Journal of Mathematical Economics 109 (2023) 102894

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Only time will tell: Credible dynamic signaling☆

Egor Starkov

Department of Economics, University of Copenhagen, Øster Farimagsgade 5, 1353 København K, Denmark

ARTICLE INFO

Article history:
Received 22 August 2022
Received in revised form 7 August 2023
Accepted 17 August 2023
Available online 7 September 2023
Manuscript handled by Editor Ying Chen

Keywords:
Dynamic signaling
Repeated signaling
Reputation
Attrition

ABSTRACT

This paper characterizes informational outcomes in a model of dynamic signaling with vanishing commitment power. It shows that contrary to popular belief, informative equilibria with payoff-relevant signaling can exist without requiring unreasonable off-path beliefs. The paper provides a sharp characterization of possible separating equilibria: all signaling must take place through attrition, when the weakest type mixes between revealing own type and pooling with the stronger types. The framework explored in the paper is general, imposing only minimal assumptions on payoff monotonicity and single-crossing. Applications to bargaining, monopoly price signaling, and labor market signaling are developed to demonstrate the results in specific contexts.

1. Introduction

In his seminal contribution, Spence (1973) argued that economic agents' actions can *signal* their private information, giving an example of schooling as a signal of ability in labor markets. In the years since, researchers have extensively studied signaling models, describing the fundamental driving forces driving them and identifying signaling patterns in a wide spectrum of applications: from bargaining (Vincent, 1990) and limit pricing (Milgrom and Roberts, 1982a,b) to corporate finance (Leland and Pyle, 1977) and advertising (Milgrom and Roberts, 1986).[1]

While most signaling models explore static interactions, dynamic signaling models may be better suited to explore some applications. For example, the choice of price to signal product quality or the choice of education/effort to signal worker's ability are both inherently dynamic problems, since the sender must in both cases repeatedly reaffirm their action choice. However, signaling in dynamic settings has a salient conceptual problem: if the agent cannot commit to an action for an extended period of time, then separation is not possible in equilibrium. For example, in Spence's story of signaling ability through education, high-ability students are supposed to pursue a college degree, while low-ability students, for whom education is more costly, would then supposedly find it optimal to enter the labor market straight out of school, skipping college. But if neither students can commit to finishing a degree, nor firms can commit to not hire undereducated workers, then a single day spent in college would imply that a student's ability is not low, and such a student could expect a correspondingly high wage in a competitive labor market — too high to actually deter low-ability workers from enrolling in college only to drop out soon thereafter. In the words of Admati and Perry (1987): "Once a high ability worker has gone to school long enough to distinguish himself from a worker of lower ability, the firms would offer wages appropriate to a high ability worker *before* enough time has elapsed to present an effective screen" (p. 363).[2]

The literature has responded to this conceptual challenge by searching for aspects of such dynamic interactions which would neutralize the argument above, thus enabling variations of the dynamic signaling model in which static separation is possible. Proposed solutions include: altering the payoffs to add intrinsic motivation for signaling (Weiss, 1983), tacit collusion on the receivers' side to generate instrumental commitment (Nöldeke and van Damme, 1990a; Swinkels, 1999),[3] evolving sender's type to

☆ This paper is based on chapter 3 of the author's Ph.D. thesis. The author thanks Nemanja Antić, Eddie Dekel, Jeffrey Ely, Yingni Guo, Nicolas Inostroza, Johan Lagerlöf, Alexey Makarin, Wojciech Olszewski, Marco Schwarz, Ludvig Sinander, Peter Norman Sørensen, Bruno Strulovici and seminar participants at Northwestern University, University of Copenhagen, and DICE for valuable feedback and helpful comments.

E-mail address: egor@starkov.email.

[1] See Riley (2001) for an excellent survey of the early literature on signaling.

[2] One possible explanation for the lack of commitment is the possibility of renegotiation; see Beaudry and Poitevin (1993).

[3] In the setting of Nöldeke and van Damme (1990a) and Swinkels (1999), a worker is signaling ability via years of education, and firms can at any point offer the worker a job. In a perfectly separating equilibrium, able workers choose to acquire college education, while the less able workers enter the job market right after high school. As soon as a worker enters college, they are revealed as able, hence firms have incentives to offer them a position with a high wage immediately, without waiting for the worker to finish their degree. In a tacit collusion equilibrium, as soon as one firm makes such an offer, other firms immediately start a bidding war for this worker, thereby eliminating any gains that could be accrued by the deviating firm by hiring this worker. This construction violates our monotonicity condition introduced in Section 2.3.

https://doi.org/10.1016/j.jmateco.2023.102894

create the need for maintaining reputation as opposed to establishing it once (Roddie, 2012a,b), or receivers observing sender's actions with noise (Dilmé, 2017; Heinsalu, 2018).[4] The basic case without any of the above is implicitly perceived as one in which informative signaling is impossible — if an equilibrium even exists, that is. However, the impossibility argument of Admati and Perry (1987) only applies to perfectly separating outcomes — it does not preclude partial separation, meaning the outcome where certain actions act as *suggestive* rather than *conclusive* evidence of the sender's information. The limits of such suggestive signaling in dynamic settings have, to our best knowledge, not been carefully investigated in the literature. We aim to fill this gap.

This paper shows that the scope for informative signaling, while limited, does in fact exist in dynamic settings without commitment, contrary to the intuition above and the ideas of Admati and Perry (1987) and Swinkels (1999). In particular, this paper explores a general signaling model, in which a sender is privately informed of his type and engages in a repeated interaction with a receiver, where the periods are vanishingly short. The receiver makes inferences about the sender's type from his action choices, and the sender's payoff is increasing in his reputation with the receiver. We show that under the appropriate monotonicity and single-crossing conditions on the sender's payoff function, payoff-relevant signaling is possible in this setting via what is effectively a war of attrition between different types of the sender. In such an attrition equilibrium, all sender types pool on the same action, with the lowest type mixing between pooling with the rest and separating to a myopically optimal action. Beyond such attrition, actions are as informative as cheap talk — meaning that attrition is the only way in which payoff-relevant signaling can proceed in this class of models. The contribution of this paper is both in showing the existence of a wedge between signaling and cheap talk in the setting under consideration (i.e., that signaling *is* possible), and characterizing the equilibrium outcomes.

The fact that informative equilibria of the attrition form exist in dynamic signaling models has been observed in applied models before. In particular, similar equilibria in specific settings have been obtained by Vincent (1990), Deneckere and Liang (2006), Daley and Green (2012), Lee and Liu (2013), Dilmé and Li (2016), Dilmé (2017), Kaya and Kim (2018) in the context of bargaining; Strebulaev, Zhu, and Zryumov (2016) in corporate finance; Vettas (1997), Aköz, Arbatli, and Celik (2020), Gryglewicz and Kolb (2022), Smirnov and Starkov (2022) in industrial organization/marketing; Smirnov and Starkov (2019) and Vong (2021) in cheap talk games; Gul and Pesendorfer (2012) in disclosure games; De Angelis, Ekström, and Glover (2022) in Dynkin games; Pei (2021) in trust games. The contribution of this paper is in setting up a general model that nests many of the models above and in identifying the sufficient conditions that yield uniqueness of attrition as the only informative equilibrium outcome, as well as providing a simple characterization of the set of equilibrium payoffs and a construction that implements them.

The main model explored in this paper is intended to serve as a framework suitable for many applications. Namely, the main model setup in Section 2 provides weak (and admittedly cumbersome) sufficient conditions for our results, while Section 5.1 then provides a more restrictive yet simple collection of sufficient conditions. To illustrate the applicability of the results, Section 5 then develops applications to price signaling, bargaining, and labor market signaling. In the price signaling example, a privately informed seller tries to signal the quality of his product by setting the price faced by generations of competitive buyers. In the bargaining example, an informed seller bargains with a single potential buyer over the price of the sale. In the labor market example, a student privately informed of his ability repeatedly chooses whether to invest in education, which may signal this ability to firms on the job market. Both problems represent classical applications of signaling theory. We show that the results from the general model apply, so all informative equilibria in each of these settings must take the attrition form.

It is worth noting that even though the attrition structure is restrictive, it allows for nontrivial equilibrium multiplicity. In addition to various possible combinations of informative and uninformative periods, multiple attrition outcomes can be sustained in equilibrium in any given period, which differ in terms of when and with what probability the lowest type separates. Section 6 characterizes the set of payoffs attainable in both pooling and attrition equilibria and provides both simple conditions for verifying which equilibria exist in a given game, and the constructive recipes for obtaining such equilibria.

Finally, this paper provides a takeaway regarding modeling assumptions that would be valuable to applied theorists investigating whether the receiver perfectly learns the sender's type (equivalently, whether social learning occurs) in a given setting in the limit as $t \to \infty$. In particular, if one adopts the simplifying assumption of there being only two types, then they could plausibly arrive at a conclusion that asymptotic learning is perfect. Yet, as this paper implies, this conclusion would not extend to the setting with finitely many types: learning can only occur regarding the lowest type, but cannot distinguish any of the higher types. It is worth noting that this impossibility result for finite types does not necessarily extend to the setting with an interval of types, where asymptotic learning is possible again (Fuchs and Skrzypacz, 2010 provide an example of such model and equilibrium in the bargaining context). The latter observation implies that the order of approximations matters: if one is interested in a continuous-type, continuous-time signaling game, discretizing types may not be an innocuous approximation and may produce results different from the fully continuous model, as well as from a discrete-time approximation.

Our analysis relies on the restriction of off-equilibrium path beliefs to be "reasonable". In particular, we adopt the assumption of non-increasing belief supports or, as labeled by Osborne and Rubinstein (1990), *NDOC* ("Never Dissuaded Once Convinced") assumption. As the name suggests, it implies that once the receiver has ruled out some type of the sender as impossible, the receiver stands by this belief and never again assigns positive probability to that type, including off the equilibrium path. Kaya (2009) and Roddie (2012a,b) have shown that in the absence of NDOC full instantaneous separation is possible in dynamic settings, since the sender's behavior can be disciplined by strong reputational threats in case of deviations. While the approach can be justified when the sender's type may change over time and hence needs constant re-verification, in other settings it is susceptible to a critique of using unreasonable off-path threats to sustain an equilibrium — a practice typically reproved in the literature on equilibrium refinements for static signaling games, as well as equilibrium concepts for dynamic games.[5] In other words, if any tremble away from the prescribed strategy can ruin all of the sender's acquired reputation, then what is such reputation worth? In contrast, the NDOC assumption allows us to explore the limits of *credible* or *robust* signaling — that which is not reversed by future deviations and is robust to trembles in the sender's strategy. See Section 2.3 for further discussion of and motivation behind NDOC.

4 See Whitmeyer (2021) for a discussion regarding the sender-optimal amount of noise in signaling.

5 C.f. Banks and Sobel (1987) and Cho and Kreps (1987) for signaling and Chapter 4 in Myerson (1997) for extensive-form games respectively.

The remainder of this paper is organized as follows. Section 2 sets up the general model and introduces the two assumptions that serve as sufficient conditions for our results: payoff monotonicity and NDOC. We then proceed to analyze two versions of this model. The two-type version in Section 3 can be seen as an illustrative example. The version with finitely many types, which requires an additional single-crossing assumption, is then explored in Section 4. Section 5 considers applications to price signaling, bargaining and labor market signaling, setting up the respective models and verifying that the required assumptions hold. Section 6 presents the equilibrium existence results. Section 7 concludes. The proofs of most statements are relegated to Appendix.

2. Model

2.1. Primitives

We will be looking at a continuous limit of a discrete-time infinite-horizon game. Time is indexed by $t \in \mathcal{T} \equiv \{0, dt, 2dt, \ldots\}$. We will be interested in the limit $dt \to 0$, which is needed to remove any implicit commitment power the agent may have. In particular, in discrete time the agent can effectively commit to not revise the action choice for the duration of the period, whereas we are interested in the case when the agent has no commitment power. At the same-time, analyzing the discrete-time model helps avoid a lot of the complexity associated with continuous-time models.

