Optimal Stopping and Reputation

Sara Shahanaghi*

February 27, 2026

Abstract

An agent decides when to take an irreversible action while learning dynamically whether it will succeed. He is reputation-driven, seeking to signal his ability to learn. In equilibrium, the agent engages in *contrarian* behavior: he is more willing to act on limited information when the prior probability of success is lower. Reputation induces systematic distortions in decision timing: the agent acts prematurely when it is highly disadvantageous ex ante, and with delay when it is advantageous. Distortions can dynamically reverse under moderate priors, with premature action at the beginning of the research process and excess caution later.

*Toulouse School of Economics, 1 Esplanade de l'Université 31000 Toulouse, France. Email contact: sara.shahanaghi@tse-fr.eu. I would like to thank Sarah Auster, Matthieu Bouvard, Alexander Guembel, Marina Halac, Johannes Hörner, Raphaël Levy, Jan Knoepfle, Erik Madsen, Thomas Mariotti, Ignacio Monzón, Konuray Mutluer, Harry Pei, Ali Shourideh, Bruno Strulovici, Maria Titova, and Juuso Välimäki for helpful feedback. Sara Shahanaghi acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program).

1. Introduction

Decision makers often decide how long to acquire information before taking an irreversible action. This is typically modeled as a Wald problem, where the stopping decision is driven by an exogenous tradeoff between accuracy and timeliness.

In practice, however, such decisions are frequently made by agents who face reputational concerns and therefore have incentives to signal their competence. For example, a politician deciding whether to implement a new policy may consider how this decision and its timing affect public perceptions of his decision-making ability, as this can influence his re-election prospects. Similarly, managers are often tasked with deciding whether to make irreversible investments in new projects and technologies, and may be concerned with how the speed and success of such decisions reflect their ability to identify profitable ventures. Indeed, the influence of CEO reputational concerns on corporate investment decisions is empirically documented (Nadeem, Zaman, Suleman, and Atawnah, 2021). In such settings, accuracy and timeliness are of *endogenous* value to the agent in signaling his competence.

I study an agent who takes an irreversible action under uncertainty. This agent seeks a reputation for learning, strategically stopping to signal a high learning ability. I pose two questions. First: what is the optimal stopping behavior of a reputation-driven agent? I consider how this behavior differs qualitatively from that of a standard Wald problem. Second, I ask how distortions in decision timing depend on the nature of the decision problem.

I present a model of optimal stopping with reputational motives. An agent (e.g., manager) dynamically learns whether some irreversible action will succeed or fail (i.e., the state), and decides when to stop learning and act before some exogenous deadline. The agent may be either good, receiving an informative signal in every period, or bad, receiving no information. He faces a combination of two objectives: a *decision payoff* which is an exogenous payoff from accurate and timely decision making as in a standard Wald problem, and a *reputation payoff* which is the belief of an outside observer (e.g., the market) that he is of high ability. This belief is formed ex post with knowledge of both the agent's stopping behavior and whether the action was a success or failure. I seek to understand how this reputation payoff affects the agent's behavior.

In equilibrium, the agent engages in *contrarian* behavior: he is more willing to act, and to act quickly, when success is ex ante less likely. This stands in contrast to a standard Wald problem, where behavior is dictated solely by the posterior belief. Contrarian behavior is driven by the endogenous impact of speed on reputation: speed is reputationally rewarded conditional on success but penalized under failure, and crucially, the relative magnitudes of this reward and penalty depend on the prior. When success is a priori unlikely, acting quickly and succeeding strongly signals learning ability, yielding a large reputational benefit, while acting quickly and failing is relatively uninformative and carries only a small reputational cost. This incentivizes early action. Meanwhile, when success is likely, acting quickly on a success does little to signal learning ability, but acting quickly on a failure carries a large reputational cost. The agent therefore waits until he is highly confident that success will occur before acting, leading to delay.

I then more generally consider the distortionary effects of reputation on the timing of action. I find that reputation can cause either premature or delayed action, and which type of distortion arises depends on the nature of the decision problem. Furthermore, distortions can reverse over the course of the research process.

When acting is highly disadvantageous from an ex ante perspective, reputation causes the agent to act prematurely at every stage of the learning process. This arises not only when the prior of success is low, reflecting the agent's contrarian motive, but also when the upside of success is small or the downside of failure is large. Intuitively, this is because in such decision problems, acting is always highly sub-optimal for the bad agent's decision payoff. However, a good agent may potentially learn over time that success will indeed occur, making acting optimal. In equilibrium, acting quickly thus serves as a signal that the agent is good at learning, one that is costly for the bad agent's decision payoff and thus credible. This induces premature action in every period. Conversely, when acting is highly advantageous a priori (due to a high prior of success, high upside of success, or low downside of failure), the opposite effect arises: reputation induces delay or excess caution. In this case, acting quickly is always the optimal course of action for the bad agent's decision payoff, while the good agent may acquire information leading him to believe that abstention is optimal. Thus, abstaining serves as a costly signal that the agent is good, leading to delay.

I then show that when the prior probability of success is moderate, distortions can dynamically reverse: the agent switches from premature action early in the research process to a reluctance to act later on. Such reversals occur when discounting is low (i.e., the agent is patient). In such decision problems, the optimal rule in the absence of reputation entails acting in early periods only if very confident that success will occur, as the option value of waiting is high. However, the reputational reward from speed induces the agent to act prematurely. As there is increasingly less time left for research, the option value of waiting declines, but late action indicates a lack of earlier informative signals. Late action can thus be reputationally damaging, making the agent hesitant to act.

In the context of irreversible investment, these results indicate that reputational motives should induce over-investment in projects that are long shots and under-investment in ones that are relatively safe or promising. The results also speak to distortions in the research process: managers conduct insufficient research before investing in speculative ventures while simultaneously exhibiting excess caution, and thus prolonging the research process, for ventures that are unlikely to fail. Meanwhile, for ventures with an intermediate probability of success, managers can be too eager to invest at the beginning of the research process but become overly hesitant towards its end. These results suggest that reputational motives can cause agents to respond inadequately to their informational environment by either under-experimenting or hesitating to act on information gathered during the research process.

This paper contributes to the literature on real options models of investment, in which a decision maker chooses when to invest under uncertainty. In canonical settings (Dixit and Pindyck (1994), McDonald and Siegel (1986)), this uncertainty concerns future flow payoffs, but not the underlying data generating process. More closely related is the subset of this literature that incorporates dynamic learning about the profitability of investment (Bernanke (1983), Cukierman (1980), Décamps, Mariotti, and Villeneuve (2005)). While these papers characterize the optimal stopping rule, I consider how a desire to signal learning ability can cause deviations from optimal behavior.

Indeed, there is an extensive literature studying an agent who uses his timing decision to influence beliefs. In Bobtcheff and Levy (2017) and Bouvard (2014), an entrepreneur decides how long to experiment before investing in a project, where investment timing signals project quality and thus affects the provision of funding. Bouvard (2014) finds that investment is delayed under the optimal contract, while Bobtcheff and Levy (2017) find that investment is hasty when learning is fast and delayed when it is slow. Meanwhile, in Grenadier and Malenko (2011), Thomas (2019), and Halac and Kremer (2020), an agent derives utility directly from market perceptions of project quality. In Thomas (2019) and Halac and Kremer (2020), the agent learns dynamically about project quality and affects public beliefs via a decision to abandon it. Both papers document inefficiently delayed abandonment. Thomas (2019) finds that timing distortions can be avoided if reputation concerns are limited. As in this paper, Halac and Kremer (2020) find that the prior belief influences distortions: the magnitude of delay is increasing in the prior that the project is profitable. I contribute to this literature by considering an agent who seeks to influence beliefs about his *ability to identify* project quality, rather than project quality itself. This different objective drives both the contrarian motive and dynamic reversals in distortions, phenomena which to my knowledge do not arise elsewhere in this literature.

Thus, this paper connects to the literature on reputation for learning. Introduced by Scharfstein and Stein (1990), Ottaviani and Sørensen (2006a) presents a general model of reputational cheap talk, finding that truthful communication is generically impossible. Regarding the nature of such distortions, Ottaviani and Sørensen (2006b) shows that reputational motives lead to less precise messages, while Gentzkow and Shapiro (2006) show low-ability senders bias towards the prior. Meanwhile, Prendergast and Stole (1996) and Dasgupta and Prat (2008) study dynamic investment and trading, respectively, under reputation for learning. My main contribution to this literature is incorporating reputation for learning into an optimal stopping problem. I find that reputation can induce reversals in distortions, which is also a key result in Prendergast and Stole (1996). The result, however, differs in several ways. First, while the distortions in Prendergast and Stole (1996) concern the responsiveness to new information, I document distortions in the agent's willingness to experiment. Moreover, I show that distortions depend critically on the decision problem: while reversals occur for intermediate priors, for extreme priors the agent persistently distorts away from the prior. In other dynamic work, Deb, Pai, and Said (2018) considers agents who learn dynamically, characterizing the optimal contract to screen for ability, whereas Shahanaghi (2025) considers a sender who learns via conclusive Poisson signals and chooses whether to make uninformed reports. Finally, Vong (2025) establishes conditions under which the sender can influence the receiver in a static setting where, as in this paper, signals are non-parametric.

Finally, the notion that the prior itself distorts behavior appears in the literature on pandering (Maskin and Tirole (2004), Canes-Wrone, Herron, and Shotts (2001), Che, Dessein, and Kartik (2013)) where agents distort towards behavior that is a priori optimal. I document a converse phenomenon: the agent distorts *away* from prior-optimal behavior. Similarly, Canes-Wrone et al. (2001) find that under certain parameters, electoral incentives can lead a politician to favor an ex ante suboptimal policy. In their setting, this occurs because the accuracy of said policy is less likely to be revealed to voters. However, I find that contrarianism is sustained by the endogenous effect of speed and accuracy on reputation.

The rest of the paper proceeds as follows. Section 2 presents the model. Sections 3 and 4 characterize equilibrium strategies and reputation, respectively, establishing the agent's contrarian motive. Section 5 characterizes equilibrium timing distortions. Section 6 presents extensions of the model. Finally, Section 7 concludes. All proofs are relegated to the appendices.

2. Model

There is one agent and one principal. Time $t \in \{1, ..., T\}$ is discrete, with a finite horizon $T < \infty$.[1] The state $\theta \in \{0, 1\}$ denotes whether some action will be a success ($\theta = 1$) or failure ($\theta = 0$). The agent and principal have a common prior $p_0 \equiv Pr(\theta = 1) \in (0, 1)$. The agent is of type $i \in \{G, B\}$ (*good* or *bad*), which is time-invariant and independent of θ. He knows his type, but the principal does not and holds a prior $R_0 \equiv Pr(i = G) \in (0, 1)$.

[1] Section 6 extends the model to an infinite-horizon setting ($T = \infty$).

Learning and Acting The type determines the agent's ability to learn about θ. At the beginning of each t, a good agent observes some signal $z_t \in (0, \infty)$, distributed according to conditional density $f(\cdot|\theta)$. The signals z_t are labeled as their likelihood ratios, i.e., $z_t = \frac{f(z_t|\theta=1)}{f(z_t|\theta=0)}$. The z_t are i.i.d. across t given θ and $f(\cdot|\theta)$ is full support on $(0, \infty)$. Meanwhile, a bad agent has no ability to learn: he observes no signals.[2]

The agent chooses if and when to act. Specifically, at each t (after observing z_t if $i = G$), the agent chooses $a_t \in \{\emptyset, 1\}$. $a_t = 1$ denotes *act* while $a_t = \emptyset$ denotes *abstain*. Acting is irreversible: if $a_t = 1$, then $a_s = 1$ for all $s > t$. Otherwise the agent chooses freely. Let $\tau \in \{1, 2, ..., T, \emptyset\}$ denote the time at which the agent acts (i.e., the first t where $a_t = 1$), with $\tau = \emptyset$ denoting that the agent never acts.

Reputation and Payoffs The principal observes the time of action τ and the state θ and forms a posterior belief about the agent's type, which is the agent's *reputation*. This belief is given by a reputation function $R : \{1, ..., T, \emptyset\} \times \{0, 1\} \to [0, 1]$, where $R(\tau, \theta)$ is the principal's belief that $i = G$ after observing τ and θ.[3]

The agent is motivated by two objectives: (1) accurate and timely decision-making and (2) reputation. His payoff is given by:

$$U(\tau, \theta) = (1 - X)D(\tau, \theta) + XR(\tau, \theta),$$

where $D(\tau, \theta)$ is the *decision payoff* and $X \in [0, 1]$. The decision payoff is given by

$$D(\tau, \theta) \equiv \beta^{\tau} K_{\theta} \mathbb{I}(\tau \neq \emptyset),$$

where $K_1 > 0 > K_0$ and $\beta \in (0, 1)$. The decision payoff captures the exogenous benefit from accurate and timely decision-making, as in a standard Wald problem. Namely, acting is beneficial if and only if $\theta = 1$ and the payoff from acting is discounted at rate β. In the context of irreversible investment, the decision payoff can be interpreted as the profit from investment. The parameter X specifies the relative weight the agent places on reputation.

[2] The assumption that the bad agent is unable to learn is made for simplicity. In Section 6, I present an extension where the bad agent observes signals with bounded likelihood ratios.

[3] One could alternatively assume that the state θ is observed by the principal only if the agent acts, and the below characterization would remain qualitatively unchanged.

2.1. Equilibrium

A Markov[4] strategy for the good agent is a sequence of functions $\{A_t^G\}_{t=1}^T$, where $A_t^G : [0,1] \to [0,1]$ and $A_t^G(p)$ denotes the probability of acting at time t (choosing $a_t = 1$) given that he has not yet acted and holds belief $p = Pr(\theta = 1)$. A strategy for the bad agent $A_t^B \in [0,1]$ denotes the probability of acting at t under belief p_0. For any signal history $(z_1, ..., z_t)$, $P(z_1, ..., z_t)$ is the good agent's posterior belief that $\theta = 1$.

