Dynamic Signaling in Wald Options*

Doruk Cetemen[†] Chiara Margaria[‡]

August 18, 2025

Abstract

A sender engages in costly signaling to influence a decision maker, who observes a biased noisy signal and decides when to take an irreversible action to match the binary state. We characterize Markov equilibria in terms of a two-dimensional boundary value problem for fixed discount rates and present sufficient conditions on the primitives for the two types of sender to choose an action exceeding their myopic optimum in all equilibria. We sharply characterize equilibrium behavior when either player becomes arbitrarily patient. In a dynamic limit pricing game, a potential entrant infers profitability of entry from the price that is affected by the incumbent's action. With demand uncertainty, a patient incumbent produces at capacity, and consumers can be hurt because the entrant strategically delays its entry. If the uncertainty pertains to the incumbent's cost, in the limit, the price conveys no information to the entrant, and consumers never benefit from limit pricing.

Keywords: Dynamic Signaling, Optimal Stopping, Continuous-time, Wald Problem.

JEL Codes: C73, D21, D43, D82, D83

*We would like to thank our discussants Laura Doval, and Mehmet Ekmekci for their insightful comments and suggestions. We thank Alp Atakan, Alessandro Bonatti, Paolo Barelli, Simon Board, Gonzalo Cisternas, Martin Cripps, Tiziano De Angelis, Toomas Hinnosaar, Johannes Hörner, Ilwoo Hwang, Hari Govindan, Edoardo Grillo, Emin Karagözoğlu, Ayca Kaya, Aaron Kolb, David Levine, Bart Lipman, Lucas Maestri, Arina Nikandrova, Juan Ortner, Pietro Ortoleva, Jérôme Renault, Patrick Rey, Klaus Ritzberger, Larry Samuelson, Bruno Strulovici, Can Urgun, and Leeat Yariv for their comments. We also thank seminar participants at Bristol University; Boston University; City, University of London; Collegio Carlo Alberto; Duke University; Durham University; Oxford University; the Parisian Seminar of Game Theory; Princeton University; Royal Holloway, University of London; Rotman School of Management; SAET 2022; The 33rd Stony Brook International Conference on Game Theory; Seoul National University; SITE 2023; Toulouse School of Economics; University College London; The University of North Carolina at Chapel Hill; University of Rochester; University of Warwick; the 2023 North American Summer Meeting of the Econometric Society; 20th Annual IO Theory Conference; the CEPR Virtual IO Seminar; 6th University of Bergamo IO Winter Symposium; University of Luxembourg; Microsoft Research New England; the Micro-CRC-Seminar at the University of Mannheim. We thank Ayush Gupta and Beixi Zhou for excellent research assistance. This work was funded by a grant from the European Research Council ERC 101165373.

[†]Luiss and EIEF, dcetemen@luiss.it.

[‡]Boston University, margaria@bu.edu.

1 Introduction

Wald's problem of sequential information acquisition has been extensively studied in economics, for example, in the context of R&D dynamics (Moscarini and Smith, 2003) or drug approval (Henry and Ottaviani, 2019). In the classical formulation, the decision maker can acquire multiple i.i.d. signals of exogenous informativeness. However, in strategic settings, interested parties may be able to affect the signal that the decision maker acquires before taking action. For example, venture capitalists wish to invest only in successful projects, but start-ups can try to affect their own periodic performance reports; policy makers would like to approve the most effective policies, but are advised by lobbies and interest groups. On the other hand, in light of insights from the literature on strategic information transmission, it is natural to wonder whether a decision maker could benefit from interacting with a sender who is informed about the state of the world and can potentially manipulate the signal.

To analyze these situations, we study a dynamic game in continuous time between two long-lived players, a decision maker (DM) and a sender. As in Wald (1945), the DM observes a public signal about a binary payoff-relevant state of the world and acts when sufficiently convinced of one state. A sender, who is privately informed about the state of the world, and has state-independent preferences over the DM's action, can engage in costly effort to affect the public signal observed by the DM. We model the public signal as a diffusion process whose drift depends on the true state of the world and the sender's action. Following Orlov, Skrzypacz, and Zryumov (2020), we call this class of games strategic Wald option games.

The continuous-time framework allows us to obtain a tractable characterization of Markov equilibria of Wald option games without the need to focus on linear equilibria, common in dynamic signaling models with Gaussian information structure. Our characterization complements the traditional adverse selection approach to reputation, providing a framework to conduct policy and welfare analysis, because the equilibrium predictions do not depend on the exogenously specified behavioral types.

For the case in which the sender is sufficiently patient, we give a sharp answer to the question of whether the DM prefers to interact with a strategic sender, who can manipulate the signal, or she would rather interact with a sender who takes a (state-dependent) constant action, resulting in a signal of exogenous informativeness. We provide necessary and sufficient conditions for the DM to benefit from interacting

with a strategic sender. In particular, our conditions imply that signal manipulation is beneficial to the DM when the difference in conditional drifts is increasing in action.

The tools we developed to study Wald option games allow us to revisit the classic model of dynamic limit pricing, which captures how an incumbent with private information may try to deter entry, and to uncover novel strategic sources of entry delay with implications for welfare and policy. The dynamic limit pricing model illustrates the analytical traction of the continuous-time setup, but the strategic Wald option game framework has far-reaching applications. Our results could be applied to study the signaling dynamics of drug approval, as in Henry and Ottaviani (2019); or, in political economy, to study lobbying by interest groups or efforts to influence the public opinion when an agent (e.g., parliament, incumbent government) decides the moment at which consultations must stop and a decision has to be reached, in the spirit of Brocas and Carrillo (2007) and Salas (2019); or, in organizational economics, to study moral hazard problem when the employer is incentivizing effort using both pay-per-performance and promotion, as in Fairburn and Malcomson (2001).

In the model, the DM observes a public signal, a diffusion process, whose drift depends on the sender's action as well as the underlying state of the world, privately known to the sender. In the leading example of dynamic limit pricing, the public signal is the market price, which the incumbent can affect by its choice of output. The binary underlying uncertainty captures either the state of the demand, which can be strong or weak, or the incumbent's marginal cost. The DM decides when if ever, to stop observing the signal and take an irreversible action to match the state—in the dynamic limit pricing, to pay an entry cost to become an incumbent's competitor or take an outside option. The sender has state-independent preferences in that it always prefers the DM to take one of the two action—the incumbent prefers the potential entrant to take the outside option, regardless of the state of the demand or its marginal cost. Affecting the public signal is costly, because it involves taking an action other than the (state-dependent) myopically optimal action, e.g., the monopoly quantity.

We provide a characterization of Markov equilibria for fixed discount rates in terms of a system of non-linear second-order ordinary differential equations. In equilibrium, both the value functions of the two types of sender and their actions are determined by a solution to a boundary value problem. A key step of our proof of equilibrium existence is showing that this multidimensional boundary value problem has a bounded solution, as there is no general existence theory for such problems.

This technical result, based on the method of upper and lower solutions, is of independent interest, as it can be readily applied, for instance, to prove the existence of Markov perfect equilibria in multi-player games with a publicly observable state variable (as in Georgiadis, 2014), or in games where some of the agents are jointly learning about an underlying state of the world (as in Bolton and Harris, 1999). Crucially, the approach is general enough to accommodate asymmetric agents and does not rely on specific assumptions about payoff functions or signal structures.

It is intuitive, and we verify that it is true, that in any equilibrium of the limit pricing example, both types of incumbent have an incentive to put downward pressure on the price. In general, however, players' incentives to choose an action higher or lower than the myopically optimal action depend on the sensitivity of the continuation value to public signals, an equilibrium object. We also provide sufficient conditions in terms of the primitives of the game to guarantee that in any equilibrium both types of sender choose an action higher than the myopically optimal action.

Using our characterization, we investigate equilibrium outcomes as the players become arbitrarily patient. First, for a fixed level of patience of the DM, if the sender's discount rate is sufficiently low, both types of sender find it optimal to choose the same extremal action at any belief, foregoing short-run gains. As a result, if the sender is sufficiently patient, the equilibrium is unique and involves a constant action, which simplifies the ranking of the informativeness of the public signal, as discussed above.

Second, we fix the sender's discount rate and look at the limit as the DM becomes arbitrarily patient. In the limit, the cutoffs at which the DM acts shift closer to the extreme values (i.e., 0 and 1): the DM acts only when the uncertainty has vanished and takes a perfectly informed action. As a result, in equilibrium, both types of sender forfeit manipulating the DM's belief and choose the myopically optimal action at any point on the equilibrium path, because persuading the DM to take the sender's favorite action is too costly for an impatient sender.

Third, we consider the case when both discount rates converge to zero at the same speed. As in the previous case, in the limit, the DM acts only at extreme beliefs; however, unlike before, the cost of manipulation does not increase unboundedly. Because the sender is becoming patient at the same speed as the DM, even if persuading

the DM takes more time, in the limit, both types of sender choose the same extremal action.

In the context of the application, we first show that when the incumbent has private information about the state of the demand, the answer to the question we posed in the general framework—does the potential entrant (i.e., the DM) benefit from extreme limit pricing (i.e., from signal manipulation by a patient incumbent firm)?—depends on the characteristics of the demand function. For instance, when demand is highly elastic, equilibrium overproduction of the incumbent results in a less informative signal for the potential entrant.

Second, while the trade-off between lower prices and entry delay is frequently mentioned by practitioners, the existing literature has focused on the welfare costs due to inefficient entry, but, to the best of our knowledge, has overlooked the intensive time margin, that is, the entry delay. In our model, on the one hand, if the incumbent adopts an aggressive strategy consumers will be better off because of the price cuts. On the other hand, depending on the effect of this aggressive pricing on the informativeness of the signal observed by the potential entrant, entry may be delayed, ultimately hurting consumers. Our results shed light on the efficacy of output restriction rules, which have been proposed, for example, by Williamson (1977) and Edlin (2002) as an antitrust policy tool to mitigate predatory behavior. Depending on the objective of the antitrust authority—whether minimizing delay or maximizing the probability of entry—and the characteristics of the demand function, it may be optimal to implement a more or a less stringent output restriction rule.

The results are quite different when the incumbents' private information concerns its marginal cost, as in Milgrom and Roberts (1982) or Matthews and Mirman (1983). In a two-period model, the incomplete cost information model and incomplete demand information model are equivalent. However, when the horizon is infinite, the two models predict starkly different equilibrium outcomes in the limit as the players become arbitrarily patient. If the incumbent has private information about the demand, the potential entrant's value function converges to the full-information value function, and consumers may benefit from a long period of aggressive limit pricing; in the case of information about cost, the potential entrant's value function converges

¹For example, in the recent Intel antitrust case, the EU Commission cited "a direct and immediate negative impact on those customers who would have had a wider price and quality choice," while Intel argued that "price declines brought large gains to the ultimate consumers who purchase computers." (Case COMP/C-3/37.990–Intel, Commission Decision, 2009, OJ C 227, 13–17).

to the no-information value function, and consumers never benefit from limit pricing. Hence, with cost uncertainty, not only the potential entrant (i.e., the DM) never benefits from limit pricing (i.e., from signal manipulation by a patient incumbent firm), but also the trade-off between delay and consumer surplus vanishes, rendering output restriction rules ineffective.

1.1 Related Literature

The paper belongs to the growing literature on dynamic signaling games with stopping decisions. In Daley and Green (2012), Kolb (2015), Kolb (2019), Dilmé (2019), and Gryglewicz and Kolb (2022), the informed player, unlike in our paper, takes the stopping decision; further, in all these papers but Dilmé (2019), the informed player cannot directly manipulate the signal. Unlike our paper, in Dilmé (2019), the action of one of the types is exogenously specified, which makes it closer to the reputation literature. Information manipulation is costless in the dynamic persuasion game of Orlov, Skrzypacz, and Zryumov (2020) when players' option exercise times are misaligned. Tangentially related, Henry and Ottaviani (2019) and McClellan (2022) analyze different versions of Wald persuasion games allowing the DM to commit.

The paper also contributes to the literature on dynamic signaling games with continuous-time Gaussian uncertainty. Bonatti, Cisternas, and Toikka (2017) and Cetemen (2020) study two-sided signaling models in continuous time, and Cisternas and Kolb (2024) analyze signaling in the presence of private monitoring. In contrast to our paper, these papers as well as the majority of papers in this literature typically focus on linear Markov equilibria. Our methodology allows us to move beyond the linear-quadratic-Gaussian framework, accommodating signals that are not additively separable in type and action, binary types, and irreversible actions.

At the intersection between the literature on dynamic signaling and the reputation literature, Ekmekci, Gorno, Maestri, Sun, and Wei (2022) study signal manipulation incentives in a dynamic principal-agent model in the presence of a commitment type.² Similarly to Pei (2021), and in contrast to the standard reputation models, we investigate reputation-building behavior as an equilibrium phenomenon in an incomplete information game without commitment types. This difference also sets our paper apart from the reputation literature with long-run players, such as Cripps and

²See also Ekmekci and Maestri (2022) for a discrete-time analogue of the model considered in Ekmekci et al. (2022).

Thomas (1997), Celetani, Fudenberg, Levine, and Pesendorfer (1996) and Atakan and Ekmekci (2012, 2015).