There are two players: a sender (agent) and a receiver. The agent has some persistent privately known *type* $\theta \in \Theta$, where $\Theta \subseteq \mathbb{R}$ is finite. Equivalently, θ can be the state of the world that the agent is privately informed of. In every period t, a Stackelberg-type sequential game is played between the agent and the receiver.[6] First, the agent chooses an *action* $a_t \in A$ from a compact action set A. This action choice affects the realization of an *outcome* $x_t \in X$, the distribution of which at time t may depend on the agent's type θ and the past history h_t. We assume that outcomes never allow to perfectly identify θ, yet they may well be *informative* of θ.[7,8] The receiver then observes both a_t and x_t and selects an action $b_t \in B$ from a compact set B. A time-t *history* is defined as $h_t \equiv ((a_0, x_0, b_0), \ldots, (a_{t-dt}, x_{t-dt}, b_{t-dt}))$, i.e., as the record of past actions and outcomes up to, but not including, time $t \in \mathcal{T}$ (so $h_0 = \varnothing$). Let $\mathcal{H}_t$ denote the set of all such time-t histories, and $\mathcal{H} \equiv \cup_{t\in\mathcal{T}}\mathcal{H}_t$ be the set of all such histories. For any two periods $s > t$, we write $h_s \succ h_t$ if history h_s succeeds history h_t, i.e., h_s and h_t coincide on $(0, \ldots, t - dt)$.

The receiver begins the game with a commonly known prior *belief* $p_0 \in \Delta(\Theta)$ about the agent's type. In every period t the receiver observes the agent's latest action a_t and the realized outcome x_t, and updates her belief p_t upon this information. Specifically, we use $p(h_t, a_t, x_t)$ to denote the resulting posterior belief. Hereinafter, belief p_t is also referred to as the agent's *reputation*. Further, we use $p(h_t)$ to denote the belief conditioned on h_t — i.e., at the beginning of period t — as a shorthand for $p(h_{t-dt}, a_{t-dt}, x_{t-dt})$, the receiver's belief at her previous action node. For any p_t, let $S(p_t) \subseteq \Theta$ be the support of belief p_t, i.e., the set of types to which p_t assigns positive weight. With abuse of notation, let $S(h_t) \equiv S(p(h_t))$.

At the end of every period t, the agent receives stage payoff $u(a_t, b_t, \theta)$, and the receiver obtains payoff $w(a_t, b_t, x_t, \theta)$. For simplicity, we assume that the agent's payoff u does not depend on the realized outcome x_t except through the effect it has on the receiver's action b_t. Both functions are assumed to be upper semi-continuous in the respective player's action: $\lim_{a\to a_t} u(a, b_t, \theta) \leq u(a_t, b_t, \theta)$ and $\lim_{b\to b_t} w(a_t, b, x_t, \theta) \leq w(a_t, b_t, x_t, \theta)$ for all a_t, b_t, x_t, θ. Further, assume that function u is bounded.

To focus on the signaling concerns (as opposed to the repeated game effects), we will be working under the assumption that the receiver behaves myopically and only maximizes the current stage payoff. This assumption can be justified on its own merit in some settings, e.g., when "the receiver" is a proxy for a competitive market of receivers or a sequence of short-lived players.[9] This assumption is not strictly necessary, although it greatly simplifies the exposition. Footnote 12 describes the extent to which the results can be applied in a model with a strategic receiver.

The *receiver's strategy* is $\mathbf{b}: \mathcal{H} \times A \times X \to B$. Strategy $\mathbf{b}^*$ is *optimal* for the receiver at a given history h_t if the action it prescribes maximizes the receiver's current stage payoff given the agent's strategy and the receiver's belief at that history

$$\mathbf{b}^*(h_t, a_t, x_t) \in \arg\max_b \mathbb{E}[w(a_t, b, x_t, \theta) \mid p(h_t)], \tag{1}$$

where the expectation is taken over θ. For simplicity, assume *no sunspots*: for any a, x, and any two histories h_t, h_s, if $p(h_t) = p(h_s)$ then $\mathbf{b}^*(h_t, a, x) = \mathbf{b}^*(h_s, a, x)$. Moving on to the agent, define a *bliss* (myopically optimal) action set for the agent of type θ at history h_t given receiver's strategy $\mathbf{b}$ as

$$A^*(h_t, \mathbf{b}, \theta) \equiv \arg\max_{a\in A} \{\mathbb{E}[u(a, \mathbf{b}(h_t, a, x_t), \theta)]\},$$

where the expectation is taken over x_t. Note that $\mathbf{b}^*$ exists and A^* is non-empty due to the upper semi-continuity of the respective utility functions w and u. A *pure strategy* for the agent of type θ is $\mathbf{a}_\theta : \mathcal{H} \to A$. Given some belief system p and the receiver's strategy $\mathbf{b}$, let $U(\mathbf{a}_\theta|h_t, \mathbf{b}, \theta)$ denote the expected discounted continuation utility of type θ from following strategy $\mathbf{a}_\theta$ starting from $h_t \in \mathcal{H}$:

$$U(\mathbf{a}_\theta|h_t, \mathbf{b}, \theta) \equiv \mathbb{E}\left[\sum_{s\in\mathcal{T}, s\geq t} e^{-r(s-t)} u\left(\mathbf{a}_\theta(h_s), \mathbf{b}_s, \theta\right) dt \mid h_t, \theta\right],$$

where r is the agent's discount rate, and the expectation is taken over future outcomes x_s. The agent is assumed to maximize his expected discounted sum of utilities. Strategy $\mathbf{a}_\theta$ is *optimal* for the agent of type θ given belief system p and receiver's strategy $\mathbf{b}$ if it maximizes his continuation payoff at every history $h_t \in \mathcal{H}$, i.e., if

$$U(\mathbf{a}_\theta|h_t, \mathbf{b}, \theta) = V(h_t, \mathbf{b}, \theta) \equiv \max_{\mathbf{a}} \{U(\mathbf{a}|h_t, \mathbf{b}, \theta)\}, \tag{2}$$

6 The results of the paper apply (after some relabeling) to simultaneous-move games as well, but for illustrative simplicity we focus on Stackelberg games in the analysis.

7 Formally, let $\zeta : \mathcal{H} \times A \times \Theta \to \Delta(X)$ denote the outcome distribution conditional on the agent's type, played action, and the history of all past actions. We then assume that for any $\theta', \theta'' \in \Theta$ and any $h_t \in \mathcal{H}, a_t \in A$, the distributions $\zeta(h_t, a_t, \theta')$ and $\zeta(h_t, a_t, \theta'')$ and absolutely continuous with respect to each other.

8 One could call x_t a "public signal"; we avoid this phrasing so as to not create confusion with the process of signaling through actions. The analysis does *not* rely on the presence of such outcomes (and indeed, the applications in Section 5 and the results in Section 6 assume no informative outcomes). Rather, these outcomes are introduced to demonstrate that the main results in Sections 3 and 4 hold *even if* some noisy public information is being revealed in addition to the agent's actions, as long as it does not allow identifying individual types. This includes both exogenous public news and public signals produced as a result of the agent's actions.

9 For example, in the context of the labor market signaling, a worker is continuously signaling his ability to a population of competitive firms, which, in an attempt to get the worker, bid wages up to the worker's expected productivity, thereby eliminating any strategic element in wage offers. Alternatively, in the context of price signaling, a firm uses its price (and price history) to signal its product's value to changing generations of short-lived consumers.

where $V(h_t, \mathbf{b}, \theta)$ is hereinafter referred to as the value function. With a slight abuse of notation we let

$$V(a|h_t, \mathbf{b}, \theta) \equiv \mathbb{E}_{x_t}\left[u(a, \mathbf{b}(h_t, a, x_t), \theta)dt + e^{-rdt}V(h_{t+dt}, \mathbf{b}, \theta) \mid h_t, a, \mathbf{b}, \theta\right]$$

denote the highest expected continuation utility that type θ can achieve conditional on taking action a at history h_t. The outer expectation is taken with respect to the realization of period-t outcome x_t, which affects the receiver's belief p_t and, thus, her action b_t and the agent's contemporaneous utility $u(a_t, b_t, \theta)$. The $t+dt$-history is $h_{t+dt} = (h_t, (a_t, x_t, b_t))$.

Finally, a *behavioral strategy* for the agent of type θ is $\alpha_\theta : \mathcal{H} \to \Delta(A)$. By the Kuhn's Theorem (Aumann, 1964), behavioral strategies are equivalent to mixed strategies in this setting. Let $\alpha_\theta(a|h_t)$ denote the probability with which action a should be played by type θ after history h_t according to strategy $\alpha_\theta(h_t)$. A behavioral strategy α_θ is then optimal for θ if there exists an equivalent mixed strategy (i.e., a probability distribution over pure strategies), such that all pure strategies in its support are optimal.

2.2. Equilibrium concept

Introduced above is a dynamic game of incomplete information. The greatest common factor among the solution concepts used for this class of games (and requiring belief consistency) is Perfect Bayesian Equilibrium (PBE). In such an equilibrium, all players maximize their expected continuation payoffs given their beliefs about other players' actions and beliefs, and these beliefs must be consistent on path with the players' knowledge of the game.

Definition 1. A *Perfect Bayesian Equilibrium* of the game is given by an agent's strategy profile $\alpha = \{\alpha_\theta\}_{\theta\in\Theta}$ with $\alpha_\theta : \mathcal{H} \to \Delta(A)$, a receiver's strategy $\mathbf{b}$: $\mathcal{H} \times A \times X \to B$, and a belief system $p : \mathcal{H} \times A \times X \to \Delta(\Theta)$ such that:

1. for all θ: strategy profile α_θ is optimal for the agent of type θ;
2. strategy $\mathbf{b}$ is optimal for the receiver at all histories $h_t \in \mathcal{H}$;
3. belief p is updated using Bayes' rule whenever possible.

Hereinafter, the PBE are referred to as simply "equilibria". We impose additional equilibrium refinements and restrictions in the following subsection and further in the text as they become relevant. Further, as mentioned previously, we explore equilibria for small but positive dt, and we are interested in the properties of these equilibria as $dt \to 0$. In order to explore those, we also need to define what exactly an "equilibrium as $dt \to 0$" means. Let $\mathcal{H}_c \equiv \cup_{t\in\mathbb{R}}\mathcal{H}_t$ denote the set of all possible histories in the continuous-time limit.[10] We can then define a limit equilibrium as follows.

Definition 2. A *Limit Equilibrium* of the game is given by an agent's strategy profile $\alpha = \{\alpha_\theta\}_{\theta\in\Theta}$ with $\alpha_\theta : \mathcal{H}_c \to \Delta(A)$, a receiver's strategy $\mathbf{b} : \mathcal{H}_c \times A \times X \to B$, and a belief system $p : \mathcal{H}_c \times A \times X \to \Delta(\Theta)$ such that there exists a sequence of equilibria $\{(\alpha_n, \mathbf{b}_n, p_n)\}_{n=1}^{\infty}$ for a respective sequence of $\{dt_n\} \xrightarrow[n\to\infty]{} 0$ such that $(\alpha_n, \mathbf{b}_n, p_n) \xrightarrow[n\to\infty]{} (\alpha, \mathbf{b}, p)$.

2.3. Restrictions

This section describes two significant equilibrium restrictions that will be imposed throughout the analysis and which are sufficient for the results in the two-type model: Monotonicity and Never Dissuaded Once Convinced. The version of the model with more than two types requires a third condition, Single Crossing, which is introduced and discussed separately in Section 4.1, and yet another weak restriction is adopted in the discussion of equilibrium existence in Section 6.

To introduce these restrictions, a few extra bits of notation would prove useful. Firstly, let δ_θ denote the Dirac delta: given some $\theta \in \Theta$, $p(h_t) = \delta_\theta$ is equivalent to $S(h_t) = \{\theta\}$. Secondly, let $\tilde{u}_t$ denote the agent's induced stage payoff function given some fixed strategy $\mathbf{b}$ of the receiver:

$$\tilde{u}_t(a_t, p(h_t, a_t, x_t), \theta) \equiv u(a_t, \mathbf{b}(h_t, a_t, x_t), \theta). \tag{3}$$

This reduction of the sender's preferences to depend on the receiver's beliefs, as opposed to actions, is a standard technique used both in signaling literature (Kaya, 2009; Roddie, 2012a,b) and other literatures such as Bayesian Persuasion (see Kamenica and Gentzkow, 2011). The equilibrium restrictions can then be phrased as follows (with the discussion following afterwards):

> **(MON)** Flow payoff function $\tilde{u}_t(a_t, p_t, \theta)$ is weakly increasing in p_t (with respect to the FOSD order) for all t, h_t, a_t, θ and all optimal $\mathbf{b}$.[11] Further, if $p_t >_{FOSD} \delta_\theta$ then $\tilde{u}(a_t, p_t, \theta) > \tilde{u}(a_t, \delta_\theta, \theta)$, and if $p_t <_{FOSD} \delta_\theta$ then $\tilde{u}(a_t, p_t, \theta) < \tilde{u}(a_t, \delta_\theta, \theta)$.
>
> **(NDOC)** Belief process $\{p_t\}_{t\in\mathcal{T}}$ is such that belief supports are non-increasing: for any $h_s \succ h_t$, $S(h_s) \subseteq S(h_t)$.

