent can choose to acquire additional information

until the agent either succeeds or chooses to quit. Once the agent quits, his continuation payoff is determined by the market's evaluation of his vision.

Our problem of "signaling vision" is qualitatively different from the standard signaling problem of Spence (1973), which we call the "standard setup." In our model, visionary agents are *ex post* turned into different informational types, depending on the information they receive over time. These different types naturally have different incentives for experimentation, generating some behavioral diversity among visionary agents. As non-visionary agents can mimic any one of these types to acquire the same reputation, the presence of behavioral diversity among visionary agents generates complicated strategic interactions between visionary and non-visionary agents and shapes the reputation dynamics. A hallmark of our analysis is the potential existence of a "mimicking phase," during which non-visionary agents randomize and quit continuously to pool with visionary agents who quit as they observe bad news.

We show that the equilibrium reputation dynamics may follow a U-shaped non-monotone path with respect to the quitting time. Our analysis thus offers a unified framework that explains how both decisiveness (quitting early) and perseverance (persisting through adversities) can be perceived as signs of competence. Under a mild condition, the equilibrium reputation is initially decreasing, suggesting that exhibiting decisiveness by quitting early is generally an effective way to signal one's vision. In contrast, the reputation value of perseverance is more nuanced: the equilibrium reputation may strictly decrease over time, in which case persistence is interpreted negatively as a sign of "stubbornness" rather than competence. The key question is, therefore, whether and when the equilibrium reputation becomes non-monotone, eventually reversing course to increase at a later stage.

The reputation dynamics of our model depends crucially on how the incentives of visionary and non-visionary agents evolve over time, conditional on having received no news.[2] Since visionary agents have additional sources of information, their incentive to persist with the project is initially stronger, starting from the same initial prior. As time passes, though, the beliefs of visionary and non-visionary agents diverge from each other. If the underlying learning process is such that visionary agents stay relatively more optimistic than non-visionary agents, they are always willing to persist longer, and the stan-

about project quality at a cost while working on the project. In our model, such additional information arrives exogenously to one type of agent but not to another type.

[2] Non-visionary agents never receive news by assumption. Here, we mainly refer to visionary agents who have received no news.

dard single-crossing property holds. In this case, visionary agents hold out until all non-visionary agents quit to achieve full separation, but to support this as an equilibrium, the equilibrium reputation must turn increasing towards the end. Otherwise, the incentives flip at some point, and the double-crossing property, as defined in Chen et al. (2022), emerges instead. Visionary agents then choose to quit midway through and settle for pooling with non-visionary agents, limiting the reputation value of perseverance.

Aside from the learning process, we also demonstrate that the agent's prior reputation—the prior belief held by the market that the agent is visionary—plays a crucial role in shaping the reputation dynamics and welfare. This insight stands in stark contrast to the predictions of the standard setup which, under standard refinements, always selects the least-cost separating equilibrium (LCSE) for any given prior belief, leading to a disturbing implication that signaling-based reputation models cannot explain why reputation is valuable in the first place. In our model, visionary agents drop out when they observe bad news, which provides opportunities for non-visionary agents to mimic them. In equilibrium, the demand and supply of mimicking opportunities must balance out, where each equilibrium of our model is indexed by what we call the "shadow value of mimicking," a measure that directly corresponds to the equilibrium payoff of non-visionary agents. When the prior reputation is low, indicating that the agent is less likely to be visionary, the supply of mimicking opportunities is limited relative to the demand. This imbalance makes the competition for mimicking more intense and forces non-visionary agents to start mimicking inefficiently early; as a consequence, the mimicking phase prolongs, and the shadow value diminishes. Our setting therefore yields an intuitive prediction that the agent's expected payoff is increasing in his prior reputation, giving rise to a mechanism through which the value of reputation emerges endogenously.

In our setting, the presence of the mimicking phase yields novel efficiency and policy implications by generating a signaling distortion that involves both types and runs in both directions. This is again in sharp contrast to the standard setup where the signaling distortion is generally unidirectional and affects only higher types: in the simplest case with two types, for instance, the low type chooses the complete-information optimum (i.e., the optimal action in the absence of information asymmetry), and the high type chooses an action just high enough to separate from the low type. To illustrate the welfare impact of equilibrium mimicking, we consider a simple scheme where a principal (e.g., a government) imposes a "tax" at some predetermined date. Such a simple tax scheme can be quite effective and achieve the efficient allocation in the standard setup, as it induces low-type agents to exit at the right time, freeing high-type agents from the need to separate by over-investing. We argue, however, that this type of intervention may backfire in the context of signaling vision, because imposing a tax effectively reduces the supply of mimicking opportunities and diminishes the equilibrium payoff of (low-type) non-visionary agents. Anticipating this, non-visionary agents start abandoning their project even earlier, further aggravating inefficient mimicking. We relate this argument to "the valley of death" that has been discussed extensively in the context of venture financing and argue that its welfare impact depends heavily on the agent's motive for signaling, i.e., what type of attribute or ability he wishes to convey to the market.

From the theoretical point of view, this paper connects two strands of the signaling literature—signaling under behavioral diversity and signaling under double-crossing preferences—that have garnered attention lately but have evolved independently. Recent works introduce behavioral diversity among the same type of agents into an otherwise standard signaling model (Dilme and Li, 2016; Ishida and Suen, 2024). In those works, some agents are randomly selected to be "behavioral" and assigned to non-equilibrium actions.[3] The remaining strategic agents may choose to mimic the choices of those behavioral agents, leading to pooling patterns that are different from standard pooling equilibria in signaling models. This mechanism is also at work in our analysis, as it underpins the mimicking phase during which non-visionary agents continuously quit. However, unlike in the existing literature where behavioral diversity is exogenously imposed, it is an endogenous response to stochastic news in our model, providing a natural signaling environment that entails behavioral diversity in a fully optimized setting.

Because different types may acquire information at different rates, experimentation models with reputation concerns often produce preferences that violate the single-crossing property (Bobtcheff and Levy, 2017; Chen et al., 2021). Chen et al. (2021) consider an environment where high-ability types have higher project success rates with good quality projects but also become pessimistic with their projects more quickly, giving rise to double-crossing preferences.[4] In contrast, visionary agents in the current model may become privately informed and decide to stay in the game or quit, depending on the information they

[3] Dilme and Li (2016) consider a situation where some agents are forced to drop out at some random time due to liquidity shocks. Ishida and Suen (2024) take a more abstract approach, abstracting away from the fine details that lead to behavioral diversity, and assume that some agents are randomly assigned to non-equilibrium actions.

[4] A general analysis of signaling under double-crossing preferences is provided by Chen et al. (2022, 2024).

receive. The rates of learning good news or bad news play an essential role in determining the preference structure—especially whether the single-crossing property holds or not—as they affect the differential signaling incentives of visionary and non-visionary agents. By combining this preference structure with the aforementioned behavioral diversity, our analysis underscores the intrinsic connection between them in the context of signaling vision and elucidates insights that arise from this combination.

While the notion of vision has been discussed extensively in the leadership and management literature (Burns, 1978; Bass, 1985; Collins, 2001), there are also some formal economic analyses of vision. Many of these analyses portray visionary agents as individuals with strong beliefs. Rotemberg and Saloner (2000) consider a visionary CEO who is consistently biased in favor of a certain kind of project and against others. Such bias can commit the firm to a narrow strategic direction and raise the frequency with which employees in the favored activities see their innovations implemented, thus increasing their incentive to seek them. Van den Steen (2005) defines vision as a strong belief held by the manager about the right course of action. He shows that hiring a manager with strong beliefs can be optimal as he will attract employees with similar beliefs, causing an alignment of beliefs within the firm.[5] Our interpretation of vision——as the ability to foresee the future course of action——is fundamentally connected to these notions, as visionary agents in our model tend to have strong and extreme beliefs stemming from their superior information.

2. Model

Environment. We consider a continuous-time model of experimentation in which an agent with reputation concerns conducts a project of unknown quality while also attempting to signal his ability type to the market. At each point in time, the agent decides whether to continue the project or quit. If he continues, the project succeeds at a rate that depends on the project quality (the state). Let $x \in \{G, B\}$ denote the project quality, which is either good ($x = G$) or bad ($x = B$). The prior probability of the project being good is p_0. If the quality is good and the agent continues for $[t, t + \mathrm{d}t)$, the project succeeds with probability $\lambda \mathrm{d}t$; if the quality is bad, the project never succeeds. The game ends either when the project succeeds or when the agent quits.

[5] Similarly, Bolton et al. (2013) considers a leader who overweights his initial information—a trait they refer to as *resoluteness*—and argue that the leader's resoluteness facilitates coordination within the organization.

Information. Let $a \in \{h, l\}$ denote the agent's ability type, which is either high ($a = h$) or low ($a = l$). The agent's ability type is his private information, where the prior probability that the agent is a high type is given by μ_0. Although the project success rate λ is the same for both types, the high type has greater learning ability than the low type, in that the former can discern more quickly whether the project he is conducting is good or bad. Specifically, in addition to experimentation outcomes, the high type may also privately observe news that reveals the true project quality while conducting the project. Given the true state x and type a, if the agent continues experimenting for $[t, t + \mathrm{d}t)$, he finds out the true state with probability $k_a \theta_x \mathrm{d}t$, where $\theta_B, \theta_G > 0$, $k_h = 1$ and $k_l = 0$. We interpret news as an event that is informative about the true state but can be recognized only by the high type. The difference in the ability to find out the true state, or "vision," is the only difference between the two types. For clarity, we say that the agent observes good (bad) news when he finds out that the state is good (bad). Define $\Delta_\theta := \theta_B - \theta_G$, which can be positive, negative, or equal to zero.

Payoffs. The agent receives a flow payoff w as long as he continues the project. The flow payoff can be either positive (the agent receives a wage) or negative (the agent incurs part of the experimentation expenses), depending on the nature of the project at hand. If the project succeeds at some random time, the agent receives a lump-sum payoff B, and the game ends. If the agent decides to quit before the project succeeds, he receives a payoff that depends on his market reputation at the time. Let $\mu(t)$ denote the agent's reputation (i.e., the market's belief that the agent is a high type) when he quits at time t. The reputation payoff the agent receives is given by $\hat{R}(\mu(t))$, where $\hat{R} : [0, 1] \to \mathbb{R}$ is a differentiable and strictly increasing function that captures the reward from being perceived as the high type. Let $R(t) := \hat{R}(\mu(t))$ represent the reputation payoff at time t, and use $\overline{\rho} := \hat{R}(1)$ and $\underline{\rho} := \hat{R}(0)$ to represent the highest and lowest possible reputation payoffs. Also, let $\rho_0 := \hat{R}(\mu_0)$ denote the reputation payoff at the prior belief. The agent maximizes the present value of payoffs with discount rate $r > 0$.

Equilibrium concept. We adopt the standard notion of a signaling equilibrium and require that each type of agent's (possibly stochastic) quitting time is a best response to the reputation dynamics $R(\cdot)$ and to the news received, if any. The market's belief $\mu(\cdot)$ about the agent's type—and hence the reputation dynamics $R(\cdot)$—must be consistent with equilibrium strategies and Bayes' rule whenever applicable. In addition, we use the D1 refinement (Cho and Kreps, 1987) to discipline off-equilibrium beliefs. Specifically, if t' is not in the support of the equilibrium quitting time of any type of agent, and if the set of values of

$R(t')$ that would induce a type-a agent to deviate to t' strictly contains the set of values of $R(t')$ that would induce type-a' ($a' \neq a$) to deviate to t', then the market imputes probability one that the deviating agent is a type-a agent. In what follows, we simply use "equilibrium" to refer to signaling equilibrium that satisfies the D1 criterion.

3. Analysis

3.1. Preliminaries

If the payoff from project success B is small, the agent may be reluctant to experiment even if he has enough confidence in the project. For the model to be interesting, we require B to be large enough so that it is always optimal for the agent to continue indefinitely until he succeeds if the project is known to be good. The payoff from continuing indefinitely is given by

$$V := \int_0^\infty e^{-(\lambda+r)s}(\lambda B + rW)\,\mathrm{d}s = \frac{\lambda B + rW}{\lambda + r}.$$

where $W := w/r$ is the present value of continuing indefinitely. Since the maximum payoff the agent can earn by quitting is $\overline{\rho}$, the agent who knows that the project is good always chooses to continue if $V > \overline{\rho}$. On the flip side, if W is positive and large, the agent would never quit even if the project is known to be bad. We thus assume that the outside option is large enough to rule out trivial equilibria in which the agent never quits.