The seminal work by Faingold and Sannikov (2011) studies reputation dynamics in a continuous-time game between a population of small players and a long-lived player, who can be a "normal" (i.e., strategic) type or a "commitment" type;³ in contrast, we consider a Bayesian game with two long-run strategic players. Strategic types introduce significant complications, as both pooling and separating incentives are present and need to be considered in the equilibrium characterization. Technically, as compared to Faingold and Sannikov (2011), not only the equilibrium behavior of the sender is now characterized by a system of ODEs (instead of a single ODE), but also the presence of a long-run player taking a stopping decision introduces a fixed point problem, absent in their paper.⁴

In our leading application, we analyze a continuous-time dynamic version of a limit pricing game à la Milgrom and Roberts (1982). Our framework is closer to Matthews and Mirman (1983) not only because the incumbent's price is observed by the entrant with some noise, but also because the incumbent's private information concerns the state of the demand, rather than its production cost. The vast majority of the empirical and theoretical literature builds upon the framework of Milgrom and Roberts (1982), which is effectively static and not suitable for analyzing the effects of entry delay. Saloner (1984) and Toxvaerd (2017) extend the two-period model of Milgrom and Roberts (1982) to multiple periods. Unlike these papers, we do not endogenously impose a finite end to the game, and show how taking delay into account changes the equilibrium welfare properties. Within the non-equilibrium limit pricing literature Kamien and Schwartz (1971) use continuous-time tools to investigate the optimal pricing strategy of the incumbent when the potential entrant's behavior is exogenously specified. Recently, Gryglewicz and Kolb (2025) study entry deterrence in a stopping game in which the incumbent, as in Milgrom and Roberts (1982), has

³Bohren (2024) shows how their results extend to a more general class of stochastic games.

⁴Anderson and Smith (2013) and Dilmé (2025) rely on the tractability of continuous-time techniques to study a signaling game between a long-run informed player and a sequence of short-run players. In both papers, the Gaussian signal only depends on the informed player's actions while we allow it to depend both on his action and his type.

⁵Sweeting, Roberts, and Gedge (2020) build a finite-horizon analytically tractable model of dynamic limit pricing to structurally investigate the reduced-form evidence from Goolsbee and Syverson (2008).

private information about its costs and can choose to imitate a committed strong type via its pricing strategy.

Lastly, we contribute to the growing literature on continuous-time methods in dynamic games. Recent work has extended classical tools to strategic settings; for example, Barilla and Gonçalves (2024) and Escudé and Sinander (2023) adapt arguments from the viscosity literature, while Durandard and Strulovici (2022) introduce a weaker solution concept for differential equations that can be used to establish the existence of equilibria in dynamic games; Cisternas and Kolb (2024) show the existence of linear Markov equilibria in a dynamic signaling game with private monitoring by solving a boundary value problem. In the same spirit, we prove equilibrium existence by tackling a multidimensional, nonlinear boundary value problem for which known methods do not apply, using a new approach that combines the method of upper and lower solutions with a suitable Nagumo-type condition.

2 Model

A sender and a decision maker (DM) interact over time. Time is continuous and potentially infinite, $t \in [0, \infty)$. A persistent state of the world θ determines the payoffs. The sender knows the state of the world $\theta \in \{H, L\} \subset \mathbf{R}$. At each time $t \geq 0$, the sender chooses an action $a_t \in A \subset \mathbf{R}$ from a compact interval, where $A := [\underline{a}, \overline{a}]$.

The DM decides when to take an irreversible action; that is, he chooses a stopping time τ , together with an action $b_{\tau} \in \{h, l\}$ to take at that time. The DM is uninformed about the state but observes at each point in time a signal which evolves according to

$$dX_t = \mu(\theta, a_t) dt + \sigma dZ_t, \qquad X_0 = 0,$$

where $\sigma^2 > 0$, and Z_t is a standard Brownian motion which is independent of θ . We assume that $\mu : \{H, L\} \times A \to \mathbf{R}$ is Lipschitz continuous and non-increasing in its second argument, the action of the sender, and for $a \in A$, $\mu(H, a) \neq \mu(L, a)$. That is, the two types are statistically distinguishable when they take the same action. Nevertheless, there may exist a pair of feasible actions $a', a'' \in A$ that make the two types statistically indistinguishable, that is, $\mu(H, a') = \mu(L, a'')$.

At each time t before the DM acts, the sender receives a flow payoff $\pi(\theta, a) \geq \underline{\pi}$, for some $\underline{\pi} \leq 0$. We assume that for any θ , and any a, $\pi(\theta, a)$ is a Lipschitz continuous and strictly concave function of the action. Let $a_{\theta}^* := \arg \max_a \pi(\theta, a)$ denote the myopic optimal action for the sender of type θ . We denote by $\pi^*(\theta) := \pi(\theta, a_{\theta}^*)$ the myopic optimal payoff.

The DM and the sender discount the future at a rate $r_{DM} > 0$ and $r_S > 0$, respectively. If the DM acts at τ , his realized payoff is

$$\begin{cases} e^{-r_{DM}\tau}G(\theta, h) & \text{if } b_{\tau} = h, \\ e^{-r_{DM}\tau}G(\theta, l) & \text{if } b_{\tau} = l. \end{cases}$$

We assume that G(H,h) > G(H,l) and G(L,l) > G(L,h), and min $\{G(H,h), G(L,l)\} > 0$ so the DM always wants to match the state.⁶ At the time the DM acts, the sender collects a lump-sum payoff $\Pi(\theta,b_{\tau})$, which depends on the state and the DM's terminal action. The sender always prefers the DM to choose action l, $\pi^*(\theta) = \Pi(\theta,l) > \Pi(\theta,h) \geq \underline{\pi}$. The assumption $\Pi(\theta,l) = \pi^*(\theta)$ captures the idea that once the sender obtains his favorite action, signaling concerns disappear, and he can achieve his myopic payoff in the (unmodelled) continuation game.⁷

A public strategy for the sender is a square-integrable process $(a_t)_{t\geq 0}$ that is progressively measurable with respect to the filtration generated by $(\theta, (X_t)_{t\geq 0})$. A strategy for the DM specifies a stopping time and an action to take when stopping that are progressively measurable with respect to the filtration generated by $(X_t)_{t\geq 0}$.

We denote by $(\phi_t)_{t\geq 0}$ the process of posterior belief that the DM attaches to $\theta = H$, where $(\phi_t)_{t\geq 0}$ is a progressively measurable process with respect to the filtration generated by $(X_t)_{t\geq 0}$, taking values in [0,1]. Hence, given a (public) strategy profile for the two types of sender, $(a_{t,H}, a_{t,L})$, by Liptser and Shiryaev (2001), the belief evolves according to

$$d\phi_{t} = \frac{\phi_{t}(1 - \phi_{t}) (\mu(H, a_{t,H}) - \mu(L, a_{t,L}))}{\sigma} \cdot \frac{dX_{t} - (\phi_{t}\mu(H, a_{t,H}) + (1 - \phi_{t}) \mu(L, a_{t,L})) dt}{\sigma}.$$
(1)

⁶In Section 4, we show that the model can be generalized to the case when the DM can take only one action but may want to prefer never to act.

⁷The assumption can be relaxed to $\Pi(\theta,l) > \pi^*(\theta)$. Details are available upon request.

The innovation process on the second line is a standard Brownian motion from the point of view of the DM. We define the speed of learning $\gamma(a_H, a_L, \phi)$ as the volatility of the DM's belief, that is,

$$\gamma(a_H, a_L, \phi) := \frac{\phi(1 - \phi) \left(\mu(H, a_H) - \mu(L, a_L)\right)}{\sigma}.$$

It is determined by the DM's expectation about the action of each type of sender and the current belief ϕ , and it converges to 0 as ϕ approaches 0 or 1. Along the path of play, when contemplating a deviation, the sender anticipates that he cannot directly affect the speed of learning, because this speed is based on the DM's conjecture rather than on the actual action of the sender. The instantaneous choice of the sender can affect only the inference that the DM draws from the public signal by affecting the actual drift of the public signal. Higher $\gamma(a_H, a_L, \phi)$ implies that the belief ϕ_t reacts more to the public signal.

Given a strategy profile, the expected discounted payoff of the type θ sender can be written as

$$\mathbf{E}_{\theta} \left[\int_{0}^{\tau} r_{S} e^{-r_{S}s} \pi(\theta, a_{s}) \, \mathrm{d}s + e^{-r_{S}\tau} \left(\Pi(\theta, h) \mathbf{1}_{b_{\tau} = h} + \Pi(\theta, l) \mathbf{1}_{b_{\tau} = l} \right) \right].$$

Similarly, the expected discounted payoff of the DM given a strategy profile can be written as

We focus on equilibria that are Markovian in the posterior belief ϕ_t . A strategy profile for the sender is Markovian in ϕ_t if $(a_{t,H}, a_{t,L}) = (a_H(\phi_t), a_L(\phi_t))$ for some measurable function $a_{\theta} : [0,1] \to A$, for $\theta \in \{H, L\}$. A strategy for the DM is Markovian in ϕ_t if $\tau = \inf\{t : \phi_t \notin \mathcal{D} \subset [0,1]\}$ a.s., and $b_{\tau} \in \arg\max\{\phi_{\tau}G(H,h) + (1-\phi_{\tau})G(L,h), \phi_{\tau}G(H,l) + (1-\phi_{\tau})G(L,l)\}$. Without loss of generality, we can assume that $\mathcal{D} = [\phi, \overline{\phi}]$.

A Markov strategy profile together with a belief process $(\phi_t)_{t\geq 0}$ is a pure-strategy Markov equilibrium if at any time, along any public history,

- (i) the DM's strategy solves his optimal stopping problem given the sender's strategy;
- (ii) the sender's strategy maximizes his expected continuation payoff at any $\phi \in (\phi, \overline{\phi})$;
- (iii) the belief process $(\phi_t)_{t\geq 0}$ evolves according to (1) for $(a_{t,H}, a_{t,L}) = (a_H(\phi_t), a_L(\phi_t))$, given the initial prior ϕ_0 .

3 Equilibrium Characterization

3.1 Equilibrium Existence

Our equilibrium characterization relies on two conditions. The first condition guarantees that the volatility of beliefs is bounded away from zero for any belief. In their continuous-time reputation model, Faingold and Sannikov (2011) show that one can guarantee this by imposing an appropriate continuous-time equivalent of the identifiability condition in Cripps et al. (2004): they assume that when the (strategic type of the) sender behaves myopically, his behavior is statistically distinguishable from the behavior of the commitment type. We show that in the absence of commitment types, it is enough to require that when different types of sender behave myopically optimal, they are statistically distinguishable. Specifically, we impose the following condition:

Condition 1.
$$\mu(H, a_H^*) \neq \mu(L, a_L^*)$$
.

Intuitively, for the two types to play a pair of observationally equivalent actions, the sender must be given intertemporal incentives. This requirement implies that in equilibrium, the volatility of the public belief is bounded above zero for any interior belief level.

In any equilibrium, the sender faces a stochastic optimal control problem, because his action affects the drift of the belief process ϕ_t . By using standard techniques, assuming that the value function $U_H:(0,1)\to \mathbf{R}$ is twice continuously differentiable, we can write the Hamilton-Jacobi-Bellman (HJB) equation for the problem of the type H sender as

$$r_{S}U_{H}(\phi) = \max_{a \in A} \left\{ r_{S}\pi(H, a) + \gamma \left(\hat{a}_{H}(\phi), \hat{a}_{L}(\phi), \phi \right) \frac{\mu(H, a)}{\sigma} U'_{H}(\phi) \right\}$$

$$- \gamma \left(\hat{a}_{H}(\phi), \hat{a}_{L}(\phi), \phi \right) \frac{\phi \mu(H, \hat{a}_{H}(\phi)) + (1 - \phi)\mu(L, \hat{a}_{L}(\phi))}{\sigma} U'_{H}(\phi)$$

$$+ \frac{1}{2} U''_{H}(\phi) \left(\gamma \left(\hat{a}_{H}(\phi), \hat{a}_{L}(\phi), \phi \right) \right)^{2},$$
(2)

where $\hat{a}_{\theta}:[0,1]\to A, \theta\in\{H,L\}$ is the conjectured strategy profile used by the DM.

When best replying, the sender trades off instantaneous payoffs (the first term in the parenthesis on line (2)) and the effect that the sender's action has on the continuation payoff (the second term in the parenthesis). In a Markov equilibrium, the expected impact of today's action on the continuation payoff depends on its effect on the belief. The expected drift of the belief, from the point of view of the sender, is proportional to the drift of the public signal, which his action affects directly, and to the speed of learning. Finally, the sensitivity of the continuation payoffs to the belief is captured by the derivative of the value function (the last term in the parenthesis).

In any Markov equilibrium, the DM's conjectured strategy profile is correct, and for any $\phi \in (\phi, \bar{\phi})$, $(a_H(\phi_t), a_L(\phi_t)) = (a_H, a_L)$ solves the following system:

$$a_{H} \in \operatorname*{argmax}_{a' \in A} \pi(H, a') + \frac{\mu(H, a_{H}) - \mu(L, a_{L})}{\sigma} \frac{\mu(H, a')}{\sigma} z_{H},$$

$$a_{L} \in \operatorname*{argmax}_{a' \in A} \pi(L, a') + \frac{\mu(H, a_{H}) - \mu(L, a_{L})}{\sigma} \frac{\mu(L, a')}{\sigma} z_{L},$$

$$(3)$$

for $z_{\theta} = \phi(1-\phi)U'_{\theta}(\phi)/r_S$, $\theta \in \{H, L\}$. Intuitively, for any $(z_H, z_L) \in [0, 1] \times \mathbf{R} \times \mathbf{R}$, any solution to the system identifies a Bayes Nash equilibrium of an auxiliary (one-shot) signaling game in which the sender is of type H with probability ϕ , and flow payoffs are perturbed by a "continuation-game term" weighted by (z_H, z_L) . Let $\mathcal{N}: (z_H, z_L) \mapsto (a_H, a_L)$ denote the equilibrium correspondence, mapping payoff perturbations to action profiles that solve the system above.

The second condition that we impose on the primitives ensures the existence of at most one interior equilibrium. This, in turn, guarantees that the equilibrium correspondence \mathcal{N} , admits a continuous selector. (We discuss an alternative more general condition in Section 3.5.)

Condition 2.