The first condition, (MON), requires that the agent's stage payoff function is increasing in his reputation p_t. This captures the core idea of signaling models: the agent would like to signal that his type is high because that induces a favorable reaction from the receiver. For example, a firm with a reputation for quality product is more likely to sell more units, an able worker is more likely to be offered a job, and a strong bargainer is more likely to see the opponent conceding to a demanding offer. Monotonicity is primarily an assumption on the model primitives, namely the utility functions: given the receiver's preferences w, her optimal strategy $\mathbf{b}$ is unique up to indifference for any a and $p(h_t)$. This makes $\tilde{u}(a, p, \theta)$ a well-defined function given some tie-breaking rule for the receiver. Hence given w, (MON) is a condition on the agent's utility function $u(a, b, \theta)$.[12] While the condition is phrased using weak monotonicity, strict preferences relative to degenerate reputation are required to guarantee the presence of signaling effects: any type must always be strictly willing to pool with the higher types and to separate away from the lower types.

The second condition, (NDOC), is the refinement of the equilibrium beliefs that drives our analysis. In particular, it says that if $p(\theta|h_t) = 0$ then $p(\theta|h_s) = 0$ for any pair of histories $h_s \succ h_t$ in $\mathcal{H}$. Note that this applies both on and off the equilibrium path. In other words, once the receiver is convinced that a given type of the agent is inconsistent with the evidence (the observed history),

[10] Note that any Borel-measurable function f defined on $\mathcal{H}$ can be extended to a Borel-measurable function on $\mathcal{H}_c$.

[11] Monotonicity with respect to the FOSD order on p_t is understood in the usual way: if p' first-order stochastically dominates p'', then $\tilde{u}_t(a_t, p', \theta) \geq \tilde{u}_t(a_t, p'', \theta)$.

[12] The assumption that the receiver behaves myopically was introduced to render (MON) expressible in terms of model primitives, but is not necessary for our results to hold. Instead, any other set of assumptions on the players' payoffs and/or the receiver's behavior that yields (MON) would work equally well. Developing such assumptions for the case of strategic receiver is not trivial, since the folk theorem dictates that any individually rational payoff for the sender can be sustained in equilibrium regardless of his reputation, which means (MON) does not hold without some additional assumptions on the receiver's behavior. However, Section 5.4 demonstrates that myopy is not a necessary assumption by applying this paper's results to a bargaining problem, in which both the sender and the receiver are strategic and forward-looking.

she can never be dissuaded from this conviction. This restriction appears reasonable, since over time, only more evidence is collected, but the existing evidence is never forgotten — including the evidence that lead the receiver to rule out certain types of the agent at the time.[13]

Alternatively, (NDOC) can be seen as a weak form of renegotiation-proofness in some settings, as suggested by Ely and Välimäki (2003). For example, in the context of labor market signaling, suppose that a firm offers a contract to a high school graduate, according to which it would hire the worker as soon as he obtains a college degree. Suppose further that in equilibrium such a contract is only accepted by able workers (whose cost of learning is low), while less able workers reject it in favor of getting a job immediately. Then if such a contract is accepted, the firm knows the worker is able, and it is in the best mutual interest of the firm and the worker to renegotiate the contract to start the job immediately, since the delay to obtain education is wasteful for both parties.

Finally, one can motivate (NDOC) through strategic attention allocation.[14] Suppose for a second that processing signals and updating beliefs is costlier for the receiver than the inertia of sticking to a previously established belief. Then it becomes quite intuitive that the receiver would only ever choose to exert this cost (however small it might be) and to pay attention to the signals if she expects there to be anything to learn — i.e., if her belief admits any uncertainty, $p_t \in \operatorname{int}\Delta(\Theta)$ (or enough uncertainty if the cost is non-trivial). If, on the other hand, her belief assigns probability one to some state/receiver type, then she would not choose to pay attention to any future signals, since those are not expected to carry any additional information, so their marginal benefit is zero. This motivation applies as well, if not better, when a receiver is myopic, since myopy can be interpreted as inattention towards future payoffs. Alternatively, in the interpretation that the receiver represents generations of short-lived players, one could argue that inertia (asking some of the past players about their beliefs) is still cheaper than paying attention to signals, past or current. In this latter interpretation, it is sufficient that a public belief is transmitted between generations (as opposed to the whole history).

The (NDOC) condition has been originally introduced by Osborne and Rubinstein (1990). It has been validated through its wide use in applied dynamic models with asymmetric information: one can find analogs of (NDOC), often labeled differently, in Rubinstein (1985), Grossman and Perry (1986), LeBlanc (1992), Vettas (1997), Kraus, Wilkenfeld, and Zlotkin (1995), Sen (2000), Ely and Välimäki (2003), Feinberg and Skrzypacz (2005), Lai (2014), Bond and Zhong (2016), Gryglewicz and Kolb (2022), Smirnov and Starkov (2019, 2022). (NDOC) is nonetheless a strong assumption and has been criticized as leading to possible equilibrium nonexistence (see Madrigal, Tan, and Werlang, 1987 and Nöldeke and van Damme, 1990b). This paper responds by showing that the set of equilibria is larger than what these papers consider due to the existence of mixed-strategy equilibria. We primarily characterize equilibria conditional on existence, but Section 6 provides an easy characterization of whether equilibria of a given type exist in a given game. The papers mentioned above also present numerous examples of settings in which equilibria exist.

[13] As argued by Osborne and Rubinstein (1990), p. 97: "[I]f we allow a player in a game of incomplete information to change his mind after he has been persuaded that he is playing with certainty against a given type, then why we do not do so in a game of complete information?"

[14] Gabaix (2019) argues that most behavioral biases explored in the behavioral economics literature can be represented as various forms of inattention. Handel and Schwartzstein (2018) survey empirical evidence of inattention in economic interactions.

In the analysis, we strengthen (NDOC) by rendering the receiver pessimistic off the equilibrium path — her beliefs off path must put all weight on the lowest type among those she has not yet ruled out. This strengthening simplifies the analysis greatly and, as discussed further, is without loss of generality. We label this stronger condition as (NDOC-P), and it is defined as follows:

> **(NDOC-P)** The off-equilibrium-path beliefs are such that after any action a that is not on equilibrium path at $h_t \in \mathcal{H}$: $p(h_t, a, x_t) = \delta_{\min S(h_t)}$ for any $x_t \in X$.[15]

Given (MON), this condition imposes the strongest possible punishment on the sender for any deviation among those punishments that satisfy (NDOC). This means if an equilibrium satisfies (NDOC), we can take its on-path actions and beliefs, replace any off-path beliefs with those prescribed by (NDOC-P) and a respective profile of off-path actions. Since the agent had no profitable deviations in the original (NDOC) equilibrium, he will not have any in the modified (NDOC-P) equilibrium, since the modifications render all deviations less appealing. This claim is formalized by the following lemma, with the proof available in Appendix.

Definition 3. We say that two equilibria $(\alpha', \mathbf{b}', p')$ and $(\alpha'', \mathbf{b}'', p'')$ are:

- *payoff-equivalent* if for all $h_t \in \mathcal{H}$ and all $\theta \in \Theta$, $V(h_t, \mathbf{b}, \theta) = V(h_t, \mathbf{b}, \theta)$;
- *strategy-equivalent* if for all $h_t \in \mathcal{H}$: $\alpha'_\theta(h_t) = \alpha''_\theta(h_t)$ for all $\theta \in \Theta$ and $\mathbf{b}'(h_t) = \mathbf{b}''(h_t)$;
- *on-path payoff-equivalent* and/or *on-path strategy-equivalent* if either of the respective definitions above hold for all histories $h_t \in \mathcal{H}$ that are on path in either equilibrium.

Lemma 1. *For any equilibrium $(\alpha', \mathbf{b}', p')$ that satisfies (MON) and (NDOC), there exists a payoff-equivalent and on-path strategy-equivalent equilibrium $(\alpha'', \mathbf{b}'', p'')$ that satisfies (MON) and (NDOC-P).*

3. Two types

This section explores the version of the model with only two types: $\Theta = \{L, H\}$. Here we show that signaling must take the form of attrition regardless of the sender's payoffs, as long as they are monotone in reputation p_t. The first part of Theorem 1 states that *perfect* separation cannot occur at any history in equilibrium: if a given action is on path for $\theta = H$ then it is also on path for $\theta = L$. This statement captures the idea of Admati and Perry (1987) and Nöldeke and van Damme (1990a). We also observe that there may effectively be only one such pooling action in any period, in the sense of all pooling actions must be payoff-equivalent for all types of the agent. This follows trivially from the fact that both types must be indifferent between playing any such action if there are more than one.

The insight that is novel to this paper (in the general setting) is that the converse to the first statement is not necessarily true: if $\alpha_L(a|h_t) > 0$ then $\alpha_H(a|h_t)$ may or may not be positive. In other words, there may exist actions which perfectly identify the low type, even if there do not exist any that identify the high type. It is immediate that the low type must be mixing for this to be possible, which is summarized by the second part of the theorem. Theorem 1 does not claim existence of any such separating actions, since they, as previously mentioned, need not exist in any given case — but Sections 5 and 6 present examples of specific settings in which such semi-separating equilibria exist and discuss equilibrium existence more generally.

[15] On-pathness is defined in the usual way; see Section 4.2 for a formal definition.

Theorem 1. *Suppose $\Theta = \{L, H\}$. In any limit equilibrium $(\alpha, \mathbf{b}, p)$ such that (MON) and (NDOC-P) hold, at any $h_t \in \mathcal{H}$ with $S(h_t) = \{L, H\}$, and for any $a' \in A$:*

1. *if $\alpha_H(a'|h_t) > 0$ then $\alpha_L(a'|h_t) > 0$. Further, all such a' are payoff-equivalent in the sense that $V(a'|h_t, \mathbf{b}, \theta)$ is the same across such a' for both types θ.*
2. *if $\alpha_H(a'|h_t) = 0$ and $\alpha_L(a'|h_t) > 0$ then $a' \in A^*(h_t, \mathbf{b}, L)$ and $V(a'|h_t, \mathbf{b}, L) = V(a''|h_t, \mathbf{b}, L)$ for any a'' such that $\alpha_H(a''|h_t) > 0$.*

Note that the attrition structure of signaling imposes strong restrictions on actions that can be played in limit equilibria. Firstly, any separating action that perfectly identifies the low type must be myopically optimal for him, since the low type does not have any strategic incentives to play anything else. Secondly, if the low type mixes between pooling and separating, then he must be indifferent between the two: the gains from pooling (higher reputation) are exactly offset by the cost of taking suboptimal actions in current and/or future periods.

It is worth emphasizing that the result holds under very minimal assumptions on payoffs and signals: the only requirements imposed on the model are that the sender's payoff is increasing in p (which, in fact, is only required for the low type) and that the outcomes x are not perfectly revealing. If the setting of interest fits this framework, then attrition is *the only* informative limit equilibrium structure that can arise in this setting, unless one is willing to allow for NDOC-nonconformant beliefs off the equilibrium path. Under attrition, the high type is playing some pooling action, while the low type mixes between that and a separating action.

One important case, which lies beyond the scope of our model, but is nonetheless worth mentioning, is that with a behaviorally committed type of the sender, and a strategic type, who prefers to mimic the committed type. For example, in the bargaining model of Abreu and Gul (2000), a player may be either committed to rejecting all offers that give him anything less than the whole surplus, or fully strategic. Similarly, in the (static) cheap talk model of Chen (2011), the sender may either be committed to truthful communication, or communicate strategically. Theorem 1 applies to such problems (with the exception of the payoff equivalence part of statement 1), since its proof only relies on the incentives of the low type — which in these settings is the strategic type. It follows that if the committed type is unable to verifiably demonstrate his commitment, perfect separation is impossible in limit equilibria, and all limit equilibria feature either attrition of the strategic type as in Abreu and Gul (2000), or full pooling.