I seek a Markov perfect equilibrium. This consists of strategies $\{A_t^i\}_{t=1}^T$, a reputation function R and beliefs P such that the strategies maximize the agent's payoff at all (t, p) and P, R are consistent with Bayes' Rule given the strategies.

Selection Because reputation is an equilibrium object, there exist equilibria one may deem unintuitive whenever the reputational motive (X) is sufficiently large. These can entail babbling or other perverse behavior in which the good agent signals his ability by acting on failures rather than successes.

I thus impose two selection criteria. First, the equilibrium must be informative about the agent's type in every period. Precisely, the good and bad agent cannot have the same joint distribution of action and state conditional on not yet having acted, which allows the principal to make inferences about the agent's type. This rules out babbling, and is stated as *informativeness* (Assumption 1).

Assumption 1 (Informativeness). *At any t, if $A_s^B \neq 1$ for all $s < t$, then for some θ:*

$$Pr(\tau = t | i = G, \theta, \tau \not< t) \neq Pr(\tau = t | i = B, \theta, \tau \not< t).$$

Next, the agent cannot be rewarded for minimizing the decision payoff. This rules out "reverse signaling" equilibria where the agent only acts if sufficiently sure that $\theta = 0$. This is stated as Assumption 2.

Assumption 2 (No reverse signaling). *In any t the following do not simultaneously hold for any $\tau \in \{t+1, ..., T, \emptyset\}$:*

$$R(t, \theta = 0) > R(\tau, \theta = 0) \text{ and } R(t, \theta = 1) < R(\tau, \theta = 1).$$

[4] It is without loss to restrict attention to Markov strategies within the class of equilibria that satisfy the selection criteria specified below.

One may ask how these criteria relate to more standard criteria in signaling games. While a generalization of D1[5] could be used instead of informativeness, it would not eliminate reverse signaling equilibria, as such equilibria exist even in the absence of off-path behavior. Rather, one can interpret these criteria as ruling out behavior that is not robust to a moderate concern for the decision payoff.[6]

3. Equilibrium behavior and prior beliefs

I now characterize the agent's decision-making behavior in equilibrium. I begin by establishing equilibrium existence and showing that strategies take a simple form. I then present and compare two benchmark cases to clearly illustrate the impact of reputation on decision making: one where the agent places no weight on reputation ($X = 0$) and the other where the agent is purely reputation-driven ($X = 1$). In this comparison, I begin by considering the static environment ($T = 1$) and later the general, dynamic one.

3.1. Equilibrium structure

Now I establish equilibrium existence and that the good agent plays cutoff strategies in equilibrium. To this end, I establish two useful properties of the equilibrium value function. Let $V_t^i(p, a)$ denote the type-i agent's time-t value from playing $a_t = a \in \{\emptyset, 1\}$, given he has not yet acted (i.e., $a_s = \emptyset$ for all $s < t$). Since the value of acting does not depend on the agent's type, I will often drop the i index and refer to it as $V_t(p, 1)$.

I first establish that the value of abstaining must be convex in the belief about the state. This follows from Blackwell's Theorem, and is stated as Lemma 1.

Lemma 1. *$V_t^G(p, \emptyset)$ is convex in p for all $t \in \{1, ..., T\}$.*

Next, I establish that an agent who knows that $\theta = 1$ strictly prefers acting, while an agent who knows that $\theta = 0$ strictly prefers abstaining. That is, at extreme beliefs, the agent behaves in a way that is consistent with maximizing the decision payoff. This property is stated as Lemma 2.

[5] D1 must be appropriately generalized, as the agent's payoff is directly a function of the principal's belief, rather than his action.

[6] Indeed, if X lies below some threshold, both assumptions must be satisfied in equilibrium.

Lemma 2. *In any equilibrium, in all t:*

$$V_t(1,1) > V_t^G(1,\emptyset) \text{ and } V_t^G(0,\emptyset) > V_t(0,1). \tag{1}$$

This follows from the above two selection assumptions, and can most easily be illustrated under a pure reputational motive ($X = 1$). Informativeness implies that reputation, and thus the agent's value, depends on his decision behavior. Meanwhile, the no reverse signaling assumption ensures that reputation cannot be maximized via behavior that minimizes the decision payoff. Thus, reputation is maximized via behavior which, at extreme beliefs, maximizes the decision payoff.

I now establish equilibrium existence, and show that the good agent must play cutoff strategies. This is stated as Proposition 1.

Proposition 1. *There exists an equilibrium. In any equilibrium, the good agent plays cutoff strategies in every period: for all t, there exists $p_t^* \in (0,1)$ such that*

$$A_t^G(p) = \begin{cases} 0 & \text{for all } p < p_t^* \\ 1 & \text{for all } p > p_t^*. \end{cases}$$

Existence follows by applying the Kakutani Fixed Point Theorem to a constrained space of strategy profiles for the two types of agent. Proposition 1 further establishes that at all times, there exists an interior, time-dependent cutoff belief such that the good agent acts if and only if his belief lies above this cutoff.

This result can be illustrated by a geometric argument. Figure 1 plots, for any given t, $V_t(p,1)$ and $V_t^G(p,\emptyset)$, i.e., the good agent's value from acting and abstaining, respectively. Now, let us make two observations. First, $V_t(p,1) = pR(t,1) + (1-p)R(t,0)$ is linear in the belief p, while $V_t^G(p,\emptyset)$ is convex in p (Lemma 1). Second, Lemma 2 states that $V_t(p,1)$ lies strictly above $V_t^G(p,\emptyset)$ when $p = 1$ and strictly below $V_t^G(p,\emptyset)$ when $p = 0$. These two facts imply that $V_t(p,1)$ intersects $V_t^G(p,\emptyset)$ at a unique interior point p_t^*, and thus $V_t(p,1) > (<)\ V_t^G(p,\emptyset)$ to the right (left) of this point.

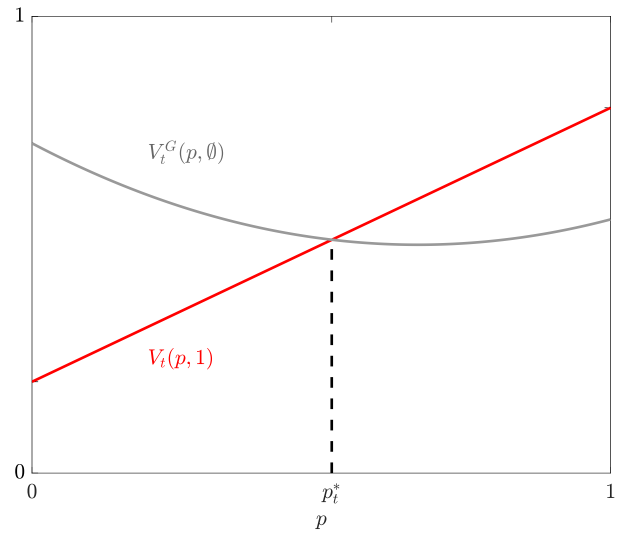

Figure 1: The good agent's value of acting ($V_t(p,1)$) and abstaining ($V_t(p,\emptyset)$).

3.2. Benchmark: Decision-optimal rule

As a benchmark, I characterize the optimal stopping rule without any reputational motive (i.e. , when $X = 0$). I refer to this as the *decision-optimal rule*. This is stated as Proposition 2.

Proposition 2. *The decision-optimal rule is the following:*

- *The good agent plays a cutoff rule $\hat{p}_t \in (0,1)$ in t, where the $(\hat{p}_t)_{t=1}^T$ are strictly decreasing in t and decreasing in K_0 and K_1 but independent of p_0.*
- *The bad agent acts in any period if and only if $p_0 > \hat{p} \equiv \frac{-K_0}{K_1 - K_0}$ for all t.*

Proposition 2 states that the good agent employs a cutoff rule in every period, where the cutoffs strictly decrease over time. The decreasing nature of these cutoffs is due to the fact that, as the deadline T approaches, the agent has less time left to learn, and hence it will be optimal to act for a wider range of beliefs as time passes. Furthermore while these cutoffs depend on the agent's learning process and payoff parameters (K_0 and K_1), they are independent of the prior belief p_0. This is because given any exogenous payoff function (such as D), the prior impacts the agent's optimal decision only via the posterior belief at any time. That is, the prior is intrinsically irrelevant to the agent's decision-making. As I will soon illustrate, this is not true for a reputation-maximizing agent.

Meanwhile, the bad agent acts immediately if acting has a positive net present value under the prior, and otherwise never acts. This is because there is discounting and the bad type is unable to learn: if his prior is such that acting is optimal, he will do so immediately and otherwise abstain indefinitely.

3.3. Static characterization under pure reputational motive

I now characterize the equilibrium when the agent is purely reputation-driven (X = 1). I begin with the static environment ($T = 1$), where the qualitative departures from the decision-optimal benchmark are most simply illustrated[7] and then consider the dynamic environment. Throughout this subsection, I omit the time index from all objects.

Claim 1 characterizes the equilibrium in the case where $X = 1$ and $T = 1$. It states that the bad agent must mix between acting and abstaining, while the good agent acts if and only if his posterior exceeds the prior.

Claim 1 (Static contrarianism). *When $X = 1$ and $T = 1$, there exists a unique equilibrium. Under this equilibrium, the bad agent mixes ($A^B \in (0, 1)$) and the good agent plays the cutoff strategy $p^* = p_0$.*

Let us begin by considering why the bad agent mixes. Proposition 1 established that the good agent plays an interior cutoff strategy $p^* \in (0, 1)$. Since the good agent's signals are full support over the likelihood ratios, it follows that the good agent both acts and abstains with positive probability. If the bad agent were to, for instance, only act in equilibrium ($A^B = 1$) then only the good type would abstain with positive probability, and thus the reputation function would assign a perfect reputation to an agent who abstains regardless of the realization of the state. A purely reputation-driven agent could then profitably deviate by always abstaining. By analogous reasoning, if the bad agent were to only abstain in equilibrium ($A^B = 0$), acting would be a profitable deviation. Thus, the bad agent must mix.

Meanwhile, the good agent's strategy follows from the fact that, in a static setting, the good and bad agent have identical continuation values: $V_t^G(p, a) = V_t^B(p, a)$ for all beliefs p and $a \in \{\emptyset, 1\}$. Since there is a unique belief at which the

[7] This static case is also a special case of Vong (2025)'s static model, which characterizes the informativeness of communication rather than the strategies of the players, as I do here.

good agent is indifferent between acting and abstaining, and the bad agent must be indifferent at p_0 to sustain mixing, the good agent's cutoff must equal the prior.

This result suggests a fundamental, qualitative difference between the decision-making behavior of a reputation maximizer compared to an agent maximizing an exogenous payoff function: while the prior does not impact the behavior of an agent maximizing decision payoffs, it is crucial to the reputation-driven agent's behavior. Namely, the reputation-driven agent is *contrarian*: he is more willing to act (formally, will act for a larger range of beliefs) when success is ex ante less likely. This holds even though the prior itself provides no payoff-relevant information beyond what is captured by the posterior p. Rather, the prior impacts the agent's equilibrium behavior via the reputation function: lowering the prior leads the reputation function to further reward acting (rather than abstaining), thus lowering the agent's cutoff. The relationship between the prior and the equilibrium reputation function will be more precisely explored in Section 4.

3.4. Dynamic characterization under pure reputational motive

I now consider a dynamic environment. In the dynamic equilibrium, the prior belief serves as a lower bound on the cutoffs the good agent employs in every period. This result is stated as Proposition 3.

Proposition 3. *Suppose $X = 1$. In any equilibrium,*

- $A_t^B \in (0, 1)$ *for all t.*
- $p_T^* = p_0$ *and* $p_t^* > p_0$ *for all $t < T$.*

The proposition first establishes that the bad agent mixes in every period between acting and abstaining, which follows from the same reasoning as in the static case. Meanwhile the good agent employs a dynamic cutoff strategy. In the final period T, this cutoff equals the prior p_0, by the same reasoning as in the static case. However, for periods prior to the deadline ($t < T$), the cutoff p_t^* strictly exceeds the prior. This is because the good agent observes informative signals about the state in future periods and thus enjoys a strictly higher continuation value than the bad agent.

This result can more precisely be illustrated by a geometric argument. Figure 2 plots the good agent's value as a function of his beliefs, as in Figure 1, for some

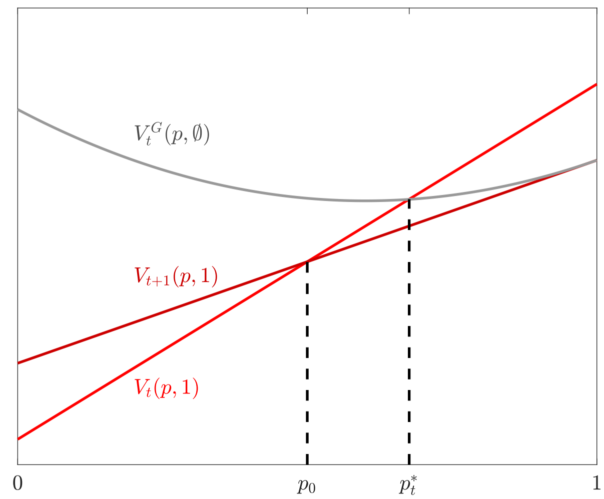

Figure 2: The good agent's cutoff, p_t^*, as compared to the prior, p_0.

$t < T$. Figure 2 also plots the value from acting in the next period, $V_{t+1}(p, 1)$. This value lies strictly below the good agent's continuation value at time t, $V_t^G(p, \emptyset)$, for all $p < 1$. This is because a good agent who continues in $t+1$ can at least obtain the value from acting in $t + 1$, and a strictly higher value by updating his belief using his new signal z_{t+1} and only acting when it is optimal to do so. Meanwhile, $V_t(p, 1)$ and $V_{t+1}(p, 1)$ must intersect at p_0: since the bad agent mixes, he must be indifferent between acting in t and $t + 1$. It follows that the good agent's point of indifference p_t^*, which is where $V_t^G(p, \emptyset)$ and $V_t(p, 1)$ intersect, lies strictly to the right of p_0.