- 1. $\mu(\theta, a)$ is linear in a, $\pi(\theta, a)$ is quadratic in a, a and $\pi_{aa}(L, a)(\mu(H, \overline{a}) \mu(L, \overline{a})) + \pi_a(L, \overline{a})\mu_a(L, a) \leq 0$.
- 2. Either of the following holds:

(a) i.
$$\Pi(H, l) - \Pi(H, h) > \Pi(L, l) - \Pi(L, h)$$
,
ii. $a_H^* > a_L^*$ and $\pi_{aa}(H, a) \ge \pi_{aa}(L, a)$,
iii. $\mu_a(H, a) \le \mu_a(L, a)$,
iv. $\mu(H, a_H^*) - \mu(L, a_L^*) > 0$,
(b) $\mu(H, a_H^*) - \mu(L, \underline{a}) > 0$.

In light of the pervasive issue of equilibrium multiplicity in signaling games—whether static (Spence, 1973) or dynamic (see Noldeke and van Damme, 1990)—it is not surprising that equilibrium uniqueness in the one-shot signaling game requires additional conditions. It is remarkable, however, that simple conditions on the primitives not only guarantee that for any payoff perturbation the equilibrium has a unique interior equilibrium, but also that this equilibrium has sufficient stability properties to allow us to construct a continuous mapping $\mathcal{N}:(z_H,z_L)\mapsto(a_H,a_L)$ with the property that $\mathcal{N}(z_H,z_L)=(\bar{a},\bar{a})$ for (z_H,z_L) low enough. That is, if the signaling incentives are strong enough, in the unique equilibrium of the auxiliary one-shot signaling game, both types of sender play the same extremal action. (The proof of this implication is in Section OA.4.6 in the Appendix.)

Armed with the continuous selector, we can characterize the sender's pseudo-best reply (i.e., a mapping from the DM's strategy to the action profile for the two types of sender). As explained above, in principle, when best replying, the sender must also take into account the conjecture used by the DM. We call these functions pseudo-best reply, rather than best reply, because in constructing them, we impose that the DM's conjecture is correct. The existence of a Markov equilibrium then follows from a fixed-point argument. In fact, given the conjectured equilibrium behavior of the two types of sender, the best-reply problem of the DM is standard. That is, the DM engages in sequential testing of two hypotheses on the mean of a Wiener process.

While simple, Condition 2 is stronger than necessary as it implies Condition 1. In light of this observation, we first state our main results under Condition 2, which we shall relax in Section 3.5.

⁸By the first part, $\pi_{aa}(L,a)$ and $\mu_a(L,a)$ are independent of a.

Theorem 1. Assume Condition 2 is satisfied. Then, a Markov equilibrium exists. In any equilibrium, the value functions of the sender solve the following system of second-order ordinary differential equations over the interval $(\phi, \bar{\phi})$,

$$U_{H}''(\phi) = -2\frac{U_{H}'(\phi)}{\phi} + \frac{2r_{S}(U_{H}(\phi) - \pi(H, a_{H}(\phi)))}{(\gamma(a_{H}(\phi), a_{L}(\phi), \phi))^{2}},$$

$$U_{L}''(\phi) = 2\frac{U_{L}'(\phi)}{1 - \phi} + \frac{2r_{S}(U_{L}(\phi) - \pi(L, a_{L}(\phi)))}{(\gamma(a_{H}(\phi), a_{L}(\phi), \phi))^{2}},$$
(4)

subject to the boundary conditions $U_{\theta}(\bar{\phi}) = \Pi(\theta, H), U_{\theta}(\underline{\phi}) = \Pi(\theta, L), \text{ and } (a_H(\phi), a_L(\phi)) = \mathcal{N}(\phi(1-\phi)U'_H(\phi)/r_S, \phi(1-\phi)U'_L(\phi)/r_S) \text{ for any } \phi \in [\phi, \bar{\phi}].$

In proving the result, we need to show existence of a bounded solution to the multidimensional second-order non-linear boundary value problem characterizing the sender's pseudo-best reply problem. As there is no off-the-shelf result that guarantees the existence of a solution to such a problem, our proof of existence, of independent interest, is based on the method of upper and lower solutions and leverages the monotonicity of the system of differential equations. See Theorem OA.1 in the Online Appendix.

Given the DM's strategy $\bar{\phi}$ and $\underline{\phi}$, equation (4) characterizes the behavior of the sender on the equilibrium path, that is, for beliefs $\phi \in [\underline{\phi}, \bar{\phi}]$. In proving the existence of a Markov equilibrium, we must complete the strategy profile outside the interval $[\underline{\phi}, \bar{\phi}]$. A natural specification of the behavior off the equilibrium path is to assume that both types of sender choose the myopically optimal action at any $\phi \notin [\underline{\phi}, \bar{\phi}]$. In the proof of Theorem 1, we prove the existence of an equilibrium by requiring that for $\phi \notin [\underline{\phi}, \bar{\phi}]$, $a_{\theta}(\phi) = a_{\theta}^*$ for $\theta \in \{H, L\}$.

Intuitively, in a discrete-time approximation of our game, the sender's action at time t affects the belief of the DM at time $t+\Delta$. As the time between periods shrinks, starting from a history off the equilibrium path, with high probability, the belief will not leave the set $[0, \underline{\phi}) \cup (\overline{\phi}, 1]$ regardless of the action of the sender. In discrete-time dynamic games, sequential rationality would then imply that off the path, the sender plays the myopically optimal action.⁹

On the one hand, the refinement has some bite in that it is possible to construct spurious equilibria in which the DM is induced to act as soon as the belief enters

⁹One way of adapting this notion to continuous time was formalized by Kuvalekar and Lipnowski (2020), who suggested an instantaneous sequential rationality refinement.

some region in the anticipation that the two types of sender will adopt a strategy that makes the difference in conditional drift nil. For example, if for any $\phi \in [0, \phi_0 - \varepsilon) \cup (\phi_0 + \varepsilon, 1]$, $\mu(H, a_H(\phi)) = \mu(L, a_L(\phi))$, for some $\varepsilon > 0$ sufficiently small, the DM finds it optimal to act as soon as the belief leaves the interval $[\phi_0 - \varepsilon, \phi_0 + \varepsilon]$. On the other hand, because the DM's best-reply problem satisfies a smooth pasting condition, the equilibria we construct could be sustained by specifying an alternative off-path behavior. 11

As discussed above, Condition 1 guarantees that, in equilibrium, the variance of the belief process never vanishes. If Condition 1 fails and $\mu(H, a_H^*) = \mu(L, a_L^*)$, the game still has a Markov equilibrium which is characterized by (4), but at one (or both) of the two cutoffs, the equilibrium behavior of the two types of the sender may converge to the myopic play so that the posterior no longer updates.

We conclude the section with a standard square root law of substitution between the discount rate and the volatility of the signals. We shall refer to this result in Section 3.3 when we study the limit as the players get patient.

Corollary 1. Multiplying the discount rate of both the sender (r_S) and the DM (r_{DM}) by a factor of $\alpha > 0$ has the same effect on the equilibrium values and equilibrium behavior as rescaling the volatility parameter σ by a factor of $\sqrt{\alpha}$.

The corollary immediately follows from the observation that the discount rates and the volatility parameter enter the DM's and the sender's problems only through the products $r_{DM}\sigma^2$ and $r_S\sigma^2$, respectively.

3.2 Signaling Incentives

First, it is intuitive and indeed true that the value function of the sender is decreasing: both types benefit from the DM holding a lower belief, that is, attaching a lower probability to $\theta = H$. In fact, in proving existence of a solution to the system of ODEs in Theorem 1, we also show that any solution is monotone.

¹⁰Technically, the best-reply problem is well-defined only when the difference in conditional variance is bounded away from zero, but a limit argument can be provided to formally justify the claim.

¹¹Ekmekci et al. (2022) circumvent the need to specify the off-path belief by introducing friction in the stopping problem, that is, the agent who takes the irreversible stopping decision can do so only upon receiving an opportunity that arrives according to a Poisson process.

Proposition 1. In any Markov perfect equilibrium, the value functions of both types of sender are weakly decreasing, i.e., $U'_{\theta}(\phi) \leq 0$ for all $\phi \in [\phi, \overline{\phi}]$.

The sender has incentives to deviate from his myopically optimal action to manipulate the public signal and induce a lower belief. However, whether a higher action increases or decreases the DM's belief depends on the sign of γ , which is pinned down by the conjectured equilibrium behavior.

Because the sender always prefers the DM to take action l as compared to action h, the type H of the sender always wants to pool with the type L of the sender, while the latter, in turn, wants to separate. These strategic considerations, absent in a model with a "commitment" type, can lead to two types of equilibrium, depending on the sign of γ . When γ is positive, in equilibrium, a lower public signal is interpreted as evidence in favor of $\theta = L$. As a result, the type H of the sender tries to put downward pressure on the signal by choosing an action higher than the myopically optimal action. At the same time, by assumption, if the type L of the sender were to choose the same action as the type H, he would induce an even lower signal, on average; in equilibrium, his effort to separate also translates in an action which is higher than his myopically optimal action. The opposite dynamics ensues when γ is negative.

While in general γ can take either sign (see Section 3.5) the restrictions in the second part of Condition 2 imply that it is always positive in equilibrium, as formalized by the following proposition.

Proposition 2. In any equilibrium, for all $\phi \in [\underline{\phi}, \overline{\phi}]$, both types of sender choose an action higher than their myopically optimal action, i.e., $a_{\theta}(\phi) \geq a_{\theta}^*$, and an unexpectedly higher signal increases the public belief, i.e., $\gamma(a_H(\phi), a_L(\phi), \phi) > 0$.

3.3 Patience Limits

Using our equilibrium characterization, we investigate signaling incentives and equilibrium outcomes as the players become arbitrarily patient. To conduct the limit exercise, we first derive a uniform lower bound on the speed of learning which holds across discount rates (see Lemma 2 in the Appendix).

¹²This follows from $\mu(H, a) \neq \mu(L, a)$ and $\mu_a(H, a) < 0$, which we maintain throughout together with $\mu(H, a_H^*) - \mu(L, \underline{a}) > 0$, which is part of Condition 2.

In Section 3.5 we discuss the additional identifiability condition required to extend Lemma 2 when Condition 2 is relaxed.

Theorem 2. Assume Condition 2 is satisfied.

- 1. Fix the discount rate of the DM, $r_{DM} > 0$. If the sender is sufficiently patient, that is, for r_S low enough, in the unique Markov equilibrium, both types of sender play the same extremal feasible actions at any belief: $a_H(\phi) = a_L(\phi) = \bar{a}$ for all $\phi \in [\phi, \bar{\phi}]$.
- 2. Fix the discount rate of the sender, $r_S > 0$. In the limit, as the DM becomes arbitrarily patient, i.e., $\lim_{n\to\infty} r_{DM,n} = 0$, the strategy profile of both types of sender converges pointwise to the myopically optimal action, i.e., $\lim_{n\to\infty} a_{H,n}(\phi) = a_H^*$ and $\lim_{n\to\infty} a_{L,n}(\phi) = a_L^*$ for all $\phi \in [\lim_{n\to\infty} \underline{\phi}_n, \lim_{n\to\infty} \overline{\phi}_n]$. Moreover, $\lim_{n\to\infty} \underline{\phi}_n = 0$ and $\lim_{n\to\infty} \overline{\phi}_n = 1$.
- 3. Let $\{r_{S,n}\}_{n\geq 1}$ and $\{r_{DM,n}\}_{n\geq 1}$ be two sequences converging to zero such that $\lim_{n\to\infty} r_{S,n}/r_{DM,n} = k \in (0,\infty)$. Then along any sequence of Markov equilibria, the strategy profile of both types of sender converges pointwise to the same extremal feasible action, i.e., $\lim_{n\to\infty} a_{\theta}(\phi) = \bar{a}$ or $\lim_{n\to\infty} a_{\theta}(\phi) = \bar{a}$ for all $\phi \in [\lim_{n\to\infty} \phi_n, \lim_{n\to\infty} \bar{\phi}_n]$, and $\theta \in \{H, L\}$. Moreover, $\lim_{n\to\infty} \phi_n = 0$ and $\lim_{n\to\infty} \bar{\phi}_n = 1$.

The first two results in Theorem 2 can be understood in light of the trade-off faced by the sender. First, as the sender becomes arbitrarily patient, his incentives to manipulate the DM's belief are stronger because short-term considerations become less salient, see (2). Notably, when both types choose the same extremal action, there is still information revelation in that the difference in expected conditional drifts is not zero.

The result also sheds light on the difference between our model and models with a "commitment" type. Consider, for example, the case when the type H of the sender is committed to an action a^* such that $\mu(L, \overline{a}) < \mu(H, a^*) < \mu(L, \underline{a})$. In this case, even in the limit, the equilibrium is not unique, as both having type L playing \overline{a} at any belief and playing \underline{a} at any belief are equilibria. The committed type's interior action serves as an anchor and leads to equilibrium multiplicity: depending on whether an unexpectedly higher signal increases or decreases the public belief, the sender has incentives to play either extremal action.

Second, if the DM adopts an extreme strategy, that is, in the limit as the cutoff approaches 0 and 1, it is too costly for an impatient sender to try to engage in signaling, because such a strategy would have to involve a long period of belief manipulation. As the DM becomes arbitrarily patient, the marginal cost of waiting for additional information decreases, and the equilibrium cutoffs converge to 0 and 1. As a result, in the limit, neither type of sender engages in signaling and the equilibrium action profile converges to the myopically optimal action.

The third result combines the first two. To understand the intuition, consider a variation of our model in which the sender can only choose his action at time zero and cannot revise it at t>0. For simplicity, assume that the sender can be of either types with equal probability. In this case, the best-reply problem of type $\theta \neq \vartheta$ of the sender can be written as

$$\max_{a \in A} \left(1 - \mathbf{E}[e^{-r_S \tau}] \right) \pi(\theta, a) + \mathbf{E}[e^{-r_S \tau}] \left(\left(1 - \Phi\left(\frac{\mu(\theta, a) - \mu(\vartheta, a_{\vartheta})}{2\sigma}\right) \right) \Pi(\theta, h) + \Phi\left(\frac{\mu(\theta, a) - \mu(\vartheta, a_{\vartheta})}{2\sigma}\right) \Pi(\theta, l) \right).$$

The cutoffs used by the DM do not appear directly in the equation above, ¹³ but they determine the distribution of the stopping time τ . In the proof, we show that in the joint limit as both the sender and the DM become patient at comparable rates, optimality of the DM's behavior implies that the expected discount factor $\mathbf{E}[e^{-r_S\tau}]$ converges to 1. Inspection of the equation above reveals that the limit equilibrium must involve an extremal action for both types. While in the actual game, the sender has the ability to revise his action at any $t \in (0, \tau]$, in the limit, any gain from any such a revision is of a lower order.