4. Finite types

We now move to exploring the setting with more than two but finitely many types. In this section we show that the insight of Theorem 1 can be extended to this case, although allowing for many types does raise a number of additional issues and calls for extra assumptions.

4.1. Single-crossing

In order to secure the result in case of many types, we need to impose the following new equilibrium restriction:

(SC) For any $\mathbf{a}', \mathbf{a}'' \in \cup_{h_t \in \mathcal{H}} \cup_{\theta \in S(h_t)} \arg\max_{\mathbf{a}} U(\mathbf{a}|h_t, \mathbf{b}, \theta)$, all optimal $\mathbf{b}$, and all $h_t \in \mathcal{H}$, the function defined by $\mathcal{U}(\theta) \equiv U(\mathbf{a}''|h_t, \mathbf{b}, \theta) - U(\mathbf{a}'|h_t, \mathbf{b}, \theta)$ either crosses zero at most once, or is identically zero.

This condition belongs to a family of single-crossing conditions widely encountered in the literature on signaling, monotone comparative statics, and mechanism design.[16] The purpose of our condition is standard: to ensure that the agent's preferred strategy is, in some sense, monotone with respect to his type. There are, however, some distinctive features that differentiate it slightly from other single-crossing conditions in the literature (which may be evident from it being framed as an equilibrium restriction, as opposed to an assumption on the payoff primitives).

Firstly, (SC) is a condition on the expectation of a discounted sum $\mathbb{E}\sum_t e^{-rt}u(a_t, b_t, \theta)$ rather than on the stage utility $u(a, b, \theta)$. While the latter would be more preferable, aggregating single-crossing is not a trivial problem. Quah and Strulovici (2012) discuss this problem and offer possible solutions, but none of them apply to our setting.[17] Secondly, (SC) is more demanding than might appear initially. The dependence of $U(\mathbf{a}|h, \mathbf{b}, \theta)$ on $\mathbf{a}$ realizes not only directly — through the effect of agent's own action a_t on his stage utility $u(a_t, b_t, \theta)$ — but also via an indirect reputation channel. The receiver's response b_t depends on agent's reputation p_t, which is, in general, affected by the agent's action choice a_t. Further, this reputation effect is persistent, with the choice of a_t affecting not only the contemporaneous response b_t, but also the continuation payoff: for a given fixed path $\{a_s, x_s\}_{s>t}$, reputation $\{p(h_s)\}_{s>t}$ will be persistently shifted by a_t, meaning the receiver's responses $\{b_s\}_{s>t}$ are affected.

All of the above means that (SC) is a non-trivial condition and may be difficult to verify in some settings. If anything, verifying (SC) might as well be the main impediment to exploiting this paper's results in applied models. However, this task is far from impossible, with Section 5 providing an easy sufficient condition, as well as a number of examples of applied models, where (SC) can be easily verified.

4.2. Attrition structure of equilibrium signaling

Theorem 2 that we gradually build up to is the analog of Theorem 1 for the case when $|\Theta| > 2$, in the sense of characterizing the actions available in limit equilibria at any history. We begin, however, by stating a weaker result which, by looking at strategies rather than actions, provides a clearer characterization of the attrition structure of equilibrium signaling with $|\Theta| > 2$. Proposition 1 below establishes that as long as (SC) and other previously stated assumptions hold, strategies played in an arbitrary limit equilibrium of the game can be split into two classes. The first class consists of pooling strategies played by all types. While a nominal multiplicity of such strategies may arise, they must all be payoff-equivalent, so this class is, in a sense, degenerate. The second class is that of separating strategies employed by the lowest type — these may vary in which pooling strategies they mimic and for how long. However, any separating strategy is only played by the lowest type.

To state this and other results we need to introduce some additional notation and definitions. Firstly, denote the two boundaries of the belief support at a given history h_t as $\bar{S}(h_t) \equiv \max S(h_t)$ and $\underline{S}(h_t) \equiv \min S(h_t)$ respectively. Furthermore, in a

[16] See Laffont and Martimort (2002) from a contract theory perspective (e.g., Ch. 2.2.3). Classic references on MCS, in turn, include Milgrom and Shannon (1994) and Athey (2002).

[17] Quah and Strulovici (2012) operate in a setting where $\mathbf{a}'$ and $\mathbf{a}''$ are comparable with respect to some order. This requirement is too strong in our setting, since even if we could impose some order on A, there is no straightforward way to extend it to $A^{\mathcal{H}}$ apart from pointwise dominance ($\mathbf{a}'' \geq \mathbf{a}' \iff \forall h_t \in \mathcal{H} : \mathbf{a}''(h_t) \geq \mathbf{a}'(h_t)$), which creates an ordering that is not useful in practice due to the vanishingly small number of pairs $\mathbf{a}', \mathbf{a}'' \in A^{\mathcal{H}}$ it can rank.

manner similar to type support S, given an equilibrium strategy profile $\{\alpha_\theta\}$, let us define action support as

$$A(h_t) \equiv \cup_{\theta \in S(h_t)} \{a \in A \mid \alpha_\theta(a|h_t) > 0\}.$$

We say that given the receiver's strategy $\mathbf{b}$, a pure strategy $\mathbf{a}$ *arrives at history* h_t — and denote it as $\mathbf{a} \curlywedge h_t$ — if $\mathbf{a}(h_\tau) = a_\tau(h_\tau)$ for all h_τ s.t. $h_t \succ h_\tau$. Further, say that $\mathbf{a}$ is *on path for* θ *at* h_t if $\mathbf{a} \curlywedge h_t$ and $\mathbf{a}$ is on path according to type θ's equilibrium strategy α starting from h_t: $\alpha_\theta(\mathbf{a}(h_t)|h_t) > 0$. Say that $\mathbf{a}$ is *on path at* h_t if it is on path at h_t for some $\theta \in S(h_t)$.

We proceed by defining payoff equivalence of strategies in a straightforward manner.

Definition 4. Fix an equilibrium $(\alpha, \mathbf{b}, p)$ and history h_t. Any two pure strategies $\mathbf{a}', \mathbf{a}'' \curlywedge h_t$ are:

- *payoff-distinct at* h_t if there exists $\theta \in S(h_t)$ such that $U(\mathbf{a}'|h_t, \mathbf{b}, \theta) \neq U(\mathbf{a}''|h_t, \mathbf{b}, \theta)$;
- *payoff-equivalent at* h_t if they are not payoff-distinct at h_t.

The result can then be stated as follows.

Proposition 1. *Fix a limit equilibrium* $(\alpha, \mathbf{b}, p)$ *such that (MON), (SC), and (NDOC-P) hold. Fix some history* $h_t \in \mathcal{H}$ *and define* $\underline{\theta} \equiv \underline{S}(h_t)$*. Then for any pure strategy* $\bar{\mathbf{a}}'$ *on path at* h_t *for some type* $\theta' \in S(h_t)\backslash\underline{\theta}$*, the following hold:*

1. $\bar{\mathbf{a}}'$ *is optimal for all* $\theta \in S(h_t)$ *at* h_t*;*
2. *any* $\bar{\mathbf{a}}''$ *optimal for any* $\theta'' \in S(h_t)\backslash\underline{\theta}$ *is payoff-equivalent at* h_t *to* $\bar{\mathbf{a}}'$*;*
3. *there exists* $\bar{\mathbf{a}}'''$ *that is payoff-equivalent at* h_t *to* $\bar{\mathbf{a}}'$ *and is on path for* $\underline{\theta}$ *at* h_t*;*
4. *any* $\underline{\mathbf{a}}$ *that is on path at* h_t *and payoff-distinct at* h_t *from* $\bar{\mathbf{a}}'$ *is only on path for* $\underline{\theta}$*.*

To understand this proposition, it is illustrative to ignore payoff equivalence for a second and treat any pair of payoff-equivalent strategies as the same strategy. In this reading, the proposition implies that any pure strategy $\mathbf{a}$ on path for some type $\theta \in S(h_t)$ is on path for *all* types $\theta \in S(h_t)$, including the currently-lowest type $\underline{\theta}$. Therefore, no type of the agent can ever conclusively separate from $\underline{\theta}$. At the same time, there may exist strategies that separate $\underline{\theta}$ away from the remaining types. The weight that the receiver's belief assigns to $\underline{\theta}$ may thus decrease over time along the pooling path of play, it may even converge to zero asymptotically as $t \to \infty$, but it may never become exactly zero. However, the interpretation above is overly strong, since payoff-equivalent strategies do not need to coincide at all histories. In other words, it is a statement about *strategies*, whereas we would like to have a result about *actions*.

4.3. From strategies to actions

The question that remains unanswered by Proposition 1 is whether payoff-equivalence implies strategy-equivalence (i.e., that any two payoff-equivalent strategies must prescribe the same actions at either all, or at least some histories) or, if not, whether payoff-equivalent strategies at least produce equivalent belief paths for all types. Unfortunately, the answer to both of the above is negative: in general, not only may there be multiple payoff-equivalent strategies, but they may even induce different beliefs. This is demonstrated by the following example.

Example 1. Suppose $\Theta = \{0, 1, 2\}$, types are ex ante equiprobable, $A = \mathbb{R}_+$, and outcomes are uninformative. Suppose the sender's reduced-form utility function (3) is given by $\tilde{u}(a, p, \theta) = \mathbb{E}_p(\theta)$ (so agent's actions are cheap talk; (SC) holds trivially in this scenario). Then the following strategies constitute an equilibrium together with respective beliefs: type $\theta = 2$ plays strategy $\mathbf{a}'' = (a'', 0, 0, \ldots)$, while types $\theta = 1, 3$ play $\mathbf{a}' = (a', 0, 0, \ldots)$, where $a' \neq a''$ are arbitrary. In this PBE some information about type is conveyed in period zero — namely, type $\theta = 2$ separates from $\theta = 1, 3$. However, all types of the sender are indifferent between the two strategies, hence information revealed by a_0 is not relevant to the sender's payoff — although it may be relevant for the receiver.

However, we are arguably more interested in *payoff-relevant* signaling, which relies on the heterogeneity of the agent's preferences across types to convey information, as opposed to the agent's utmost indifference. Narrowing the focus to such payoff-relevant information revelation allows to carry the insight of Proposition 1 over from strategies to actions. We begin by stating the formal definitions of payoff-relevant and irrelevant signaling in our setting.

Definition 5. Fix an equilibrium $(\alpha, \mathbf{b}, p)$ and history $h_t \in \mathcal{H}$. *Payoff-relevant signaling* happens at h_t if there exist $a', a'' \in A(h_t)$ and $\theta \in S(h_t)$ such that $V(a'|h_t, \mathbf{b}, \theta) \neq V(a''|h_t, \mathbf{b}, \theta)$.

In other words, payoff-relevant signaling implies that at a given history h_t there are two distinct actions on path, a' and a'', and there is some type of the agent for which the choice between these two actions has payoff consequences. Note that since both actions are on path, it cannot be the case that all types prefer one over another — both a' and a'' must be optimal for some types of the agent. Payoff-relevance of this action choice is then defined as some type $\theta \in S(h_t)$ having strict preference between the two. To clarify, the example above features an instance of a *payoff-irrelevant signaling*, which happens (in a given equilibrium at a given h_t) if there exist $a', a'' \in A(h_t)$ such that $p(h_t, a', x) \neq p(h_t, a'', x)$ for some $x \in X$ but $V(a'|h_t, \mathbf{b}, \theta) = V(a''|h_t, \mathbf{b}, \theta)$ for all $\theta \in S(h_t)$.

We are now ready to state the theorem that characterizes payoff-relevant signaling in terms of actions, making the implications of Proposition 1 more explicit. The result below expands the message obtained in Theorem 1 to the case of finitely many types.

Theorem 2. *Fix a limit equilibrium* $(\alpha, \mathbf{b}, p)$ *such that (MON), (SC), and (NDOC-P) hold. Fix some history* $h_t \in \mathcal{H}$*. If payoff-relevant signaling happens at* h_t *then, defining* $\underline{\theta} \equiv \underline{S}(h_t)$*, the following hold:*

1. *any on-path action* $a \in A(h_t)$ *is on path for* $\underline{\theta}$ *at* h_t*;*
2. $A(h_t) \cap A^*\left(h_t, \mathbf{b}, \underline{\theta}\right)$ *is nonempty, and any* $\underline{a}$ *in the intersection is on path only for* $\underline{\theta}$ *at* h_t*;*
3. *any action* $\bar{a} \in A(h_t)\backslash A^*\left(h_t, \mathbf{b}, \underline{\theta}\right)$ *is optimal at* h_t *for all* $\theta \in S(h_t)$*.*

What the theorem says is that in any limit equilibrium with payoff-relevant signaling, there are effectively at most two types of *actions* — as opposed to *strategies* in Proposition 1 — on path at any history: pooling actions (typical element $\bar{a}$) and separating actions (typical element $\underline{a}$). The latter are only ever played by the currently-lowest type $\theta = \underline{S}(h_t)$ and separate him from the remaining types. As in Theorem 1, any separating action must be myopically optimal for the lowest type given that he is revealed.