Assumption 1. $V > \overline{\rho} > \underline{\rho} > W$.

3.2. Informed and uninformed types

At any time $t > 0$, the agent is either *informed* about the true state or *uninformed*. Conditional on no project success and no news up to time t, the belief of a type-a agent ($a = h, l$) that the project is good is

$$p_a(t) = \frac{p_0 e^{-\lambda t}}{p_0 e^{-\lambda t} + (1-p_0)e^{-k_a \Delta_\theta t}}. \tag{1}$$

The belief of a low-type agent always goes down with time because $k_l = 0$, but the belief of a high-type agent may go up in the absence of news if $\Delta_\theta > \lambda$. Note that $p_h(t) > p_l(t)$ for all $t > 0$ if and only if $\Delta_\theta > 0$.

An informed agent must be a high type because a low-type agent never receives any news. The following statement establishes important properties that allow us to substantially simplify the subsequent analysis. The proofs of Proposition 1 and of other results are presented in the Appendix.

Proposition 1. *The following properties hold in any equilibrium:*

(a) the high type who has received good news never quits until he succeeds;

(b) the high type quits immediately if he receives bad news.

Proposition 1 allows us to focus on the problems of the high type who has received no news and the low type; we will refer to the former simply as the high type for brevity whenever it is not confusing. Upon receiving bad news, a high-type agent knows that it is the time to quit. Because high-type agents know when to abandon a futile project, low-type agents may try to mimic high-type agents by strategically choosing their quitting times, while high-type agents who are still uninformed have the incentive to separate from low-type agents. The main difference between our model and a standard signaling model is that the behavior of the high type is not completely determined by his equilibrium strategy—the random arrival of bad news may also cause him to quit even if he intends to continue otherwise. The remainder of the analysis studies how such random behavior induced by news distorts the strategies chosen by the high and low types

3.3. Marginal rates of substitution and indifference curves

To pin down the optimal strategies of the two types, we need to clarify how their incentives evolve over time. Define $u_a(t;R)$ as the expected payoff of a type-a agent who plans to quit at time t and receives a reputation payoff of $R(t)$ upon quitting. By Proposition 1, if he receives good news at time $s < t$, his payoff at that point becomes V, and if he receives bad news at time $s < t$, he quits immediately to obtain a payoff of $R(s)$. Therefore, his expected payoff is given by

$$u_a(t;R) = \int_0^t \left[p_0 e^{-(\lambda+k_a\theta_G+r)s}(\lambda B + k_a\theta_G V + rW) + (1-p_0)e^{-(k_a\theta_B+r)s}(k_a\theta_B R(s) + rW)\right] \mathrm{d}s \\ + e^{-rt}\left[p_0 e^{-(\lambda+k_a\theta_G)t} + (1-p_0)e^{-k_a\theta_B t}\right]R(t). \tag{2}$$

Assuming the differentiability of $R(\cdot)$, we have $\dot{u}_a(t;R) = 0$ if and only if

$$p_a(t)\left[\lambda(B - R(t)) + k_a\theta_G(V - R(t))\right] + r(W - R(t)) = -\dot{R}(t). \tag{3}$$

When the agent has continued up to time t and obtained no success and no news, his belief that the project is good is $p_a(t)$. If the state is good, the capital gain from continuing a little longer is $\lambda(B-R(t))+k_a\theta_G(V-R(t))$, while the gain in flow payoff is $r(W-R(t))$. These two effects must be balanced against the appreciation or depreciation in the reputation payoff from exit, shown on the right-hand-side of equation (3).

It is useful to define the *marginal rate of substitution* as the increase in reputation payoff at some time t that is needed to compensate an agent for continuing a little longer beyond time t. For $a=h,l$, we let

$$MRS_a(t,\rho) := -p_a(t)[\lambda(B-\rho)+k_a\theta_G(V-\rho)]-r(W-\rho).$$

We also define an *indifference curve* of a type-a agent, denoted $\rho_a(\cdot)$, as the locus of points $(t,\rho_a(t))$ such that the agent is indifferent between quitting and continuing along this locus. By equation (3), this requires the indifference curve to satisfy the differential equation given by

$$\dot{\rho}_a(t) = MRS_a(t,\rho_a(t)), \tag{4}$$

which corresponds to the property that the slope of the indifference curve is equal to the marginal rate of substitution. If a type-a agent quits continuously on an interval (t',t'') and the reputation dynamics is given by $R(\cdot)$, the first-order condition requires $\dot{u}_a(t;R)=0$ on this interval, which is equivalent to requiring $\dot{R}(t)=MRS_a(t,R(t))$ for all $t\in(t',t'')$. In other words, the indifference curve $\rho_a(\cdot)$ of type-a must coincide the reputation dynamics $R(\cdot)$ on the interval (t',t'').

An agent with a greater $MRS_a(t,\rho)$ has more incentive to quit at time t when the reputation payoff is ρ. It is important to study how the marginal rate of substitution varies with type. The fact that an indifference curve of the low type is quasi-convex (i.e., decreasing then increasing in t) is particularly important, as it dictates the equilibrium reputation dynamics.

Lemma 1. *An indifference curve $\rho_l(\cdot)$ of the low type is quasi-convex. An indifference curve $\rho_h(\cdot)$ of the high type is quasi-convex if $\Delta_\theta<\lambda$. Further,*

(a) if $\Delta_\theta \geq 0$, then $MRS_h(t,\rho) < MRS_l(t,\rho)$ for all (t,ρ);

(b) if $\Delta_\theta < 0$, there exists a continuous and decreasing function $D(\cdot)$ such that

$$MRS_h(t,\rho)\begin{cases} < MRS_l(t,\rho) & \text{if and only if } t<D(\rho)\\ > MRS_l(t,\rho) & \text{if and only if } t>D(\rho),\end{cases}$$

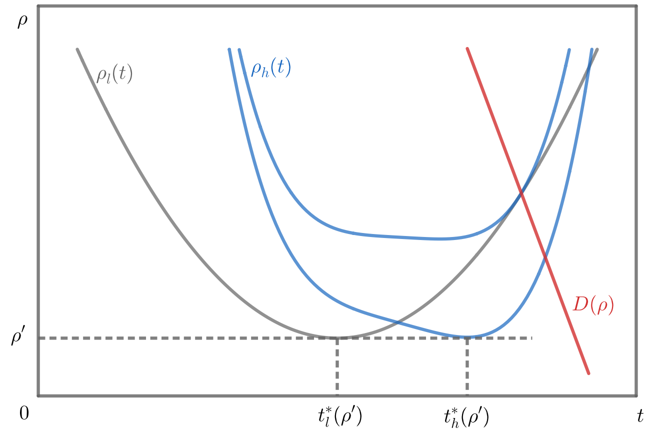

Figure 1. *Indifference curves are U-shaped. The bottom of these curves indicates the optimal quitting time if the reputation payoff is constant. The indifference curve of a high-type agent is more convex than that of a low-type agent when their indifference curves are tangent to one another.*

and $MRS_h(t,\rho_l(t)) - MRS_l(t,\rho_l(t))$ *is single-crossing from below in* t.

When $\Delta_\theta \geq 0$, no news is good news; so the high type is always more optimistic than the low type. Lemma 1 states that the standard single-crossing property prevails, with the high type always having less incentive to quit than does the low type. When $\Delta_\theta < 0$, no news is bad news, and a double-crossing property (Chen et al., 2022) obtains as a consequence. We call $D(\cdot)$ the *dividing line* because it separates the standard single-crossing domain from the reverse single-crossing domain. For (t,ρ) to the left of the dividing line, the high type has less incentive to quit than does the low type, because the high type's payoff from continuing the project includes the option value from the possibility of receiving news. The incentive comparison reverses to the right of the dividing line, because the high type becomes pessimistic more quickly upon receiving no news if $\Delta_\theta < 0$.[6] Lemma 1 also establishes that, in the double-crossing case, the difference in the marginal rate of substitution is single-crossing from below along an indifference curve. This implies that when the indifference curves of the two types are tangent to one another (i.e., for (t,ρ) on the dividing line), the high type's indifference curve is "more convex" than that of the low type. See Figure 1 for an illustration.

[6]For consistency of notation, we use the convention to denote $D(\rho) = \infty$ in the single-crossing case such that $D(\rho) > t$ holds for any finite t.

When $\Delta_\theta < \lambda$, the beliefs of both types strictly decrease over time. Lemma 1 shows that their indifference curves are U-shaped, as depicted in Figure 1. If the reputation payoff is exogenously fixed at some ρ' and is constant over time, our model reduces to a standard bandit experimentation problem in which the agent stops experimenting when his belief crosses a threshold and becomes pessimistic enough. Denote the optimal stopping time for such a problem by $t_a^*(\rho')$ for $a = h, l$. The optimal stopping time satisfies $MRS_a(t_a^*(\rho'), \rho') = 0$, which corresponds to the bottom of the relevant indifference curve. In particular, $t_a^*(\rho)$ solves

$$p_a(t_a^*(\rho)) = \frac{r(\rho - W)}{\lambda(B-\rho) + k_a\theta_G(V-\rho)}. \tag{5}$$

If $\Delta_\theta < \lambda$ and the right-hand-side of (5) is strictly between 0 and p_0, then $t_a^*(\cdot)$ is interior and is strictly decreasing. If the right-hand-side of (5) is greater than or equal to p_0, we let $t_a^*(\rho) = 0$. If $\Delta_\theta \geq \lambda$, the high type's belief is weakly increasing over time, and his optimal choice is either to quit immediately ($t_h^*(\rho) = 0$) or never quit ($t_h^*(\rho) = \infty$).

When the prior belief p_0 is too low, the low type may choose to quit immediately at time 0 under complete information. As this situation is uninteresting, we restrict attention to the case where there is an interior complete-information optimum for the low type, i.e., $t_l^*(\underline{\rho}) > 0$. From (5), we obtain the following condition.

Assumption 2. $p_0 > [r(\underline{\rho} - W)]/[\lambda(B - \underline{\rho})]$.

As will be seen later, this assumption also implies that the high type will not quit immediately when the reputation payoff is very low. It does not imply, however, that the low type never quits at time 0 in equilibrium, because the reputation payoff $R(0)$ is endogenously determined and may be higher than $\underline{\rho}$. Assumptions 1 and 2 are maintained throughout this paper.

4. Equilibrium

4.1. Optimal stopping strategies

The key strategic feature of our model arises from the fact that the high type will quit immediately as he receives bad news (Proposition 1). Because only the high type can receive bad news, mimicking this behavior may allow an agent to gain a higher reputation even if this agent is uninformed. Moreover, the opportunity to mimic the high type is not entirely exogenous, because the quitting strategies of the two types determine the unconditional probability that an agent who remains in the game will receive news in the next interval of time. The standard analysis of using exit as a signal must be be adjusted to account for these strategic features.

For $a = h, l$, let Q_a be the support of the equilibrium stopping times for type a. When the equilibrium reputation function is $R(\cdot)$, an equilibrium stopping time must be a best response to $R(\cdot)$, i.e., $q_a \in Q_a \implies q_a \in \operatorname{argmax}_t u_a(t;R)$. We allow the possibility that it is optimal for the high type to continue indefinitely, in which case we define $\operatorname{argmax}_t u_h(t;R) = \infty$. The following result shows that Q_h must be a singleton set (with either a finite or infinite optimal quitting time).

Proposition 2. *In any equilibrium, the high type adopts a pure strategy—he either continues indefinitely or quits once and for all at some finite time.*

By Proposition 2, we can use q_h to represent the unique stopping time of the high type. If $\Delta_\theta < \lambda$, then $p_h(t)$ approaches 0 as t approaches infinity, and the expected payoff of continuing indefinitely converges to W. Because $\overline{\rho} > W$, there must exist some t' such that $R(t)$ is strictly bounded away from $\overline{\rho}$ for all $t > t'$, but this is infeasible in equilibrium because this would require the low type to quit at a rate strictly bounded away from 0 indefinitely. Thus q_h must be finite when $\Delta_\theta < \lambda$. On the other hand, if $\Delta_\theta \geq \lambda$, the high type becomes more optimistic about the project upon receiving no news. It is possible that the high type may want to continue the project indefinitely until either he succeeds or receives bad news. That is, q_h may be finite or infinite when $\Delta_\theta \geq \lambda$. As we will illustrate below, the high type's stopping time q_h is a key factor that determines the equilibrium reputation dynamics.