This result implies that a lower prior will yield a strictly lower cutoff in some periods (e.g., T), but not necessarily all. However, one can show that a sufficiently large decrease in the prior yields strictly lower cutoffs in all periods. This is stated formally as Proposition 4.

Proposition 4 (Dynamic contrarianism)**.** *Assume $X = 1$. Fixing all other parameters, for any p_0 there exists $\overline{p} \in (0, 1)$ such that if $p_0' < \overline{p}$,*

$$p_t^{*\prime} < p_t^*$$

for all t and any equilibrium strategies $(p_t^*)_{t=1}^T$ and $(p_t^{*\prime})_{t=1}^T$ under p_0 and p_0', respectively.

Proposition 4 shows that contrarianism has dynamic implications: the agent not only acts more often, but more hastily when success is ex ante less likely. Put differently, a lower probability of success yields the agent less cautious, more willing to prioritize early action over further experimentation. In the following section, I illustrate the equilibrium forces that induce contrarianism.

4. Reputation: speed and accuracy

This section characterizes the equilibrium reputation function. This elucidates the equilibrium incentives at play and why the agent engages in contrarian behavior. To most clearly illustrate these incentives, I again restrict attention to a purely reputation-driven agent ($X = 1$).

I begin by characterizing the impact of speed on the agent's reputation. Formally, *speed* refers to acting in an earlier period, holding fixed the state. I show that speed improves reputation conditional on success ($\theta = 1$) but is reputationally damaging if failure occurs ($\theta = 0$). That is, early errors (acting when $\theta = 0$) are worse for reputation than late ones. This is stated as Corollary 1.

Corollary 1 (Speed)**.** *In any equilibrium when $X = 1$, for all t:*

- *$R(t, \theta = 1)$ is strictly decreasing in t,*
- *$R(t, \theta = 0)$ is strictly increasing in t.*

The reason for this conditional effect of speed lies in the fact that the bad agent mixes in equilibrium: if speed is reputationally beneficial conditional on success, it must be costly conditional on failure. Otherwise, the bad agent could profitably deviate by always acting sooner.

While Corollary 1 establishes that speed has a positive effect on reputation when $\theta = 1$ and a negative effect when $\theta = 0$, it does not speak to their magnitudes. In fact, the relative magnitudes of these effects are dictated by the prior p_0: the higher the prior, the lower the benefit of speed when $\theta = 1$ compared to its cost when $\theta = 0$. I argue that this relationship explains why the agent engages in contrarianism. The result is formalized as Corollary 2.

Corollary 2 (Speed and the prior)**.** *In equilibrium when $X = 1$, for all $t < T$:*

$$\frac{R(t,1) - R(t+1,1)}{R(t+1,0) - R(t,0)} = \frac{1-p_0}{p_0}.$$

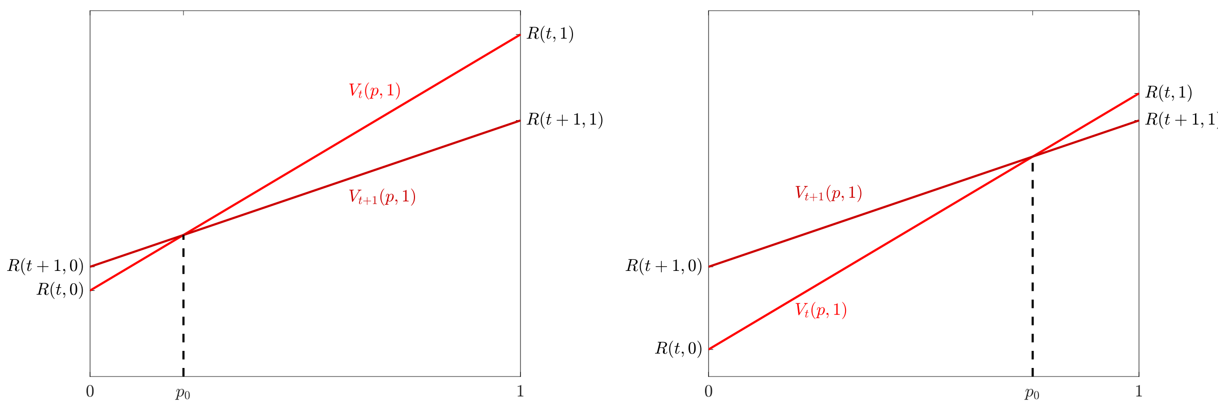

Figure 3: The value from acting in t and $t+1$ under a low prior (left panel) and high prior (right panel).

There is a concise explanation for this result: if the prior increases, $\theta = 1$ is more likely to realize ex ante. To preserve the bad agent's indifference between acting in t and $t+1$, the benefit of speed when $\theta = 1$, $R(t,1) - R(t+1,1)$, must decrease compared to the cost of speed when $\theta = 0$, $R(t+1,0) - R(t,0)$, to compensate for the fact that speed is more likely to be beneficial. It can also be illustrated graphically: Figure 3 plots the agent's value from acting in period t and $t+1$, under a low prior (left panel) and high prior (right panel). These lines have endpoints that correspond to realizations of reputation and must intersect at the prior. When the prior is low, speed is highly beneficial for reputation when acting correctly (under $\theta = 1$) but minimally harmful when acting erroneously (under $\theta = 0$). Conversely, when the prior is high, speed is minimally beneficial when acting correctly, but highly harmful when acting erroneously.

Such a relationship incentivizes contrarian behavior. When the prior is high, speed provides only a minor reputational benefit if success occurs, but can inflict substantial reputational harm in the event of failure. Thus, the agent is incentivized to gather more information to be sufficiently sure that success will occur before acting, leading to delay. Meanwhile, when the prior is low, speed provides a substantial reputational benefit if success occurs, but minimal harm if failure occurs. This incentivizes the agent to act quickly even if lacking confidence that success will occur. Together, these incentives induce contrarian behavior.

While we have thus far considered speed, one may ask what effect accuracy has on reputation. Formally, *accuracy* refers to acting on a success ($\theta = 1$) rather than a failure ($\theta = 0$), fixing the time of action. While accuracy always benefits the agent, it is more beneficial in earlier periods.

Corollary 3. *In equilibrium,* $R(t, 1) - R(t, 0) > 0$ *and is strictly decreasing in* t.

This follows immediately from Corollary 1, and can be seen from either panel of Figure 3: acting is more reputationally beneficial conditional on success and reputationally damaging conditional on failure in early periods. Thus, the effect of accuracy on reputation decreases over time.

5. Timing distortions

I now consider how reputation distorts the timing of action compared to the decision-optimal benchmark. I find that reputation can induce both premature and delayed action depending on the nature of the decision problem. The agent acts prematurely when acting is particularly disadvantageous ex ante. Conversely, the agent acts with delay when acting is favorable ex ante. At intermediate priors, reputation can induce dynamic reversals in distortions wherein the agent acts prematurely at the beginning of his learning process and with excess caution towards its end.

5.1. Premature action

I begin by characterizing the conditions under which the agent acts prematurely throughout his learning process. Precisely, *premature action* occurs at t if $p_t^* < \hat{p}_t$, i.e., the good agent acts under a strictly larger range of beliefs in equilibrium than the decision-optimal rule. Throughout this characterization, I relax the assumption that the agent is purely reputation-driven, and instead consider an arbitrary X. It is thus helpful to decompose the equilibrium value function into two components:

$$V_t^i(p, a) = (1 - X)V_t^{D,i}(p, a) + XV_t^{R,i}(p, a),$$

where $V_t^{D,i}$ denotes the agent's *decision value* and $V_t^{R,i}$ denotes his *reputation value*:

$$V_t^{D,i}(p, a) \equiv E_{\tau,\theta}[D(\tau, \theta)|(A_s^i)_{s=t+1}^T, a_t = a, i]$$
$$V_t^{R,i}(p, a) \equiv E_{\tau,\theta}[R(\tau, \theta)|(A_s^i)_{s=t+1}^T, a_t = a, i],$$

and $(A_s^i)_{s=1}^T$ denotes the agent's equilibrium strategy. Notably, both $V_t^{D,i}$ and $V_t^{R,i}$ condition on playing the equilibrium strategy in the continuation game. Because the agent chooses their equilibrium strategy to maximize a convex combination of his decision value and reputation value, this strategy does not necessarily maximize either his decision payoff or reputation payoff alone. This implies that while the equilibrium decision value and decision-optimal value of acting are equivalent, the decision value of abstaining may in general be less than that of the decision-*optimal* value of abstaining. I.e., letting $\hat{V}$ denote the value function under $X = 0$ (i.e., the decision-optimal value):

$$V_t^{D,G}(p,1) = \hat{V}_t^G(p,1) \text{ and } V_t^{D,G}(p,\emptyset) \leq \hat{V}_t^G(p,\emptyset). \tag{2}$$

Before proceeding, let the *net decision (reputation) value of acting* $W_t^{D,i}(p)$ ($W_t^{R,i}(p)$) denote the difference in value from acting and abstaining:

$$W_t^{D,i}(p) \equiv V_t^{D,i}(p,1) - V_t^{D,i}(p,\emptyset) \qquad W_t^{R,i}(p) \equiv V_t^{R,i}(p,1) - V_t^{R,i}(p,\emptyset).$$

I first show that premature action occurs whenever acting has a positive net reputation value at the decision-optimal cutoff. This is stated as Lemma 3.

Lemma 3. *Suppose $X > 0$. In any equilibrium, for any t: if $W_t^{R,G}(\hat{p}_t) > 0$, then the good agent acts prematurely in t.*

At the decision-optimal cutoff $\hat{p}_t$, the expected value from acting equals the decision-optimal value of abstaining: $\hat{V}_t^G(p,1) = \hat{V}_t^G(p,\emptyset)$. Thus by (2), the net decision value of acting must be weakly positive in equilibrium: $W_t^{D,G}(\hat{p}_t) \geq 0$. If the agent has a strictly positive net reputation value at $\hat{p}_t$, then for any positive weight on reputation, there must be a strictly positive net *total value* of acting at $\hat{p}_t$. Thus, the equilibrium point of indifference lies to the left of $\hat{p}_t$: $p_t^* < \hat{p}_t$.

I now show that the agent acts prematurely whenever acting is sufficiently sub-optimal ex ante. This could be driven by a low prior, meaning that the agent distorts the timing of their action in accordance with the contrarian motive established above. However, other factors can also drive such suboptimality: fixing a prior, the agent will also act prematurely whenever the payoff parameters align with acting being sub-optimal. This can entail a low benefit from acting on a success (K_1) or a large cost from acting on a failure (K_0). This is formalized as Theorem 1.

Theorem 1 (Premature action)**.** *Suppose $X > 0$. There exist thresholds $\overline{p} \in (0,1)$, $\overline{K}_0 < 0$, and $\overline{K}_1 > 0$ such that the agent acts prematurely in all t when any one of the following holds:*

$$p_0 < \overline{p} \qquad K_0 < \overline{K}_0 \qquad K_1 < \overline{K}_1,$$

holding fixed all other parameters.

The intuition behind this result lies in the bad agent's incentives. In equilibrium, any behavior that is reputationally rewarded —formally, any (τ, θ) pair yielding a reputation exceeding the prior reputation R_0—must be exhibited less often by the bad agent than the good agent. However, the bad agent faces two objectives: optimal decision-making and reputation maximization.[8] So, behavior that is reputationally rewarded must be costly *in expectation* for his decision payoff; otherwise he would always engage in such behavior, and it would not serve as signal that the agent is good. When acting is highly suboptimal ex ante (due to a low p_0, low K_1 or low K_0), the reputation function thus rewards acting in every period regardless of the realized state θ. This induces the good agent to act at beliefs below the decision-optimal cutoffs, leading to premature action.

This result holds not only when the agent is highly reputation-driven, but for any weight $X > 0$ on reputation. Thus, it can be explained more precisely by considering two different cases: one where reputational motives are small (low X) and the other where they are large (high X). When the reputational motive is small and acting is sufficiently suboptimal ex ante for the decision payoff, the bad agent will never act in equilibrium. This is because the negative decision value from acting at t, $W_t^{D,B}(p_0)$, outweighs any possible gain in reputation value, $W_t^{R,B}(p_0)$. The good agent, however, may observe signals strongly indicating $\theta = 1$, making acting optimal with positive probability. Since only the good agent acts, the reputation function must assign a perfect reputation to acting regardless of θ. By Lemma 3, this implies premature action.

When the reputational motive is large, the bad agent does not exclusively abstain in equilibrium but instead mixes (as was true for the purely reputation-driven agent in Section 3). Thus, the bad agent is indifferent between acting

[8] While this intuition applies to the case where $X \in (0, 1)$, this result holds even when $X = 1$.

and abstaining at any t even though acting has a negative net decision value ($W_t^{D,B}(p_0) < 0$). This implies the bad agent must enjoy a positive net reputation value from acting: $W_t^{R,B}(p_0) > 0$. Moreover, if acting is reputationally beneficial at the prior, it must be even more beneficial at higher beliefs, since accuracy is rewarded in equilibrium. So, the good agent must also enjoy a positive net reputation value from acting at the decision-optimal cutoffs, i.e., $W_t^{R,G}(\hat{p}_t) > 0$, since these cutoffs lie well above the prior when acting is sufficiently sub-optimal ex ante. It again follows that the good agent acts prematurely.

5.2. Delayed action

The previous section established that reputation induces premature action when acting is highly sub-optimal ex ante. One might expect the opposite effect when acting is highly *optimal* ex ante, namely that the agent delays action. In this subsection, I establish that this is indeed true. However, because the agent's stopping problem is inherently one-sided (acting is irreversible, abstaining is not), the result is not symmetric to that of premature action.