Pooling actions, on the other hand, are optimal for all types. Further, if no payoff-irrelevant signaling takes place, then any pooling action is, in fact, on path for all $\theta \in S(h_t)$ — i.e., all types do actually pool on the pooling action(s). Notably, both payoff-relevant and payoff-irrelevant signaling may occur simultaneously at a given history. In that case there will be more than one pooling action, and while all of them are necessarily on path for $\underline{\theta}$, the higher types may vary in their action choices, despite all types being indifferent between all of these pooling actions.

The corollary below relates to the situations when payoff-relevant signaling occurs at successive histories. It states that the pooling action in the earlier history must then be such that the low type is indifferent between separating and pooling — meaning that *stage payoffs* the low type gets from the separating and pooling actions must be the same. This is because the low type must be indifferent between separating at t and $t + dt$ and the continuation payoffs after pooling or separating at t are the same (since after pooling, separating at $t + dt$ is, by assumption, optimal). One period of pooling must then be exactly as attractive as one period of being identified as $\underline{\theta}$. In practice, this means that pooling action must be costlier for $\underline{\theta}$ than the separating action, since the former yields higher reputation payoff.

Corollary 1. *Suppose the conditions in Theorem 2 hold. Suppose payoff-relevant signaling occurs also at $h_{t+dt} \equiv (h_t, (\bar{a}, x, \mathbf{b}_t))$ for some $\bar{a}$ and all x in the support. Then such $\bar{a}$ must satisfy $\mathbb{E}_x \left[\tilde{u}(\bar{a}, p(h_t, (\bar{a}, x)), \underline{\theta}) \mid \underline{\theta}\right] = \tilde{u}(\underline{a}, \delta_{\underline{\theta}}, \underline{\theta})$, where $\tilde{u}(a, p, \theta)$ is defined by (3).*

Finally, Theorem 2 applies to all histories, including those off the equilibrium path. Applying it inductively starting from the root history, we obtain Corollary 2 below, which states that in the absence of payoff-irrelevant signaling, only the lowest type $L \equiv \min \Theta$ can ever separate from the rest, while the remaining ones can never separate from one another.

Corollary 2. *Suppose the conditions in Theorem 2 hold. In any limit equilibrium in which no payoff-irrelevant signaling happens, for any on-path history h_t, one of the following must hold:*

1. $S(h_t) = \Theta$;
2. $S(h_t) = \{\min \Theta\}$.

It is worth noting that there may be histories h_t at which $p(h_t)$ assigns arbitrarily small weight to the lowest type. So while this type can never be ruled out completely along the pooling path, asymptotically the receiver's belief may assign arbitrarily low weight to it. Notably, this implies that the mechanism of attrition of the lowest type can yield full separation asymptotically if there are only two types of the sender but not if there are more (but finitely many), which is an important takeaway, since many applied papers treat two-type models as proxies for more general settings.

At the same time, the discreteness of the type space is crucial to our analysis and, specifically, to the conclusion regarding the impossibility of full information revelation, even in the limit. If one were to adopt an interval type space instead, then full asymptotic revelation is possible again, at least in principle. An example of such outcome in the context of bargaining is presented by Fuchs and Skrzypacz (2010). Their equilibrium resembles the attrition equilibria of this paper, except in their equilibrium, an interval of lowest types separates away in every period instead of the single lowest type mixing between that and pooling. This discrepancy serves to illustrate the nontrivial implications of model discretization: if one attempts to compute equilibria of a continuous-time signaling model with an interval of types by approximating it with a discrete-time finite-type model, different approximations may yield qualitatively different results. In particular, if an interval type space is approximated by a grid that is too coarse relative to time discretization, the researcher would conclude that no asymptotic learning takes place in the discretized model, whereas it could take place in the continuous model.

Theorem 2 and its corollaries effectively provide a cookbook on how to construct an equilibrium with payoff-relevant signaling only. Suppose we want signaling to occur during the time interval $[0, T]$. Then in every period, along the pooling path we shall have two actions available to the sender: a separating action $\underline{a} \in A^* \left(h_t, \mathbf{b}, \underline{\theta}\right)$ only taken by the lowest type $\underline{\theta} \equiv \min \Theta$ and a pooling action $\bar{a}$ that satisfies the condition in Corollary 1 — the latter action will be played by $\underline{\theta}$ with some probability and by all other types for sure. Note that we have a degree of freedom in this construction: reputation from taking a pooling action depends on the probability with which type $\underline{\theta}$ separates in that given period. Hence by changing these probabilities we will be able to sustain different pooling actions $\bar{a}$ in equilibrium. To complete the construction, we need to verify that from time T onwards, the pooling strategy is such that $\underline{\theta}$ is exactly indifferent at T (or the last period before T) between separating and following this pooling path, and to verify that all other types always weakly prefer the pooling action to the optimal deviation. Section 5.2 provides an example of an equilibrium constructed using this cookbook, in the context of the price signaling model developed in the following section.

5. Applications

This section presents examples of applied models in different settings that fit our framework. It is meant to demonstrate some instances of models yielding additively and/or multiplicatively separable payoff functions that allow (MON) and (SC) to be verified with little effort. We begin by presenting in Section 5.1 a specific separable framework and show that there exist simple sufficient conditions for (MON) and (SC) within this framework. We then shift to exploring more specific applications. Section 5.2 sets up a simple price signaling model and verifies (both directly and using results from Section 5.1) that our results apply to it; it also constructs an example of an informative equilibrium to verify that they exist and to show more concretely how they can look. We show that in an extremely simple model, both high and low prices can be (suggestive, but not conclusive) signals of high product quality. Sections 5.3 and 5.4 then explore other applications, to labor market signaling and bargaining, respectively. The presentation in these sections is limited to verifying that (MON) and (SC) hold in the respective models; relation to the respective literatures is also discussed, but no equilibria are constructed and no novel results are claimed beyond those from Sections 3 and 4.

The examples also demonstrate different approaches that can be taken towards verifying (MON) and interpreting the assumption regarding the receiver's myopic behavior that underlies it. The price signaling in Section 5.2 uses "the receiver" to proxy a steady flow of short-lived (and hence myopic) consumers. The labor market signaling model in Section 5.3 relies on a *competitive* market of (potentially forward-looking) receivers to produce an effectively myopic response of the receiver-market to the sender's actions. Finally, the bargaining model is Section 5.4 considers a fully strategic receiver, showing that (MON) and (NDOC) do not require myopy or competitive markets and can be easily verified in some games with non-myopic receivers, if, e.g., non-trivial actions (acceptance of a proposed split) effectively end the game.

5.1. Separable settings

This section shows that if the agent's stage utility function $\tilde{u}(a, p, \theta)$ is separable in a specific way, then (MON) and (SC) can be easily verified. While this form of the utility function may appear restrictive, the remainder of Section 5 shows that it captures a wide range of settings, including classic signaling and bargaining models.

In particular, suppose the agent's stage utility function $\tilde{u}$ defined in (3) can be represented as

$$\tilde{u}(a,p,\theta) = \phi_0(a,p) + \phi_1(a,p)\psi(\theta) \tag{4}$$

for some collection of functions ϕ_0, ϕ_1, ψ. Then we can derive simple conditions on these three functions that are sufficient for (MON) and (SC) to hold. These conditions are given by the two respective propositions below.

Proposition 2. *If representation (4) applies, with $\psi(\theta) \geq 0$ and $\phi_0(a,p), \phi_1(a,p)$ weakly increasing in p, then $\tilde{u}$ satisfies (MON).*

As in the rest of the paper, monotonicity in p is understood with respect to the FOSD order on p (see Footnote 11). Note that $\phi_1(a,p)\psi(\theta)$ can be negative for some or all a, p, θ; without loss we let ϕ_1 absorb the negative sign in this case.

Proposition 3. *If representation (4) applies and outcomes x_t are uninformative at all h_t, then (SC) holds if $\psi(\theta)$ is strictly monotone in θ.*

We now continue to the more specific applications that use these results.

5.2. Price signaling

This section looks at a simple model of price signaling with product reviews. Price signaling is a phenomenon that is widespread in the real world — high-quality products may be priced at a premium to signal quality, or, conversely, they may offer more free trials or giveaways to help consumers learn about the product. Models exist that support both kinds of behavior. For example, Bagwell and Riordan (1991) show that if some consumers are initially informed of product quality while others learn from repeated purchases, then high and declining prices signal product quality. They also refer to empirical cases which support their conclusions. On the other hand, Vettas (1997) demonstrates that in the presence of social learning, high-type firm prices low on entry, gradually increasing the price afterwards, which is another pattern commonly observed in reality. While this apparent contradiction – that both high and low introductory prices can serve to signal quality – has been recognized in the literature (see, e.g., Kirmani and Rao, 2000), we are not aware of a theory that addresses it. The simple dynamic model below amends this and shows that both high and low prices are equally fit to serve as (suggestive yet inconclusive) signals of high quality.

Consider a long-lived firm i that faces a continuum of consumers $j \in [0,1]$ every period $t \in \mathcal{T} \equiv \{0, dt, 2dt, \dots\}$. The firm offers for sale a single product of privately known quality $\theta \in \Theta$, which is commonly known ex ante to be distributed according to $p_0 \in \mathtt{int}\Delta(\Theta)$. The marginal costs of production are zero. In every period the firm sets price a_t of its product and the consumers decide whether to purchase it or not. There is a unit continuum of short-lived consumers $j \in [0,1]$, renewed every period. Consumer j's payoff from buying the product is given by $\theta v_j - a_t$, where $v_j \sim i.i.d.U[0,1]$ is the consumer's value for quality. Payoff from not buying the product is zero. Consumer j then buys the item if and only if $\mathbb{E}[\theta|h_t] \geq \frac{a_t}{v_j}$; suppose hereinafter that all consumers follow this strategy and denote it as $\mathbf{b}$. With probability $1 - e^{-\phi dt} \approx \phi dt$ for $\phi \in [0,1)$ the population of period-t consumers generates an informative review $x_t = \theta$ which perfectly reveals the firm's quality and is observable by all future consumers. With complementary probability $e^{-\phi dt}$ no review is generated: $x_t = \varnothing$.[18] The period-t consumers observe the whole price path $\{a_s\}_{s\leq t}$ and product reviews as described below.

The stage profit of a firm of type θ in period t from setting price a_t is then given by

$$\tilde{u}(a_t, p_t, \theta) \equiv a_t \left(1 - \frac{a_t}{\mathbb{E}[\theta|p_t]}\right)_+. \tag{5}$$

Note that this profit function is independent of θ and is increasing in $\mathbb{E}[\theta|p]$ and hence increasing in p, so (MON) is satisfied.[19] This example features informative outcomes, hence Proposition 3 does not apply, but we can verify (SC) directly as follows.

Whenever a review arrives – which happens with probability $1 - e^{-\phi dt}$ in every period – the continuation play is trivial (cf. Lemma 2 in the Appendix): every type θ sets myopically optimal price $a^*(p,\theta) = a^*(\delta_\theta, \theta) = \frac{\theta}{2}$ and obtains payoff $\frac{\theta}{4}$ per period. Continuation payoffs at such histories are then given by $U(\mathbf{a}^*|h_t, \mathbf{b}, \theta) = \frac{\theta}{4r}$, which satisfies (SC). Now consider a history h_t with a non-degenerate belief $p(h_t)$. The type-θ firm's continuation payoff from following a given strategy $\mathbf{a}$ starting from any such history h_t in the reduced game is given by

$$U(\mathbf{a}|h_t, \mathbf{b}, \theta) = \sum_{s\in\mathcal{T}, s\geq t} e^{-(r+\phi)(s-t)} \left[e^{-\phi dt} \cdot \mathbf{a}_t \left(1 - \frac{\mathbf{a}_t}{\mathbb{E}[\theta|h_s]}\right)_+ dt + (1 - e^{-\phi dt}) \cdot \frac{\theta}{4r} \right],$$

For any $\mathbf{a}', \mathbf{a}''$ we then have

$$\begin{aligned} \mathcal{U}(\theta) &\equiv U(\mathbf{a}''|h_t, \mathbf{b}, \theta) - U(\mathbf{a}'|h_t, \mathbf{b}, \theta) \\ &= \sum_{s\in\mathcal{T}, s\geq t} e^{-(r+\phi)(s-t) - \phi dt} \left[\mathbf{a}''_t \left(1 - \frac{\mathbf{a}''_t}{\mathbb{E}[\theta|h''_s]}\right)_+ - \mathbf{a}'_t \left(1 - \frac{\mathbf{a}'_t}{\mathbb{E}[\theta|h'_s]}\right)_+ \right] dt, \end{aligned}$$

where histories h'_s and h''_s correspond to $\mathbf{a}'$ and $\mathbf{a}''$, respectively. The expression above is independent of θ, and hence satisfies (SC).