In light of Corollary 1, the joint limit is equivalent to the limiting case as noise vanishes, i.e., $\sigma \to 0$, as monitoring becomes perfect. In the proof, we also show that in the limit, the value function of the DM converges to his full-information value function. The expected time before the DM acts diverges as the volatility parameter vanishes.

The DM takes action h (action l) whenever $\phi_{\tau} \geq \phi_0$ ($\phi_{\tau} \leq \phi_0$), and by assumption $\phi_0 = 1/2$.

3.4 Patient Sender vs Myopic Sender

Theorem 2 offers a simple answer to our initial question: does the DM prefer to interact with a strategic sender, who can manipulate the signal, or she would rather interact with a myopic sender, who always takes the (state-dependent) myopically optimal action? The DM prefers a strategic patient sender to a myopic sender if and only if $\mu(H, a_H^*) - \mu(L, a_L^*) < \mu(H, \bar{a}) - \mu(L, \bar{a})$.

As a result, if $a_H^* > a_L^*$ and the drift exhibits increasing differences, as explained below, the DM always benefits from facing a patient sender, as compared to a myopic sender.

Corollary 2. Assume $a_H^* > a_L^*$ and $\mu(H, a) - \mu(L, a)$ is increasing in a. Then, if the sender is patient enough, in the unique equilibrium, the DM achieves a higher payoff as compared to the case when the sender is myopic.

Because the corollary follows from Part 1 of Theorem 2, the DM benefits from interacting with a strategic sender as long as the sender discount rate belongs to interval $[0, \underline{r}_S]$, as the game has a unique equilibrium for a positive range of discount rate.

The result is in stark contrast with the existing results in the reputation literature. For instance, Ekmekci and Maestri (2022) address this question in a discrete-time reputation game between a sender and a DM who decides when to take an irreversible action. They examine how much useful information the DM can elicit when one of two types of sender is a commitment type. They show that when the sender wants the DM to stop as late as possible, in the limit as both players become arbitrarily patient, no information is revealed in equilibrium.

3.5 Beyond the Case of Linear Drift

Condition 2 in Theorem 1 could be replaced by the following condition, which is more general, but seemingly, harder to check.

Condition 2'. $\mathcal{N}(z_H, z_L)$ is a non-empty single-valued correspondence for each $(z_H, z_L) \in \mathbf{R} \times \mathbf{R}$. Moreover, \mathcal{N} is continuous on every bounded subset of $\mathbf{R} \times \mathbf{R}$.

Continuity suffices to ensure that, as before, there exists $(\underline{z}_H, \underline{z}_L) \in \mathbf{R}_- \times \mathbf{R}_-$ such that for any $\phi \in [0, 1]$, $\mathcal{N}(z_H, z_L) \subset \{(\underline{a}, \underline{a}), (\bar{a}, \bar{a})\}$ whenever $(z_H, z_L) \in (-\infty, \underline{z}_H] \times$

 $(-\infty, \underline{z}_L]$. That is, if the signaling incentives are strong enough, in the unique equilibrium of the auxiliary one-shot signaling game, both types of sender play the same extremal action.

As we discussed, Condition 2 is somewhat stronger than necessary, as its sole purpose is to show that Condition 2'. Theorem 1 can be restated with no change under Condition 1 and Condition 2'—it is now necessary to require Condition 1, as it is not implied by Condition 2'. However, under Condition 2' in equilibrium, γ is not necessarily positive, that is, an unexpected high public signal is not necessarily interpreted as evidence in favor of $\theta = H$, even if γ never changes sign in equilibrium. In the online Appendix, we generalize Proposition 2 by providing sufficient conditions for γ to be positive in any equilibrium under Condition 2'.

We also derive sufficient conditions on the primitives to determine the sign of γ in the limit as the sender becomes sufficiently patient, as again, in this case the equilibrium is unique. However, to derive a uniform lower bound on the speed of learning which holds across discount rates we need to strengthen the identifiability Condition 1 as follows.

Condition 3. Either of the following holds: $\mu(H, \overline{a}) > \mu(L, a_L^*)$ or $\mu(H, \underline{a}) < \mu(L, a_L^*)$.

As we assumed throughout that either $\mu(H, a) > \mu(L, a)$ or $\mu(H, a) < \mu(L, a)$ for all $a \in A$, the two inequalities in Condition 3 are mutually exclusive. The theorem below generalizes Part 1 of Theorem 2. It is more general in that, in the limit, the unique equilibrium involves both types of the sender playing either the highest or the lowest feasible action at all beliefs, depending on the primitives.

Theorem 2.1'. Assume Condition 2', and Condition 3. Fix the discount rate of the DM, $r_{DM} > 0$. If the sender is sufficiently patient, that is, for r_S low enough, in the unique Markov equilibrium, both types of sender play the same extremal feasible actions at any belief. If $\mu(H, a) - \mu(L, a) > 0$, then $a_H(\phi) = a_L(\phi) = \bar{a}$ for all $\phi \in [\underline{\phi}, \bar{\phi}]$; if $\mu(H, a) - \mu(L, a) < 0$, then $a_H(\phi) = a_L(\phi) = \underline{a}$ for all $\phi \in [\underline{\phi}, \bar{\phi}]$.

4 Dynamic Limit Pricing

As our leading application, we consider a dynamic model of entry deterrence à la Matthews and Mirman (1983). The objective of the section is two-fold. First, it illustrates the tractability of the general model, as well as the applied relevance of

the results in Section 3.4. Second, it demonstrates through a parsimonious framework that policy instruments can affect the informativeness of the signal when the incumbent engages in aggressive entry-deterrence behavior.

The sender is an incumbent firm that has private information about the state of the demand: it knows the demand shifter θ .¹⁴ We interpret the action of the sender as its choice of output level, which results in a stochastic inverse demand,

$$dX_t = \theta (1 - b \cdot a_t) dt + \sigma dZ_t, \qquad X_0 = 0, \tag{5}$$

where $\sigma^2 > 0$ is the variance, b > 0, and Z_t is a standard Brownian motion that is independent of θ . We assume that the set of feasible actions is $A = [0, \bar{a}]$, for some $\bar{a} > 0$, which can be interpreted either as a capacity constraint or as an output restriction rule, as in Williamson (1977) and Edlin (2002). We will discuss the policy implication of this interpretation in Section 4.2.1.

The DM is a potential entrant who is uninformed about the state of the demand and decides when to take an irreversible action: becoming a competitor by entering the incumbent's market or abandoning all entry prospects and taking an outside option.

The potential entrant observes the prevailing price at each point in time dX_t/dt , a linear demand perturbed by an additive i.i.d. noise.

In line with this interpretation, at each time t, before the potential entrant acts, given the incumbent's output choice and the realization of the inverse demand, the resulting increment in the incumbent's profit is

$$dX_t a_t - (c/2)a_t^2 dt, \qquad c > 0.$$

For tractability, we assume quadratic production costs.¹⁵ We leave the continuation game following the entrant's decision unmodelled, since our focus is on entry deterrence rather than potential collusion in the contestable industry. We assign continuation payoffs capturing the idea that the potential entrant would like to enter only if demand is strong ($\theta = H$), while the incumbent is better off when its monopoly remains unchallenged.

¹⁴We conjecture that our equilibrium construction extends to the case when the demand fluctuates over time (i.e., the demand shifter changes over time), as in Keller and Rady (1999), provided that the incumbent can observe the prevailing state.

¹⁵Hence, the myopic optimal action is the optimal monopoly quantity $a_{\theta}^* = \theta/(c+2b\theta)$.

That is, if the potential entrant acts at τ , its payoff is

$$\begin{cases} e^{-r_E \tau} o & \text{if } b_\tau = l \\ e^{-r_E \tau} (D_\theta - F) & \text{if } b_\tau = h \end{cases},$$

where $D_H - F > o > D_L - F$, and until Section 4.2.1 o > 0. The incumbent's payoff in the game is

$$r_I \int_0^{\tau} e^{-r_I t} a_t \, dX_t - r_I \int_0^{\tau} e^{-r_I t} (c/2) (a_t)^2 \, dt + e^{-r_I \tau} \left(M_{\theta} \mathbf{1}_{b_{\tau}=l} + D_{\theta} \mathbf{1}_{b_{\tau}=h} \right),$$

where $M_{\theta} > D_{\theta}$, for $\theta \in \{H, L\}$, $M_H > M_L$, and $D_H > D_L$. We set M_{θ} and D_{θ} equal to the profit of a firm with quadratic cost facing a linear demand curve with intercept θ and slope $b \cdot \theta$, that operates as a monopolist, or competes à la Cournot with a symmetric firm, respectively.^{16,17}

It may be argued that upon entry, the competitor is likely to not have access to the same information about the demand as the incumbent; or that even if the potential entrant takes the outside option, the incumbent's monopoly may still be threatened by a future competitor. In principle, it is possible to capture alternative information structures or the threat of future entry in the specification of the expected discounted payoffs that the firms collect at τ . While both the analytical and the numerical results we derive rely on the specific payoff assumptions, we believe that the insights generalize to these variations of the baseline model.

Discussion of the Assumptions. Because the demand shifter θ changes both the demand intercept and the slope, the difference in conditional drifts, $\mu(H, a) - \mu(L, a)$ is decreasing in a. To put it differently, if firms were to choose the same output, the informativeness of the price would be decreasing in output, see Figure 1.

¹⁶That is, $M_{\theta} = \theta^2/(4b\theta + 2c)$ and $D_{\theta} = (2b\theta + c)\theta^2/(2(3b\theta + c)^2)$.

¹⁷Asymmetric payoffs in the duopoly continuation game can easily be accommodated, to capture either asymmetric cost structure or a different equilibrium selection.

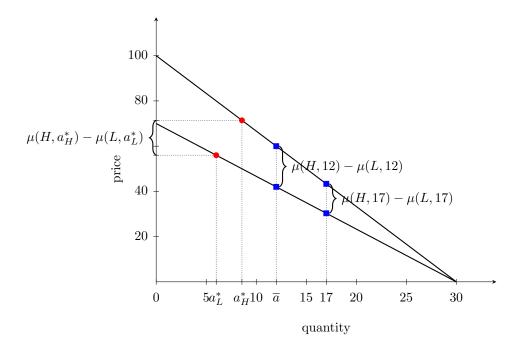


Figure 1: Illustration of the demand function and the information structure, $(H, L, b, c, \bar{a}) = (100, 70, 1/30, 5, 12)$.

Traditional theories of test market predation, ¹⁸ as well as the most cited cases, ¹⁹ and anecdotal evidence, ²⁰ concern situations in which the incumbent seeks to prevent the potential entrant from obtaining reliable information about the profitability of the market. In principle, however, aggressive pricing may allow the entrant to learn about the relevant segments of the market, or the demand at the relevant prices. ²¹

While parsimonious, our model captures both the scenario in which aggressive pricing decreases the informativeness of the signal, and the scenario in which the opposite is true. Compared to the case when the incumbent produces the (state-dependent) monopoly quantity, when the incumbent produces at capacity, the price

¹⁸Traditional industrial organization theories distinguish between aggressive pricing to deter entry or induce exit (limit pricing or predatory pricing, respectively), and aggressive pricing to reduce the information garnered by a new firm to deter its expansion. The latter is referred to as *test market predation*. (See, for example, Viscusi et al., 2018, Ch. 8.)

¹⁹ "In the 1980s P&G tried to get into the bleach business. [...] We went to test-market in Portland, Maine. [...] Do you know what Clorox did? They gave every household in Portland, Maine, a free gallon of Clorox bleach—delivered to the front door." (Dillon, 2011)

²⁰Smiley's 1988 surveys of firms find that firms attempt to limit entry by masking product-specific data on profitability within a multiproduct firm.

²¹For example, in the airline industry, the low-end market had been generally ignored by major carriers up until the entry or threat of entry of low-cost carriers (Brady and Cunningham, 2001).

can be more or less informative, depending on parameters, as summarized by the following observation.

Observation 1. The price is more informative when both types produce \bar{a} , as compared to when both produce the monopoly quantity, if and only if $\bar{a} < (Ha_H^* - La_L^*)/(H-L)$.

Thus, the informativeness of the price when the firm floods the market depends on the magnitude of \bar{a} . As discussed below, we interpret this parameter as an output restriction rule.

Lastly, we focus on the case of linear demand, however, we believe that the insights generalize not only to the case of nonlinear demand but also to the case of entry-deterrence strategies other than limit pricing. For example, Ellison and Ellison (2011) suggest that pharmaceutical firms reduce their advertising on drugs that have lost their patent in the face of the threat of generic entry. Broadly speaking, an incumbent is likely to have superior information about the demand, and any entry-deterrence strategy it engages in affects the information available to the potential entrant.

4.1 Equilibrium Properties

The following parametric assumptions are sufficient to guarantee that the baseline assumptions Condition 2 hold in the dynamic limit pricing game.

Assumption 1. The following conditions hold:

(i)
$$a_H^* < \bar{a} < (H - L)/(bH) + H/L \cdot a_L^*$$
,

(ii)
$$(b(H - L) - bL + c)H > cL$$
.

Assumption 1 requires that the wedge between the demand functions be sufficiently large; it fails, for example, when there is little demand uncertainty, that is, when the difference in demand intercepts is small. Assumption 1 does not rule out the existence of a range of prices that both types can induce. Crucially, it implies Condition 2.1 and Condition 2.2(b), which in turn imply Condition 1 and Condition 3.