Formally, *delayed action* occurs at t when $p_t^* > \hat{p}_t$, i.e., the good agent acts for a strictly smaller range of beliefs than the decision-optimal rule. I now establish that when acting is sufficiently optimal ex ante, the agent acts with delay for at least some time. When the agent's reputational motive is sufficiently large, then they act with delay in all periods. This is stated as Theorem 2.

Theorem 2 (Delayed action)**.** *If $p_0 > \hat{p}_t$ for all t, then action is delayed in at least some t for any $X > 0$, and action is delayed in all t if X is sufficiently large.*

This proposition establishes that the agent acts with delay for at least some periods when the prior exceeds the decision-optimal cutoffs in all t. While this may seem to establish only that action is delayed for a sufficiently high prior, it also implies that holding fixed a prior, action is delayed whenever K_1 is sufficiently large, or K_0 is insignificant. This is because, per Proposition 2, the decision-optimal cutoffs are decreasing in both payoff parameters K_1 and K_0.

The basic intuition is analogous to the case of premature action. When acting is highly optimal ex ante, delaying action (or never acting) is costly for the bad agent's decision payoff. So in equilibrium, abstaining is a costly signal that the agent is good and a reputation-driven agent responds to this by delaying action.

To more precisely convey its reasoning, it is again helpful to separately consider the cases of a small and large reputational motive. Regardless of X, when acting is highly optimal ex ante, abstaining in any period is damaging to the bad agent's decision value. When X is large, the bad agent nonetheless mixes in every period, so to ensure indifference, he must benefit reputationally from abstaining. This in turn implies an expected reputational benefit from abstaining at the good agent's decision-optimal cutoffs $\hat{p}_t$, leading to delayed action. When X is small, the damage to the bad agent's decision value from abstaining can outweigh any reputational benefit from doing so, leading him to act with certainty by some $s \leq T$. In all periods $t \leq s$, abstaining remains reputationally beneficial, leading to delayed action. However, an agent who continues past s will have revealed that he is good, and thus enjoys a perfect reputation regardless of his behavior thereafter. So from this time onwards, he will not distort his behavior and instead employs the decision-optimal cutoffs $\hat{p}_t$. It is for this reason that Theorem 2 includes the caveat that action may not be delayed in all periods when X is small.

Discussion Let us interpret the previous two results in the context of irreversible investment. Theorem 1 suggests that for projects with a low ex ante value — because success is unlikely, the downside potential is high, or the upside potential is modest— reputational motives lead to premature investment. This implies over-investment in long-shot projects since investment, and especially early investment, in such projects serves as a costly signal of learning ability. Moreover, because the agent invests prematurely in every period, he also engages in an inefficiently low level of experimentation. Meanwhile, Theorem 2 suggests that for projects with a high ex ante value, the opposite phenomenon arises: investment is delayed. This implies that the agent will both underinvest and experiment for too long in promising projects, only waiting until he is extremely sure that the project will succeed before investment.

Together, these results imply that reputation distorts the research process by either shading up or shading down the agent's option value. Furthermore, the agent distorts in a way that is opposed to the ex ante profitability of a project. That is, reputation incentivizes behavior that is diametrically opposed to the optimal behavior of an uninformed decision maker.

5.3. Dynamics in distortions

Theorem 1 and Theorem 2 showed that when action is highly (dis)advantageous ex ante, reputation induces distortions that are time-invariant in direction: the agent either acts prematurely in every period or with delay in every period. Here, I show that for decision problems where the value of acting takes an intermediate value, the direction of distortions can dynamically reverse. Specifically, the agent may act prematurely in early periods, and with delay in later ones. This phenomenon arises whenever the agent is sufficiently reputationally motivated and discounting is slow (i.e., β is high). The result is formalized as Theorem 3.

Theorem 3. *Suppose $p_0 \in (\hat{p}_T, \hat{p}_1)$. Fixing all other parameters, there exist $\overline{X} \in (0,1)$ and $\overline{\beta} \in (0,1)$ such that if $\beta > \overline{\beta}$ and $X > \overline{X}$ the good agent:*

- *acts prematurely in early periods (i.e., there exists $\underline{t}$ such that $p_t^* < \hat{p}_t$ for all $t \leq \underline{t}$),*
- *acts with delay in later periods (i.e., there exists $\overline{t}$ such that $p_t^* > \hat{p}_t$ for all $t \geq \overline{t}$).*

This can be understood by considering how the decision value and reputation value interact over time. In early periods, the net decision value of acting is low: discounting is minimal and there is a high option value of waiting from learning in future periods. Accordingly, the decision-optimal rule entails acting early only under very high beliefs that $\theta = 1$. However, as the deadline T approaches, this option value falls because there is little time left to learn, and so the decision-optimal cutoff falls. Meanwhile, the reputation payoff evolves differently. There is always a reputational benefit to speed, and thus cost to delay, when $\theta = 1$. Importantly, this benefit from speed is sizeable (i.e., is bounded below) in early periods even when discounting is slow. Because the decision-optimal cutoffs are high in these early periods, speed is reputationally beneficial in expectation at these cutoffs, inducing premature action. As the deadline approaches and the decision-optimal cutoff falls, speed is no longer expected to be rewarded at these cutoffs. Instead, delay becomes increasingly reputationally beneficial, as abstaining forever is costly for the bad agent's decision payoff (since $p_0 > \hat{p}_T$) and serves as a credible signal of learning ability. This induces a shift to delayed action in later periods.

Discussion This result illustrates that reputational motives can give rise to distortions that reverse over the course of the research process. More precisely, Theorem 3 indicates that when discounting is sufficiently slow, the agent is initially too eager to act on limited information. I.e., he under-experiments. However, as he approaches the end of the research process, he becomes overly conservative, choosing to either continue experimenting excessively or to avoid acting altogether even though doing so is not decision-optimal. This reversal arises because the reputational benefit of speed incentivizes early action, but the reputational cost from late action incentivizes conservatism late in the research process.

6. Extensions

I now present two extensions of the baseline model, allowing for: (1) learning by the bad type and (2) an infinite horizon.

6.1. Learning by the bad type

Formally, let us now suppose that B observes some signal z_t^B in every period t, distributed according to conditional density $f^B(\cdot|\theta)$. As with G, I suppose that B's signals are i.i.d. given θ and are labeled as their likelihood ratios: $z_t^B \equiv \frac{f^B(z_t^B|\theta=1)}{f^B(z_t^B|\theta=0)}$. I now impose two assumptions to formalize the notion that, while B has access to an informative signal, G has a greater ability to learn. First, B's signals have bounded support on the likelihood ratios: there exists $\underline{z} < \overline{z} \in (0, \infty)$ such that the support of f^B is $(\underline{z}, \overline{z})$. This stands in contrast to G's signal, which is full support on the likelihood ratios. Second, G's signal is Blackwell more informative than B's signal. Otherwise, I maintain all assumptions of the baseline model.

I first show that when B has access to such a signal, he does not mix but rather plays a cutoff strategy in equilibrium. Proposition 5 characterizes this strategy.

Proposition 5. *In equilibrium, the bad and good agents play cutoff strategies* $(p_t^{*,B})_{t=1}^T$ *and* $(p_t^{*,G})_{t=1}^T$*, respectively, where* $p_t^{*,B} < p_t^{*,G}$ *for all* $t < T$*, and* $p_T^{*,B} = p_T^{*,G}$.

Proposition 5 establishes that B acts for a strictly larger range of beliefs than G in all periods prior to the deadline. This is due to the fact that G has access to a more informative signal than B in every period, granting him a strictly higher continuation value in $t < T$, and thus a greater willingness to abstain.

Next, I show that Theorem 1 is robust to allowing the bad agent the ability to learn.[9] That is, the good agent will act prematurely when acting is disadvantageous ex ante. This result is formalized as Proposition 6. For analytic simplicity, I assume the agent is not purely reputationally motivated ($X < 1$).

Proposition 6. *If $X \in (0, 1)$ and the bad agent observes signals drawn from f^B, then Theorem 1 continues to hold.*

Theorem 1 relied on the notion that the reputation function rewards behavior which is always sub-optimal from the bad agent's perspective, but sometimes optimal given the good agent's information. When the bad agent is unable to learn, his perspective is always the ex ante one. This notion extends to a setting where the bad agent observes signals of bounded informativeness, as specified here, because this ensures that the bad agent's perspective never drastically departs from the ex ante one, whereas the good agent's perspective does so with positive probability.

6.2. Infinite horizon

The baseline model assumed an exogenous deadline $T < \infty$. I now extend the baseline model to an infinite horizon. I show that even under an infinite horizon, the agent behaves as if there is a deadline in equilibrium. For simplicity, I restrict attention to a purely reputation-driven agent ($X = 1$).

Formally, let us assume $T = \infty$ (i.e., $t \in \mathbb{N}^+$). I otherwise maintain all assumptions of the baseline model, except for informativeness (Assumption 1). I relax this assumption because, as I now show, informativeness cannot hold in every period under an infinite horizon. This is stated as Proposition 7.

Proposition 7. *When $T = \infty$, Assumption 1 must fail in equilibrium.*

This follows from the fact that the bad agent's indifference between acting and abstaining, which must hold in equilibrium, cannot be satisfied indefinitely. Under an infinite horizon, the reputational benefit of speed must vanish as $t \to \infty$ (due to the boundedness of the reputation function), leading the good agent's equilibrium cutoffs to approach 1. I.e., in the limit the good agent only acts if he is arbitrarily certain that $\theta = 1$. This in turn implies acting when $\theta = 0$ reveals that the agent

[9] While I do not state and prove the result formally, Theorem 2 can analogously be generalized.

is the bad type, making acting sub-optimal for the bad agent in the limit and thus leading to a failure of indifference.

Since informativeness cannot hold indefinitely, I now define a relaxed version of the informativeness assumption that accommodates the agent stopping acting in finite time. I call this relaxed assumption *T-informativeness*. Precisely, it imposes that informativeness applies for all $t \leq T$, but that the agent ceases acting after T.

Definition 1. For any $T \in \mathbb{N}^+ \cup \infty$, an equilibrium satisfies *T-informativeness* if for any $t \leq T$, for some θ, $Pr(\tau = t | i = G, \theta, \tau \nless t) \neq Pr(\tau = t | i = B, \theta, \tau \nless t)$, and for all $t > T$ and θ, $Pr(\tau = t | i = G, \theta) = 0$.

One can show that there exists a T-informative equilibrium for any $T < \infty$. This is formalized as Corollary 4.

Corollary 4. *For any $T < \infty$, there exists a T-informative equilibrium.*

Together, Proposition 7 and Corollary 4 imply that the deadline of the baseline model need not be exogenously imposed: under an infinite horizon, a deadline must arise endogenously.

7. Conclusion

This paper studies how a motive to signal learning ability can distort the timing of action. This reputational motive gives rise to contrarian behavior, driven by the endogenous role speed plays in reputation formation. In equilibrium, the agent's action may be either premature or delayed, with premature action occurring when the ex ante value of action is low, and delayed action occurring when it is high. When the probability of success takes an intermediate value, the agent's distortionary behavior can change over time, with premature action at the beginning of the research process but excess caution towards its end. More generally, this paper illustrates how reputational motives can distort the value of experimentation driven by the endogenous, signaling value of speed.

In this paper, I have studied the behavior of a single agent in isolation, i.e., in the absence of competition. Not only would competition give rise to new strategic motives, it would allow agents to condition their investment behavior on that of their opponents. How competition between reputation-driven agents could give rise to unique dynamics in behavior is a question that warrants investigation.

References

Ben S Bernanke. Irreversibility, Uncertainty, and Cyclical Investment. *The Quarterly Journal of Economics*, 98(1):85–106, 1983.

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

Matthieu Bouvard. Real Option Financing Under Asymmetric Information. *The Review of Financial Studies*, 27(1):180–210, 2014.

Brandice Canes-Wrone, Michael C Herron, and Kenneth W Shotts. Leadership and Pandering: A Theory of Executive Policymaking. *American Journal of Political Science*, pages 532–550, 2001.

Yeon-Koo Che, Wouter Dessein, and Navin Kartik. Pandering to Persuade. *American Economic Review*, 103(1):47–79, 2013.

Alex Cukierman. The Effects of Uncertainty on Investment under Risk Neutrality with Endogenous Information. *Journal of Political Economy*, 88(3):462–475, 1980.

Amil Dasgupta and Andrea Prat. Information Aggregation in Financial Markets with Career Concerns. *Journal of Economic Theory*, 143(1):83–113, 2008.

Rahul Deb, Mallesh M Pai, and Maher Said. Evaluating Strategic Forecasters. *American Economic Review*, 108(10):3057–3103, 2018.

Jean-Paul Décamps, Thomas Mariotti, and Stéphane Villeneuve. Investment Timing Under Incomplete Information. *Mathematics of Operations Research*, 30 (2):472–500, 2005.

Avinash K Dixit and Robert S Pindyck. *Investment under Uncertainty*. Princeton university press, 1994.

Aaron S Edlin and Chris Shannon. Strict monotonicity in comparative statics. *Journal of Economic Theory*, 81(1):201–219, 1998.

Matthew Gentzkow and Jesse M Shapiro. Media Bias and Reputation. *Journal of Political Economy*, 114(2):280–316, 2006.

Steven R Grenadier and Andrey Malenko. Real Options Signaling Games with Applications to Corporate Finance. *The Review of Financial Studies*, 24(12):3993–4036, 2011.

Marina Halac and Ilan Kremer. Experimenting with Career Concerns. *American Economic Journal: Microeconomics*, 12(1):260–288, 2020.

Eric Maskin and Jean Tirole. The Politician and the Judge: Accountability in Government. *American Economic Review*, 94(4):1034–1054, 2004.