We conclude that the results of Theorems 1 and 2 therefore apply: the firm cannot conclusively signal its quality through prices (if it unable to commit to a price). In the remainder of this section we show, however, that partial signaling is still possible, and there exist informative attrition equilibria. In such an attrition equilibrium, type $\underline{\theta} \equiv \min \Theta$ mixes at every on-path history h_t s.t. $p(h_t) \neq \delta_{\underline{\theta}}$ between some pooling price profile $\mathbf{a}^p_t$ and separating to bliss price $a^*(\delta_{\underline{\theta}}, \theta) = \frac{\delta_{\underline{\theta}}}{2}$, whereas all other types $\theta \neq \underline{\theta}$ set the pooling price $\mathbf{a}^p_t$. We then show that Theorem 2 applies as well, meaning that the result translates to the case $|\Theta| > 2$.

We will be looking for an equilibrium of the game that satisfies (NDOC-P). In particular, let us construct an equilibrium in which prices are informative at every history where the firm's type is not perfectly known (i.e., $\nexists\theta : p(h_t) = \delta_\theta$). This would be the most informative equilibrium, since in the remaining histories informative signaling is trivially impossible. If the firm is believed to be bad ($p(h_t) = \delta_{\underline{\theta}}$) then there are no signaling motives — it sets price $a^*(\delta_{\underline{\theta}}, \underline{\theta}) \equiv \frac{\underline{\theta}}{2}$ and earns profit $u\left(a^*(\delta_{\underline{\theta}}, \underline{\theta}), \delta_{\underline{\theta}}, \underline{\theta}\right) = \frac{\underline{\theta}}{4}$.

We begin by deriving the pooling price a^p_t that renders the low type indifferent between separating and pooling. The former

[18] The analysis carries over fully to the case when the review arrival rate depends on the number of consumers who purchased the product in a given period.

[19] It is easier to see this property directly, but Proposition 2 applies as well.

yields the continuation value equal to

$$U(\mathbf{a}(\delta_{\underline{\theta}}, \underline{\theta}) \mid h_t, \mathbf{b}, \underline{\theta}) = \frac{dt}{1-e^{-rdt}} \frac{\underline{\theta}}{4} \approx \frac{\underline{\theta}}{4r},$$

since $1 - e^{-rdt} \approx rdt$ is a valid approximation when dt is small enough. Pooling, in turn, yields

$$U\left(\mathbf{a}^p \mid h_t, \mathbf{b}, \underline{\theta}\right) = \mathbf{a}_t^p \left(1 - \frac{\mathbf{a}_t^p}{\mathbb{E}[\theta|p_t]}\right) dt + e^{-rdt}\mathbb{E}_{x_t}\left[U\left(\mathbf{a}^p \mid h_{t+dt}, \mathbf{b}, \underline{\theta}\right)\right], \quad (6)$$

where $h_{t+dt} = (h_t, (\mathbf{a}_t^p, \mathbf{b}_t, x_t))$.

After a bad review $x_t = \underline{\theta}$ the consumers are sure that the firm is bad: $p(h_t, (\mathbf{a}_t^p, \mathbf{b}_t, \underline{\theta})) = \delta_{\underline{\theta}}$. After no review $x_t = \varnothing$ the consumers' belief is inconclusive, hence in our construction signaling should continue. This means type $\underline{\theta}$ must be indifferent between pooling and separating once again, so $U\left(\mathbf{a}^p \mid h_{t+dt}, \mathbf{b}, \underline{\theta}\right) = \frac{\underline{\theta}}{4r}$ as well. Finally, the indifference of type $\underline{\theta}$ at h_t implies that $U\left(\mathbf{a}^p \mid h_t, \mathbf{b}, \underline{\theta}\right) = U(\mathbf{a}(\delta_{\underline{\theta}}, \underline{\theta}) \mid h_t, \mathbf{b}, \underline{\theta}) = \frac{\underline{\theta}}{4r}$. Plugging all of these into (6), we obtain that the stage payoff from pooling must coincide with that from separating:

$$\frac{\underline{\theta}}{4} = \mathbf{a}_t^p \left(1 - \frac{\mathbf{a}_t^p}{\mathbb{E}[\theta|p_t]}\right).$$

The solution to the above is given by

$$\mathbf{a}_t^p = \frac{\mathbb{E}[\theta|p_t]}{2}\left[1 \pm \sqrt{1 - \frac{\underline{\theta}}{\mathbb{E}[\theta|p_t]}}\right]. \quad (7)$$

Therefore, for a fixed $\mathbb{E}[\theta|p_t]$ we have two equivalid candidates for the pooling price $\mathbf{a}_t^p$. The negative root corresponds to signaling by setting a low price — below L's preferred price (for any reputation). Such pooling price can be seen as low entry pricing à la (Vettas, 1997). The positive root, conversely, corresponds to the price well above the myopic optimum and signaling through exclusivity, in the spirit of Bagwell and Riordan (1991). Further, the pooling output $\mathbf{a}_t^p$ (whichever root to (7) we choose) is a function of the seller's reputation

$$p(h_{t+dt}) = \frac{p(h_t)}{1 - \lambda_t(1 - p(h_t))}, \quad (8)$$

where λ_t is the probability with which type $\underline{\theta}$ separates at h_t. In particular, we can choose these probabilities freely and construct an equilibrium for arbitrary λ_t, so long as an appropriate pooling action $\mathbf{a}_t^p$ is available to guarantee the indifference. Therefore, in general the pooling price path is indeterminate given the market conditions (seller's reputation), so inferring whether price signaling is taking place in a given market by looking at price data is a daunting task. This point was originally raised by Kaya (2013) in relation to advertising expenditures.

To complete the equilibrium description we only need to argue that the pooling price $\mathbf{a}_t^p$ is optimal for the seller of type $\theta \neq \underline{\theta}$. His continuation value from following the pooling strategy is

$$U\left(\mathbf{a}^p \mid h_t, \mathbf{b}, \theta\right) = \frac{\underline{\theta}}{4}dt + e^{-rdt}\left[(1 - e^{-\phi dt})\frac{\theta}{4r} + e^{-\phi dt} U\left(\mathbf{a}^p \mid h_{t+dt}, \mathbf{b}, \theta\right)\right].$$

Indeed, his stage payoff is the same as for $\underline{\theta}$ (the two types only differ in the reviews they get), and in case a good review is generated at h_t, he will be receiving $\frac{\theta}{4}$ in every future period. The same applies to any history h_{t+dt} with $p(h_{t+dt}) \in \mathtt{int}\Delta(\Theta)$, which allows us to conclude that $U\left(\mathbf{a}^p \mid h_t, \mathbf{b}, \theta\right) = \frac{1}{4r} \cdot \frac{r\underline{\theta}+\phi\theta}{r+\phi}$. Setting any price other than a_t^p results in $p(h_{t+dt}) = \delta_{\underline{\theta}}$ (by (NDOC-P)), and hence yields value of at most $\frac{\underline{\theta}}{4r} < \frac{1}{4r} \cdot \frac{r\underline{\theta}+\phi\theta}{r+\phi}$. Therefore, pooling is indeed optimal for θ. All of the above proves the following proposition.

Proposition 4. *The following constitutes an equilibrium of the price signaling game for any profile of $\{\lambda_t\}_{t\in\mathcal{T}}$. At any history $h_t \in \mathcal{H}$:*

1. *if $p(h_t) = \delta_\theta$ then type θ' plays $a_t = \frac{\theta'}{2}$, and $p(h_{t+dt}) = p(h_t)$ for all $h_{t+dt} \succ h_t$;*
2. *if $p(h_t) \in \mathtt{int}\Delta(\Theta)$ then:*
   - *(a) type $\theta = \underline{\theta}$ plays $\mathbf{a}_t^p$ as given by (7) w.p. $1 - \lambda_t$ and $a^*(\delta_{\underline{\theta}}, \underline{\theta}) = \frac{\underline{\theta}}{2}$ w.p. λ_t;*
   - *(b) all types $\theta \neq \underline{\theta}$ play $\mathbf{a}_t^p$ with probability one;*
   - *(c) the consumers' belief $p(h_{t+dt})$ is updated according to (8) after $\mathbf{a}_t^p$ and $x_t = \varnothing$, to $p(h_{t+dt}) = \delta_\theta$ after $x_t = \theta$, and $p(h_{t+dt}) = \delta_{\underline{\theta}}$ for all other $h_{t+dt} \succ h_t$.*

To summarize, the price signaling model presented in this section does, despite being highly stylized, demonstrate that informative price signaling is possible without commitment, cost advantages, and with or without consumers learning from experiences. The signaling price may be either low (e.g., in the form of free trials or frequent sales) or inefficiently high (excluding most consumers) as a result of sunspots. Further, multiple informative equilibria exist that differ in the speed of separation. Due to the multiplicity issues above, the empirical identification of price signaling in a given market is a daunting task.

5.3. Labor market signaling

In this section we revisit the classic labor market signaling model (Spence, 1973), which sparked the original discussion around dynamic signaling (Nöldeke and van Damme, 1990a; Swinkels, 1999). In the dynamic version of this model, a long-lived candidate of privately known ability $\theta \in \Theta \subseteq \mathbb{R}_+$ acquires costly and, w.l.o.g., unproductive education in an attempt to signal her ability to potential employers. A high-ability worker is more productive on the job and can thus bargain for a higher wage, while also having lower cost of education than a low-ability worker. In every period $t \in \mathcal{T} \equiv \{0, dt, 2dt, \dots\}$ she chooses education intensity $e \in E \subset \mathbb{R}$. The stage (ongoing) cost of education is given by $c(e|\theta) \equiv l(e) \cdot m(\theta)$, where $l(e)$ is increasing in e with $l(0) = 0$, and $m(\theta)$ is strictly decreasing in θ.

There is a population of homogeneous competitive employers, who observe the full history of the candidate's education choices and grades. In every period they simultaneously offer employment contracts to the candidate.[20] After observing all contracts, the candidate may accept at most one of them. If a contract is accepted, in every future period the candidate receives wage $w \cdot dt$, where w is as specified in the contract. Let $d \in \{0, 1\}$ denote the worker's acceptance decision — whether she chooses to accept an offer in a given period or not. If the candidate chooses to accept, she would trivially find it optimal to choose the highest-wage contract.

W.l.o.g., let θ be equal to the candidate's on-the-job productivity (so her output is $\theta \cdot dt$ per period). This means that at any history h_t, all competitive firms will offer the same wage $w(h_t) = \mathbb{E}[\theta|h_t, d(h_t) = 1]$. A history here consists of the candidate's past actions: $h_t = \{d_s, e_s\}_{s\in\mathcal{T},s<t}$. The implied timing in the stage game at any history h_t is:

20 Suppose that the offers are made privately and so are not observed by other firms. Nöldeke and van Damme (1990a) and Swinkels (1999) show that otherwise – if the offers are public, – a tacit collusion equilibrium with perfect separation can be sustained, see footnote 3 for details.

1. if the candidate has accepted an offer with wage w in the past, she receives the contracted wage $w \cdot dt$, and the game proceeds to the next period. Otherwise,
2. firms make wage offers $w(h_t)$ to the candidate;
3. the candidate decides $d(h_t)$ whether to accept the highest-wage contract. If $d(h_t) = 1$ then the game continues to the next period. Otherwise,
4. the candidate chooses $e(h_t)$, her education effort in the current period, and the game continues to the next period.