Although the high type always adopts a pure strategy, the low type may not adopt a pure strategy. Consider the simple case where the single-crossing property holds (i.e., $\Delta_\theta \geq 0$), so that the high type always has more incentive to persist longer. Suppose $q_l = t_l^*(\underline{\rho})$ and $q_h > q_l$ is such that the low type is indifferent between quitting at $(q_l, \underline{\rho})$ and quitting at $(q_h, \overline{\rho})$. Even though a low-type agent has no strict incentive to pool with high-type agents by quitting at time q_h, he has an incentive to pool with high-type agents who receive bad news by, say, quitting at some time slightly after q_l and to get a reputation payoff of $\overline{\rho}$. While this is a form of pooling, it is different from its conventional use (in the sense of pooling equilibrium under the standard setup); we thus specifically call it *mimicking* to distinguish it from the standard form of pooling. The fact that the high type quits continuously as he receives bad news before time q_h makes such a mimicking strategy possible and breaks the standard fully separating equilibrium.

4.2. Characterization

Although our analytical focus is on the possibility of mimicking, our model admits an equilibrium that does not entail any mimicking behavior. This is the case if the high type becomes pessimistic much faster than the low type, so that the dividing line is located far to the left. More precisely, let $\underline{\rho}_l(\cdot)$ represent the indifference curve that passes through the point $(t_l^*(\underline{\rho}), \underline{\rho})$, and define $t^{\text{sep}} < t_l^*(\underline{\rho})$ such that

$$\begin{cases} \underline{\rho}_l(t^{\text{sep}}) = \overline{\rho} & \text{if } \underline{\rho}_l(0) > \overline{\rho}, \\ t^{\text{sep}} = 0 & \text{if } \underline{\rho}_l(0) \leq \overline{\rho}. \end{cases} \tag{6}$$

Because indifference curves are quasi-convex, there are typically two values of t such that the low type is indifferent between quitting at $(t, \overline{\rho})$ and quitting at $(t_l^*(\underline{\rho}), \underline{\rho})$, and t^{sep} is defined to be the smaller of these two values of t.[7] Given this, if $t^{\text{sep}} \geq D(\overline{\rho})$, there is a fully separating equilibrium in which the low type quits at $q_l = t_l^*(\overline{\rho})$, and the high type quits at $q_h = \min\{t^{\text{sep}}, t_h^*(\overline{\rho})\} < q_l$. However, as neither the incentive compatibility constraint of the high type nor that of the low type binds, this type of equilibrium is technically trivial and adds little economic insight, aside from the fact that it is somewhat unrealistic to have $D(\overline{\rho})$ so close to 0. In the remainder of the analysis, therefore, we make the following assumption to rule out this possibility and focus our attention on more relevant cases.

Assumption 3. $t^{\text{sep}} < D(\overline{\rho})$.

Assumption 3 states that the incentives of the two types may flip but not too early. The following lemma establishes that under this assumption, there is always an interval (q_l^1, q_l^2) during which the low type quits continuously, which we refer to as the *mimicking phase* throughout the analysis.

Lemma 2. *If $t^{\text{sep}} < D(\overline{\rho})$, the low type adopts a mixed strategy in any equilibrium. The support of the low type's strategy takes one of the following forms:*

(a) $Q_l = (q_l^1, q_l^2)$ *for some* $q_l^1 < q_l^2$*;*

[7] If the low type strictly prefers quitting at $(t_l^*(\underline{\rho}), \underline{\rho})$ to quitting at $(0, \overline{\rho})$, then t^{sep} is defined to be equal to 0.

(b) $Q_l = (q_l^1, q_l^2) \cup \{q_l^2\}$*;*

(c) $Q_l = (q_l^1, q_l^2) \cup \{q_l^2\} \cup \{t_l^*(\underline{\rho})\}$*.*

In each of these cases, $q_l^2 \leq q_h$*.*

Observe that if the low type quits continuously, we must have $\dot{u}_l(t;R) = 0$ on this interval, which is equivalent to $\dot{R}(t) = MRS_l(t, R(t))$. Therefore, the reputation payoff $R(\cdot)$ offered by the market must coincide with an indifference curve $\rho_l(\cdot)$ of the low type on (q_l^1, q_l^2). The indifference curve of the low type can be obtained by solving the differential equation (4) for $a = l$. The solution to this differential equation is obtained as

$$\rho_l(t; C) = p_l(t)V + (1 - p_l(t))W + \frac{p_l(t)}{p_0 e^{-(\lambda+r)t}} C, \tag{7}$$

where $C > 0$ is a constant of integration. If the low type's current belief is $p_l(t)$, the payoff from continuing indefinitely is $p_l(t)V + (1 - p_l(t))W$. The last term in this solution represents the option value from quitting, and this option value is positive. The indifference curves given by (7) are indexed by C, with a higher C indicating a higher payoff to the low type. If a low-type agent earns a payoff above what is ensured under complete information, it is precisely because of the possibility of mimicking; we thus refer to the value of C as the *shadow value of mimicking*. To determine the equilibrium reputation dynamics $R(\cdot)$ requires pinning down the shadow value of mimicking.

We denote by $\underline{C}$ the value of C such that the indifference curve $\rho_l(\cdot; C)$ passes through the point $(t_l^*(\underline{\rho}), \underline{\rho})$, i.e., $\rho_l(\cdot; \underline{C})$ is the same as $\underline{\rho}_l(\cdot)$ defined earlier. This value $\underline{C}$ constitutes the lower bound of the shadow value of mimicking, because the low type can always choose to quit at time $t_l^*(\underline{\rho})$ and achieve a reputation payoff of at least $\underline{\rho}$. Similarly, denote by $\overline{C}$ the value of C such that the indifference curve $\rho_l(\cdot; C)$ passes through the point $(t_l^*(\overline{\rho}), \overline{\rho})$, which provides the upper bound of the shadow value. For each shadow value $C \in [\underline{C}, \overline{C}]$, define

$$(t_{\min}(C), t_{\max}(C)) := \{t : \rho_l(t; C) < \overline{\rho}\}. \tag{8}$$

By construction, for any t in this open interval, the point $(t, \overline{\rho})$ lies strictly above the indifference curve $\rho_l(\cdot; C)$; and for any $t < t_{\min}(C)$ or $t > t_{\max}(C)$, the point $(t, \overline{\rho})$ lies strictly below the same indifference curve. Since $\rho_l(\cdot; C)$ is U-shaped and reaches a minimum at the point $(t_l^*(\rho), \rho)$ for some $\rho \in [\underline{\rho}, \overline{\rho}]$, the interval in (8) is well defined. In particular, $t_{\min}(\underline{C})$ is the same as t^{sep} defined in equation (6). It is easy to see that $t_{\min}(\cdot)$ is increasing while $t_{\max}(\cdot)$ is decreasing, and $t_{\min}(\overline{C}) = t_{\max}(\overline{C}) = t_l^*(\overline{\rho})$.

For $C \in [\underline{C}, \overline{C}]$, if $D(\overline{\rho})$ belongs to the interval $[t_{\min}(C), t_{\max}(C)]$, then the indifference curve $\rho_l(\cdot; C)$ must intersect the dividing line $D(\cdot)$ at some point in the (t, ρ)-space. We label this intersection point as $(t_{\text{div}}, \rho_{\text{div}})$. If $D(\overline{\rho})$ is outside the interval $[t_{\min}(C), t_{\max}(C)]$, then ρ_{div} would exceed $\overline{\rho}$. In this case we define $\rho_{\text{div}}(C) = \overline{\rho}$ and define $t_{\text{div}}(C)$ accordingly to maintain continuity in C:

$$t_{\text{div}}(C) = \begin{cases} t_{\min}(C) & \text{if } D(\overline{\rho}) < t_{\min}(C) \\ t_{\max}(C) & \text{if } D(\overline{\rho}) > t_{\max}(C) \\ D\left(\rho_l\left(t_{\text{div}}(C); C\right)\right) & \text{otherwise.} \end{cases} \tag{9}$$

Because the indifference curve $\rho_l(\cdot; C)$ crosses the dividing line once and from below, and because $D(\cdot)$ is downward-sloping, $t_{\text{div}}(C)$ is continuous and decreasing in C. It follows that $\rho_{\text{div}}(C)$ is continuous and increasing in C.

The dividing line $D(\cdot)$ is an important construct of our model because it is the locus of points where the high type's indifference curve is tangent to the low type's indifference curve. Recall that Lemma 2 requires that, in any equilibrium, the low type must continuously quit over an interval of time. As a result, the market reputation payoff $R(\cdot)$ must coincide with the low type's indifference curve $\rho_l(\cdot; C)$ over this interval. Refer to Figure 1 again. The high type's optimal quitting decision can be solved by finding his highest indifference curve that is feasible given $\rho_l(\cdot; C)$, and the solution will lie on the locus of tangency points. This is why the intersection points $t_{\text{div}}(C)$ are good candidates for the high type's equilibrium quitting time. Moreover, because the high type's indifference curve $\rho_h(\cdot)$ is more convex than the low type's $\rho_l(\cdot; C)$ at $t_{\text{div}}(C)$, we have $\rho_h(t) > \rho_l(t; C)$ for all $t \neq t_{\text{div}}(C)$. If the equilibrium support sets Q_l and Q_h both contain $t_{\text{div}}(C)$, and if there is an off-equilibrium deviation to $t' \neq t_{\text{div}}(C)$, the D1 refinement would attribute such a deviation to the low type.

To find an equilibrium, we need to find the shadow value of mimicking that satisfies all the equilibrium conditions (to be discussed in more detail in the next section). In what follows, we let C^* denote the equilibrium shadow value, and denote $t^*_{\min} := t_{\min}(C^*)$, $t^*_{\max} := t_{\max}(C^*)$, $t^*_{\text{div}} := t_{\text{div}}(C^*)$, and $\rho^*_{\text{div}} := \rho_{\text{div}}(C^*)$. The following result provides a characterization of the equilibrium strategies and reputation dynamics under the maintained assumptions.

Proposition 3. *Suppose $t^{\text{sep}} < D(\overline{\rho})$. In any equilibrium, the high type quits with probability 1 at time q_h and the low type quits for $t \in Q_l$, where*

(a) *if* $D(\overline{\rho}) \geq t^*_{\max}$*, then* $q_h = \max\{t^*_{\max}, t^*_h(\overline{\rho})\}$ *and* $Q_l = (t^*_{\min}, t^*_{\max})$*;*

(b) *if* $D(\overline{\rho}) < t^*_{\max}$ *and* $C^* > \underline{C}$*, then* $q_h = t^*_{\text{div}}$ *and* $Q_l = (t^*_{\min}, t^*_{\text{div}}) \cup \{t^*_{\text{div}}\}$*;*

(c) *if* $D(\overline{\rho}) < t^*_{\max}$ *and* $C^* = \underline{C}$*, then* $q_h = t^*_{\text{div}}$ *and* $Q_l = (t^*_{\min}, t^*_{\text{div}}) \cup \{t^*_{\text{div}}\} \cup \{t^*_l(\underline{\rho})\}$*.*

The equilibrium reputation dynamics $R(\cdot)$ *follows:*

$$R(t) = \begin{cases} \overline{\rho} & \text{if } t < q^1_l \\ \rho_l(t; C^*) & \text{if } t \in [q^1_l, q^2_l), \end{cases}$$

with $R(t) = \overline{\rho}$ *for* $t \in [q^2_l, q_h]$ *in case (a), or* $R(q^2_l) = \rho^*_{\text{div}}$ *and* $R(t) = \underline{\rho}$ *for* $t > q^2_l$ *in cases (b) and (c).*

Panel (a) of Figure 2 illustrates case (a) of Proposition 3. In this panel, the dividing line $D(\cdot)$ is located relatively far to the right, and the standard single-crossing property holds to the left of the dividing line. We have $Q_h = \{t^*_{\max}\}$ and $Q_l = (t^*_{\min}, t^*_{\max})$. Although the two equilibrium support sets Q_h and Q_l do not overlap, this is not a conventional separating equilibrium because the low type adopts a mixed strategy to mimic the high type and obtains a reputation payoff that exceeds the payoff lower bound corresponding to $\underline{C}$. We may call this a *mimicking equilibrium*. In a mimicking equilibrium, the high type reveals his type upon quitting and obtains a reputation payoff of $\overline{\rho}$ in the absence of any news. In Figure 2(a), the low type is just indifferent between following his equilibrium strategy and mimicking the high type. But in a mimicking equilibrium, it is also possible that $q_h = t^*_h(\overline{\rho}) > t^*_{\max}$, in which case the high type's quitting decision is not distorted because the low type's incentive constraint is not binding.