Proposition 3. Assume Assumption 1 holds. In any Markov perfect equilibrium of the limit pricing game,

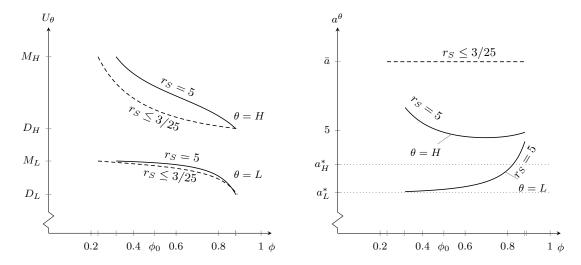


Figure 2: Left: Incumbent's value functions. Right: Incumbent's equilibrium actions. $(H, L, b, c, F, o, \sigma, r_{DM}, \bar{a}) = (100, 70, 1/20, 10, 100, 20, 4, 3/2, 7).$

- (i) both types of the incumbent engage in limit pricing, that is, they produce a quantity higher than myopically optimal;
- (ii) in any equilibrium, for any belief the expected price is always higher when the demand is strong, i.e., $H(1 b \cdot a_H(\phi)) \ge L(1 b \cdot a_L(\phi))$, for any $\phi \in [\phi, \overline{\phi}]$.

Figure 2 illustrates the equilibrium value functions and the quantities for different discount rates. First, as proved in Proposition 1, the value functions are decreasing. Second, the incumbent has incentives to overproduce to induce a lower price and deter entry. As explained in Section 3.2, when best replying, the incumbent takes the potential entrant's conjecture and, hence, the speed of learning, as given: the higher the speed of learning and the steeper the value function, the stronger the incentives to overproduce.

As stated in (ii), in equilibrium, the expected price is always higher when the demand is strong. To put it differently, in equilibrium, γ is always positive and both types have an incentive to overproduce to put pressure on the price.

The expected price path can never cross in equilibrium: even if the firm facing a strong demand wants to overproduce so to induce a price as low as the firm facing a weak demand, by assumption the weak firm has the ability to push the price to levels unattainable to the strong firm. In the limit as the two firms become arbitrarily patient, this logic "unravels" and both firms produce at capacity.

4.2 Welfare

One can decompose the welfare effect of limit pricing into three components. First, the welfare gains or losses before the potential entrant take the irreversible action; with regard to consumer surplus, aggressive limit pricing is always beneficial. Second, the welfare gains or losses once the potential entrant takes the irreversible action; with regard to consumer surplus, the lower the probability of entry, the larger the loss. Third, conditional on the potential entrant taking the decision that maximizes welfare, the delay entails a cost.

The welfare analysis of limit pricing has been vastly influenced by Milgrom and Roberts (1982) and has mostly focused on the first component, namely, the gains from price cuts. For example, the estimates of the welfare effects of limit pricing in Sweeting et al. (2020) rely on the fact that, if one focuses on the separating equilibrium of Milgrom and Roberts (1982)—as commonly done in empirical applications—entry decisions would be the same under either complete or asymmetric information.

In Matthews and Mirman (1983) and in the pooling equilibrium of Milgrom and Roberts (1982) the welfare effects can be negative because entry may be successfully threatened. However, unless it reduces the probability of entry, limit pricing is always beneficial. This section uncovers a novel welfare trade-off that emerges once the intensive time-margin delay is taken into account.

Compared to Milgrom and Roberts (1982) and Matthews and Mirman (1983)'s two-period models, in our model the potential entrant faces a trade-off between a more informed decision (that is, waiting and gathering more data about the market conditions before acting) and discounting (that is, foregoing profits by not acting early on). Further, the incumbent pricing strategy affects the precision of the information that the potential entrant has access to before taking its decision.

We divide our analysis of the welfare trade-offs in two sections. First, we present analytical results regarding the probability of entry and the delay. Second, we illustrate numerically when limit pricing can hurt consumers.

4.2.1 Deterrence Probability and Entry Delay

As mentioned in the introduction, one can interpret \bar{a} as an output restriction rule, in the spirit of the rule proposed by Williamson (1977) and Edlin (2002) and recently analyzed by Rey et al. (Forthcoming), to mitigate predatory behavior. While restrict-

ing output reduces the consumer surplus by curtailing the price cuts, the effects on the probability of entry and the delay are a priori unclear.

To understand how the informativeness of the price ultimately affects welfare, we first consider the case in which the potential entrant only decides when to enter, if ever, or equivalently, o = 0, and focus on the case of a patient incumbent. In this case, in the unique equilibrium, both types of incumbent produce at capacity and the potential entrant always enters when the demand is strong. As a result, the expected discounted consumer surplus is a function of two statistics: the unconditional probability of entry, $\Pr[b_{\tau} = h]$ and the unconditional expected discounted entry time, a measure of the cost of delay. Specifically, define the cost of delay as $1 - \mathbf{E}[e^{-r_{DM}\tau}]$.

Proposition 4. Suppose the regulator can choose $\bar{a} \in \bar{A}$, for some closed $\bar{A} \subset (a_H^*, (H-L)/(bH) + H/L \cdot a_L^*))$ and the incumbent is patient enough.

- (i) Assume $\phi_0 = 1/2$, $D_H + D_L < 0$. The output restriction rule \bar{a} that minimizes the cost of delay is min \bar{A} .
- (ii) The output restriction rule \bar{a} that maximizes the unconditional probability of entry is max \bar{A} .

The proposition sheds light on the trade-offs faced by a regulator designing an output restriction rule. On the one hand, when delay is likely to harm consumers, either because of impatience or via (unmodeled) lower product quality and more limited choice offered in monopoly, a conservative output restriction rule is beneficial. Because the price is more informative, entry occurs sooner. On the other hand, a patient regulator may be willing to tolerate the delay caused by a more lax output restriction rule to increase the probability of entry, inducing an equilibrium in which the price is a less informative signal. While both the expected discounted consumer surplus and the expected discounted welfare in equilibrium are only a function of these statistics, the exact optimal policy depends on parameters, as it must weigh the benefit of the price cuts and the cost of delay.

These results hold exactly only for r_S low enough, but we believe that the insights are robust to perturbations in the discount rate of the incumbent. While we do not

 $^{^{22}\}mathrm{Our}$ equilibrium characterization generalizes to the case of one-sided action. Because the right-hand side of (4) diverges in the limit $\phi \to 0$, following Faingold and Sannikov (2011) we construct the solution as the limit of a sequence of solutions on expanding closed subintervals, as we show in the Online Appendix. An alternative approach would have been using the results in Durandard and Strulovici (2022)

prove continuity of the set of equilibria with respect to r_S , upper-hemicontinuity is expected. Moreover, the trade-off between informativeness of the signal and delay would be present also in a model in which the incumbent engages in entry deterrence via product proliferation, advertising, or some other means that reduces the amount of information collected by the potential entrant.

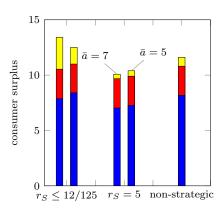
While Proposition 4 lays bare a novel trade-off, a complete analysis of the welfare effect of limit pricing must take into account also the short-term gain coming from lower prices. Furthermore, while analytically tractable, the one-sided model has some undesirable built-in asymmetries. In the next section, we derive some insights into the welfare effect of limit pricing by simulating equilibria of the two-sided model.

4.2.2 Numerical Simulations

Figure 3 plots the consumer welfare, expected discount factor, and the ex ante probability of entry for different parameters of the model. In the figure, the consumer welfare is decomposed into three components: in blue and red, the expected discounted consumer welfare after the potential entrant takes its irreversible action, conditional on the demand being strong or weak, respectively; in yellow, the expected discounted consumer welfare before the potential entrant acts. On all panels, we plot on the left the case when the incumbent is sufficiently patient, so that the unique equilibrium involves production at capacity for both types; in the center, the case of an intermediate discount rate; on the right, the case when the incumbent does not engage in limit pricing.

When the discount rate of the incumbent is intermediate, the equilibrium involves interior actions, and limit pricing hurts consumer surplus. Perhaps surprisingly, the welfare is larger when the capacity constraint is lower. As in Proposition 4, (i), when the capacity constraint is lower, in equilibrium, the price is more informative. However, in contrast to Proposition 4, a more informative price results in a higher probability of entry (because the problem is two-sided), albeit a larger cost of delay, as measured by the potential entrant's expected discount factor.

The effect of the capacity constraint is reversed if we look at the case where the incumbent is sufficiently patient. In both cases, in the unique equilibrium, both types of incumbent produce the maximum feasible quantity and the consumers gain from limit pricing. Because the gains are mostly coming from lower pricing, the higher the maximum feasible quantity, the larger the gains in consumer surplus.



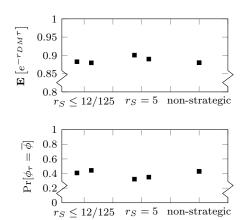


Figure 3: Consumer surplus discounted at rate 3/2, expected discount factor, and probability of entry. In yellow, the expected discounted consumer surplus between $[0,\tau)$. In blue (red) the expected discounted (lump-sum) consumer surplus at τ when $\theta = H$ ($\theta = L$). $(H, L, b, c, F, o, \sigma, r_{DM}) = (100, 70, 1/20, 10, 100, 20, 4, 3/2)$.

The conclusion we draw in this section resonates with the empirical evidence. For example, in the 1970s, Folger (owned by Procter & Gamble) delayed by several years entry in the Eastern United States regular coffee market because General Foods reacted to this threat by sharply reducing the price of its Maxwell House. ^{23,24} Procter & Gamble had a practice of conducting market tests before undertaking large-scale entry but Maxwell's aggressive pricing behavior affected the informativeness of early tests. While the FTC dismissed the complaint, ²⁵ the case is frequently cited as an example of test market predation (see, Viscusi et al., 2018, Ch. 8).

In line with the insights provided by existing models, the FTC based its decision in that case on the finding that the alleged predator did not have a dangerous probability of success (see, Hilke and Nelson, 1989). We believe that our model provides a new theoretical perspective for reexamining the court's approaches toward predatory pricing and attempted monopolization.

²³See Hilke and Nelson (1989) and Ross and Scherer (1990).

²⁴Traditional industrial organization distinguishes between limit pricing, that is, pricing strategy to discourage entry, and predatory pricing, that is pricing strategy to encourage exit. However, our limit pricing model resonates with the theories of test market predation, according to which an incumbent pricing strategy affects the information acquired by a competitor and deters expansion.

²⁵ General Foods Corporation, 103 F.T.C. 204.

²⁶See Footnote 18.

4.3 Cost Uncertainty vs Demand Uncertainty

In the previous section, we studied a dynamic version of the model proposed by Matthews and Mirman (1983), who assume that the incumbent has private information about the market demand. However, Milgrom and Roberts (1982) workhorse model of limit pricing assumes that the incumbent's private information is about its marginal cost. In this section, we compare our example to the case in which the two types of incumbent differ only in their marginal cost.

We depart minimally from the framework we introduced before. The sender is an incumbent firm that chooses an output level, which results in a stochastic inverse demand as in (5). In contrast to the previous section, the two types of incumbent face the same demand function, and we normalize the shifter θ to 1. However, the two types differ in their marginal cost; the cost of producing quantity a is $c_{\theta}a^{2}/2$, where $c_{H} > c_{L} > 0$.

The assumption $\mu(a, H) \neq \mu(a, L)$ for all $a \in A$ fails in this case but we show in the online Appendix that Condition 1 and Condition 2' hold, so Theorem 1 generalizes to this case, guaranteeing equilibrium existence.

The key difference between the two models becomes apparent when we examine the limit as the two players become arbitrarily patient. In the baseline model, namely when the incumbent has private information about the demand, we derived a lower bound on the speed of learning which holds for any discount rate. In contrast, when the incumbent has private information about the marginal cost, but the drift functions are the same for both types, in the patience limit, the volatility of beliefs must vanish along any sequence of equilibria, as shown in the following proposition.

Proposition 5.

Let $\{r_{S,n}\}_{n\geq 1}$ and $\{r_{DM,n}\}_{n\geq 1}$ be two sequences converging to zero such that $\lim_{n\to\infty} r_{S,n}/r_{DM,n} = k \in (0,\infty)$. Then along any sequence of Markov equilibria, volatility of belief converges to zero and the potential entrant's value function converges to the no-information value function $\max\{o, \phi D_H + (1-\phi)D_L - F\}$.

One immediate implication of Proposition 5 is that a potential entrant facing an arbitrarily patient incumbent acts without delay, as the equilibrium price does not reveal any information about the incumbent's type. In fact, in sharp contrast with Observation 1, when the two types differ only in their production cost, there is no ambiguity about the informativeness of the price when the incumbent engages in

aggressive limit pricing: when the two types choose the same action, and in particular, the highest possible output, the potential entrant does not learn anything from observing the price.

By Corollary 1, the joint limit in Proposition 5 is equivalent to the limit as $\sigma \to 0$ while keeping the discount rates fixed, allowing us to derive a few simple comparative statics, which were elusive in the case in which the incumbent's private information was about the state of the demand. Specifically, the expected time before the potential entrant acts is a non-monotone function of σ .²⁷ In fact, both when the noise vanishes, $\sigma \to 0$, and when the price becomes arbitrarily noisy, $\sigma \to \infty$, the potential entrant acts with no delay. The result resonates with Gonçalves (2024) who proves that in a standard Wald's problem the expected stopping time is inverse U-shaped in the signal-to-noise ratio.²⁸ The key difference between our result and his is that in the limit as the noise vanishes, in the strategic setting, the potential entrant's value function does not converge to the full-information value function: decisions are faster not because the signal is very informative but, to the contrary, because waiting has no value.