Once a candidate has accepted an offer, at all future histories set $d(h_t) = e(h_t) = 0$.

In such a game, the candidate's payoff from following some given strategy $\mathbf{a} = \{d(h), e(h)\}_{h\in\mathcal{H}}$ conditional on some history $h_t \in \mathcal{H}$ at which she has not yet accepted an offer is given by

$$U(\mathbf{a}|h_t, \theta) \equiv \sum_{s\in\mathcal{T}, s\geq t} e^{-r(s-t)} \left[d(h_s)\frac{w(h_s)}{r} - (1 - d(h_s)) \cdot c\,(e(h_s)|\theta)\, dt \right] dt,$$

since accepting an offer at h_t is equivalent to receiving a lumpsum payoff of $w(h_t)/r$. The stage utility function can thus be framed as (4) using functions $\phi_0((d, e), p) = d\frac{w(p,d)}{r}$, $\phi_1((d, e), p) = -(1 - d)l(e)$, and $\psi(\theta) = m(\theta)$. Here $m(\theta)$ is strictly monotone, $\phi_1((d, e), p)$ is independent of p, and $\phi_0((d, e), p)$ is weakly increasing in p, since $w(p, d) = \mathbb{E}[\theta|p, d = 1] = \mathbb{E}\big[\mathbb{E}[\theta|p] \mid d = 1\big]$ by the law of iterated expectations, and $\mathbb{E}[\theta|p]$ is strictly increasing in p. Therefore, by Propositions 2 and 3, (MON) and (SC) hold.

Verifying (MON) directly is not much different from invoking Proposition 2 above, same as in Section 5.2. It is easy to verify (SC) directly as well: for any pair of strategies $\mathbf{a}'$, $\mathbf{a}''$, the function $\mathcal{U}(\theta) \equiv U(\mathbf{a}''|h_t, \theta) - U(\mathbf{a}'|h_t, \theta)$ can be written as

$$\begin{aligned}\mathcal{U}(\theta) = &\sum_{s\in\mathcal{T}, s\geq t} e^{-r(s-t)} \Bigg[\left(d'(h_s) - d''(h_s)\right) \frac{w(h_s)}{r} - m(\theta)dt \\ &\cdot \left[\left(1 - d'(h_s)\right) \cdot l\left(e'(h_s)\right) - \left(1 - d''(h_s)\right) \cdot l\left(e''(h_s)\right) \right]\Bigg], \\ =&C_1 + C_2 m(\theta)\end{aligned}$$

for some C_1, C_2 that depend on $\mathbf{a}'$, $\mathbf{a}''$, h_t. Since $m(\theta)$ is strictly decreasing in θ, $\mathcal{U}(\theta)$ is either strictly monotone, or constant, depending on whether the coefficient at $m(\theta)$ is positive, negative, or zero. Since strategies $\mathbf{a}'$, $\mathbf{a}''$ and history h_t were arbitrary, this means that payoff function $U(\mathbf{a}|h, \theta)$ satisfies (SC) and our results apply. The candidate will only be able to signal her ability via attrition: in an informative limit equilibrium – if it exists – the lowest type would drop out of education and take up a job at a random date, while all other types go through the whole education process prescribed by the equilibrium and only accept the job afterwards. The latter path would also be taken by some low-ability candidates. Swinkels (1999), however, shows that pooling is the only equilibrium in this model, i.e., even the attrition-form informative equilibria do not exist.

5.4. Bargaining

In this section we consider a simple model of bilateral bargaining, in which one party's valuation is commonly known, while another party's valuation is their private information. Consider two players, a buyer B and a seller S, who interact repeatedly in every period $t \in \mathcal{T} \equiv \{0, dt, 2dt, \ldots\}$. There is a unit of indivisible good of some quality $\theta \in \Theta \subseteq \mathbb{R}_+$ privately known by the seller. The buyer has some commonly known prior belief $p_0 \in \Delta(\Theta)$ about quality θ, but does not observe the realization of θ. The seller is initially in possession of the item and values it at $c(\theta)$, assumed to be either constant, or strictly increasing. The buyer's valuation is $v(\theta)$. In every period one of the players $x_t \in \{B, S\}$ is chosen as the proposer (the choice rule can be random or deterministic, and/or history-dependent) and can offer a price $y_t \in \mathbb{R}_+$. The other player then decides $z_t \in \{0, 1\}$ whether to accept the offer. If the offer is accepted, the item is traded at that price and the game ends (which can be emulated by setting $y = z = 0$ for both players at all subsequent histories). Otherwise the game continues to the next period. Both players discount the future at rate r.

History $h_t = \{x_s, y_s, z_s\}_{s\in\mathcal{T}, s<t}$ in this game is given by the players' past offers and responses, as well as identities of the proposing player. The players' pure strategies are given by $\mathbf{a} = \{y_S(h), z_S(y, h)\}_{h\in\mathcal{H}, y\in\mathbb{R}_+}$ for the seller (conditional on type θ) and $\mathbf{b} = \{y_B(h), z_B(y, h)\}_{h\in\mathcal{H}, y\in\mathbb{R}_+}$ for the buyer respectively, where $y_i(h)$ is the player's proposal if they are selected, and $z_i(y, h)$ is their response to the opponent's proposal y.[21] Since the buyer is neither short-lived, nor myopic, this setting does not fall within the scope of the model defined in Section 2. However, as mentioned in Footnote 12, those assumptions are only needed to shut down the folk theorem effects, so that (MON) is a non-vacuous assumption. In the context of this bargaining model, this can be achieved by assuming instead that the buyer follows a *Markov* strategy $\{y_B(p), z_B(y_S, p)\}$, which treats belief p_t as a sufficient statistic of the whole history h_t, and that both $y_B(p)$ and $z_B(y_S, p)$ are increasing in p (where monotonicity is interpreted in the same semi-strong sense with respect to the F.O.S.D. order on p as in the (MON) assumption).

Then denoting the probability with which the buyer is selected to propose at h_t as $\chi(h_t) \equiv \mathbb{P}\left(x(h_t) = B \mid h_t\right)$, the seller's expected payoff from following some strategy $\mathbf{a}$ conditional on the buyer's strategy $\mathbf{b}$ and some history $h_t \in \mathcal{H}$, by which no offer had been accepted, is given by

$$\begin{aligned}&U(\mathbf{a}|h_t, \theta, \mathbf{b}) \\ &\equiv \mathbb{E}\Bigg[\sum_{s\in\mathcal{T}, s\geq t} e^{-r(s-t)} \left[\chi(h_s) z_S(y_B(p_s), h_s)\left(y_B(p_s) - c(\theta)\right) \right.\\ &\left.+ \left(1 - \chi(h_s)\right) z_B(y_S(h_s), p_s)\left(y_S(h_s) - c(\theta)\right) \right]\Bigg].\end{aligned} \tag{9}$$

It is easy to see that $U(\mathbf{a}|h_t, \theta, \mathbf{b})$ is linear in $c(\theta)$, hence so is $\mathcal{U}(\theta) \equiv U(\mathbf{a}''|h_t, \theta, \mathbf{b}) - U(\mathbf{a}'|h_t, \theta, \mathbf{b})$ for any pair of seller's strategies $\mathbf{a}'$, $\mathbf{a}''$ and any given buyer's strategy $\mathbf{b}$. Therefore, (SC) holds in this model as long as $c(\theta)$ is constant or strictly increasing, regardless of $v(\theta)$. Since both $y_B(p)$ and $z_B(y_S, p)$ are assumed to be increasing in p, (MON) holds as well.[22]

Alternatively, we could once again invoke Propositions 2 and 3, since (9) can, given the buyer's equilibrium strategy $\{y_B(p), z_B(y_S, p)\}$, be represented in terms of (4) with

$$\phi_0((y_S, z_S), p|\chi) = \chi z_S y_B(p) + (1 - \chi)z_B(y_B, p)y_S,$$
$$\phi_1((y_S, z_S), p|\chi) = -\chi z_S - (1 - \chi)z_B(y_B, p),$$

[21] Note that regardless of x_t, the seller in this example always acts as "the sender" and the buyer as "the receiver". In particular, the setup of Section 2 implies that in this setting, the seller chooses the time-t action $(y_S(h_t), z_S(y, h_t))$ before knowing time-t role allocation x_t, and the buyer chooses $(y_B(h_t), z_B(y_S(h_t), h_t))$ after learning both x_t and the seller's action. However, this deviation from the standard offer-response procedure is not meaningfully impactful when the buyer is restricted to Markov strategies.

[22] Representation (3) assumes stage utility only depends on x_t via $p(h_t, a_t, x_t)$, whereas the stage utility in (9) explicitly depends on x_t — or, after taking expectations, on its distribution $\chi(h_t)$. However, in common settings of buyer-proposing ($\chi(h_t) \equiv 1$), seller-proposing ($\chi(h_t) \equiv 0$), and even alternating-offer ($\chi_0 \in \{0, 1\}$, $\chi(h_{t+dt}) = 1-\chi(h_t)$) bargaining, we can treat $\chi(h_t)$ as an exogenous parameter.

both of which are monotone in p, and $\psi(\theta) = c(\theta)$, which is also strictly monotone.[23]

The main implication of Theorems 1 and 2 for this bargaining model is that the perfectly efficient allocation is unattainable for any finite delay. However, our results make no statements on whether the Coase conjecture is realized and all seller types trade at $t = 0$ at the lowest acceptable price, or delay can be used as an (imperfect but effective) screening device.[24] The exact shape that the equilibria can take depends on $c(\theta)$, $v(\theta)$, and $x(h_t)$. For example, if $c(\theta) = c$ for some constant c and $v(\theta) \geq c + g$ for some "gap" value g and the uninformed buyer makes all the offers — then the Coase conjecture realizes, and in the unique equilibrium, all types of the seller pool on accepting the lowest price (Gul, Sonnenschein, and Wilson, 1986). In the alternating-offer scenario with the same assumptions, on the other hand, a delay equilibrium is possible, where attrition occurs for some finite time. During that time, the lowest-type seller sells with some probability, and after that all the remaining sellers sell at the same price (Ausubel and Deneckere, 1998, as presented in Ausubel, Cramton, and Deneckere, 2002). If $v(\theta) = v \geq c(\theta)$ for some constant v and all θ, then in the seller-proposing game pooling is, again, the only equilibrium, but not in the Coase conjecture sense. In that equilibrium all types of the seller propose v in every period and obtain surplus $v - c(\theta)$, with their offer being accepted straight away (Ausubel and Deneckere, 1989). For a review of these and other results on bargaining under incomplete information, see Ausubel et al. (2002). As applied to more recent literature, our framework also subsumes the model of Daley and Green (2020), who explore a setting in which public news arrive during the bargaining process.

Our results further imply that any non-pooling ("non-Coasian") equilibrium would be of the form that the buyer's offers gradually increase, the high-type seller rejects all initial low offers, whereas the low-type seller randomly accepts. The seller's expected type (reputation) conditional on the no-sale path is then gradually increasing over time. This is in line with the bargaining/dynamic adverse selection literature (see, e.g., Vincent, 1998; Deneckere and Liang, 2006; Daley and Green, 2012 for some examples). Conversely, Kaya and Kim (2018) present a model that does not fit our framework. They show that if the informed strategic seller faces a sequence of short-lived buyers with noisy private signals, the opposite dynamic may arise, with time on the market becoming a negative signal about the seller's type. Private signals on the receivers' side are the main driving force behind this reversal and the part of their model not allowed in our framework.

6. Equilibrium existence

This section provides the conditions under which pooling and informative equilibria exist in our setting. While Madrigal et al. (1987) and Nöldeke and van Damme (1990b) claim that (NDOC) may lead to equilibrium nonexistence, their argument concerns perfectly separating equilibria, and hence aligns with our message well. Their constructions involve (different) dominant strategies for every type of the sender, meaning that no pooling equilibrium can exist. Our main focus, in turn, is on the mixed-strategy attrition equilibria, which may or may not exist regardless of pooling equilibria.

The results below provide a simple geometric characterization of the set of payoffs achievable in equilibrium. While they do not guarantee that this set is nonempty in any given game, they provide a simple characterization of attainable equilibria. This characterization can then be used to test whether a given type of equilibrium exists in a given game. This characterization is only provided for the case of outcomes x being uninformative, since, while informative outcomes have no effect on equilibrium properties described above, they create substantial difficulties for the characterization below.