In Figure 2(b), the dividing line $D(\cdot)$ intersects the low type's equilibrium indifference curve $\rho_l(\cdot; C^*)$. Case (b) of Proposition 3 shows that there is pooling of the two types at the intersection point t^*_{div}. We may call this a *pooling equilibrium*. In a pooling equilibrium, the low type is indifferent between mimicking by quitting on the time interval $(t^*_{\min}, q_h)$ and pooling by quitting at time q_h. By the double-crossing property, this implies that quitting at q_h is optimal for the high type only if $\rho_h(t)$ is tangent to $\rho_l(t; C^*)$ at $t = q_h$. This requires that q_h must lie on the dividing line $D(\cdot)$, and therefore $q_h = t^*_{\text{div}}$.

Finally, in case (c), the equilibrium shadow value is at its lower bound, $C^* = \underline{C}$. As this equilibrium is qualitatively similar to the fully separating equilibrium discussed above, we call it a *semi-separating equilibrium*. Observe that the fully separating equilibrium is a

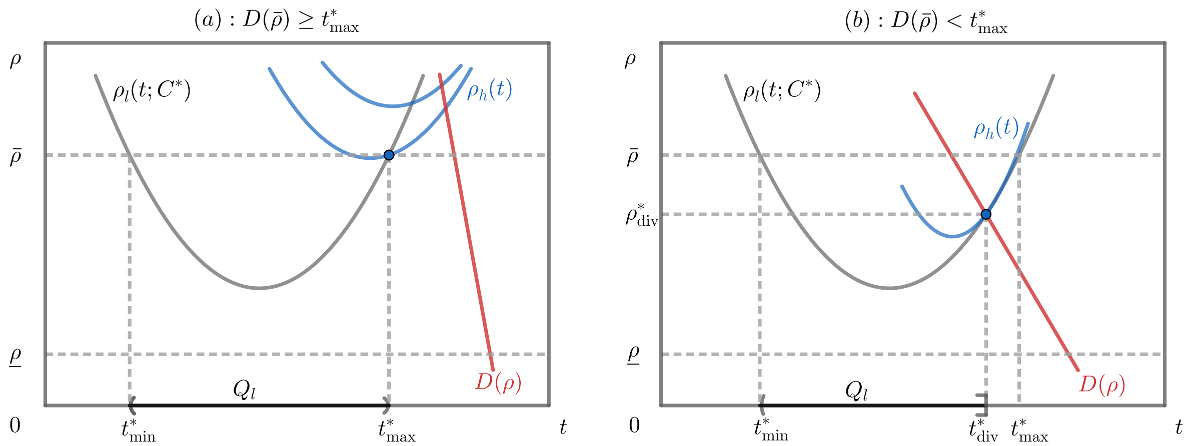

Figure 2. *A mimicking equilibrium in panel (a) and a pooling equilibrium in panel (b).*

limiting case of the semi-separating equilibrium. Suppose $D(\overline{\rho})$ is just slightly larger than t^{sep}. Then the dividing line crosses $\rho_l(\cdot;\underline{C})$ at some $t_{\text{div}}(\underline{C})$ which is just slightly larger than t^{sep}. The low type quits continuously for just a very short interval of time on $(t^{\text{sep}}, t_{\text{div}}(\underline{C})]$. Therefore, the probability of the low type quitting in the mimicking phase is close to 0, while that of quitting at $t_l^*(\underline{\rho})$ is close to 1. The high type pools with the low type at $t_{\text{div}}(\underline{C})$, which is close to t^{sep}, and obtains a reputation payoff $\rho_{\text{div}}(\underline{C})$, which is close to $\overline{\rho}$. Therefore, the high type's payoff is almost the same as his payoff in the fully separating equilibrium when $D(\overline{\rho}) = t^{\text{sep}}$.

4.3. Existence

Proposition 3 provides a characterization of equilibrium taking the equilibrium shadow value C^* as given. In this subsection, we establish that such a C^* always exists.[8] To construct an equilibrium, we need to pin down the high type's equilibrium stopping time q_h and the low type's quitting strategy. We define the low type's strategy by the unconditional distribution of stopping times. Specifically, let $G(\cdot; C)$ be the low type's strategy consistent with $\rho_l(\cdot; C)$, with the interpretation that the low type draws a stopping time q_l from $G(q_l; C)$ at the outset and quits at time q_l if he has not achieved success at that point. As we will see below, $G(\cdot; C)$ is differentiable, and there is a well-defined density $g(\cdot; C)$ almost everywhere. Cases (b) and (c) of Proposition 3(b) show that it is possible that the

[8] When $t^{\text{sep}} \geq D(\overline{\rho})$, it is straightforward to construct a fully separating equilibrium and hence establish its existence.

low type may quit with positive probability mass at certain points, in which case $G(\cdot;C)$ jumps up.

An overview of our equilibrium construction goes as follows. Once we fix the shadow value C, we can pin down the reputation dynamics from the indifference curve $\rho_l(\cdot;C)$. Given $\rho_l(\cdot;C)$, equation (8) gives the interval $(t_{\min}(C), t_{\max}(C))$ such that $\rho_l(t;C) < \overline{\rho}$ if and only if t belongs to that interval, and equation (9) gives the intersection point $(t_{\text{div}}(C), \rho_{\text{div}}(C))$ such that $\rho_l(\cdot;C)$ crosses the dividing line $D(\cdot)$. Proposition 3 shows that the high type must quit at

$$q_h = \begin{cases} \max\{t_{\max}(C), t_h^*(\overline{\rho})\} & \text{if } D(\overline{\rho}) \geq t_{\max}(\underline{C}), \\ \min\{t_{\text{div}}(C), t_{\max}(C)\} & \text{if } D(\overline{\rho}) < t_{\max}(\underline{C}). \end{cases}$$

The quitting time q_h is a continuous function of C. Similarly, the low type quits continuously on the open interval

$$\left(q_l^1, q_l^2\right) = \begin{cases} (t_{\min}(C), t_{\max}(C)) & \text{if } D(\overline{\rho}) \geq t_{\max}(\underline{C}), \\ (t_{\min}(C), \min\{t_{\text{div}}(C), t_{\max}(C)\}) & \text{if } D(\overline{\rho}) < t_{\max}(\underline{C}), \end{cases}$$

which also varies continuously with C. The support $Q_l(C)$ includes the open interval (q_l^1, q_l^2) and possibly with atoms at q_l^2 and $t_l^*(\overline{\rho})$. The strategy $G(\cdot;C)$ must be consistent with the equilibrium reputation dynamics given by $\rho_l(\cdot;C)$ and must integrate to 1 on the support $Q_l(C)$. This allows us to pin down the equilibrium shadow value C^*.

Proposition 4. *There exists an equilibrium as characterized in Proposition 3.*

While there always exists an equilibrium in our model, the equilibrium shadow value C^* need not be unique in the case of $\theta_B < \theta_G$. When equilibrium is not unique, we adopt an equilibrium selection rule that chooses the largest value of C^* that can constitute an equilibrium. Because the shadow value C^* can serve as a utility index, this is equivalent to selecting the best equilibrium for the low type. As multiple equilibria can only occur when the high type pools with the low type at $(t_{\text{div}}(C^*), \rho_{\text{div}}(C^*))$ on the dividing line, and as the high type's indifference curve is tangent to the low type's along the dividing line, the best equilibrium for the low type is also the best equilibrium for the high type. We will assume this equilibrium selection rule when we discuss the comparative statics of this model.

5. Discussion

5.1. Reputation dynamics and the value of perseverance

Because $\rho_l(\cdot; C^*)$ is quasi-convex, the equilibrium reputation $R(\cdot)$ may be decreasing, increasing, or decreasing and then increasing in the mimicking phase. Observe that $R(\cdot)$ can be strictly increasing only if $t^*_{\min}$ is on the upward-sloping part of the equilibrium indifference curve $\rho_l(\cdot; C^*)$. A necessary condition for this to hold is that $t^*_l(\overline{\rho}) = 0$. If $t^*_l(\overline{\rho}) > 0$, or equivalently,

$$p_0 > \frac{r(\overline{\rho} - W)}{\lambda(B - \overline{\rho})},$$

then the equilibrium reputation $R(\cdot)$ must be either strictly decreasing or non-monotone on $(q^1_l(C^*), q^2_l(C^*))$. Note that this condition is likely to be satisfied if the agent is sufficiently forward-looking (a low r), suggesting that quitting early is perceived as a sign of competence under mild conditions.

In contrast, the reputation value of perseverance is more nuanced. Specifically, we say that perseverance has a positive reputation value if the equilibrium reputation is non-monotone and eventually turns increasing at some point. The following statement establishes a necessary and sufficient condition for perseverance to have a positive reputation value. Let $\hat{\rho} \in [\underline{\rho}, \overline{\rho}]$ be the reputation payoff such that $D(\rho) \geq t^*_h(\rho) \geq t^*_l(\rho)$ if and only if $\rho \geq \hat{\rho}$, and $D(\rho) \leq t^*_h(\rho) \leq t^*_l(\rho)$ if and only if $\rho \leq \hat{\rho}$.[9]

Proposition 5. *Suppose $t^*_l(\overline{\rho}) > 0$. The equilibrium market reputation $R(\cdot)$ is strictly decreasing on $(q^1_l(C^*), q^2_l(C^*))$ if and only if $\rho^*_{\text{div}} \leq \hat{\rho}$, and is non-monotone on $(q^1_l(C^*), q^2_l(C^*))$ if and only if $\rho^*_{\text{div}} > \hat{\rho}$.*

Proposition 5 is stated in terms of an endogenous variable ρ^*_{div}. Since ρ^*_{div} is the point at which the equilibrium indifference curve and the dividing line intersect, its value depends crucially on where the dividing line is located. Roughly speaking, when the dividing line is located far to the right, it is not binding, and the standard single-crossing property prevails in the relevant space. In this case, the high type is always more willing to persist longer than the low type, and we have a mimicking equilibrium as in Figure 3(a). In this type of equilibrium, the equilibrium reputation follows a non-monotone path, and the reputation value of perseverance is maximized. As the dividing line shifts to the left, however,

[9]Lemma 3 in the Appendix establishes that such $\hat{\rho}$ exists. If $D(\rho) > t^*_h(\rho) > t^*_l(\rho)$ for all $\rho \in [\underline{\rho}, \overline{\rho}]$, we define $\hat{\rho} = \underline{\rho}$; if $D(\rho) < t^*_h(\rho) < t^*_l(\rho)$ for all $\rho \in [\underline{\rho}, \overline{\rho}]$, we define $\hat{\rho} = \overline{\rho}$.

the reputation value of perseverance gradually diminishes and eventually dissipates as in Figure 3(b).

Recall that the dividing line is the locus of points at which the incentives of the two types flip. The argument thus ultimately boils down to how quickly the high type becomes pessimistic by receiving no news, which is determined by the news arrival rates (θ_G, θ_B).