Robert McDonald and Daniel Siegel. The Value of Waiting to Invest. *The Quarterly Journal of Economics*, 101(4):707–727, 1986.

Paul Milgrom and Ilya Segal. Envelope Theorems for Arbitrary Choice Sets. *Econometrica*, 70(2):583–601, 2002.

Muhammad Nadeem, Rashid Zaman, Tahir Suleman, and Nader Atawnah. CEO Ability, Career Concerns, Firms' Lifecycle and Investments in Intellectual Capital. *International Review of Economics & Finance*, 75:237–251, 2021.

Marco Ottaviani and Peter Norman Sørensen. Reputational Cheap Talk. *The RAND Journal of Economics*, 37(1):155–175, 2006a.

Marco Ottaviani and Peter Norman Sørensen. The Strategy of Professional Forecasting. *Journal of Financial Economics*, 81(2):441–466, 2006b.

Canice Prendergast and Lars Stole. Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning. *Journal of Political Economy*, 104(6):1105–1134, 1996.

David S Scharfstein and Jeremy C Stein. Herd Behavior and Investment. *The American economic review*, pages 465–479, 1990.

Sara Shahanaghi. Reputation and Misreporting in News Media. 2025.

Caroline Thomas. Experimentation with reputation concerns–Dynamic signalling with changing types. *Journal of Economic Theory*, 179:366–415, 2019.

Allen Vong. Reputational Cheap Talk: Influentialness and Prior Information, 2025.

Appendix

Before proceeding, let us define two different conditional distributions. To this end, let F denote the unconditional distribution of likelihood ratios:

$$F(s) \equiv \int_0^s p_0 f(s'|\theta = 1) + (1 - p_0) f(s'|\theta = 0) ds'.$$

Now, let $G_t(\cdot|p_{t-1})$ denote the good agent's distribution of time t beliefs given that his time $t-1$ belief was p_{t-1}. It follows from the definition of F that:

$$G_t(p_t|p_{t-1}) = F\left(\left(\frac{1 - p_{t-1}}{p_{t-1}}\right)\left(\frac{p_t}{1 - p_t}\right)\right).$$

Second, let $H_t(\cdot|p_{t-1})$ denote the distribution of time-t beliefs given $\tau \notin \{1, ..., t\}$ and that the time-$t-1$ belief was p_{t-1}. Finally, let $H_t(\cdot)$ denote the good agent's distribution of time-t beliefs given $\tau \notin \{1, ..., t\}$, conditional on p_0 (namely, not conditional on the time-$t-1$ belief). It is computed recursively as follows:

$$H_1(p_1) = H_1(p_1|p_0)$$

$$H_t(p_t) = \int_0^1 H_t(p_t|p_{t-1}) dH_{t-1}(p_{t-1}).$$

Further, note that we refer to $\tau = \emptyset$ and $\tau = T + 1$ interchangeably. We will also frequently drop the G superscript when referring to the good agent's value function. Finally, we will frequently refer to an equilibrium object, which is the agent's *interim reputation*. It is defined as follows.

Definition 2 (Interim reputation)**.** The agent's time t interim reputation is the principal's belief $i = G$ given that they did not report at or before t:

$$R_t \equiv Pr(i = G|\tau \notin \{1, ..., t\}).$$

Proof of Proposition 1: necessity of cutoff strategies. Here, we show that in any equilibrium, for any t, G plays an interior cutoff strategy (we later prove existence of an equilibrium). To this end, fix any equilibrium. Fix a t. By Lemma 2, $V_t(1,\emptyset) < V_t(1,1)$ and $V_t(0,\emptyset) < V_t(0,1)$. Because $V_t(p,\emptyset)$ is convex in p (Lemma 1) and $V_t(1,p) = pR(t,1) + (1-p)R(t,0)$ is linear in p, there exists a unique $p_t^* \in (0,1)$ such that

$$V_t(p,1) > V_t(p,\emptyset) \text{ for all } p > p_t^*$$

$$V_t(p,1) < V_t(p,\emptyset) \text{ for all } p < p_t^*.$$

Thus, the good agent's strategy must be an interior cutoff strategy. □

We now seek to establish existence of an equilibrium. To this end, let us define the correspondence Φ as follows. First, let us define R^x. Let R^x denote the reputation function that is consistent with the strategy profile $x = (p_1^*, ..., p_T^*, A_1^B, ..., A_T^B) \in [0,1]^{2T}$. Formally, whenever Bayes' Rule applies, $R^x(t,\theta)$ is given by

$$R^x(t,\theta) = \frac{1}{1 + (\frac{1-R_{t-1}}{R_{t-1}})(\frac{Pr(\tau=t,\theta|\tau\geq t, i=B)}{Pr(\tau=t,\theta|\tau\geq t,i=G)})}, \tag{3}$$

where the probabilities, including R_{t-1}, are those that obtain given the strategy profile x. The only case in which Bayes Rule does not apply is when $p_t^* = 1$ and $A_t^B = 0$ for some t, and in this case we impose $R^x(t,\theta) = 1$ for all θ.

Now, let $V_{s-1}^{G,x}(p, (\hat{p}_t)_{t=s}^T)$ denote G's value, under belief p at time $s-1$, from playing cutoff strategies $(\hat{p}_t)_{t=s}^T$ in periods $s, ..., T$, respectively, given reputation function R^x and that the agent did not act in $s-1$. Now, define the $\Phi_s^G(x)$ recursively as follows:

$$\Phi_s^G(x) \equiv \min_{\overline{p}_s \in [0,1]} \arg\max_{\overline{p}_s \in [0,1]} [V_{s-1}^{G,x}(p_0, (\overline{p}_t)_{t=s}^T)],$$

where $\overline{p}_t \equiv \Phi_t^G(x)$ for all $t > s$. Let $\Phi^G(x) \equiv (\Phi_t^G(x))_{t=1}^T$. Note that the value could have been taken at any interior belief (not necessarily p_0) and the analysis that follows would remain unchanged.

Next, let $V_s^{B,x}((b_t)_{t=s}^T)$ denote B's value from playing strategy $A_t^B = b_t$ for all

$t \geq s$, given reputation function R^x and that the agent did not act before s. Now, define the $\Phi_s^B(x)$ recursively as follows:

$$\Phi_s^B(x) \equiv \arg\max_{b_s \in [0,1]} V_s^{B,x}((b_t)_{t=s}^T),$$

where $b_t \in \Phi_t^B(x)$ for all $t > s$. Define $\Phi^B(x) \equiv \Phi_1^B(x) \times ... \times \Phi_T^B(x)$, and finally $\Phi(x) \equiv \Phi^G(x) \times \Phi^B(x)$.

I wish to show that any fixed point of Φ is an equilibrium. To this end, I begin by establishing two lemmas, the proofs of which are relegated to the Online Appendix.

Lemma 4. *In any fixed point of Φ, $p_t^* < 1$ for all t.*

Lemma 5. *Any fixed point x of Φ, together with R^x, is an equilibrium that satisfies Assumption 1 and Assumption 2.*

Proof of Proposition 1: equilibrium existence. I now establish existence of a fixed point to Φ. It follows from Lemma 5 that this is an equilibrium. To this end, for each $\varepsilon > 0$, I define a constrained correspondence Φ^ε and show that for some ε, there exists a fixed point of Φ^ε which is also a fixed point of Φ. I proceed in a number of steps, as outlined below.

1. **Define constrained correspondence:** For any $\varepsilon \in (0,1)$, we define the following correspondence:

$$\Phi^\varepsilon : [0, 1-\varepsilon]^T \times [0,1]^T \to [0, 1-\varepsilon]^T \times (2^{[0,1]})^T,$$

where $\Phi^\varepsilon(x) \equiv \Phi^{G,\varepsilon}(x) \times \Phi^B(x)$,

$$\Phi_s^{G,\varepsilon}(x) \equiv \min_{\overline{p}_s \in [0,1-\varepsilon]} \arg\max_{\overline{p}_s \in [0,1-\varepsilon]} [V_{s-1}^{G,x}(p_0, (\overline{p}_t)_{t=s}^T)], \tag{4}$$

and $\Phi^B(x)$ is defined as before.

2. **Existence of fixed point for Φ^ε:** I now claim that for any $\varepsilon < 1$, Φ^ε has a fixed point. To prove this, I invoke the Kakutani Fixed Point Theorem. To this end, I show that Φ^ε satisfies the following properties:

(a) $\Phi^\varepsilon(x)$ **is non-empty for all** x**.** This follows from the fact that $[0, 1-\varepsilon]$ and $[0,1]$ are compact and $R^x(\tau, \theta)$ is bounded, implying by the Extreme Value Theorem that both $\Phi_t^B(x)$ and $\Phi_t^{G,\varepsilon}(x)$ are non-empty for all t, x.

(b) $\Phi^\varepsilon(x)$ **is convex and closed for all** x**.** $\Phi_t^{G,\varepsilon}(x)$ is a singleton by definition for all x, t. Now, fix an (x,t) and consider $\Phi_t^B(x)$. Define $\underline{b}_t = 0$, $\bar{b}_t = 1$, $\bar{b}_s = \underline{b}_s = A_s^B$ for all $s > t$. It follows that

$$\Phi_t^B(x) = \begin{cases} 1 & \text{if } V_t^{B,x}((\underline{b}_s)_{s=t}^T) < V_t^{B,x}((\bar{b}_s)_{s=t}^T) \\ 0 & \text{if } V_t^{B,x}((\underline{b}_s)_{s=t}^T) > V_t^{B,x}((\bar{b}_s)_{s=t}^T) \\ [0,1] & \text{if } V_t^{B,x}((\underline{b}_s)_{s=t}^T) = V_t^{B,x}((\bar{b}_s)_{s=t}^T), \end{cases} \tag{5}$$

so $\Phi_t^B(x)$ is convex and closed, and thus $\Phi^\varepsilon(x)$ is convex and closed.

(c) Φ^ε **is UHC.** I will show that for all t, Φ_t^B and $\Phi_t^{G,\varepsilon}$ are UHC everywhere on the domain. It follows that their Cartesian product Φ^ε is also UHC. Fix an $x \in X$ and a t. Let us begin with Φ_t^B. Now note that because $\varepsilon > 0$, $R^x(t, \theta)$ is continuous in x, and thus both $V_t^{B,x}((\underline{b}_s)_{s=t}^T)$ and $V_t^{B,x}((\bar{b}_s)_{s=t}^T)$ are continuous in x. Thus, it follows from (5) that $\Phi_t^B(x)$ is UHC at x. Next, consider $\Phi_t^{G,\varepsilon}$. It again follows from the continuity of $R^x(t,\theta)$ that $V_t^{G,x}$ is continuous in x, and thus by (4), $\Phi_s^{G,\varepsilon}(x)$ is continuous in x.

Thus, by the Kakutani Fixed Point Theorem Φ^ε has a fixed point.

3. **Show that for some** $\varepsilon > 0$**,** Φ^ε **has an interior fixed point:** I now claim that for some $\varepsilon > 0$, Φ^ε has a fixed point that lies within $[0, 1-\varepsilon)^T \times [0,1]^T$ (i.e., a fixed point such that $p_t^* < 1-\varepsilon$ for all t). Suppose not, by contradiction. Then, there exists t^* and a sequence $\{\varepsilon_n\}_{n=1}^\infty$ such that $\varepsilon_n > 0$ for all n, $\lim_{n\to\infty} \varepsilon_n = 0$ and there exists a sequence $\{x_n\}_{n=1}^\infty$ where x_n is a fixed point of Φ^{ε_n} such that $p_{t^*}^* = 1 - \varepsilon_n$. Hereafter, let t^* refer to the last such period.

   I first claim that there exists an N such that if $n > N$, $A_{t^*}^B \in (0,1)$ under x_n. Suppose not, by contradiction. First consider the case where for some infinite subsequence $\{x_m\}_{m=1}^\infty$ of $\{x_n\}_{n=1}^\infty$, $A_{t^*}^B = 0$ for all m. It follows that $R^{x_m}(t^*, \theta) = 1$ for all θ. Thus, for m sufficiently large, $p_{t^*}^* = 1 - \varepsilon_m$ cannot be optimal. Contradiction. Next, consider the case where for some infinite subsequence $\{x_m\}_{m=1}^\infty$, $A_{t^*}^B = 1$ for all m. Since for all m, $p_{t^*}^* = 1-\varepsilon_m$ is optimal and $p_0 < p^*_{t^*}$ for m sufficiently large, it follows that for some m $A^B_{t^*} = 1$ is not optimal. Contradiction.

Hereafter, let us thus assume that the sequence $\{x_n\}_{n=1}^{\infty}$ is such that $A^B_{t^*} \in (0,1)$ for all n. First, suppose $X < 1$. That $p^*_{t^*} = 1 - \varepsilon_n$ for all n implies that for all $D_1 > -K_1(\beta^{t^*} - \beta^{t^*+1})\frac{1-X}{X}$, there exists an N such that for $n > N$,

$$R^{x_n}(t^*, 1) - R^{x_n}(t^* + 1, 1) < D_1.$$

It follows from Bayes Rule (given that G plays cutoff strategies) that

$$R^{x_n}(t^*, 0) - R^{x_n}(t^* + 1, 0) < D_1.$$

Thus, $A^B_{t^*} = 0$ is a profitable deviation. Contradiction.