Interpreting Proposition 5 as the limit as $\sigma \to 0$, the game approaches the perfect monitoring framework of Milgrom and Roberts (1982), where the entrant learns the incumbent's cost perfectly in the separating equilibrium. From a game-theoretic perspective, the failure of upper hemicontinuity of the equilibrium correspondence as we approach perfect monitoring is not surprising. However, from a policy standpoint, the result highlights the fragility of conclusions drawn from the complete-information game—the idea that limit pricing is beneficial to consumers.

On the one hand, the trade-off between delay and lower prices discussed in Section 4.2 disappears when the incumbent has private information about the marginal cost instead of the state of the demand. On the other hand, because the result in Proposition 5 is independent of the magnitude of the capacity constraint \bar{a} , the output restriction rules we discussed in Section 4.2.1 have no effect on the informativeness of the price or on the equilibrium delay, and become completely ineffective policy instruments.

²⁷In the general model, we conjecture that the expected time before the DM takes an action is monotone in σ ; in fact, Theorem 2, Part 3, it diverges to infinity as $\sigma \to 0$, while clearly it converges to zero as $\sigma \to \infty$.

 $^{^{28} \}mbox{Bobtcheff}$ and Levy (2017) and Cetemen and Margaria (2024) prove a similar result in a Poisson learning setup.

In this case, a policymaker could increase the probability of entry by introducing entry subsidies, paid to the potential entrant when it decides to become a competitor of the incumbent, effectively lowering the entry cost.²⁹ Moreover, the effect of such a policy on the probability of entry is ambiguous when the incumbent's private information concerns the state of the demand because the informativeness of the signal is determined in equilibrium, and a monotonic change in the best reply of the potential entrant does not necessarily lead to a monotonic change in the equilibrium outcome. However, when the incumbent has private information only about its marginal cost, there is no information revealed so the dynamic information channel is muted, and the potential entrant only acts based on its prior.

5 Conclusion

While our analysis is mainly theoretical, our two applications illustrate how our tractable characterization can be leveraged to conduct comparative static exercises and inform policy interventions.

In the paper, we focused on a canonical setup, but the framework can be easily extended in a few directions. Generalizing the result to allow for multidimensional Gaussian signals or for additional conclusive Poisson signals is straightforward.

We assumed that the DM collects payoffs only upon taking the irreversible action. The extension to the case in which the DM also collects payoffs before does not add conceptual difficulties but is not immediate especially when the payoffs depend on the state and the action of the sender. In the context of dynamic competition, this extension would allow us to talk about predation, that is, about a dominant firm trying to induce the exit of a competitor.

A more sophisticated predation model, however, would allow for investment in capacity expansion. For example, airlines can invest in their fleet to increase their capacity, or workers can invest in education to improve their skills, leading to endogenously evolving types. This could be captured, for example, by allowing the sender to affect the evolution of the state, as in Board and Meyer-ter Vehn (2013). We are pursuing this in ongoing work.

²⁹Such subsidies are common for example in the airline industry, see for example Ryerson (2016) and Sweeting et al. (2020).

References

- Amster, Pablo (2007) "Nagumo and Landesman-lazer Type Conditions for Nonlinear Second Order Systems," Nonlinear Differential Equations and Applications NoDEA, 13, 699–711.
- Anderson, Axel and Lones Smith (2013) "Dynamic Deception," American Economic Review, 103, 2811–47.
- Atakan, Alp E. and Mehmet Ekmekci (2012) "Reputation in Long-run Relationships," The Review of Economic Studies, 79, 451–480.
- Atakan, Alp E. and Mehmet Ekmekci (2015) "Reputation in The Long-run with Imperfect Monitoring," *Journal of Economic Theory*, 157, 553–605.
- Barilla, César and Duarte Gonçalves (2024) "The Dynamics of Instability," *Theoretical Economics*, 19, 365–405.
- Board, Simon and Moritz Meyer-ter Vehn (2013) "Reputation for Quality," *Econometrica*, 81, 2381–2462.
- Bobtcheff, Catherine and Raphaël Levy (2017) "More Haste, Less Speed? Signaling Through Investment Timing," American Economic Journal: Microeconomics, 9, 148–86.
- Bohren, J. Aislinn (2024) "Persistence in a Dynamic Moral Hazard Game," *Theoretical Economics*, 19, 449–498.
- Bolton, Patrick and Christopher Harris (1999) "Strategic Experimentation," *Econometrica*, 67, 349–374.
- Bonatti, Alessandro, Gonzalo Cisternas, and Juuso Toikka (2017) "Dynamic Oligopoly with Incomplete Information," *The Review of Economic Studies*, 84, 503–546.
- Brady, Stephan P. and William A. Cunningham (2001) "Exploring Predatory Pricing in the Airline Industry," *Transportation Journal*, 41, 5–15.
- Brocas, Isabelle and Juan D. Carrillo (2007) "Influence Through Ignorance," *The RAND Journal of Economics*, 38, 931–947.
- Celetani, Marco, Drew Fudenberg, David K. Levine, and Wolfgang Pesendorfer (1996) "Maintaining a Reputation against a Long-lived Opponent," *Econometrica*, 64, 691–704.
- Cetemen, Doruk (2020) "Efficiency in Repeated Partnerships," Working Paper.

- Cetemen, Doruk and Chiara Margaria (2024) "Exit Dilemma: The Role of Private Learning on Firm Survival," *American Economic Journal: Microeconomics*, 16, 110–154.
- Cisternas, Gonzalo and Aaron Kolb (2024) "Signalling with Private Monitoring," The Review of Economic Studies, 92, 909–953.
- Cripps, Martin W., George J. Mailath, and Larry Samuelson (2004) "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72, 407–432.
- Cripps, Martin W. and Jonathan P. Thomas (1997) "Reputation and Perfection in Repeated Common Interest Games," Games and Economic Behavior, 18, 141–158.
- Daley, Brendan and Brett Green (2012) "Waiting for News in the Market for Lemons," *Econometrica*, 80, 1433–1504.
- Dillon, Karen (2011) "I Think of My Failures As A Gift," *Harvard Business Review*, 89, 86–89.
- Dilmé, Francesc (2019) "Dynamic Quality Signaling with Hidden Actions," Games and Economic Behavior, 113, 116–136.
- Dilmé, Francesc (2025) "Repeated Trade with Imperfect Information about Previous Transactions," *Theoretical Economics*, 20, 209–254.
- Durandard, Théo and Bruno Strulovici (2022) "Smoothness of Value Functions in General Control-Stopping Diffusion Problems," Working Paper.
- Edlin, Aaron S. (2002) "Stopping Above-cost Predatory Pricing," *The Yale Law Journal*, 111, 941–991.
- Ekmekci, Mehmet, Leandro Gorno, Lucas Maestri, Jian Sun, and Dong Wei (2022) "Learning From Manipulable Signals," *American Economic Review*, 112, 3995–4040.
- Ekmekci, Mehmet and Lucas Maestri (2022) "Wait Or Act Now? Learning Dynamics in Stopping Games," *Journal of Economic Theory*, 205, 105541.
- Ellison, Glenn and Sara Fisher Ellison (2011) "Strategic Entry Deterrence and the Behavior of Pharmaceutical Incumbents Prior to Patent Expiration," *American Economic Journal: Microeconomics*, 3, 1–36.
- Escudé, Matteo and Ludvig Sinander (2023) "Slow persuasion," Theoretical Economics, 18, 129–162.
- Faingold, Eduardo and Yuliy Sannikov (2011) "Reputation in Continuous-time Games," *Econometrica*, 79, 773–876.

- Fairburn, James A. and James M. Malcomson (2001) "Performance, Promotion, and the Peter Principle," *The Review of Economic Studies*, 68, 45–66.
- Georgiadis, George (2014) "Projects and Team Dynamics," The Review of Economic Studies, 82, 187–218.
- Gonçalves, Duarte (2024) "Speed, Accuracy, and Complexity," Working Paper.
- Goolsbee, Austan and Chad Syverson (2008) "How Do Incumbents Respond to the Threat of Entry? Evidence From the Major Airlines," *The Quarterly Journal of Economics*, 123, 1611–1633.
- Gryglewicz, Sebastian and Aaron Kolb (2022) "Dynamic Signaling with Stochastic Stakes," *Theoretical Economics*, 17, 539–559.
- Gryglewicz, Sebastian and Aaron Kolb (2025) "Strategic Pricing in Volatile Markets," Operations Research, 73, 444–460.
- Henry, Emeric and Marco Ottaviani (2019) "Research and The Approval Process: The Organization Of Persuasion," *American Economic Review*, 109, 911–55.
- Hilke, John C. and Philip B. Nelson (1989) "Strategic Behavior and Attempted Monopolization: The Coffee (general Foods) Case," in *The Antitrust Revolution* ed. by Jr. John E. Kwoka and Lawrence J. White. Glenview, Ill.: Scott, Foresman.
- Kamien, Morton I. and Nancy L. Schwartz (1971) "Limit Pricing and Uncertain Entry," *Econometrica*, 39, 441–454.
- Keller, Godfrey and Sven Rady (1999) "Optimal Experimentation in A Changing Environment," The Review of Economic Studies, 66, 475–507.
- Kolb, Aaron M. (2015) "Optimal Entry Timing," Journal of Economic Theory, 157, 973–1000.
- Kolb, Aaron M. (2019) "Strategic Real Options," Journal of Economic Theory, 183, 344–383.
- Kuvalekar, Aditya and Elliot Lipnowski (2020) "Job Insecurity," American Economic Journal: Microeconomics, 12, 188–229.
- Liptser, Robert S. and Albert N. Shiryaev (2001) Statistics of Random Processes: I. General Theory, Berlin, Heidelberg: Springer.
- Matthews, Steven A. and Leonard J. Mirman (1983) "Equilibrium Limit Pricing: The Effects of Private Information and Stochastic Demand," *Econometrica*, 51, 981–996.

- McClellan, Andrew (2022) "Experimentation and Approval Mechanisms," *Econometrica*, 90, 2215–2247.
- Milgrom, Paul and John Roberts (1982) "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis," *Econometrica*, 50, 443–459.
- Moscarini, Giuseppe and Lones Smith (2003) "The Optimal Level of Experimentation," *Econometrica*, 69, 1629–1644.
- Noldeke, Georg and Eric van Damme (1990) "Signalling in A Dynamic Labour Market," The Review of Economic Studies, 57, 1–23.
- Ok, Efe A. (2007) Real Analysis with Economic Applications, Princeton: Princeton University Press.
- Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2020) "Persuading the Principal to Wait," *Journal of Political Economy*, 128, 2542–2578.
- Pei, Harry (2021) "Trust and Betrayals: Reputational Payoffs and Behaviors Without Commitment," *Theoretical Economics*, 16, 449–475.
- Rey, Patrick, Yossi Spiegel, and Konrad Stahl (Forthcoming) "A Dynamic Model of Predation," The RAND Journal of Economics.
- Ross, David and Frederic M. Scherer (1990) Industrial Market Structure and Economic Performance, Boston: Houghton Mifflin, 3rd edition.
- Ryerson, Megan S. (2016) "Incentivize It and They Will Come? How Some of the Busiest U.S. Airports Are Building Air Service With Incentive Programs," *Journal of the American Planning Association*, 82, 303–315.
- Salas, Christian (2019) "Persuading Policy-makers," Journal of Theoretical Politics, 31, 507–542.
- Saloner, Garth (1984) "Dynamic Equilibrium Limit-pricing in An Uncertain Environment," Working Paper.
- Smiley, Robert (1988) "Empirical Evidence on Strategic Entry Deterrence," *International Journal of Industrial Organization*, 6, 167–180.
- Spence, Michael (1973) "Job Market Signaling," The Quarterly Journal of Economics, 87, 355–374.
- Stokey, Nancy L. (2009) The Economics of Inaction: Stochastic Control Models with Fixed Costs: Princeton University Press.

Sweeting, Andrew, James W. Roberts, and Chris Gedge (2020) "A Model of Dynamic Limit Pricing with An Application to the Airline Industry," *Journal of Political Economy*, 128, 1148–1193.

Toxvaerd, Flavio (2017) "Dynamic Limit Pricing," *The RAND Journal of Economics*, 48, 281–306.

Viscusi, W. Kip, Joseph E. Harrington, and David E. M. Sappington (2018) *Economics of Regulation and Antitrust*, Cambridge, MA: MIT Press, 5th edition.

Wald, Abraham (1945) "Sequential Tests of Statistical Hypotheses," The Annals of Mathematical Statistics, 16, 117–186.

Williamson, Oliver E. (1977) "Predatory Pricing: A Strategic and Welfare Analysis," *The Yale Law Journal*, 87, 284–340.

6 Appendix

6.1 Proof of Theorem 1

The proof of Theorem 1 is organized as follows: first, we prove the statement under Condition 1 and Condition 2'. In the Online Appendix, we prove that Condition 2 implies Condition 2'.

6.1.1 Bounds on Coefficient γ

This section proves a bound on γ which is used in the subsequent analysis.

Lemma 1. Assume Condition 1 and Condition 2'. There exists a C > 0 such that for all $(a_H, a_L, \phi, z_H, z_L) \in A \times A \times (0, 1) \times \mathbf{R} \times \mathbf{R}$, if $(a_H, a_L, \phi, z_H, z_L)$, then

$$(1+|z_H|+|z_L|)\frac{|\gamma(a_H,a_L,\phi)|}{\phi(1-\phi)} \ge C.$$

Proof. Suppose that such a constant does not exist. Then, there exists a sequence $\{(a_{H,n}, a_{L,n}, \phi_n, z_{L,n}, z_{H,n})\}_{n\geq 1}$ with $\phi_n \in (0,1)$ and $(a_{H,n}, a_{L,n}) \in \mathcal{N}(z_{H,n}, z_{L,n})$, such that for both $\theta = H$ and $\theta = L$ the following hold:

$$|z_{\theta,n}| \frac{|\gamma(a_{H,n}, a_{L,n}, \phi_n)|}{\phi_n(1 - \phi_n)} \to 0, \quad \text{and} \quad \frac{|\gamma(a_{H,n}, a_{L,n}, \phi_n)|}{\phi_n(1 - \phi_n)} \to 0.$$
 (6)

By compactness, there exists a sub-sequence converging to some $(a_H, a_L, \phi, z_L, z_H) \in A \times A \times [0, 1] \times \mathbf{R} \times \mathbf{R}$. By continuity, (a_H, a_L, ϕ) must be a Bayes Nash equilibrium of the auxiliary signaling game with prior ϕ , i.e., $(a_H, a_L) \in \mathcal{N}(z_L, z_H)$.