Corollary 1. *Suppose $t_l^*(\overline{\rho}) > 0$. For each $\theta = \theta_B$, there exists $\hat{\theta}_G^2 > \hat{\theta}_G^1 > \theta_B$ such that if $\theta_G \geq \hat{\theta}_G^2$ then $R(\cdot)$ is decreasing in the mimicking phase, and if $\theta_G < \hat{\theta}_G^1$ then $R(\cdot)$ is decreasing then increasing in the mimicking phase.*

Whenever we are in a "good-news environment," in which the arrival rate of good news is much higher than the arrival rate of bad news (i.e., when $\theta_G > \hat{\theta}_G^2$), $R(\cdot)$ is strictly decreasing during the mimicking phase. In such an environment, the high type becomes pessimistic very quickly and more reluctant to continue further. Perseverance is then interpreted as a sign of incompetence or "stubbornness" by the market. In contrast, for $\theta_G < \hat{\theta}_G^1$ (which always holds if $\theta_G \leq \theta_B$), bad news is relatively more frequent than good news. The high type becomes pessimistic less quickly (or even becomes optimistic) by having found "no flaws" in the project. As a consequence, the high type is willing to persist longer than the low type. This prolongs the mimicking phase as the high type holds out until the low type completely exits, making the equilibrium reputation swing upward at some point.[10]

5.2. Prior reputation and the shadow value of mimicking

The news arrival rates (θ_G, θ_B) affect the equilibrium dynamics largely by inducing a shift in the dividing line. Aside from the dividing line, the form of equilibrium depends also on the shadow value of mimicking (which corresponds to the low type's expected payoff). In general, the shadow value is determined by the relative demand and supply of mimicking opportunities: the shadow value is high when the supply is scarce relative to the demand and decreases as it gets more abundant. Since mimicking opportunities are supplied by high-type agents who receive bad news, the prior type distribution, captured by μ_0, plays an essential role in determining the shadow value. This implication stands in contrast to the standard setup that always selects the LCSE whose allocation is independent of the prior belief, leading to a disturbing implication that signaling-based reputation models cannot explain why and how reputation matters in the first place.

[10]If $\theta_G \in (\hat{\theta}_G^1(\theta_B), \hat{\theta}_G^2(\theta_B))$, then $\hat{\rho} \in (\underline{\rho}, \overline{\rho})$. Whether $R(\cdot)$ is decreasing or non-monotone depends on whether the equilibrium value of ρ^*_{div} is less than or greater than $\hat{\rho}$.

To see the role of the prior reputation, suppose that μ_0 is small and arbitrarily close to 0. Since there are almost no high-type agents providing mimicking opportunities, $g(\cdot; C)$ becomes arbitrarily small for any C. In the limit, the low type quits in the neighborhood of $t_l^*(\underline{\rho})$ with probability close to 1, and his payoff is either equal to $\underline{C}$ (when $D(\underline{\rho}) < t_l^*(\underline{\rho})$) or converges to $\underline{C}$ (when $D(\underline{\rho}) \geq t_l^*(\underline{\rho})$). As μ_0 increases, however, more high-type agents receive bad news and quit over time, allowing the low type to quit at a faster rate in the mimicking phase to maintain the same reputation. Then, to satisfy the equilibrium condition that the strategy must integrate to 1, the equilibrium indifference curve must shift up. As $\mu_0 \to 1$, $g(\cdot; C)$ diverges to infinity for any C, and the mimicking phase must degenerate to a single point to satisfy the equilibrium condition. The form of equilibrium is then determined by $D(\overline{\rho})$. If $D(\overline{\rho}) \in (t^{\text{sep}}, t_l^*(\overline{\rho})]$, there is a pooling equilibrium in which the mimicking phase degenerates to $D(\overline{\rho})$, with the high type also quitting at $q_h(C^*) = D(\overline{\rho})$. If $D(\overline{\rho}) > t_l^*(\overline{\rho})$, there is a mimicking equilibrium in which the mimicking phase degenerates to $t_l^*(\overline{\rho})$, with the high type quitting at his complete-information optimum $q_h(C^*) = t_h^*(\overline{\rho})$.[11]

This argument indicates that a higher prior reputation raises the shadow value of mimicking and shifts the equilibrium reputation to a higher indifference curve of the low type, which directly benefits the low type. The high type also benefits from this upward shift as it raises the reputation payoff when he quits before he succeeds. The following statement summarizes this observation, which states that a higher prior reputation leads to a higher expected payoff, i.e., the value of reputation emerges endogenously in this setting.[12]

Proposition 6. *The expected payoffs of both types, given by (2), are increasing in μ_0.*

5.3. Efficiency and policy implications: the valley of death

In our model, a high-type agent quits immediately when he receives bad news, while he continues indefinitely when he receives good news. Given the additional information they have, these decisions are efficient under the maintained assumptions. Our focus is thus on the decisions of those who have observed no news. In what follows, we say

[11] If $D(\overline{\rho}) \leq t^{\text{sep}}$, the equilibrium would still be a separating equilibrium for any μ_0.

[12] Here, the expected payoffs of the two types are evaluated *ex ante* at time 0 before any news is observed. Ishida and Suen (2024) make a related but distinct observation that the payoffs of strategic types are increasing in the prior belief. Note that the welfare criterion used by Ishida and Suen (2024) is evaluated *ex post* (after the realization of stochastic shocks) and excludes behavioral types. The welfare of those behavioral types cannot be meaningfully analyzed in that setting because their signaling choices are exogenously imposed with no well-defined payoff functions for these types.

that an agent quits too early (late) if he quits earlier (later) than his complete-information optimum ($t_l^*(\underline{\rho})$ for the low type and $t_h^*(\overline{\rho})$ for the high type). In the standard setup, the D1 criterion always selects the LCSE in which the low type chooses his complete-information optimum; the only distortion that arises is therefore that the high type needs to choose a higher signaling action to separate from the low type. This is not the case in our setting where the signaling distortion is a two-way interaction involving both types and runs in both directions (either too early or too late) due to the presence of the mimicking phase.

To illustrate the efficiency and policy implications stemming from mimicking behavior, it is instructive to introduce a principal (e.g., a government) to the model and consider a simple tax scheme where the principal sets a deadline and imposes a lump-sum tax if the agent continues the project past the deadline. Chen et al. (2021) argue that such a scheme is potentially welfare-improving because low-type agents may choose to quit at the deadline, freeing high-type agents from the need to separate from those low-type agents.[13] For illustration, suppose $t_l^*(\underline{\rho}) < t_h^*(\overline{\rho})$, so that the low type would quit earlier than the high type under complete information. Suppose further that the principal sets the deadline at $t_l^*(\underline{\rho})$ and imposes a lump-sum tax that is just high enough to deter low-type agents from continuing further. If this scheme works as designed, all low-type agents quit by $t_l^*(\underline{\rho})$, and we have $R(t) = \overline{\rho}$ for $t > t_l^*(\underline{\rho})$. Given this, since $t_l^*(\overline{\rho}) < t_l^*(\underline{\rho})$, the optimal deviation for the low type is to quit immediately after $t_l^*(\underline{\rho})$ and earn a payoff of $\overline{\rho}$. As such, if the tax is set at $\overline{\rho}$, the low type quits at $t_l^*(\underline{\rho})$ with probability 1, which enables the high type to quit at $t_h^*(\overline{\rho}) > t_l^*(\underline{\rho})$. If $t_{\max}^* > t_h^*(\overline{\rho})$, such that the high type is forced to persist in order to separate from the low type, this improves welfare by eliminating the over-investment incentive by the high type. Since the over-investment problem is typically the only welfare concern in the LCSE of the standard setup, this simple tax scheme would implement the efficient outcome.

Unfortunately, this type of intervention does not work, and may even backfire, in the context of signaling vision. In our setting, due to the mimicking phase, the low type earns a payoff above the complete-information level. It is straightforward to show that this is the case if $t_h^*(\overline{\rho}) > t_l^*(\underline{\rho})$ as assumed here.[14] This implies that the proposed intervention, which forces low-type agents to quit at $t_l^*(\underline{\rho})$, would surely decrease their expected payoff.

[13] Chen et al. (2021) considers a more complicated environment where the payoff of success may also depend on the market belief, and the low type's behavior is also distorted.

[14] Observe that $t_l^*(\underline{\rho}) < t_h^*(\overline{\rho})$ implies $t_l^*(\underline{\rho}) < D(\underline{\rho})$. Then, the original equilibrium (before imposing the tax scheme) must be either pooling or mimicking with $C^* > \underline{C}$.

Anticipating this, they must start quitting even earlier with the intervention than without. In effect, this type of intervention can be welfare-reducing as it limits the effective supply of mimicking opportunities by imposing an upper bound on the mimicking phase. This in turn intensifies the competition for mimicking opportunities, thereby diminishing the shadow value of mimicking. This is a reflection of the fact that the signaling distortion involves both types and runs in both directions.

In venture financing, many startups face difficulty raising follow-on ("series A") funding after initial seed funding. This funding gap is known as the "valley of death," as most startups struggle to survive beyond this phase. Formally, the valley of death can be seen as an exogenous increase in the operation cost of running a venture startup and in this sense functions quite similarly to the tax scheme discussed above. This line of argument then implies that despite its apparent cost of resulting in premature termination of startups, there can be a bright side of the valley of death, as a screening device to weed out wasteful signaling activities. Our analysis suggests, however, that the validity of such an argument depends heavily on the agent's motive for signaling, i.e., what type of attribute or ability the agent wishes to convey to the market. In markets where vision matters and is a valuable commodity, there is an additional dimension to consider, which is the possibility of inefficient mimicking, and the potential benefit of the valley of death as a screening device can easily be outweighed by its negative impact on mimicking behavior.

6. Conclusion

Economists have studied how visionary—and sometimes even overconfident—leaders can improve the performance of the organizations they lead by motivating subordinates to work harder or facilitating coordination (Rotemberg and Saloner, 2000; Van den Steen, 2005; Bolton et al., 2013). In this paper, we take a step further and model "vision" as the ability to correctly foresee the future course of action, endogenously inducing extreme beliefs and extreme actions. A visionary leader in our model should know when to persist (if his initial strategic direction is revealed to be good) and when to recede (if his initial vision turns out to be misguided). Furthermore, because there is a demand for people with vision, a leader (visionary or not) has an incentive to induce the market to believe that he is visionary, leading to a signaling game that is the focus of this paper. The signaling game analyzed here is different from the standard setup, in that the quitting decisions of informed leaders provide opportunities for leaders without vision to mimic them. This produces distortions in the timing of project abandonment in both directions and can result in non-monotone reputation dynamics. The non-monotonicity of reputation dynamics in turn provides a unified explanation for why and how both decisiveness (quitting early) and perseverance (persisting through adversities) can be perceived as signs of competence. Our analysis also provides a mechanism through which the agent benefits from a higher prior reputation.

References

Bass, Bernard M., *Leadership and Performance Beyond Expectations*, New York: Free Press, 1985.

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

Bolton, Patrick, Markus K. Brunnermeier, and Laura Veldkamp, "Leadership, Coordination, and Corporate Culture," *Review of Economic Studies*, 2013, *80* (2), 512–537.

Burns, James MacGregor, *Leadership*, New York: Harper and Row, 1978.

Chen, Chia-Hui, Junichiro Ishida, and Wing Suen, "Reputation Concerns in Risky Experimentation," *Journal of the European Economic Association*, 2021, *19* (4), 1981–2021.

__, __, and __, "Signaling under Double-Crossing Preferences," *Econometrica*, 2022, *90* (3), 1225–1260.

__, __, and __, "Signaling under Double-Crossing Preferences: The Case of Discrete Types," *Journal of Mathematical Economics*, 2024, *114*, 103046.

Cho, In-Koo and David M. Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 1987, *102* (2), 179–221.

Collins, Jim, *Good to Great: Why Some Companies Make the Leap...And Others Don't*, New York: Harper Business, 2001.

Damiano, Ettore, Hao Li, and Wing Suen, "Learning While Experimenting," *Economic Journal*, 2020, *130* (625), 65–92.

Dilme, Francesc and Fei Li, "Dynamic Signaling with Dropout Risk," *American Economic Journal: Microeconomics*, 2016, *8* (1), 57–82.

Duckworth, Angela, *Grit: The Power of Passion and Perseverance*, New York: Scribner, 2016.

Godin, Seth, *The Dip*, New York: Portfolio, 2007.

Ishida, Junichiro and Wing Suen, "Pecuniary Emulation and Invidious Distinction: Signaling under Behavioral Diversity," *Games and Economic Behavior*, September 2024, *147*, 449–459.

Merriman, Kimberly K., "Leadership and Perseverance," in J. Marques and S. Dhiman, eds., *Leadership Today*, Springer, 2017, pp. 335–350.

Rotemberg, Julio J. and Garth Saloner, "Visionaries, Managers, and Strategic Direction," *RAND Journal of Economics*, 2000, *31* (4), 693–716.

Spence, Michael, "Job Market Signaling," *Quarterly Journal of Economics*, 1973, *87* (3), 355–374.

Van den Steen, Eric, "Organizational Beliefs and Managerial Vision," *Journal of Law, Economics, and Organization*, 2005, *21* (1), 256–283.