Next, suppose $X = 1$. I claim that

$$\lim_{n\to\infty} R^{x_n}(s, 1) - R^{x_n}(s + 1, 1) = 0 \text{ and } \lim_{n\to\infty} R^{x_n}(s, 0) = 0 \tag{6}$$

for all $s \geq t^*$. Proof by induction. Begin with $s = t^*$. Note that by the contradiction assumption, for all n, $V^G_{t^*}(1 - \varepsilon_n, 1) \leq V^G_{t^*}(1 - \varepsilon_n, \emptyset)$ (where this is the value function that obtains from R^{x_n}) because otherwise $p^*_t < 1 - \varepsilon_n$ under x_n. I claim this implies $\lim_{n\to\infty} R^{x_n}(t^*, 1) - R^{x_n}(t^* + 1, 1) = 0$. Suppose not, by contradiction. Then there exists $\delta > 0$ and an infinite subsequence $\{\varepsilon_{n_k}\}_{k=1}^{\infty}$ of $\{\varepsilon_n\}_{n=1}^{\infty}$ such that $R^{x_{n_k}}(t^*, 1) - R^{x_{n_k}}(t^* + 1, 1) > \delta$ for all k. Thus, there exists k such that $V^G_{t^*}(1 - \varepsilon_{n_k}, 1) - V^G_{t^*}(1 - \varepsilon_{n_k}, \emptyset) > 0$. Contradiction.

Next, I show $\lim_{n\to\infty} R^{x_n}(t^*, 0) = 0$. Recall that by Bayes Rule, under any x_n:

$$R^{x_n}(t^*, 0) = \frac{1}{1 + (\frac{1-R_0}{R_0})(\frac{1}{Q_{t^*}(n)})} \text{ where } Q_t(n) \equiv \frac{Pr^n(\theta = 0|\tau = t, i = G)Pr^n(\tau = t|i = G)}{Pr^n(\theta = 0|\tau = t, i = B)Pr^n(\tau = t|i = B)},$$

I claim that $\lim_{n\to\infty} Q_{t^*}(n) = 0$. Suppose not, by contradiction. Since $\lim_{n\to\infty} \varepsilon_n = 0$, $\lim_{n\to\infty} \frac{Pr^n(\theta=0|\tau=t^*,i=G)}{Pr^n(\theta=0|\tau=t,^*i=B)} = 0$, and thus it suffices to show that $\frac{Pr^n(\tau=t^*|i=G)}{Pr^n(\tau=t^*|i=B)}$ does not diverge as $n \to \infty$. This is only possible if there exists a subsequence $\{\varepsilon_{n_k}\}_{k=1}^{\infty}$ of $\{\varepsilon_n\}_{n=1}^{\infty}$ such that $\lim_{k\to\infty} \frac{Pr^{n_k}(\tau=t^*|i=G)}{Prn_k(\tau=t^*|i=B)} = \infty$. This implies $\lim_{k\to\infty} R^{x_{n_k}}(t^*, \theta) = 1$ for all θ, and thus for k sufficiently large, $A^B_{t^*} = 1$ is a profitable deviation from what is specified under x_{n_k}.

Now, fix some $t > t^*$ and assume by induction that (6) holds for all s such that $t^* \leq s < t$. We want to show that it also holds for t. First, let us show that $\lim_{n\to\infty} R^{x_n}(t,0) = 0$. Since for all n, when $X = 1$, $A_t^B \in (0,1)$:

$$p_0 R^{x_n}(t,1) + (1-p_0)R^{x_n}(t,0) = p_0 R^{x_n}(t-1,1) + (1-p_0)R^{x_n}(t-1,0).$$

Thus,

$$\lim_{n\to\infty} R^{x_n}(t,0) = \frac{p_0}{1-p_0}[\lim_{n\to\infty}[R^{x_n}(t-1,1) - R^{x_n}(t,1)] - \lim_{n\to\infty} R^{x_n}(t-1,0)] = 0,$$

where the last equality follows from the inductive assumption.

Next, let us show that $\lim_{n\to\infty} R^{x_n}(t,1) - R^{x_n}(t+1,1) = 0$. Suppose not, by contradiction. Then, there exists $\delta > 0$ and subsequence $\{\varepsilon_{n_k}\}_{k=1}^{\infty}$ of $\{\varepsilon_n\}_{n=1}^{\infty}$ such that $R^{x_{n_k}}(t,1) - R^{x_{n_k}}(t+1,1) \geq \delta$ for all k. This implies that there exists $p \in (p_0, 1)$ such that $p_t^* \leq p$ under x_{n_k} for all k. However, $\lim_{n\to\infty} R^{x_n}(t,0) = 0$, and thus for all $p \in (p_0,1)$, there exits an $N \in \mathbb{N}$ such that $p_t^* > p$ under x_n for all $n > N$. Contradiction.

Now, note that for all n, $p_T^* = p_0$ under x_n. Thus, $Q_T(n)$ does not converge to 0 as $n \to \infty$. However, because $\lim_{n\to\infty} R^x(T,0) = 0$, $\lim_{n\to\infty} Q_T(n) = 0$. Contradiction.

4. **This interior fixed point of Φ^ε is also a fixed point of Φ:** Fix an $\varepsilon > 0$ such that there is a fixed point x of Φ^ε such that $x \in [p_0, 1-\varepsilon)^T \times [0,1]^T$. I claim that x is also a fixed point of Φ. This is equivalent to showing that for all t:

$$A_t^B \in \Phi_t^B(x) \text{ and } p_t^* = \Phi_t^G(x).$$

Note that $A_t^B \in \Phi_t^B(x)$ holds because this is necessary for x to be a fixed point of Φ^ε. Next, let us show that $p_t^* = \Phi_t^G(x)$ for all t. Proof by induction. Fix a t, and suppose $p_s^* = \Phi_s^G(x)$ for all $s > t$. We want to show $p_t^* = \Phi_t^G(x)$.

Since $p_t^* < 1 - \varepsilon$,

$$V_t^G(p,1) > V_t^G(p,\emptyset) \text{ for all } p > 1-\varepsilon,$$

where this is the value function that obtains given the reputation function

R^x. Thus,

$$V_t^{G,x}(p_0, (p_s^*)_{s=t}^T) > V_t^{G,x}(p_0, (\tilde{p}_s)_{s=t}^T)$$

for any $\tilde{p}_t > 1 - \varepsilon$ and $\tilde{p}_s = p_t^*$ for all $s > t$. This, combined with the fact that $p_t^* = \Phi_t^{G,\varepsilon}(x)$, implies $p_t^* = \Phi_t^G(x)$.

□

Proof of Proposition 2. Part 2 follows from the fact that the B agent holds belief p_0 in every period. For part 1, let us first show that G plays a cutoff rule and these cutoffs are unique. Proof by induction.

In period T, the agent acts if and only if his belief lies above $\hat{p} \equiv \frac{-K_0}{K_1 - K_0}$. Next, fix a t and suppose that the agent plays an interior cutoff rule in all $s > t$. Note that

$$V_t(p, 1) = \beta^t(pK_1 + (1-p)K_0).$$

Now, let us observe three facts about $V_t^G(p, \emptyset)$:

1. Since a cutoff rule is played in $t+1$,

$$V_t^G(1, \emptyset) = V_{t+1}^G(1, 1) = \beta^{t+1}K_1 < \beta^t K_1 = V_t^G(1, 1)$$

2. $V_t^G(0, \emptyset) = 0 > \beta^t K_0 = V_t^G(0, 1)$
3. $V_t(p, \emptyset)$ is convex in p.

These three facts together with the linearity of $V_t(p, 1)$ imply that there is a unique $\hat{p}_t \in (0, 1)$ such that the good agent plays cutoff rule $\hat{p}_t$ in t.

Now, it remains to show that the $\hat{p}_t$ are strictly decreasing in t. To this end, fix a $t < T$. Suppose by contradiction that $\hat{p}_t \leq \hat{p}_{t+1}$. Then

$$V_t^G(\hat{p}_t, 1) = V_t^G(\hat{p}_t, \emptyset)$$

$$V_{t+1}^G(\hat{p}_t, 1) \leq V_{t+1}^G(\hat{p}_t, \emptyset).$$

Since all these values are strictly positive,

$$\beta = \frac{V_{t+1}^G(\hat{p}_t, 1)}{V_t^G(\hat{p}_t, 1)} \leq \frac{V_{t+1}^G(\hat{p}_t, \emptyset)}{V_t^G(\hat{p}_t, \emptyset)}. \tag{7}$$

Now, let $\tilde{V}_t(p, \emptyset)$ denote the agent's value from the modified problem which is identical to the original problem except that the time horizon is $T - 1$. It follows that for all $t < T$:

1. $\tilde{V}_t^G(p, \emptyset) = \frac{V_{t+1}^G(p,\emptyset)}{\beta}$
2. $\tilde{V}_t^G(p, \emptyset) < V_t^G(p, \emptyset)$.

These two facts together imply

$$\frac{V_{t+1}^G(\hat{p}_t, \emptyset)}{V_t^G(\hat{p}_t, \emptyset)} < \beta,$$

contradicting (7).

It remains to show that $\hat{p}_t$ is strictly decreasing in K_1 (that it is also strictly decreasing in K_0 follows analogously). Fix any t. Note first that for any p_t

$$\frac{\delta \hat{V}_t^G(p_t, 1)}{\delta K_1} = \beta^t p_t.$$

Meanwhile, by the Envelope Theorem (Milgrom and Segal, 2002):

$$\frac{\delta \hat{V}_t^G(p_t, \emptyset)}{\delta K_1} = E_{\tau > t, p_\tau}[\beta^\tau p_\tau \mathbb{I}(\tau \neq \emptyset)],$$

where p_τ denotes the belief at time τ, and the expectation is take with respect to this belief and the agent's stopping time. By the Martingale property of the belief,

$$E_{\tau > t, p_\tau} = p_t.$$

where the expectation is taken over all $\tau > t$, including $\tau = \emptyset$. Since $\mathbb{I}(\tau \neq \emptyset) < 1$ and $\beta^\tau < \beta^t$,

$$\frac{\delta \hat{V}_t^G(p_t, \emptyset)}{\delta K_1} < \frac{\delta \hat{V}_t^G(p_t, 1)}{\delta K_1},$$

thus, $\hat{p}_t$ is strictly decreasing in K_1. □

I now state a lemma that will be used in the proof of Proposition 3. The proof of this lemma is relegated to the Online Appendix.

Lemma 6. *In any equilibrium, if for all $s \leq t$ there exists a $p_s^* \in (0,1)$ such that*

$$A_t^G(p) = \begin{cases} 0 & \text{for all } p < p_t^* \\ 1 & \text{for all } p > p_t^* \end{cases}$$

and $A_t^B \in (0,1)$, then $R_t \in (0,1)$.

Proof of Proposition 3. First, we want to show that in any equilibrium $A_t^B \in (0,1)$ for all t. Proof by induction. Assume by induction that $A_s^B \in (0,1)$ for all $s < t$ (this holds vacuously when $t = 1$). Assume by contradiction $A_t^B \in \{0,1\}$. Now, consider the case where $A_t^B = 0$ (the case where $A_t^B = 1$ follows analogously). It follows from Bayes Rule that

$$R(t,0) = \frac{1}{1 + (\frac{1-R_{t-1}}{R_{t-1}})(\frac{Pr(\tau=t,\theta=0|\tau\geq t,i=B)}{Pr(\tau=t,\theta=0|\tau\geq t,i=G)})} \tag{8}$$

First, note that

$$Pr(\tau = t, \theta = 0|\tau \geq t, i = B) = A_t^B = 0.$$

Meanwhile,

$$Pr(\tau = t, \theta = 0|\tau \geq t, i = G) = \int_0^1 \int_{p_t^*}^1 (1 - p_t) dG_t(p_t|p_{t-1}) dH_{t-1}(p_{t-1}) > 0,$$

where the strict inequality follows from the fact that $p_t^* \in (0,1)$. By Lemma 6, it follows from (8) that $R(t,0) = 1$. One can analogously show that $R(t,1) = 1$. Thus,

$$V_t(p_0, 1) = p_0 R(t,1) + (1 - p_0) R(t,0) = 1. \tag{9}$$

Now, by the Law of Iterated Expectations

$$\begin{aligned} R_{t-1} &= Pr(i = G, \tau = t|\tau \geq t)(1) \\ &\quad + Pr(i = G, \tau \neq t|\tau \geq t) \int_0^1 V_t^G(\emptyset, p) dH_t(p) \\ &\quad + Pr(i = B|\tau \geq t) V_t^B(p_0, \emptyset). \end{aligned} \tag{10}$$

Because R is consistent with the A^i in equilibrium, $V_t^B(p_0, \emptyset) \leq \int_0^1 V_t^G(p, \emptyset) dH_t(p)$.

Because $R_{t-1} < 1$ (Lemma 6), it follows from (10) that $V_t^B(p_0, \emptyset) < 1$. Combining this with (9) implies $V_t^B(p_0, \emptyset) < V_t(p_0, 1)$. Thus, $A_t^B(p_0) = 1$. Contradiction.

Now, we want to show that in any equilibrium, $p_T^* = p_0$. Fix any equilibrium. We have already shown $A_t^B \in (0, 1)$. Thus,

$$V_T(p_0, 1) = V_T^B(p_0, \emptyset) = V_T^G(p_0, \emptyset), \tag{11}$$

Note further that (1) both $V_T(p, 1)$ and $V_T(p, \emptyset)$ are linear in p and (2) by Lemma 2, $V_T(0, 1) < V_T^G(0, \emptyset)$ and $V_T(1, 1) > V_T^G(1, \emptyset)$. These two facts, combined with (11), imply that $V_T(p, 1) < V_T(p, \emptyset)$ for all $p < p_0$ and $V_T(p, 1) > V_T(p, \emptyset)$ for all $p > p_0$. Thus $p_T^* = p_0$.

It remains to show that for all $t < T$, $p_t^* > p_0$. To this end, fix a $t < T$. It follows from the above that B mixes between $a \in \{1, \emptyset\}$ in every t, and thus $V_t(p_0, 1) = V_{t+1}(p_0, 1)$. Now, it follows from Lemma 2 that

$$V_t(1, \emptyset) = V_{t+1}(1, 1) \qquad \text{and} \qquad V_t(0, \emptyset) > V_{t+1}(0, 1).$$

Thus, it follows that $V_t(p, \emptyset) > V_{t+1}(p, 1)$ for all $p < 1$. Because $p_0 \in (0, 1)$, then

$$V_t(p_0, \emptyset) > V_{t+1}(p_0, 1). \tag{12}$$

Since $V_t^G(p, \emptyset) > V_t(p, 1)$ if and only if $p < p_t^*$, it follows from (12) that $p_t^* > p_0$.