Hence, the first limit in (6) implies that $a_{H,n} \to a_H^*$ and $a_{L,n} \to a_L^*$. Let $\varepsilon = |\mu(H, a_H^*) - \mu(L, a_L^*)|$. By Condition 2' and the continuity of μ , for any n sufficiently high, $|\mu(H, a_{H,n}) - \mu(L, a_{L,n})| \ge \varepsilon/2 > 0$, contradicting the second limit in (6).

The following corollary implies that our system satisfies a quadratic growth condition for each type of the sender.

Corollary 3. Assume Condition 1 and Condition 2'. For all $\varepsilon > 0$, there exist a K > 0 such that for all $\phi \in [\varepsilon, 1-\varepsilon]$, $(U_H, U_L) \in [\Pi(H, l), \Pi(H, h)] \times [\Pi(L, l), \Pi(L, h)]$, and $(U'_H, U'_L) \in \mathbf{R}_-$, if $(a_H, a_L) \in \mathcal{N}(\phi(1-\phi)U'_H(\phi)/r_S, \phi(1-\phi)U'_L(\phi)/r_S)$ then

$$\left| -2\frac{U'_{H}(\phi)}{\phi} + \frac{2r_{S}\left(U_{H}(\phi) - \pi(H, a_{H}(\phi))\right)}{\left(\gamma\left(a_{H}(\phi), a_{L}(\phi), \phi\right)\right)^{2}} \right| \leq K\left(1 + \left(U'_{H}(\phi)\right)^{2} + \left(U'_{L}(\phi)\right)^{2}\right),$$

$$\left| 2\frac{U'_{L}(\phi)}{1 - \phi} + \frac{2r_{S}\left(U_{L}(\phi) - \pi(L, a_{L}(\phi))\right)}{\left(\gamma\left(a_{H}(\phi), a_{L}(\phi), \phi\right)\right)^{2}} \right| \leq K\left(1 + \left(U'_{H}(\phi)\right)^{2} + \left(U'_{L}(\phi)\right)^{2}\right).$$

The proof follows directly from the bounds derived in Lemma 1.

6.1.2 Sender's Best-Reply Problem

The following proposition characterizes of the sender's value functions in any Markov equilibrium. That is, it characterizes the sender's pseudo-best reply (i.e., a mapping from the DM's strategy to the action profile for the two types of sender), as explained in Section 3.1.

Proposition 6. Assume Condition 1 and Condition 2' are satisfied. If the DM's strategy $\bar{\phi}$ and $\underline{\phi}$ and the sender's strategy profile $(a_H(\phi), a_L(\phi))$ are part of a Markov equilibrium, then the value functions of the sender solve the following system of second-order ordinary differential equations over the interval $(\phi, \bar{\phi})$,

$$U_{H}''(\phi) = -2\frac{U_{H}'(\phi)}{\phi} + \frac{2r_{S}\left(U_{H}(\phi) - \pi(H, a_{H}(\phi))\right)}{\left(\gamma\left(a_{H}(\phi), a_{L}(\phi), \phi\right)\right)^{2}},$$

$$U_{L}''(\phi) = 2\frac{U_{L}'(\phi)}{1 - \phi} + \frac{2r_{S}\left(U_{L}(\phi) - \pi(L, a_{L}(\phi))\right)}{\left(\gamma\left(a_{H}(\phi), a_{L}(\phi), \phi\right)\right)^{2}},$$
(7)

subject to the boundary conditions $U_{\theta}(\bar{\phi}) = \Pi(\theta, H)$, $U_{\theta}(\underline{\phi}) = \Pi(\theta, L)$, and $(a_H(\phi), a_L(\phi)) = \mathcal{N}(\phi(1-\phi)U'_H(\phi)/r_S, \phi(1-\phi)U'_L(\phi)/r_S)$ for any $\phi \in [\underline{\phi}, \bar{\phi}]$. Moreover, for any $0 < \underline{\phi} < \bar{\phi} < 1$, the boundary value problem has a solution.

Proof. The proof relies on a modification of Theorem 2.1 of Amster (2007), which we state and prove in the Online Appendix. To apply Theorem OA.1, define the constant functions $\alpha : [0, 1] \to \mathbf{R}^2$, $\alpha \equiv (\underline{\pi}, \underline{\pi})$, and $\beta : [0, 1] \to \mathbf{R}^2$, $\beta \equiv (\pi^*(H), \pi^*(L))$.

Let $f:[0,1]\times\mathbf{R}^4\to\mathbf{R}^2$ be defined as

$$f_{1}(\phi, U, U') = -2\frac{U'_{1}}{\phi} + 2r_{S}\frac{U_{1} - \pi (H, \operatorname{proj}_{1} \mathcal{N} (\phi(1-\phi)U'_{1}/r_{S}, \phi(1-\phi)U'_{2}/r_{S}))}{(\gamma (\mathcal{N} (\phi(1-\phi)U'_{1}/r_{S}, \phi(1-\phi)U'_{2}/r_{S}), \phi))^{2}},$$

$$f_{2}(\phi, U, U') = 2\frac{U'_{2}}{1-\phi} + 2r_{S}\frac{U_{2} - \pi (L, \operatorname{proj}_{2} \mathcal{N} (\phi(1-\phi)U'_{1}/r_{S}, \phi(1-\phi)U'_{2}/r_{S}))}{(\gamma (\mathcal{N} (\phi(1-\phi)U'_{1}/r_{S}, \phi(1-\phi)U'_{2}/r_{S}), \phi))^{2}}.$$

Consider the following boundary value problem

$$U''(\phi) = f(\phi, U(\phi), U'(\phi)), \qquad \phi \in [\underline{\phi}, \overline{\phi}],$$

$$U(\underline{\phi}) = (\Pi(H, l), \Pi(L, l)), \qquad U(\overline{\phi}) = (\Pi(H, h), \Pi(L, h)).$$

It can be verified that α and β are a lower and an upper solution of the boundary value problem, respectively. By Corollary 3, there exist a constant K > 0 such that

$$|f_i(\phi, U_i(\phi), U_i'(\phi))| \le K(1 + (|U_1'(\phi)| + |U_2'(\phi)|)^2)$$

for i = 1, 2 and for all $(\phi, U, U') \in [\underline{\phi}, \overline{\phi}] \times [\alpha_1, \beta_1] \times [\alpha_2, \beta_2] \times \mathbf{R}^2$. Let $\psi(s) = K(1+s^2)$. Hence, one can choose M

$$\int_{r}^{M} \frac{1}{K} \frac{s}{1+s^2} \, \mathrm{d}s > 4 \max\{\alpha_H, \alpha_L\},\,$$

where $r := (\Pi(H, l) - \Pi(H, h) + \Pi(L, l) - \Pi(L, h)) / (\overline{\phi} - \underline{\phi})$, so to satisfy the Nagumotype condition in Theorem OA.1. Further, it is readily verified that condition (OA.4) in Theorem OA.1 holds with $U_i = \underline{U}$ for i = 1, 2, see (4). We can conclude that the boundary value problem has a bounded C^2 solution in the domain $[\underline{\phi}, \overline{\phi}]$ taking values in $[\underline{\pi}, \Pi(H, l)) \times [\underline{\pi}, \Pi(L, l))$.

6.1.3 Proof of Theorem 1

Lemma OA.3, relegated to the online appendix, proves the continuity of the solutions to the boundary value problem of the sender with respect to the boundary conditions.

For any ordered pair $(\underline{\phi}, \overline{\phi}) \in [0, 1]^2$, define the sender's best-reply mapping BR^S : $(\underline{\phi}, \overline{\phi}) \mapsto (a_H(\phi), a_L(\phi)) \in (A \times A)^{[0,1]}$ by pasting together a solution to the system of ordinary differential equations in Proposition 6, which specifies the value of the function for $\phi \in [\underline{\phi}, \overline{\phi}]$, and the constant functions (a_H^*, a_L^*) . By Lemma OA.3, we can assume that BR^S is continuous.

Similarly, for any regular strategy profile $(a_H, a_L) : [0, 1] \to A^2$ define the DM's best-reply mapping BR^{DM} : $(A \times A)^{[0,1]} \mapsto [0, 1]^2$ as the unique pair of cutoffs characterizing the DM's best reply.

The best-reply problem of the DM is standard and its characterization is relegated to the Online Appendix (see Section OA.4.2). The best reply of the DM is continuous: by the characterization in the proof of Proposition OA.2, and since the fundamental solutions to the ODE (OA.7) are continuous in the action profile of the sender, both the value function and the optimal cutoffs of the DM are continuous in the action profile of the sender.

Define the mapping $\Gamma:[0,1]^2\to [0,1]^2$ by combining the two best replies, that is, $\Gamma:(\underline{\phi},\overline{\phi})\mapsto \mathrm{BR}^{DM}(\mathrm{BR}^S(\underline{\phi},\overline{\phi}))$. Since the composition of the continuous functions, Γ is continuous. Therefore, by Brouwer's fixed-point theorem, Γ has a fixed point. By construction, any fixed point is a Markov Perfect equilibrium.

6.1.4 Proof of Proposition 1

Follows directly from the proof of Proposition 6, as Theorem OA.1 proves the existence of a monotone bounded solution to the boundary value problem.

6.1.5 Proof of Theorem 2

We start with two technical lemmas which are used later.

Lemma 2. Consider a sequence of $\{(r_{S,n}, r_{DM,n})\}_{n\geq 1}$, where each $r_{S,n}$ and $r_{DM,n}$ is strictly positive, together with an associated sequence of equilibria and corresponding value functions $\{(U_{H,n}, U_{L,n})\}_{n\geq 1}$ and cutoffs $\{(\underline{\phi}_n, \overline{\phi}_n)\}_{n\geq 1}$. Set $\mathcal{S} = \{\{\phi_n\}_{n\geq 1} : \phi_n \in [\underline{\phi}_n, \overline{\phi}_n]\}$. That is, \mathcal{S} is the set of sequences of beliefs such that for each n, the belief ϕ_n belongs

to the interval $[\underline{\phi}_n, \overline{\phi}_n]$. Then,

$$\Delta \mu := \inf_{\mathcal{S}} \liminf_{n \to \infty} \frac{\left| \gamma \left(a_{H,n} \left(\phi_n \right), a_{L,n} \left(\phi_n \right), \phi_n \right) \right|}{\phi_n \left(1 - \phi_n \right)}$$

is such that $\Delta \mu > 0$.

Proof. See Online Appendix.

Lemma 3. Suppose that for any ϕ , the speed of learning is equal to $\phi(1-\phi)\Delta\mu$. Then, denoting with τ the time when the DM acts, the following holds:

$$\mathbf{E}[\tau] = -\frac{4(1 - 2\phi_0) \tanh^{-1}(1 - 2\phi_0)}{(\Delta\mu)^2} + \frac{4(1 - 2\underline{\phi})(\overline{\phi} - \phi_0) \tanh^{-1}(1 - 2\underline{\phi})}{(\overline{\phi} - \underline{\phi})(\Delta\mu)^2} + \frac{4(1 - 2\overline{\phi})(\phi_0 - \underline{\phi}) \tanh^{-1}(1 - 2\overline{\phi})}{(\overline{\phi} - \phi)(\Delta\mu)^2}.$$

In addition, for any r > 0, as $(\phi, \overline{\phi}) \to (0, 1)$, $\mathbf{E}[e^{-r\tau}] \to 0$.

Proof. See Online Appendix.

Proof of Part 1: $r_{DM} > 0$, $r_S \to 0$. Consider a sequence $\{r_{S,n}\}_{n\geq 1}$ such that $\lim_{n\to\infty} r_{S,n} = 0$ together with a sequence of equilibria and corresponding value functions $\{(U_{H,n}, U_{L,n})\}_{n\geq 1}$ and cutoffs $\{(\underline{\phi}_n, \overline{\phi}_n)\}_{n\geq 1}$. We claim that $U_{\theta,n}(\phi)/r_S \to -\infty$ for $\theta \in \{H, L\}$ and for all $\phi \in [\underline{\phi}_n, \overline{\phi}_n]$. Suppose by contradiction that this is not the case. Denote with $(\underline{\phi}, \overline{\phi})$ the limit cutoffs of the DM. Since $r_{DM} > 0$, $0 < \underline{\phi} < \overline{\phi} < 1$. Using an argument similar to Lemma OA.3, one can show that $\{(U_{H,n}, U_{L,n})\}_{n\geq 1}$ must pointwise converge to $(\mathring{U}_1(\phi), \mathring{U}_2(\phi))$, where

$$\begin{split} \mathring{U}_1(\phi) &= \frac{(\Pi(H,l) - \Pi(H,h))\underline{\phi}\overline{\phi}}{\phi(\overline{\phi} - \underline{\phi})} + \frac{\overline{\phi}\Pi(H,h) - \underline{\phi}\Pi(H,l)}{\overline{\phi} - \underline{\phi}} \\ \mathring{U}_2(\phi) &= -\frac{(\Pi(L,l) - \Pi(L,h))(1 - \underline{\phi})(1 - \overline{\phi})}{(1 - \phi)(\overline{\phi} - \underline{\phi})} + \frac{(1 - \underline{\phi})\Pi(L,l) - (1 - \overline{\phi})\Pi(L,h)}{\overline{\phi} - \underline{\phi}}, \end{split}$$

is the unique solution to the boundary value problem

$$U_1''(\phi) = -2U_1'(\phi)/\phi, \qquad U_2''(\phi) = 2U_2'(\phi)/(1-\phi),$$

under the boundary conditions

$$U(\underline{\phi}) = \left(\begin{array}{c} \Pi(H,l) \\ \Pi(L,l) \end{array} \right), \qquad U(\overline{\phi}) = \left(\begin{array}{c} \Pi(H,h) \\ \Pi(L,h) \end{array} \right).$$

But this implies $\lim_{n\to\infty} U'_{H,n}(\phi) < 0$ and $\lim_{n\to\infty} U'_{L,n}(\phi) < 0$, contradicting $\lim_{n\to\infty} U'_{\theta,n}(\phi)/r_S \neq -\infty$ for at least one θ . Hence, for any sequence of discount rate $\{r_{S,n}\}_{n\geq 1}$ converging to zero, the sequence of corresponding value functions $\{(U_{H,n}, U_{L,n})\}_{n\geq 1}$ is such that $\lim_{n\to\infty} U_{H,n}(\phi)/r_S = \lim_{n\to\infty} U_{H,n}(\phi)/r_S = -\infty$, which in turns implies that for r_S low enough the pair of action solving (3) is not interior, for any belief. That is, $a_H(\phi) \to \bar{a}$ and $a_L(\phi) \to \bar{a}$ or $a_H(\phi) \to \bar{a}$ and $a_L(\phi) \to \bar{a}$ for any $\phi \in [\underline{\phi}, \overline{\phi}]$. Whether one or the other case occurs depends on the sign of $\mu(H,a) - \mu(L,a)$, which, by assumption, is independent of a. Part 2 of Condition 2 implies that $\mu(H,a) - \mu(L,a) > 0$, so $a_{\theta}(\phi) \to \bar{a}$, but under Condition 2' either cases can occur.