Appendix

Proof of Proposition 1. Part (a) follows from Assumption 1, because $V > \overline{\rho}$ means that the payoff from continuing for the informed-good type is strictly larger than the maximum payoff he can earn from quitting.

For part (b), we first establish that there is no equilibrium in which no type quits for some interval of time. Suppose on the contrary that no type quit for $t \in (t', t'')$. At time t'', if any type quits, the informed-bad type must also quit because he has the lowest return to continuing the project. Then, under D1, a deviation to quitting at time $t'' - \varepsilon$ is attributed to the informed-bad type, who must be a high type. This means that $\mu(t'' - \varepsilon) = 1$ for small $\varepsilon > 0$. Since $\overline{\rho}$ is the highest payoff the informed-bad type can earn, he would surely deviate and quit slightly earlier than at t''.

The above argument implies that there are no quitting times that are off the equilibrium path as long as the high type stays in the game with positive probability. Given this, suppose that the high type observes a bad signal at time t' but waits until time $t'' > t'$ to quit. This is optimal only if

$$\int_0^{t''-t'} e^{-rs} rW \,\mathrm{d}s + e^{-r(t''-t')} R(t'') \geq R(t'). \tag{10}$$

For this to hold, we must have $R(t') < \overline{\rho}$ by Assumption 1, implying that $\mu(t') < 1$. Hence t' must belong to the support of the low type's equilibrium quitting times. A necessary condition for this is that the low type must be weakly better off quitting at time t' than waiting until time t'':

$$\begin{aligned} R(t') \geq p_l(t') \int_0^{t''-t'} & e^{-(\lambda+r)s}(\lambda B + rW)\,\mathrm{d}s + (1 - p_l(t')) \int_0^{t''-t'} e^{-rs} rW \,\mathrm{d}s \\ & + e^{-r(t''-t')} \big(p_l(t') e^{-\lambda(t''-t')} + 1 - p_l(t')\big) R(t''). \end{aligned}$$

Adding this inequality to (10) gives

$$\begin{aligned} 0 &\geq p_l(t') \big((1 - e^{-(\lambda+r)(t''-t')})V - (1 - e^{-r(t''-t')})W - e^{-r(t''-t')}(1 - e^{-\lambda(t''-t')})R(t'')\big) \\ &\geq p_l(t') \big((1 - e^{-(\lambda+r)(t''-t')})V - (1 - e^{-r(t''-t')})W - e^{-r(t''-t')}(1 - e^{-\lambda(t''-t')})\overline{\rho}\big) \\ &= p_l(t') \big((1 - e^{-(\lambda+r)(t''-t')})(V - \overline{\rho}) + (1 - e^{-r(t''-t')})(\overline{\rho} - W)\big), \end{aligned}$$

which contradicts Assumption 1. This shows that in any equilibrium, the informed-bad type has no incentive to wait and stay in the game. ■

Proof of Lemma 1. Since $\dot{\rho}_a(t) = MRS_a(t, \rho_a(t))$, we have

$$\ddot{\rho}_a(t) = -\dot{p}_a(t)[\lambda(B - \rho_a(t)) + k_a\theta_G(V - \rho_a(t))] + [p_a(t)(\lambda + k_a\theta_G) + r]\dot{\rho}_a(t).$$

For $a = l$, $\dot{p}_a(t) < 0$. For $a = h$ and $\Delta_\theta < \lambda$, $\dot{p}_a(t) < 0$. In these two cases, $\ddot{\rho}_a(t) > 0$ when $\dot{\rho}_a(t) = 0$. This shows that the corresponding indifference curves are quasi-convex.

Let $\phi(t) := p_l(t)/p_h(t)$. The condition that $MRS_h(t, \rho) = MRS_l(t, \rho)$ is equivalent to

$$\phi(t) = \frac{\lambda(B - \rho) + \theta_G(V - \rho)}{\lambda(B - \rho)}. \tag{11}$$

The right-hand-side is strictly greater than 1 for all $\rho \in [\underline{\rho}, \overline{\rho}]$ by Assumption 1. Since $\phi(0) = 1$, the left-hand-side is smaller than the right-hand-side when t is sufficiently small. If $\Delta_\theta \geq 0$, then $\phi(t)$ is non-increasing in t and hence $MRS_h(t, \rho) < MRS_l(t, \rho)$ for any (t, ρ). If $\Delta_\theta < 0$, on the other hand, $\phi(\cdot)$ is strictly increasing in t and there is a unique $D(\rho)$ for any given $\rho \in [\underline{\rho}, \overline{\rho}]$ such that the two sides of the above are equal for $t = D(\rho)$. Moreover, since the right-hand-side of (11) decreases in ρ, $D(\cdot)$ is decreasing.

Let $\delta(t)$ represent the inverse of $D(\cdot)$. We show that the indifference curve $\rho_l(t)$ of the low type crosses the dividing line $\delta(t)$ once and from below when $\Delta_\theta < 0$. If the two curves cross at some t', this would imply that $MRS_h(t, \rho_l(t)) < MRS_l(t, \rho_l(t))$ for $t < t'$ and $MRS_h(t, \rho_l(t)) > MRS_l(t, \rho_l(t))$ for $t > t'$. In other words, $MRS_h(\cdot, \rho_l(\cdot)) - MRS_l(\cdot, \rho_l(\cdot))$ would be single-crossing from below.

The dividing line $\delta(\cdot)$ is given by the value of ρ that solves equation (11). This gives

$$\delta(t) = B - \frac{\theta_G(B - V)}{\theta_G - (\phi(t) - 1)\lambda},$$

where the denominator in the fraction is positive because $\theta_G - (\phi(t) - 1)\lambda = \theta_G[1 - (V - \delta(t))/(B - \delta(t))]$ and $\delta(t) \leq \overline{\rho}$. The slope of the dividing line is

$$\dot{\delta}(t) = -\frac{\dot{\phi}(t)\lambda\theta_G r(B - W)}{(\lambda + r)(\theta_G - (\phi(t) - 1)\lambda)^2}.$$

Since $\Delta_\theta < 0$,

$$\dot{\phi}(t) = \frac{p_l(t)}{p_h(t)}[(1 - p_h(t))(\lambda - \Delta_\theta) - (1 - p_l(t))\lambda] > p_l(t)(\phi(t) - 1)\lambda.$$

Therefore

$$\dot{\delta}(t) < -\frac{p_l(t)(\phi(t) - 1)\lambda^2\theta_G r(B - W)}{(\lambda + r)(\theta_G - (\phi(t) - 1)\lambda)^2}. \tag{12}$$

An indifference curve $\rho_l(\cdot)$ for the low type is obtained from the solution to the the differential equation $\dot{\rho}_l(t) = MRS_l(t, \rho_l(t))$, given by equation (7) in the text:

$$\rho_l(t; C) = p_l(t)V + (1 - p_l(t))W + \frac{p_l(t)}{p_0 e^{-(\lambda+r)t}} C,$$

where $C > 0$ is a constant of integration indexing the relevant indifference curve. Differentiating this solution leads to

$$\dot{\rho}_l(t; C) = -p_l(t)(1 - p_l(t))\lambda(V - W) + (r + p_l(t)\lambda)\frac{p_l(t)}{p_0 e^{-(\lambda+r)t}} C.$$

When an indifference curve meets the dividing line, we have $\rho_l(t_l; C) = \delta(t)$, and therefore

$$\begin{aligned} \frac{p_l(t)}{p_0 e^{-(\lambda+r)t}} C &= \delta(t) - W - p_l(t)(V - W) \\ &= (B - W)\left(1 - \frac{(r + p_l(t)\lambda)\theta_G}{(\lambda + r)(\theta_G - (\phi(t) - 1)\lambda)} + \frac{p_l(t)\lambda^2(\phi(t) - 1)}{(\lambda + r)(\theta_G - (\phi(t) - 1)\lambda)}\right). \end{aligned} \tag{13}$$

Therefore the slope of the indifference curve when it crosses the dividing line is

$$\begin{aligned} \dot{\rho}_l(t) &= -p_l(t)(1 - p_l(t))\lambda(V - W) + (r + p_l(t)\lambda)\left(\delta(t) - W - p_l(t)(V - W)\right) \\ &= r(B - W)\left(1 - \frac{(r + p_l(t)\lambda)\theta_G}{(r + \lambda)(\theta_G - (\phi(t) - 1)\lambda)}\right). \end{aligned}$$

Since $C > 0$, the right-hand-side of (13) is positive. This implies

$$\dot{\rho}_l(t) > -r(B - W)\frac{p_l(t)\lambda^2(\phi(t) - 1)}{(r + \lambda)(\theta_G - (\phi(t) - 1)\lambda)}. \tag{14}$$

Combining (14) and (12) gives

$$\dot{\rho}(t) - \dot{\delta}(t) > \frac{r(B - W)p_l(t)\lambda^2(\phi(t) - 1)}{(r + \lambda)(\theta_G - (\phi(t) - 1)\lambda)}\left(-1 + \frac{\theta_G}{\theta_G - (\phi(t) - 1)\lambda}\right) > 0.$$

This shows that the indifference curve of the low type $\rho(t)$ cuts the dividing line $\delta(t)$ (in the (t, ρ)-space) at most once and from below. ■

Proof of Proposition 2. We show that the high type cannot quit at two (or more) different points in time. Suppose to the contrary that both q' and q'' belong to Q_h, with $q'' > q'$. This means that the high type remains in the game with positive probability at least up to q''. At any time in the interval (q', q''), the high type may receive bad news and exit the game. This implies that the low type must quit continuously for all $t \in (q', q'')$. Otherwise there would be some t' in this interval such that $R(t') = \overline{\rho}$. But since the high type is indifferent between quitting at q' and q'', the points $(q', R(q'))$ and $(q'', R(q''))$ are on the same indifference curve $\rho_h(\cdot)$ of the high type. The fact that $\rho_h(\cdot)$ is quasi-convex then implies that $(t', \overline{\rho})$ is strictly above this indifference curve, violating the optimality of quitting at q' or q''.

Since the low type quits continuously in the interval (q', q''), the reputation dynamics $R(\cdot)$ must coincide the indifference curve $\rho_l(\cdot)$ of the low type in this interval. If $\theta_B \geq \theta_G$, the single-crossing property holds by Lemma 1. Because $MRS_h(q', R(q')) < MRS_l(q', R(q')) = \dot{R}(q')$ in this case, the high type would prefer to continue rather than quit at q', a contradiction. If $\theta_B < \theta_G$, the double-crossing property holds. Because $MRS_h(\cdot, \rho_l(\cdot))$ crosses $MRS_l(\cdot, \rho_l(\cdot))$ at most once, the fact that $MRS_a(t, R(t)) = \dot{R}(t)$ holds at both $t = q'$ and $t = q''$ for type $a = l$ implies that the same cannot hold for type $a = h$. This violates the assumption that q' and q'' are optimal quitting times for the high type. ■

Proof of Lemma 2. We first claim that Q_l is not a singleton if $t^{\text{sep}} < D(\overline{\rho})$. Suppose $t^{\text{sep}} < D(\overline{\rho})$ and, contrary to the proposition, the low type quits with probability 1 at some time q_l. By Proposition 2, the high type also quits with probability 1 at some time q_h. There are three cases to consider.

1. If $q_l = q_h = 0$, the marginal rate of substitution at the point $(0, \rho_0)$ is strictly lower for the high type than for the low type because $D(\rho_0) > 0$. Under D1, an agent who deviates by quitting a little later would be interpreted as a high type, making such deviation profitable.
2. If $0 < q_l \leq q_h$, because the uninformed-bad type quits continuously on the interval $(0, q_h)$, the reputation payoff from quitting slightly before q_l is $\overline{\rho}$, while the payoff from quitting at q_l is bounded below $\overline{\rho}$, making such deviation profitable for the low type.
3. If $q_l > q_h$, because the quitting times are fully separating, we must have $q_l = t_l^*(\underline{\rho})$. If $q_h \leq t^{\text{sep}}$, $MRS_h(q_h, \overline{\rho})$ must be negative and must be lower than $MRS_l(q_h, \overline{\rho})$. Under D1, the reputation payoff for deviation slightly later than q_h will be $\overline{\rho}$, and the high type would gain from such deviation. If $q_h > t^{\text{sep}}$, the low type would strictly prefer to quit at $(q_h, \overline{\rho})$ rather than $(q_l, \underline{\rho})$.