□

Proof of Claim 1. That $A^B \in (0, 1)$ and $p^* = p_0$ follows directly from Proposition 3. It remains to show that there is a unique equilibrium. Namely, we want to show that there exists a unique $b \in (0, 1)$ such that $(A^B = b, p^* = p_0)$ is an equilibrium strategy profile. First, define

$$W(a, b) \equiv p_0 R^b(a, 1) + (1 - p_0) R^b(a, 0)$$

where

$$R^b(a, \theta) \equiv \frac{1}{1 + \frac{1-R_0}{R_0} \frac{Pr(a,\theta|i=B,A^B=b)}{Pr(a,\theta|i=G,p^*=p_0)}}$$

is the unique reputation function that is consistent with the strategy profile $(A^B =$

$b, p^* = p_0)$. I claim that there exists a unique $b \in (0,1)$ such that $W(1,b) = W(\emptyset, b)$. First, note that

$$Pr(a, \theta | i = G, p^* = p_0) \in (0,1) \text{ for all } a, \theta. \tag{13}$$

Now, I make two observations about W:

1. $W(1, b = 0) - W(\emptyset, b = 0) > 0$ and $W(1, b = 1) - W(\emptyset, b = 1) < 0$.
   To show this, note that $Pr(1, \theta | i = B, A^B = 0) = 0$ for all θ. Thus, by (13), $R^{b=0}(1, \theta) = 1$ and $R^{b=0}(\emptyset, \theta) < 1$ for all θ. Thus, $W(1, b = 0) - W(\emptyset, b = 0) > 0$. One can analogously show that $W(1, b = 1) - W(\emptyset, b = 1) < 0$.

2. $W(1, b) - W(\emptyset, b)$ is continuous and strictly decreasing in b.
   To show this, note that

$$R^b(1,1) = \frac{1}{1 + \frac{1-R_0}{R_0} \frac{p_0 b}{1 - F(1|\theta=1)}},$$

   which is continuous and strictly decreasing in b. One can similarly show that $R^b(1, \theta)$ ($R^b(\emptyset, \theta)$) is continuous and strictly decreasing (increasing) in b for all θ. The statement then follows from the definition of W.

1. and 2. imply that there exists a unique b such that $W(1, b) = W(\emptyset, b)$. It follows immediately that $(A^B = b, p^* = p_0)$ is the unique equilibrium strategy profile. □

I now state a lemma that will be used to show that Theorem 1 holds in the special case where $X = 1$. The proof of this lemma is relegated to the Online Appendix.

Lemma 7. *Assume $X = 1$, and fix all parameters except p_0. For any $p \in (0,1)$, there exists $\overline{p} \in (0,1)$ such that if $p_0 < \overline{p}$, $p_t^* < p$ in any equilibrium under p_0.*

Proof of Proposition 4. Fix any R_0, T and f. Consider some $p_0 \in (0,1)$. By Lemma 7, there exists a $\overline{p} \in (0, p_0)$ such that if $p_0' < \overline{p}$, $p_t^* < \overline{p}$ in any equilibrium strategy $(p_t^{*\prime})_{t=1}^T$ under p_0', for all t. Furthermore, it follows from Proposition 3 that $p_r^* \geq p_0$ for all t in any equilibrium strategy $(p_t^*)_{t=1}^T$ under p_0. Thus, $p_t^{*\prime} < p_t^*$ for all t, under any such equilibria.

□

Proof of Corollary 1. Fix any $t < T$. First note

$$V_t(1, \emptyset) = V_{t+1}(1, 1) = R(t + 1, 1),$$

where the first inequality follows from Lemma 2. Next, note $V_t(1, 1) = R(t, 1)$. Finally, by Lemma 2, $V_t(1, 1) > V_t(1, \emptyset)$. Thus, $R(t, 1) > R(t + 1, 1)$. The bad agent's indifference condition then implies $R(t, 0) < R(t + 1, 0)$.

□

Proof of Lemma 3. Because G plays a cutoff strategy, it suffices to show that $V_t^G(\hat{p}_t, 1) > V_t^G(\hat{p}_t, \emptyset)$. Note that

$$V_t^{D,G}(\hat{p}_t, 1) = \hat{V}_t^G(\hat{p}_t, 1) = \hat{V}_t^G(\hat{p}_t, \emptyset) \geq V_t^{D,G}(\hat{p}_t, \emptyset),$$

where the second equality follows from the definition of $\hat{p}_t$. Thus,

$$V_t^G(\hat{p}_t, 1) = (1-X)V_t^{D,G}(\hat{p}_t, 1)+XV_t^R(\hat{p}_t, 1) > (1-X)V_t^{D,G}(\hat{p}_t, \emptyset)+XV_t^R(\hat{p}_t, \emptyset) = V_t^G(\hat{p}_t, \emptyset),$$

where the strict inequality follows from the assumptions that $V_t^R(\hat{p}_t, 1) > V_t^R(\hat{p}_t, \emptyset)$ and $X > 0$. □

Proof of Theorem 1. First, let us consider K_0 and K_1. Begin by considering the case where $X < 1$. Fixing all other parameters, we want to show that there exists $\overline{K}_1 > 0$ such that if $K_1 < \overline{K}$, then $p_t^* < \hat{p}_t$ for all t under any equilibrium (the existence of $\overline{K}_0$ follows analogously). Suppose not, by contradiction. Then, there exists a sequence $\{K_1^n\}_{n=1}^{\infty}$ such that $K_1^n > 0$ for all n and $\lim_{n\to\infty} K_1^n = 0$ such that there exists an equilibrium strategy $\{p_t^{*,n}\}_{t=1}^{T}$ associated with each K_1^n where $p_s^{*,n} \leq \hat{p}_s^n$ for all n and some $s < T$. Note $\hat{p}_s^n$ is the decision-optimal cutoff under K_1^n. In what follows, let superscript n denote objects for said equilibrium under K_1^n.

Since $\lim_{n\to\infty} \hat{p}_s^n = 1$, it must be that $\lim_{n\to\infty} p_s^{*,n} = 1$. In order for such $p_s^{*,n}$ to be optimal, it must be that:

$$\lim_{n\to\infty} [R^n(s, \theta = 1) - R^n(s + 1, \theta = 1)] = 0 \tag{14}$$

Meanwhile, since we cannot have $R^n(s, \theta = 0) = R^n(s, \theta = 1) = 1$, $A_s^{B,n} \in (0, 1)$ for all n. Now note there exists an N and $\overline{W} < 0$ such that for all $n > N$, $W_s^{D,B,n}(p_0) < \overline{W}$. Thus, to maintain B's indifference, for all $n \geq N$, $W_s^{R,B,n}(p_0) > -\overline{W} > 0$, where:

$$\begin{aligned} W_s^{R,B,n}(p_0) &= [R^n(s, 1) - R^n(s + 1, 1)]p_0 + [R^n(s, 0) - R^n(s + 1, 0)](1 - p_0) \\ &< R_n(s, 1) - R^n(s + 1, 1). \end{aligned}$$

Thus, for all $n \geq N$, $R^n(s, 1) - R^n(s + 1, 1) > -\overline{W} > 0$, which contradicts (14).

Next, suppose $X = 1$. I claim that for all t, there exists $\overline{p}_t \in (0, 1)$ such that $p_t^* < \overline{p}_t$. This proves the statement for K_0 and K_1 under $X = 1$. Suppose not, by contradiction. Let s be the last period where the statement fails. Since $p_T^* = p_0$ in any equilibrium, $s < T$. Thus, there exists a sequence of equilibria with time-s strategies $p_s^{*,n}$ such that the $p_s^{*,n}$ are strictly increasing in n and $\lim_{n\to\infty} p_s^{*,n} = 1$. This implies

$$\lim_{n\to\infty} [R_n(s, 1) - R_n(s + 1, 1)] = 0,$$

which implies

$$\lim_{n\to\infty} [R_n(s, 0) - R_n(s + 1, 0)] = 0.$$

Meanwhile, $\lim_{n\to\infty} R_n(s, 0) = 0$, because otherwise $\lim_{n\to\infty} R_n(s, 1) = 1$, which would imply that for n sufficiently large, $p_s^{*,n}$ is not optimal. Thus, $\lim_{n\to\infty} R_n(s + 1, 0) = 0$. However, for all n,

$$R_n(s + 1, 0) = \frac{1}{1 + \frac{1-R_0}{R_0} \frac{Pr^n(\tau=s+1|i=B)}{Pr^n(\tau=s+1|i=G,\theta=0)}}.$$

Since there exists $\overline{p}_{s+1} \in (0, 1)$ such that $p_{s+1}^{*,n} < \overline{p}_{s+1}$ for all n, there exists a $Q > 0$ such that $Pr^n(\tau = s + 1|i = G, \theta = 0) > Q$. Thus, we cannot have $\lim_{n\to\infty} R_n(s + 1, 0) = 0$. Contradiction.

Next, consider p_0. Fix all parameters except p_0 and first assume $X \in (0, 1)$. I show that there exists $\overline{p} \in (0, 1)$ such that if $p_0 < \overline{p}$, $p_t^* < \hat{p}_t$ for all $t \in \{1, ..., T\}$ in any equilibrium. Fix any $t \in \{1, .., T\}$. I will show there exists $\overline{p}_t \in (0, 1)$ such that if $p_0 < \overline{p}_t$, then $p_t^* < \hat{p}_t$ under any equilibrium strategy of G. Letting $\overline{p} \equiv \min_t \overline{p}_t$, this implies the statement we wish to prove.

Suppose by contradiction that for some t, there does not exist such a $\overline{p}_t$. Then, there exists a sequence $\{p_{0,n}\}_{n=1}^{\infty}$ such that $p_{0,n} \in (0, \hat{p})$ for all n and $\lim_{n\to\infty} p_{0,n} = 0$, where for all n and for some equilibrium strategy profile $(A^B_{t,n}, p^*_{t,n})_{t=1}^T$ under prior $p_{0,n}$, $p^*_{t,n} \geq \hat{p}_t$. Now, I proceed in a number of steps:

1. **Show that for all n, $A^B_{s,n} > 0$ for all $s \geq t$.** Begin with time t. Suppose by contradiction that $A^B_{t,n} = 0$ for some n. Because $p^*_{t,n} \in (0, 1)$, $R_n(t, \theta) = 1$ for all θ. Thus, $V^{R,G}_{t,n}(p, 1) = 1$ for all p, where $V^{R,G}_{t,n}$ is G's reputation value under the equilibrium. Now, let $R_{t-1,n}$ denote the interim reputation under the equilibrium. Because $A^B_{s,n} < 1$ for all $s < t$, it follows that $R_{t-1,n} < 1$. Thus, it must be that $V^{R,G}_{t,n}(\hat{p}_t, \emptyset) < 1$. Thus by Lemma 3, $p^*_{t,n} < \hat{p}_t$. Contradiction.

   Next, consider some $s > t$. Suppose by contradiction that $A^B_{s,n} = 0$ for some n. By Lemma 2, $p^*_{s,n} < 1$, and thus by Bayes Rule $R_n(s, \theta) = 1$. But then:

$$\begin{aligned} V^B_{s,n}(p_{0,n}, 1) &= (1 - X)(p_{0,n}K_1 + (1 - p_{0,n})K_0)\beta^s + X \\ &> (1 - X)(p_{0,n}K_1 + (1 - p_{0,n})K_0)\beta^t + XV^R_{t,n}(p_{0,n}, 1) = V^B_{t,n}(p_{0,n}, 1), \end{aligned}$$

   where the strict inequality follows from the fact that $p_{0,n} < \hat{p}$ by assumption, and thus $p_{0,n}K_1 + (1 - p_{0,n})K_0 < 0$. Contradiction.