Proof of Part 2: $r_S > 0$, $r_{DM} \to 0$. Consider a sequence $\{r_{DM,n}\}_{n\geq 1}$ such that $\lim_{n\to\infty} r_{DM,n} = 0$ together with a sequence of equilibria and corresponding value functions $\{(U_{H,n}, U_{L,n})\}_{n\geq 1}$ and cutoffs $\{(\bar{\phi}_n, \underline{\phi}_n)\}_{n\geq 1}$.

First, we show that $\{(\bar{\phi}_n, \underline{\phi}_n)\}_{n\geq 1} \to (0,1)$. By Lemma 2, there exists a strictly positive lower bound to the difference in conditional drifts in the sequence of games. With abuse of notation, let $\Delta\mu$ denote this lower bound when one considers the sequence of discount rates $\{(r_S, r_{DM,n})\}_{n\geq 1}$. Consider the sequence of optimal stopping problems for the DM associated to the discount rate $\{r_{DM,n}\}_{n\geq 1}$ in which the speed of learning is, for any ϕ , $\phi(1-\phi)\Delta\mu$. It follows an adaptation of Moscarini and Smith (2003)'s argument (see their proof of Proposition 5(e)), that the cutoff shift out strictly as the discount rate increases. As a result, in the limit, they converge to 0 and 1, respectively. A fortiori, $(\bar{\phi}_n, \underline{\phi}_n) \to (0, 1)$.

Second, we show that the $\lim_{n\to\infty} U_{\theta,n}(\phi) = \pi^*(\theta)$, for all $\phi \in (0,1)$. For any n, the payoff of type θ of the sender can be written as,³⁰

$$U_{\theta,n}(\phi) = \mathbf{E}_{\theta}[e^{-r_{S}\tau} \mid \phi_{\tau} = \underline{\phi}_{n}, \phi_{0} = \phi] \operatorname{Pr}_{\theta}[\phi_{\tau} = \underline{\phi}_{n} \mid \phi_{0} = \phi] \Pi(\theta, l)$$

$$+ \mathbf{E}_{\theta}[e^{-r_{S}\tau} \mid \phi_{\tau} = \overline{\phi}_{n}, \phi_{0} = \phi] \operatorname{Pr}_{\theta}[\phi_{\tau} = \overline{\phi}_{n} \mid \phi_{0} = \phi] \Pi(\theta, h)$$

$$+ \mathbf{E}_{\theta} \left[\int_{0}^{\tau} r_{S} e^{-r_{S}t} \pi(\theta, a(\phi_{t})) dt \right]. \quad (8)$$

By Lemma 3 and Lemma 2, the first two terms vanish. As a result, the limit value functions reduce to the limit of the last term in (8). By Theorem 3.7 of Stokey (2009), we have

$$\mathbf{E}_{\theta} \left[\int_{0}^{\tau} r_{S} e^{-r_{S} t} \pi(\theta, a(\phi_{t})) \, \mathrm{d}t \mid \phi_{0} = \hat{\phi} \right] = \mathbf{E} \left[\int_{\underline{\phi}}^{\overline{\phi}} r_{S} \pi(\theta, a(\phi)) \hat{\ell}(\phi; \hat{\phi}, \tau; r_{S}) \, \mathrm{d}\phi \right]$$

$$= \int_{\underline{\phi}}^{\overline{\phi}} r_{S} \pi(\theta, a(\phi)) \hat{L}(\phi; \hat{\phi}, \underline{\phi}, \overline{\phi}; r_{S}) \, \mathrm{d}\phi$$

where $\hat{\ell}$ denotes the discounted local time function evaluated at ϕ and $\hat{L}(\phi; \hat{\phi}, \underline{\phi}, \overline{\phi}; r_S) := \mathbf{E}[\hat{\ell}(\phi; \hat{\phi}, \tau; r_S)].$

As in the proof of Lemma OA.3, by Arzelà-Ascoli theorem (Ok, 2007, Chapter D.6), the sequence of pairs of value functions has a converging subsequence and the limit pair is a continuously differentiable function that solves the limit boundary value problem. To put it differently, the limit pair of value functions solve the limit best-reply problem. On the other hand, in the limit, each type of sender maximizes

$$\max_{a_{\theta}(\phi) \in A^{[0,1]}} \int_0^1 r_S \pi(\theta, a_{\theta}(\phi)) \hat{L}(\phi; \hat{\phi}, \underline{\phi}, \overline{\phi}; r_S) d\phi,$$

where the choice of action affects not only the payoff but also the expected discounted local time function. Clearly, the expected discounted local time function is affected not only by type θ 's action but also by the other type, as well as by the DM's conjecture. However, we can look at a relaxed problem in which the only constraint

The expectation in the last term in the equation below also depends on ϕ_n and $\bar{\phi}_n$ but we are omitting this dependence for notational convenience.

on the choice of the expected discounted local time function is

$$\int_0^1 \hat{L}(\phi; \hat{\phi}, \underline{\phi}, \overline{\phi}; r_S) d\phi = 1/r_S.$$

Recall that one can interpret the expected discounted local time function as a weighting function, similar to a density, with the difference, that integrates to $1/r_s$.

Since π is single-peaked, it is then immediately that is optimal to choose the myopically optimal action at any time, that is, $a_{\theta}(\phi) = a_{\theta}^*$ for any $\phi \in (0,1)$.

Proof of Part 3: $\lim_{n\to\infty} r_{S,n}/r_{DM,n} = k \in (0,\infty)$. For each $n=1,2,\ldots$, we consider an equilibrium associated with the discount factors $r_{DM,n}$ and $r_{S,n}$ together with the corresponding value functions $\{(U_{H,n},U_{L,n})\}_{n\geq 1}$ and cutoffs $\{(\underline{\phi}_n,\overline{\phi}_n)\}_{n\geq 1}$. Assume that $\lim_{n\to\infty} r_{S,n} = 0$ and $\lim_{n\to\infty} r_{DM,n} = 0$ and $\lim_{n\to\infty} (r_{S,n}/r_{DM,n}) = \kappa \in (0,\infty)$. Recall that by Lemma 2, there exists a strictly positive lower bound to the difference in conditional drifts in the sequence of games. Again, abusing notation, let $\Delta\mu$ denote this lower bound when one considers the sequence of discount rates $\{(r_{S,n},r_{DM,n})\}_{n\geq 1}$.

Lemma 4. For each n, consider the sequence of DM's optimal stopping problems when the difference in conditional drifts is, for each ϕ , $\Delta\mu\phi(1-\phi)$. Denote the value function in these decision problems with V_n^{\dagger} and the decision time with τ_n^{\dagger} . Then,

$$\begin{split} &\lim_{n \to \infty} V_n^{\dagger}\left(\phi\right) = \phi G(H,h) + \left(1 - \phi\right) G(L,l), \\ &\lim_{n \to \infty} \mathbf{E}\left[e^{-r_{DM,n}\tau_n^{\dagger}}\right] = 1, \quad and \quad \lim_{n \to \infty} \mathbf{E}\left[\tau_n^{\dagger}\right] = \infty. \end{split}$$

Proof. From the proof of Proposition OA.2, the value function admits the following characterization,

$$V_n^{\dagger}(\phi) = \tilde{A}(1-\phi)^{\left(1+\sqrt{1+8r_{DM_n}\sigma^2/\overline{\Delta\mu}^2}\right)/2} (\phi)^{\left(1-\sqrt{1+8r_{DM_n}\sigma^2/\overline{\Delta\mu}^2}\right)/2} + \tilde{B}(1-\phi)^{\left(1-\sqrt{1+8r_{DM_n}\sigma^2/\overline{\Delta\mu}^2}\right)/2} (\phi)^{\left(1+\sqrt{1+8r_{DM_n}\sigma^2/\overline{\Delta\mu}^2}\right)/2},$$
(9)

for any $\phi \in (0,1)$ such that $V_n^{\dagger}(\phi) > \max\{\phi G(H,h) + (1-\phi)G(L,h), \phi G(H,l) + (1-\phi)G(L,l)\}$. In the limit, as $r_{DM,n} \to 0$, both fundamental solutions become linear in ϕ . It follows that $V_n^{\dagger}(\phi)$, which by standard results is convex and bounded

below by $\max\{\phi G(H,h) + (1-\phi)G(L,h), \phi G(H,l) + (1-\phi)G(L,l)\}$ and above by $\phi G(H,h) + (1-\phi)G(L,l)$, converges to $\phi G(H,h) + (1-\phi)G(L,l)$. But then, smooth pasting can hold only if the cutoffs characterizing the optimal policy converge to 0 and 1. For the value function to converge to the complete information value, the cost of delay must converge to zero, that is, $\lim_{n\to\infty} \mathbf{E}\left[e^{-r_{DM,n}\tau_n^{\dagger}}\right] = 1$. To show that $\mathbf{E}\left[\tau_n^{\dagger}\right] = \infty$, one follow the same steps as Lemma 3.

Lemma 4 implies, a fortiori, that

$$\lim_{n \to \infty} V_n(\phi) = \phi G(H, h) + (1 - \phi) G(L, l),$$

$$\lim_{n \to \infty} \mathbf{E} \left[e^{-r_{DM, n} \tau_n} \right] = 1, \quad \text{and} \quad \lim_{n \to \infty} \mathbf{E} \left[\tau_n \right] = \infty,$$

In turns, by Claim OA.6 in Ekmekci et al. (2022), $\lim_{n\to\infty} \mathbf{E}\left[e^{-r_{S,n}\tau_n}\right] = 1$. Given that π is bounded, it follows that, $\mathbf{E}\left[\int_0^{\tau} e^{-r_{S,n}t}\pi\left(\theta,a\left(\phi_t\right)\right)\mathrm{d}t\right] = 0$. As a result, the value function of the sender converges to

$$\lim_{n \to \infty} U_{H,n}(\phi) = \Pi(H,h) \qquad \lim_{n \to \infty} U_{L,n}(\phi) = \Pi(L,l) \qquad \phi \in (0,1). \tag{10}$$

Next, we show that for any $\phi \in (0,1)$ $\lim_{n\to\infty} U'_{\theta,n}(\phi)/r_{S,n} = -\infty$. For each n, consider the unique solutions to the boundary value problem

$$U_1''(\phi) = -2U_1'(\phi)/\phi, \qquad U_2''(\phi) = 2U_2'(\phi)/(1-\phi),$$

under the boundary conditions

$$U_1'(\underline{\phi}_n) = U_{H,n}'(\underline{\phi}_n), \qquad U_2'(\overline{\phi}_n) = U_{L,n}'(\overline{\phi}_n),$$

$$U_1(\phi_n) = \Pi(H, l), \qquad U_2(\overline{\phi}_n) = \Pi(L, h).$$

Because of (10), for any n > 0, the boundary values are bounded away from zero, and the unique solution to the above boundary problem is

$$\mathring{U}_{1,n}(\phi) = -\frac{U'_{H,n}(\underline{\phi}_n)(\underline{\phi}_n)^2}{\phi} + \Pi(H,l) + U'_{H,n}(\underline{\phi}_n)\underline{\phi}_n,
\mathring{U}_{2,n}(\phi) = \frac{U'_{L,n}(\bar{\phi}_n)(1-\bar{\phi}_n)^2}{1-\phi} + \Pi(L,h) + U'_{L,n}(\bar{\phi}_n)(1-\bar{\phi}_n).$$

Proceeding as in Lemma OA.3, one can show that the sequences $\{(U_{H,n}-\mathring{U}_{1,n},U_{L,n}-\mathring{U}_{2,n})\}_{n\geq 1}$ and $\{(U'_{H,n}-\mathring{U}'_{1,n},U'_{L,n}-\mathring{U}'_{2,n})\}_{n\geq 1}$ converge to zero uniformly. By Lemma OA.4,

$$\lim_{n\to\infty} (1-\bar{\phi}_n)^2/r_{S,n} = \lim_{n\to\infty} (\underline{\phi}_n)^2/r_{S,n} = \infty,$$

which implies that $\lim_{n\to\infty} U'_{1,n}(\phi_n)/r_{S,n}=\lim_{n\to\infty} U'_{2,n}(\phi_n)/r_{S,n}=-\infty$, so that $\lim_{n\to\infty} U'_{\theta,n}(\phi_n)/r_{S,n}=-\infty$.