Fix the high type's strategy at time q_h. Suppose there is some $q_l \in Q_l$ such that $q_l > q_h$.

Then the reputation payoff must satisfy $R(q_l) = \underline{\rho}$. If $q_l \neq t_l^*(\underline{\rho})$, the low type could strictly improve his payoff by quitting at $t_l^*(\underline{\rho})$. This shows that the only quitting time in Q_l that can strictly exceed q_h is $q_l = t_l^*(\underline{\rho})$. Since we have already established that Q_l is not a singleton, there exists $q' \in Q_l$ such that $q' \leq q_h$.

Suppose there is only one $q' \in Q_l$ such that $q' \leq q_h$. Then $Q_l = \{q', t_l^*(\underline{\rho})\}$. The low type can be indifferent between these two quitting times only if $R(q') > \underline{\rho}$, which implies $q_h = q'$. But since the high type receives bad news continuously and quits in the interval $(0, q')$, we have $R(t) = \overline{\rho}$ for $t < q'$. This would induce the low type to deviate by quitting a little earlier than q', a contradiction. This argument shows that there exist at least two distinct elements of Q_l that are both less than or equal to q_h.

Suppose that $q' < q'' \leq q_h$ and $q', q'' \in Q_l$, but there exists $t' \in (q', q'')$ that does not belong to Q_l. Since $t' < q_h$, we have $R(t') = \overline{\rho}$. Because the low type is indifferent between quitting at q' and q'', there is an indifference curve $\rho_l(\cdot)$ connecting $(q', R(q'))$ and $(q'', R(q''))$. The quasi-convexity of $\rho_l(\cdot)$ implies that $(t', \overline{\rho})$ lies strictly above this indifference curve, violating the optimality of quitting at q' or q''. This shows that the set $\{q_l \in Q_l : q_l \leq q_h\}$ is an open interval. ■

Proof of Proposition 3. We first establish the following properties that prove to be quite useful for the subsequent analysis.

Lemma 3. *If $\Delta_\theta \geq 0$, then whenever $t_h^*(\overline{\rho})$ is interior, $t_h^*(\rho) > t_l^*(\rho)$ for all $\rho \in [\underline{\rho}, \overline{\rho}]$. If $\Delta_\theta < 0$, then whenever $t_h^*(\overline{\rho})$ is interior, one of the following properties hold:*

(a) $D(\rho) > t_h^*(\rho) > t_l^*(\rho)$ for all $\rho \in [\underline{\rho}, \overline{\rho}]$;

(b) $D(\rho) < t_h^*(\rho) < t_l^*(\rho)$ for all $\rho \in [\underline{\rho}, \overline{\rho}]$; or

(c) there exists a unique $\hat{\rho} \in [\underline{\rho}, \overline{\rho}]$ such that $D(\rho) > t_h^*(\rho) > t_l^*(\rho)$ iff $\rho > \hat{\rho}$ and $D(\rho) < t_h^*(\rho) < t_l^*(\rho)$ for $\rho < \hat{\rho}$.

Proof. Since $t_h^*(\overline{\rho})$ is either equal to zero or infinity when $\Delta_\theta \geq \lambda$, we only need to consider the case of $\Delta_\theta < \lambda$. In this case, $MRS_a(\cdot, \rho)$ is strictly decreasing. If $t_h^*(\rho) > t_l^*(\rho)$, then $MRS_l(t_l^*(\rho), \rho) > MRS_h(t_h^*(\rho), \rho) = 0$, which implies $t_h^*(\rho) < D(\rho)$ by Lemma 1. If $t_l^*(\rho) > t_h^*(\rho)$, then $MRS_l(t_h^*(\rho), \rho) > MRS_h(t_l^*(\rho), \rho) = 0$, which implies $t_h^*(\rho) > D(\rho)$.

If $\Delta_\theta \geq 0$, then the single-crossing property holds and $t_h^*(\rho) > t_l^*(\rho)$ for any ρ. If $\Delta_\theta < 0$, we show that $t_h^*(\cdot) - t_l^*(\cdot)$ is strictly increasing. This would imply that if $t_h(\underline{\rho}) >$

$t_l(\overline{\rho})$, then $t_h(\rho)$ will remain higher than $t_l(\rho)$ for any ρ, corresponding to case (a) of the lemma; and if $t_h(\underline{\rho}) \leq t_l(\underline{\rho})$, then either case (b) or case (c) obtains.

To show that $t_h^*(\cdot) - t_l^*(\cdot)$ is strictly increasing, we differentiate equation (5) to obtain:

$$\dot{p}_l(t_l^*(\rho))\frac{\partial t_l^*(\rho)}{\partial \rho} = \frac{r(B-W)}{\lambda(B-\rho)^2}.$$

Since $\dot{p}_l(t) = -p_l(t)(1-p_l(t))\lambda$, the above equation reduces to

$$\begin{aligned}\frac{\partial t_l^*(\rho)}{\partial \rho} &= -\frac{1}{p_l(t_l^*(\rho))(1-p_l(t_l^*(\rho)))\lambda}\frac{r(B-W)}{\lambda(B-\rho)^2} \\ &= \frac{-\lambda(B-W)}{\lambda(\rho-W)(\lambda(B-\rho)-r(\rho-W))}.\end{aligned}$$

Following similar steps,

$$\begin{aligned}\frac{\partial t_h^*(\rho)}{\partial \rho} &= \frac{-(\lambda(B-W)+\theta_G(V-W))}{(\lambda-\Delta_\theta)(\rho-W)(\lambda(B-\rho)+\theta_G(V-\rho)-r(\rho-W))} \\ &> \frac{-(\lambda(B-W)+\theta_G(V-W))}{\lambda(\rho-W)(\lambda(B-\rho)+\theta_G(V-\rho)-r(\rho-W))},\end{aligned}$$

where the inequality follows because $\Delta_\theta < 0$. Therefore,

$$\begin{aligned}\frac{\partial t_h^*(\rho)}{\partial \rho} - \frac{\partial t_l^*(\rho)}{\partial \rho} &> \frac{1}{\lambda(\rho-W)}\left(\frac{\lambda(B-W)}{\lambda(B-\rho)-r(\rho-W)} - \frac{\lambda(B-W)+\theta_G(V-W)}{\lambda(B-\rho)+\theta_G(V-\rho)-r(\rho-W)}\right) \\ &= \frac{\theta[r(V-W)-\lambda(B-V)]}{\lambda[\lambda(B-\rho)-r(\rho-W)][\lambda(B-\rho)+\theta_G(V-\rho)-r(\rho-W)]} = 0,\end{aligned}$$

where the last equality follows from the definition of V. ∎

Given this result, we now prove the proposition. First, Lemma 2 shows that the low type must quit continuously on some interval (q_l^1, q_l^2) when $t^{\text{sep}} < D(\overline{\rho})$. Because the low type is adopting a mixed strategy, his payoff from quitting at any time in this interval must be constant and is indexed by the equilibrium value C^*. The first-order condition $\dot{u}_l(t;R) = 0$ implies that the market reputation dynamics $R(\cdot)$ must follow the indifference curve of the low type $\rho_l(\cdot;C^*)$ on this interval. For $t < q_l^1$, only the informed-bad type quits on the equilibrium path. Therefore, we have $R(t) = \overline{\rho}$ for such t. The reputation dynamics must be continuous at $t = q_l^1$; otherwise the low type could gain by quitting slightly earlier than q_l^1. But $R(q_l^1) = \overline{\rho}$ and the fact that $(q_l^1, R(q_l^1))$ is on the indifference curve $\rho_l(\cdot;C^*)$ imply, by definition (8) of $t_{\min}(\cdot)$, that $q_l^1 = t_{\min}^*$.

If $D(\overline{\rho}) \geq t^*_{\max}$, we must have $t^*_l(\overline{\rho}) < t^*_h(\overline{\rho}) < D(\overline{\rho})$ by Lemma 3, and the single-crossing property holds at $(t^*_{\max}, \overline{\rho})$. Lemma 2 already establishes that $q_h \geq q^2_l$. If $q^2_l < t^*_{\max}$ and $q^2_l < q_h$, then $R(\cdot)$ would jump up at q^2_l and the low type would gain by quitting a bit later. If $q^2_l = q_h < t^*_{\max}$, then $R(q^2_l)$ would be strictly below $\overline{\rho}$, but quitting a bit later than q^2_l would yield a reputation payoff of $\overline{\rho}$ under D1. This argument shows that $q^2_l \geq t^*_{\max}$. But q^2_l cannot be strictly greater than $t^*_{\max}$ because, for $t > t^*_{\max}$, having the market dynamics follow the indifference curve $\rho_l(\cdot; C^*)$ would require $R(t) > \overline{\rho}$, which is infeasible. Thus we must have $q^2_l = t^*_{\max}$. If $t^*_h(\overline{\rho}) < t^*_{\max}$, we have $MRS_h(t^*_{\max}, \overline{\rho}) > 0$ and so the high type prefers quitting at $t^*_{\max}$ to quitting at a later time. We have $q_h = \overline{q}^*$. If $t^*_h(\overline{\rho}) \leq \overline{q}^*$, the opposite is true and the high type optimally quits and we have $q_h = t^*_h(\overline{\rho})$. Moreover in this case, because the informed-bad type quits before the (uninformed) high type quits, we have $R(t) = \overline{\rho}$ for $t \in [q^2_l, q_h]$. This corresponds to case (a) of the proposition.

If $D(\overline{\rho}) < t^*_{\max}$, we must also have $D(\overline{\rho}) \geq t^*_{\min}$ in equilibrium. Otherwise, $q^1_l = t^*_{\min}$ is to the right of the dividing line. The fact that the low type is indifferent between quitting and continuing at $(q^1_l, \overline{\rho})$ would imply that the high type strictly prefers to quit earlier, violating the requirement that $q_h \geq q^1_l$ in Lemma 2. Since $D(\overline{\rho})$ is between $t^*_{\min}$ and $t^*_{\max}$, the intersection point $(t^*_{\text{div}}, \rho^*_{\text{div}})$ of the indifference curve $\rho_l(\cdot; C^*)$ and the dividing line $D(\cdot)$ is well defined. Notice that q^2_l cannot exceed t^*_{div}; otherwise the fact that the low type is indifferent between quitting and continuing on $(t^*_{\text{div}}, q^2_l)$ would imply that the high type strictly prefers to quit, again violating $q_h \geq q^2_l$. If $q^2_l < t^*_{\text{div}}$ and $q_h = q^2_l$, then the high type have greater incentive than the low type to continue the project at q^2_l. Under D1, the off-equilibrium reputation from quitting a bit later would be $\overline{\rho}$, making such deviation profitable. If $q^2_l < t^*_{\text{div}}$ and $q_h > q^2_l$, then the high type would receive bad news and quit continuously on the interval (q^2_l, q_h), and the reputation payoff would again be $\overline{\rho}$ on this interval, making it profitable for the low type to quit a bit later than q^2_l. This argument shows that we must have $q^2_l = t^*_{\text{div}} < t^*_{\max}$. We must also have $q_h = t^*_{\text{div}}$, because otherwise the low type could profitably deviate by quitting a bit later to mimic the high type who receives bad news. Finally, note that $R(t^*_{\text{div}}) = \rho^*_{\text{div}} < \overline{\rho}$. This requires the low type to quit with positive probability mass at time t^*_{div}. Moreover, quitting at any $t > t^*_{\text{div}}$ would yield a reputation payoff of $\underline{\rho}$. The maximum payoff to the low type from such deviation is obtained from choosing $t^*_l(\underline{\rho})$, yielding the payoff lower bound $\underline{C}$. If $C^* > \underline{C}$, the low type has no incentive to choose $t^*_l(\underline{\rho})$ with positive probability. This corresponds to case (b) of the proposition. If $C^* = \underline{C}$, quitting at $t^*_l(\underline{\rho})$ with positive probability can be part of the low type's equilibrium strategy because it yields the equilibrium payoff C^*. This corresponds to case (c). ■

Proof of Proposition 4. For a given C, let $\mu(t;C)$ represent the market belief that an agent who quits at time t is a high type. Since the reputation dynamics is $\rho_l(\cdot;C)$, the market belief implied by the reputation dynamics satisfies

$$\mu(t;C)=\begin{cases}0 & \text{if } \rho_l(t;C)\leq \underline{\rho}\\ \hat{R}^{-1}(\rho_l(t;C)) & \text{if } \rho_l(t;C)\in(\underline{\rho},\overline{\rho})\\ 1 & \text{if } \rho_l(t;C)\geq \overline{\rho},\end{cases} \tag{15}$$

and is uniquely defined because $\hat{R}(\cdot)$ is strictly increasing. A strategy $G(\cdot;C)$ is consistent with market belief $\mu(\cdot;C)$ if $\mu(\cdot;C)$ can be derived from $G(\cdot;C)$ and from the exit strategies of the high and fully pessimistic types through Bayes' rule. For t in the open interval (q_l^1, q_l^2) for which the low type quits continuously, Bayes' rule requires

$$\frac{\mu(t;C)}{1-\mu(t;C)}=\frac{\mu_0}{1-\mu_0}\frac{(1-p_0)\theta_B e^{-\theta_B t}}{(p_0 e^{-\lambda t}+1-p_0)g(t;C)}, \tag{16}$$

In equation (16), $(1-p_0)\theta_B e^{-\theta_B t}\,\mathrm{d}t$ is the probability that a high-type agent receives bad news and quits during time $[t,t+\mathrm{d}t)$. Under strategy $G(\cdot;C)$, the probability that a low-type agent has not achieved success and quits the project during the same time interval is $(p_0 e^{-\lambda t}+1-p_0)g(t;C)\,\mathrm{d}t$. The ratio between these two is the relevant likelihood ratio for Bayesian updating.