2. **Show $\lim_{n\to\infty}[R_n(s, 0) - R_n(s + 1, 0)] > 0$ for all s such that $t \leq s < T$, and $\lim_{n\to\infty}[R_n(T, 0) - R_n(\emptyset, 0)] > 0$.**
   Fix some s such that $t \leq s < T$, By 1., $V^B_{s,n}(p_{0,n}, 1) = V^B_{s+1,n}(p_{0,n}, 1)$ for all n, where

$$V^B_{s,n}(p_{0,n}, 1) = (1-X)[p_{0,n}K_1+(1-p_{0,n})K_0]\beta^s+X[p_{0,n}R_n(s, 1)+(1-p_{0,n})R_n(s, 0)].$$

   Thus,

$$\lim_{n\to\infty} V^B_{s,n}(p_{0,n}, 1)-V^B_{s+1,n}(p_{0,n}, 1) = (1-X)(\beta^s-\beta^{s+1})K_0+X \lim_{n\to\infty}[R_n(s, 0)-R_n(s+1, 0)] = 0.$$

   So,

$$\lim_{n\to\infty}[R_n(s, 0) - R_n(s + 1, 0)] = -\frac{1 - X}{X}(\beta^s - \beta^{s+1})K_0 > 0.$$

   Next, it follows from 1. that $V^B_{T,n}(p_{0,n}, 1) = V^B_{T,n}(p_{0,n}, \emptyset)$ for all n. Thus by the same reasoning as above,

$$\lim_{n\to\infty}[R_n(T,0) - R_n(\emptyset,0)] = -\beta^T K_0(\frac{1-X}{X}) > 0.$$

3. **Show** $\lim_{\mathbf{n}\to\infty} \mathbf{R_n}(\mathbf{t},\mathbf{1}) = \mathbf{1}$. For the equilibrium $(p^*_{t,n}, A^B_{t,n})^T_{t=1}$ under $p_{0,n}$, let $H_{t,n}$ denote the good agent's distribution of time-t beliefs given $\tau \notin \{1,...,t\}$, and likewise for $G_{t,n}$. Define

$$Q_{n,t} \equiv \frac{\int_0^1 \int_{p^*_{t,n}}^1 (1-p_t) dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}{\int_0^1 \int_{p^*_{t,n}}^1 p_t dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}.$$

I begin by showing

$$\lim_{n\to\infty}(\frac{p_{0,n}}{1-p_{0,n}})Q_{n,t} = 0. \tag{15}$$

Since $\lim_{n\to\infty} p_{0,n} = 0$, it suffices to show that there exists $L \in \mathbb{R}^+$ such that $Q_{n,t} < L$ for all n. Note that for all n,

$$\int_0^1 \int_{p^*_{t,n}}^1 p_t dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1}) > \int_0^1 \int_{p^*_{t,n}}^1 p^*_{t,n} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1}).$$

Thus,

$$Q_{n,t} < \frac{\int_0^1 \int_{p^*_{t,n}}^1 dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}{\int_0^1 \int_{p^*_{t,n}}^1 p^*_{t,n} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})} - 1 = \frac{1}{p^*_{t,n}} - 1.$$

By the assumption that $p^*_{t,n} \geq \hat{p}_t$ for all n, $Q_{n,t} < \frac{1}{\hat{p}_t} - 1$, thus establishing (15). Now, recall that

$$R_n(t,0) = \frac{1}{1 + (\frac{1-R_{t-1,n}}{R_{t-1,n}})(\frac{A^B_{t,n}}{\int_0^1 \int_{p^*_{t,n}}^1 \frac{1-p_t}{1-p_{0,n}} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})})} \tag{16}$$

$$R_n(t,1) = \frac{1}{1 + (\frac{1-R_{t-1,n}}{R_{t-1,n}})(\frac{A^B_{t,n}}{\int_0^1 \int_{p^*_{t,n}}^1 \frac{p_t}{p_{0,n}} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})})}.$$

It follows from 2. that

$$\frac{A^B_{t,n}}{\int_0^1 \int_{p^*_{t,n}}^1 \frac{1-p_t}{1-p_{0,n}} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})} \text{is bounded.} \tag{17}$$

Now note

$$\lim_{n\to\infty} \frac{A^B_{t,n}}{\int_0^1 \int_{p^*_{t,n}}^1 \frac{p_t}{p_{0,n}} dG_t(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}$$
$$= \lim_{n\to\infty} \Big(\frac{A^B_{t,n}}{\int_0^1 \int_{p^*_{t,n}}^1 \frac{1-p_t}{1-p_{0,n}} dG_t(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}\Big)\Big(\frac{\int_0^1 \int_{p^*_{t,n}}^1 \frac{1-p_t}{1-p_{0,n}} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}{\int_0^1 \int_{p^*_{t,n}}^1 \frac{p_t}{p_{0,n}} dG_{t,n}(p_t|p_{t-1}) dH_{t-1,n}(p_{t-1})}\Big) = 0$$

where the final equality follows from (15) and (17). Thus, by (16), $\lim_{n\to\infty} R_n(t,1) = 1$.

4. **For some n, $\mathbf{V^{R,G}_{t,n}(\hat{p}_t, 1) > V^{R,G}_{t,n}(\hat{p}_t, \emptyset)}$.** For all n,

$$V^{R,G}_{t,n}(\hat{p}_t, 1) = \hat{p}_t R_n(t, \theta = 1) + (1 - \hat{p}_t) R_n(t, \theta = 0).$$

It thus follows from 2. and 3. above, and the fact that $\hat{p}_t < 1$, that there exists $K > 0$ and $N \in \mathbb{N}$ such that if $n > N$,

$$V^{R,G}_{t,n}(\hat{p}_t, 1) > \hat{p}_t + (1 - \hat{p}_t) R_n(t + 1, 0) + K. \tag{18}$$

Now, note that for any n,

$$V^{R,G}_{t,n}(\hat{p}_t, \emptyset) \le \hat{p}_t \max_{\tau\in\{t+1,\ldots,T,\emptyset\}} R_n(\tau, 1) + (1 - \hat{p}_t) \max_{\tau\in\{t+1,\ldots,T,\emptyset\}} R_n(\tau, 0).$$

Thus, by 2., there exists $N' \in \mathbb{N}$ such that for all $n \ge N'$

$$V^{R,G}_{t,n}(\hat{p}_t, \emptyset) < \hat{p}_t + (1 - \hat{p}_t) R_n(t + 1, 0).$$

Thus, by (18), for all $n \ge \max\{N, N'\}$,

$$V^{R,G}_{t,n}(1, \hat{p}_t) > V^{R,G}_{t,n}(\emptyset, \hat{p}_t).$$

By Lemma 3 it follows from 4. that for all $n \geq \max\{N, N'\}$, $p^*_{t,n} < \hat{p}_t$, contradicting the assumption that $p^*_{t,n} \geq \hat{p}_t$ for all n.

Next, suppose $X = 1$. It follows directly from Lemma 7 that there exists a $\overline{p} \in (0,1)$ such that if $p_0 < \overline{p}$, $p^*_t < \hat{p}_T \leq \hat{p}_t$ for all t under any equilibrium strategy $(p^*_t)_{t=1}^T$ under p_0.

□

Proof of Theorem 2. Fix any equilibrium $(p^*_t, A^B_t)_{t=1}^T$ and any $t \in \{1, .., T\}$. First, consider the case where $A^B_s = 1$ for some $s < t$. Since G plays an interior cutoff strategy at all t, the equilibrium reputation function R must be such that

$$R(\tau, \theta) = 1 \text{ for all } \tau \in \{t, .., T, \emptyset\}, \theta \in \{0, 1\}.$$

Thus, $V^{R,G}_t(p, a) = 1$ for all $a \in \{\emptyset, 1\}$, $p \in [0, 1]$. Hence, G's problem at time t is to choose a strategy which maximizes the following:

$$E_\theta[U(\tau, \theta)] = (1 - X)\beta^\tau(K_\theta \mathbb{I}(\tau \neq \emptyset)) + X.$$

This problem is equivalent to maximizing $\beta^\tau(K_\theta \mathbb{I}(\tau \neq \emptyset))$. Hence, the equilibrium strategy must be equal to the optimal cutoff rule under $X = 0$, i.e., $p^*_t = \hat{p}_t$.

Next, consider the case where $A^B_s < 1$ for all $s < t$ (this holds vacuously when $t = 1$). I claim that in this case $p^*_t > \hat{p}_t$. First, suppose that $A^B_t = 1$. The equilibrium reputation function must be such that (1) $R(s, \theta) = 1$ for all $s \in \{t + 1, ..., T, \emptyset\}$ and $\theta \in \{0, 1\}$ and (2) $R(t, \theta) < 1$ for $\theta \in \{0, 1\}$. Together, these two facts imply that $V^{R,G}_t(p, 1) < 1$ and $V^{R,G}_t(p, \emptyset) = 1$ for all p. Furthermore, by the same reasoning as above, $p^*_s = \hat{p}_s$ for all $s > t$ and thus $V^{D,G}_t(p, a) = \hat{V}_t(p, a)$ for all a, p. So, $V^G_t(\hat{p}_t, 1) < V^G_t(\hat{p}_t, \emptyset)$ and thus $p^*_t > \hat{p}_t$. Next, suppose that $A^B_t < 1$. It must also be that $A^B_t > 0$. To show this, suppose not by contradiction. Then, the reputation function must be such that $R(t, \theta) = 1$ for $\theta \in \{0, 1\}$. Thus, $V^{R,B}_t(p, 1) \geq V^{R,B}_t(p, \emptyset)$ for all p. Since $V^{D,B}_t(p_0, 1) > V^{D,B}_t(p_0, \emptyset)$, it follows that $V^B_t(p_0, 1) > V^B_t(p_0, \emptyset)$, and thus $A^B_t = 1$. Contradiction. So, $A^B_t \in (0, 1)$ which implies B must be indifferent at p_0: $V^B_t(p_0, 1) = V^B_t(p_0, \emptyset)$. Since $V^G_t(p_0, \emptyset) \geq V^B_t(p_0, \emptyset)$ and $V^B_t(p_0, 1) = V^G_t(p_0, 1)$, $V^G_t(p_0, \emptyset) \geq V^G_t(p_0, 1)$. Thus, $p^*_t \geq p_0 > \hat{p}_t$.

Further, note that there exists an $\underline{X}$ such that if $X > \underline{X}$, then $A^B_t \in (0, 1)$ for all t.

Thus, if $X > \underline{X}$, $p_t^* > \hat{p}_t$ in all t. □

To prove Theorem 3, I state and prove three lemmas. The proofs of the first two lemmas are relegated to the Online Appendix.

Lemma 8. *Fix all parameters except β and X. There exists an $\underline{X}$ such that if $X > \underline{X}$, $A_t^B \in (0,1)$ in any equilibrium, under any β.*

Lemma 9. *Fix all parameters except β and X. Suppose $X > \underline{X}$, where $\underline{X}$ is established in Lemma 8. For any t, if there exists $\overline{p}_t$ such that $p_t^* < \overline{p}_t$ in all equilibrium for any β, there exits $\underline{R} > 0$ and $\overline{R} < 1$ such that $R(t,0) \geq \underline{R}$ and $R(t,1) \leq \overline{R}$.*

Lemma 10. *Fix all parameters except β and X. There exists $\underline{X} > 0$ such that if $X > \underline{X}$, then for all t, there exists $\overline{p}_t \in (0,1)$ such that $p_t^* < \overline{p}_t$ in any equilibria.*

Proof. Suppose $X > \underline{X}$, where $\underline{X}$ is the bound established in Lemma 8. I will now prove that for all t, there exists $\overline{p}_t \in (0,1)$ such that $p_T^* < \overline{p}_t$ in all equilibria. Proof by (backwards) induction.

Base case: By Lemma 8, $p_T^* = p_0$.

Induction step: Fix a $t \in \{1,..,T\}$. Assume for all $s > t$, there exists such a $\overline{p}_s \in (0,1)$. By Lemma 9, for all $s > t$, there exists an $\underline{R}_s > 0$ and $\overline{R}_s < 1$ such that $R(s,0) > \underline{R}_s$ and $R(s,1) < \overline{R}_s$. We want to show that there exists a $\overline{p}_t$ satisfying the above condition.

Suppose not, by contradiction. Then, there exists a sequence of X $\{X_n\}_{n=1}^{\infty}$ with $\lim_{n\to\infty} X_n = 1$ such that $\lim_{n\to\infty} p_t^{*,n} \to 1$, where $p_t^{*,n}$ is some equilibrium under $p_t^{*,n}$ for some β.

First, consider the case where $\lim_{n\to\infty} R^n(t,1) = 1$. By Lemma 9, there exists a $\overline{R}_{t+1} < 1$ such that $R^n(t+1,1) < \overline{R}_{t+1}$ for any n. By the continuity of the V^n, there exists $\overline{p}_t$ such that $p_t^* < \overline{p}_t$ for all n. Contradiction.

Next, consider the case where $R^n(t,1)$ does not converge to 1. Then, there exists an infinite subsequence indexed by m and an upper bound $\overline{R}$ such that for all m, $R^m(t,1) < \overline{R}$. Recall by the definition of R^m, there exists $\underline{p}$ such that for all m,

$$\frac{Pr^m(\tau = t|B, \theta = 1)}{Pr^m(\tau = t|G, \theta = 1)} > \underline{p}.$$

Since $\lim_{m\to\infty} p_t^{*,m} = 1$,

$$\lim_{m\to\infty} \frac{Pr^m(\tau = t|G, \theta = 1)}{Pr^m(\tau = t|G, \theta = 0)} = \infty.$$

Since $Pr^m(\tau = t|B, \theta = 1) = Pr^m(\tau = t|B, \theta = 0)$, it follows that

$$\lim_{m\to\infty} \frac{Pr^m(\tau = t|B, \theta = 0)}{Pr^m(\tau = t|G, \theta = 0)} = 1,$$

and thus $\lim_{m\to\infty} R^m(t, 0) = 0$. Note that by B's indifference condition, for all $s > t$,

$$(1 - p_0)R^m(t, 0) + p_0 R^m(t, 1) = (1 - p_0)R^m(s, 0) + pR^m(s, 1),$$

Since there exists a $\underline{R}$ such that $R^m(s, 0) > \underline{R}$ for all m, it follows that there exists an M and $\Delta > 0$ such that for all $m > M$,

$$R^m(t, 1) > \Delta + R^m(s, 1) \text{ for all } s > t.$$

Thus, it again follows that from the continuity of V^m that there exists $\overline{p}_t$ such that $p_t^{*,m} < \overline{p}_t$ for all $m > M$. Contradiction. □

Proof of Theorem 3. To prove the first bullet, note that it follows from Lemma 10 that for all t, there exists a $\underline{X} \in (0, 1)$ and $\overline{p}_t \in (0, 1)$ such that $p_t^* < \overline{p}_t$ under any equilibrium for any β when $X > \underline{X}$. Fix this $\underline{X}$ and these $\overline{p}_t$. For any $t < T$, fixing all other parameters (except X and β) there exists $\underline{\beta}_t \in (0, 1)$ such that if $\beta > \underline{\beta}_t$, $\hat{p}_t > \overline{p}_t$. Thus, for any $t < T$ and for all $X > \underline{X}$ and $\beta \geq \max_{s\leq t} \underline{\beta}_t$, $\hat{p}_t > p_t^*$.

To prove the second bullet, it suffices to show $p_T^* > \hat{p}_T$. Note that there exists $\tilde{X} \in (0, 1)$ such that if $X > \tilde{X}$, $A_t^B \in (0, 1)$ for all t, in any equilibrium. Thus, by the same reasoning presented in the proof of Theorem 2, since $\hat{p}_T < p_0$, $p_T^* > \hat{p}_T$. □