When $D(\overline{\rho})<t_{\max}(\underline{C})$, if the equilibrium is a pooling equilibrium (cases (b) and (c) of Proposition 3), the strategy $G(\cdot;C)$ specifies that the low type quits with some positive probability mass $N_1(C)$ at time $t_{\text{div}}(C)$. The high type who still remains in the game must also quit with probability 1 at the same time. Bayes' rule requires that, in this case,

$$\frac{\mu(t_{\text{div}}(C);C)}{1-\mu(t_{\text{div}}(C);C)}=\frac{\mu_0}{1-\mu_0}\frac{p_0 e^{-(\lambda+\theta_G)t_{\text{div}}(C)}+(1-p_0)e^{-\theta_B t_{\text{div}}(C)}}{(p_0 e^{-\lambda t_{\text{div}}(C)}+1-p_0)N_1(C)}. \tag{17}$$

If $D(\overline{\rho})\geq t_{\max}(C)$, cases (b) and (c) will not obtain, and we define $N_1(C)=0$. This ensures that $N_1(C)$ is continuous in C. In addition, in case (c) of Proposition 3, the strategy $G(\cdot;C)$ also specifies that the low type quits with some positive probability mass $N_2(C)$ at time $t_l^*(\underline{\rho})$. The corresponding equilibrium reputation is $R(t_l^*(\underline{\rho}))=\underline{\rho}$. Since this case can occur only when $t_l^*(\underline{\rho})>q_h$, there is no high-ability agent who quits at $t_l^*(\underline{\rho})$. Therefore, any value of $N_2(C)$ is consistent with Bayes' rule.

The low type's strategy $G(\cdot;C)$ is an equilibrium strategy if (i) there is a value of C such that its density $g(\cdot;C)$ satisfies equation (16) for $t \in (q_l^1(C), q_l^2(C))$, (ii) $N_1(C)$ satisfies equation (17), and (iii) $G(\sup Q_l(C);C) = 1$. Toward establishing the last requirement, define

$$\Gamma(C) := \int_{q_l^1(C)}^{q_l^2(C)} g(t;C)\,\mathrm{d}t + N_1(C). \tag{18}$$

If there is a value $C^* \in [\underline{C}, \overline{C}]$ such that $\Gamma(C^*) = 1$, then the density $g(\cdot;C^*)$ and the probability mass $N_1(C^*)$ given by (16) and (17) completely describe the low type's equilibrium strategy. This corresponds to case (a) (if $N_1(C^*) = 0$) or case (b) (if $N_1(C^*) > 0$) of Proposition 3. If there is no $C \in [\underline{C}, \overline{C}]$ such that $\Gamma(C) = 1$ and $\Gamma(\underline{C}) < 1$, then $C^* = \underline{C}$. In this case, we set $N_2(C^*) = 1 - \Gamma(C^*)$. This corresponds to case (c) of Proposition 3.

Since $\Gamma(\cdot)$ is a continuous function, it suffices to establish $\Gamma(\underline{C}) > 1$ and $1 > \Gamma(C)$ for C sufficiently close to $\overline{C}$. If $D(\overline{\rho}) < t_l^*(\overline{\rho})$, there is some $C' < \overline{C}$ such that the interval $(q_l^1(\overline{C}), q_l^2(\overline{C}))$ degenerates to $D(\overline{\rho})$ as $C \to C'$. If $D(\overline{\rho}) \geq t_l^*(\overline{\rho})$, the interval $(q_l^1(\overline{C}), q_l^2(\overline{C}))$ degenerates to $t_l^*(\overline{\rho})$ as $C \to \overline{C}$. In either case, the integral term in (18) converges to 0. Moreover, since $\rho_{\mathrm{div}}(\overline{C}) = \overline{\rho}$, the implied market belief is 1 at the end of the interval, Bayes' rule requires $N_1(C') = 0$ in the first case and $N_1(\overline{C}) = 0$ in the second case by equation (17). Thus we have $\Gamma(C) < 1$ when C is sufficiently close to $\overline{C}$.

At $C = \underline{C}$, we consider two cases. First suppose that $D(\underline{\rho}) \geq t_l^*(\underline{\rho})$. In this case the interval $(q_l^1(\underline{C}), q_l^2(\underline{C}))$ must contain $t_l^*(\underline{\rho})$. Recall that $\rho_l(t;\underline{C})$ is quasi-convex and reaches a minimum of $\underline{\rho}$ at $t = t_l^*(\underline{\rho})$. This implies that $\mu(t;\underline{C})$ is quasi-convex and reaches a minimum of 0 at $t = t_l^*(\underline{\rho})$. Define $\xi := \mu(t_l^*(\underline{\rho}) - \varepsilon; \underline{C})/\varepsilon > 0$ for some small $\varepsilon > 0$ and a linear function $L(t;C) = \xi(C)(t_l^*(\underline{\rho}) - t)$. Observe that $L(t_l^*(\underline{\rho}) - \varepsilon; \underline{C}) = \mu(t_l^*(\underline{\rho}) - \varepsilon; \underline{C})$, $L(t_l^*(\underline{\rho}); \underline{C}) = \mu(t_l^*(\underline{\rho}); \underline{C})$, and $L(t;\underline{C}) = \mu(t;\underline{C})$ for all $t \in (t_l^*(\underline{\rho}) - \varepsilon, t_l^*(\underline{\rho}))$. From equation (16) for $g(t;C)$, it then follows that

$$\lim_{C \downarrow \underline{C}} \int_{q_l^1(C)}^{q_l^2(C)} g(t;C)\,\mathrm{d}t > \lim_{C \downarrow \underline{C}} \int_{t_l^*(\underline{\rho})-\epsilon}^{t_l^*(\underline{\rho})} \frac{\mu_0}{1-\mu_0} \frac{(1-p_0)\theta_B e^{-\theta_B t}}{p_0 e^{-\lambda t} + 1 - p_0} \frac{1 - L(t;C)}{L(t;C)}\,\mathrm{d}t = \infty.$$

This means $\lim_{C \downarrow \underline{C}} \Gamma(C) = \infty$, and there exists $C^* \in (\underline{C}, \overline{C})$ such that $\Gamma(C^*) = 1$.

For the second case, suppose that $D(\underline{\rho}) < t_l^*(\underline{\rho})$. In this case, $\Gamma(\underline{C})$ does not diverge to infinity. If $\Gamma(\underline{C}) > 1$, then continuity of $\Gamma(\cdot)$ implies the existence of $C^* \in (\underline{C}, \overline{C})$ such that $\Gamma(C^*) = 1$. If $\Gamma(\underline{C}) \leq 1$, we set $C^* = \underline{C}$ and set $N_2(C^*) = 1 - \Gamma(\underline{C})$ so that $\Gamma(C^*) = 1$. ∎

Proof of Proposition 5. By Proposition 3, if the low type is continuously quitting on (q_l^1, q_l^2), we must have $q_l^2 = t^*_{\text{div}}$. When $\rho^*_{div} = \rho_l(t^*_{\text{div}}; C^*) \leq \hat{\rho}$, we must have $t^*_l(\rho^*_{div}) \geq D(\rho^*_{\text{div}})$, and hence the intersection point $(t^*_{\text{div}}, \rho^*_{\text{div}})$ is on the downward-sloping part of the equilibrium indifference curve. Thus $R(\cdot)$ must be strictly decreasing on (q_l^1, q_l^2). On the other hand, when $\rho^*_{div} > \hat{\rho}$, we must have $t^*_l(\rho^*_{\text{div}}) < D(\rho^*_{\text{div}})$, and hence the intersection point $(t^*_{\text{div}}, \rho^*_{\text{div}})$ is on the upward-sloping part of the equilibrium indifference curve. If $t^*_l(\overline{\rho}) > 0$, $R(\cdot)$ must be first decreasing then increasing on (q_l^1, q_l^2). ■

Proof of Corollary 1. For any $\rho \in [\underline{\rho}, \overline{\rho}]$, we have $t^*_l(\rho) > t^*_h(\rho)$ if and only if

$$\frac{1}{\lambda} \ln \frac{p_0[\lambda(B-\overline{\rho}) - r(\overline{\rho}-W)]}{(1-p_0)r(\overline{\rho}-W)} > \frac{1}{\lambda-(\theta_B-\theta_G)} \ln \frac{p_0[\lambda(B-\overline{\rho}) + \theta_G(V-\overline{\rho}) - r(\overline{\rho}-W)]}{(1-p_0)r(\overline{\rho}-W)}. \tag{19}$$

Since $t^*_l(\cdot)$ is decreasing, the left-hand-side is positive when $t^*_l(\overline{\rho}) > 0$. The right-hand-side is quasi-concave in θ_G, and is greater than the left-hand-side at $\theta_G = \theta_B$ and less than the left-hand-side as θ_G tends to infinity. It follows that there is a critical value $\hat{\theta}_G(\theta_B; \rho) \in (\theta_B, \infty)$ such that condition (19) holds if and only if $\theta_G > \hat{\theta}_G(\theta_B; \rho)$. Let $\hat{\theta}^1_G := \hat{\theta}_G(\theta_B; \underline{\rho})$ and $\hat{\theta}^2_G := \hat{\theta}_G(\theta_B; \overline{\rho})$. The proof of Lemma 3 shows that $t^*_l(\cdot) - t_h(\cdot)$ is strictly decreasing. This implies $\hat{\theta}^2_G > \hat{\theta}^1_G$. For $\theta_G \geq \hat{\theta}^2_G$, we have $\hat{\rho} = \overline{\rho}$; and for $\theta_G < \hat{\theta}^1_G$, we have $\hat{\rho} = \underline{\rho}$. The corollary thus follows from Proposition 5. ■

Proof of Proposition 6. Observe that from (16), an increase in μ_0 reduces $g(t; C^*)$ and $N_1(C^*)$ (if $N_1(C^*) > 0$) for a fixed C^*. Therefore, $\Gamma(C^*)$ decreases from 1. To restore the equilibrium condition, C^* must increase, so that the equilibrium reputation function shifts to a higher indifference curve.

An increase in C^* raises the low type's expected payoff by definition. For the high type, consider $u_h(q_h; R)$ as defined in (2). An increase in C^* weakly raises $R(t)$ for all $t \in [0, q_h]$ (and strictly for some t), which directly raises $u_h(q_h; R)$. If $q_h = t^*_h(\overline{\rho})$, q_h is unaffected by a small change in C^*, and we are done. If $q_h = t^*_{\text{div}}$, q_h changes along the dividing line, but the indirect effect from this shift is canceled out by the envelope theorem. If $q_h = t^*_{\max}$, q_h moves with $t^*_{\max}$, but the indirect effect must be positive because q_h moves closer to $t^*_h(\overline{\rho})$. This proves that $u_h(q_h; R)$ is increasing in C^*. ■