Before we proceed, some new notation should be defined. With abuse of notation, let $\tilde{u}(a, p) \equiv (\tilde{u}(a, p, \theta))_{\theta\in\Theta}$ denote a vector of stage payoffs for all types from a given action a and given reputation p. Further, given a prior belief p_0 and denoting the lowest type as $\underline{\theta} \equiv \min \Theta$, let $\tilde{p}(\nu) \equiv \frac{p_0 - \nu\cdot p_0(\underline{\theta})\cdot\delta_{\underline{\theta}}}{1-\nu\cdot p_0(\underline{\theta})}$ for $\nu \in [0, 1]$ denote the receiver's belief if share ν of the lowest types decide to separate, and $P_s \equiv \{\tilde{p}(\nu) \mid \nu \in [0, 1]\}$ the set of such beliefs. This would be the set of beliefs that can be attained on path in separating equilibria satisfying the characterization in Theorem 2 (assuming no payoff-irrelevant signaling). I.e., P_s includes the prior belief p_0 and all beliefs p that assign smaller weight to $\underline{\theta}$ than p_0 but induce the same likelihood ratios for all other types: $\frac{p(\theta')}{p(\theta'')} = \frac{p_0(\theta')}{p_0(\theta'')}$ $\forall\theta', \theta'' \in \Theta\backslash\{\underline{\theta}\}$. Let $\underline{u}(\theta) \equiv \max_{a\in A} \tilde{u}(a, \delta_{\underline{\theta}}, \theta)$ denote the "outside option" of type θ, which consists of the payoff from separating from the prescribed equilibrium strategy and playing the myopically optimal action while having the worst reputation. If type θ's payoff $\tilde{u}(a, p)$ (strictly) exceeds $\underline{u}(\theta)$, we call it (strictly) *individually rational* (IR).

Our main object of interest is $W \equiv \left\{(u, \nu) \mid \forall\theta : u(\theta) \geq \underline{u}(\theta); \exists a \in A : u = \tilde{u}\left(a, \tilde{p}(\nu)\right)\right\}$, which denotes the collection of the sender's IR stage payoffs compatible with some action profile and some reputation attainable in principle in attrition equilibria. Let $a(u, \nu) \in A$ denote the action corresponding to a given $(u, \nu) \in W$. Of particular interest would be the two sections of W. First, let $W_0 \equiv \left\{(u, \nu) \in W \mid \nu = 0; \forall\theta : u(\theta) > \underline{u}(\theta)\right\}$ denote the set of strictly IR stage payoffs attainable under pooling (i.e., under the prior belief p_0 — equivalently, under $\nu = 0$). Second, let $\underline{W} \equiv \left\{(u, \nu) \in W \mid u(\underline{\theta}) = \underline{u}(\underline{\theta}); \forall\theta \neq \underline{\theta} : u(\theta) > \underline{u}(\theta)\right\}$ denote the section of the payoff set W corresponding to the outside option payoff for type $\underline{\theta}$ and strictly IR payoffs. As usual, $co(W_0)$ and $co(\underline{W})$ denote the convex hulls (i.e., the smallest convex supersets) of sets W_0 and $\underline{W}$, respectively.

One final technical restriction that we need to impose relates to the richness of the game and states that the two actions described in Corollary 1 – the pooling action that sets type $\underline{\theta}$ on the reservation utility given any reputation feasible in equilibrium and his myopically optimal separating action – should be distinct. In this case we say the game *admits separation*. This condition is formalized as follows:

> **(AS)** For all h_t and $(u, \nu) \in \underline{W}$, there exists an optimal $\mathbf{b}$ such that $A^*\left(h_t, \mathbf{b}, \underline{\theta}\right) \backslash \{a(u|\nu)\}$ is nonempty.

Worth noting that this condition is trivially satisfied if we consider a cheap talk extension of the baseline game, in which the sender has access to payoff-irrelevant messages (in addition to payoff-relevant actions). Given this new notation, we can make equilibrium existence claims formalized by Theorem 3 below. Part 1 describes pooling equilibria, while part 2 relates to potentially informative equilibria.

[23] If $c(\theta)$ is constant, Proposition 3 can be applied by setting $\tilde{u}(a, p, \theta) = \phi_0(a, p)$.

[24] There is a subset of literature exploring possible reasons for delay in bargaining (e.g., Abreu and Gul, 2000 and Feinberg and Skrzypacz, 2005). Our focus is different: instead of demonstrating sufficient conditions within the bargaining models under which *delay is the only equilibrium*, this paper considers a much more general class of models and aims to provide weaker sufficient conditions under which *attrition is the only informative (limit) equilibrium*. In particular, when applied to bargaining, our model does not rule out the possibility of instant agreement, in which all types of the sender pool on the same action, but claims instead that either instant agreement, or delay by the high types must take place in equilibrium.

Theorem 3. *Suppose outcomes x are uninformative and given some optimal* **b**, *(MON) and (SC) hold. Then the following are true:*

1. *For any $u \in co(W_0)$ and any $\epsilon > 0$, exists Δ such that if $dt < \Delta$, then there exists a pooling equilibrium that yields expected discounted utility within ϵ of $\frac{u(\theta)}{r}$ to the sender of type θ for all $\theta \in \Theta$ and satisfies (NDOC).*
2. *If (AS) holds, then for any $(u, v) \in co(\underline{W})$ and any $\epsilon > 0$, exists Δ such that if $dt < \Delta$, then there exists an equilibrium that yields expected discounted utility within ϵ of $\frac{\underline{u}(\theta)}{r}$ to the sender of type θ for all $\theta \in \Theta$ and satisfies (NDOC).*

Note that the statements relate to the PBE, meaning that the respective limit equilibria also exist. Part 1 of Theorem 3 characterizes the set of payoffs attainable in pooling equilibria, where all types of the agent pool on the same action (or a sequence/distribution of actions) on equilibrium path. Part 2 of Theorem 3 is more exciting, since it characterizes the set of payoffs attainable in semi-separating/attrition equilibria, which are ignored in the existing literature. Specifically, in the equilibria constructed in the proof, there exists at most a finite collection of periods, in which the lowest type $\underline{\theta}$ mixes between pooling with other types and separating with positive probability (and pooling for sure in all other periods).

We make no claims about the sets W_0 and $\underline{W}$ being non-empty. Indeed, Madrigal et al. (1987) and Nöldeke and van Damme (1990b) provide an example of a game with no pooling equilibria (empty W_0). Theorem 3 should hence be seen not as a proof that equilibria of the respectivy type necessarily exist, but rather as a characterization of when they do. The characterization in the theorem is meant to be straightforward enough to be immediately applicable to any game of interest.

The argument in the proof of part 2 of the theorem shows, in particular, that any separation payoff can be achieved by an equilibrium, in which (partial) separation happens in a finite number of periods (to be precise, at most $N+2$ separation periods are necessary). However, the construction presented therein applies equally well to equilibria with arbitrarily many separation periods, as evidenced by the following corollary. It shows that any path of feasible actions and reputations can be implemented in equilibrium if dt is small enough.

Corollary 3. *Suppose outcomes x are uninformative and given some optimal* **b**, *(MON), (SC), and (AS) hold. For any sequence $\{u_k, v_k, \tau_k\}_{k\in\mathbb{N}}$ such that $(u_k, v_k) \in \underline{W}$ for all k, $\tau_1 = 0$, $\exists \Delta_t$: $\min_k\{\tau_k - \tau_{k-1}\} \geq \Delta_t$, and v_k and τ_k are strictly increasing, there exists Δ such that for any $dt < \Delta$ there exists an attrition equilibrium that satisfies (NDOC) and is such that on the equilibrium path, in any period $t \in \mathcal{T}$ s.t. $\tau_k \leq t < \tau_{k+1}$:*

- *all sender types $\theta \in \Theta$ play pooling action $a(u_k, v_k)$;*
- *the receiver's belief is $\tilde{p}(v_k)$;*
- *if $t \geq \tau_k > t - dt$, then type $\underline{\theta}$ plays some separating action with positive probability.*

7. Conclusion

This paper explores a model of dynamic signaling with observable actions. In this model a single privately-informed agent takes an action every period, but cannot commit to future actions. The receiver tries to infer the agent's information from his actions, and the receiver's opinion is relevant to the agent's payoff. The existing literature has implied that signaling is impossible in such setting, unless strong assumptions about off-equilibrium-path beliefs are adopted. This paper confirms the negative result that *perfect separation* is impossible in such a setting. However, it provides a novel positive result, showing that *imperfect signaling* is possible under reasonable off-path beliefs. Further, we show that such signaling must necessarily happen through attrition of the lowest type of the agent. In this attrition scenario, all types pool on the same action (or split across a number of different yet payoff-equivalent actions), while the lowest type also plays some separating action with positive intensity. Finally, the paper characterizes the set of payoffs attainable in such equilibria.

The paper identifies sufficient conditions, under which the results hold. These include a restriction on monotonicity of the agent's preferences with respect to his reputation, and a restriction on the off-path beliefs to be reasonable. In the case of many types single-crossing of agent's preferences must also hold. The latter is, arguably, the strongest of the three assumptions and the one that would be most difficult to verify in the applied work. However, the paper presents a number of applied signaling models to demonstrate that this notion of single-crossing can be used in applied work. Future work could involve the exploration of simpler notions of single-crossing that would work for dynamic signaling games.

Importantly, the paper assumes the receiver behaves myopically. The main issue that arises when both the agent and the receiver are strategic is the folk theorem, which says that any individually rational payoff for either player can be sustained in equilibrium for dt small enough. The consequence of this equilibrium multiplicity is that (MON) and (SC) almost never hold across the whole spectrum of equilibria in a given setting. The solution, if one wishes to explore settings with a strategic long-lived receiver, is to focus on some selected equilibria — i.e., to restrict attention to some fixed strategy (or a class of strategies) **b** of the receiver and to test (MON) and, if necessary, (SC) against those strategies. This approach has been demonstrated in this paper in the bargaining application. Exploration of the specific equilibrium selection criteria that could yield favorable results lies beyond the scope of this paper, but could be another prospective direction for future research.

Declaration of competing interest

Declarations of interest: none.

Data availability

No data was used for the research described in the article.

Appendix. Proofs and supplementary results

A.1. Proofs: Preliminaries

The first observation states that once there is no need for signaling any more – i.e., when the receiver's belief assigns probability 1 to some type of the agent – there are no reasons for the agent to steer away from the myopically optimal action. Hereinafter, $h_s \succeq h_t$ denotes weak succession and means "either $s > t$ and $h_s \succ h_t$, or $s = t$ and $h_s = h_t$".

Lemma 2. *For any dt, in any equilibrium that satisfies (NDOC), at any $h_t \in \mathcal{H}$, if $|S(h_t)| = 1$ then for all θ and all $h_s \succeq h_t$: $\alpha_\theta\left(A^*(h_s, \mathbf{b}, \theta) \mid h_s\right) = 1$.*

Proof. By (NDOC), for all $h_s \succeq h_t$: $p(h_s) = \delta_{S(h_t)}$. In particular, $p(h_s)$ is independent of all actions and outcomes during $[t, s)$, hence the receiver's best response $\mathbf{b}(h_s, a_s, x_s)$ is the same at all such h_s (given a_s and x_s). Therefore, the solution to (2) is given by pointwise maximization of the stage utility. □

Lemma 2 above is the direct consequence of (NDOC): actions cannot change a degenerate belief under this assumption, hence the myopic optimum is chosen. This captures the main tension between signaling and sequential rationality: signaling requires sticking to the costly action over an extended period of time, while sequential rationality as captured by Lemma 2 pushes against that when no further signaling concerns are present. The remaining statements formalize this intuition. However, before proceeding any further, we use Lemma 2 to prove Lemma 1 from the text.

Proof of Lemma 1. Denote the original equilibrium as $(\alpha', \mathbf{b}', p')$. Construct the new equilibrium $(\alpha'', \mathbf{b}'', p'')$ by copying the strategies and beliefs the original equilibrium prescribes for all on-path histories h_t. For all off-path histories h_t, set $p''(h_t) = \delta_{\underline{S}(h)}$, where h is the last on-path history preceding h_t. For the receiver's strategy $\mathbf{b}''(h_t, \alpha''(h_t), x_t)$ at off-path h_t, take any optimal strategy that is consistent with on-path play with respect to the "no sunspots" assumption. Finally, the agent's strategy profile $\alpha''(h_t)$ for off-path histories h_t is set in conformance with Lemma 2.

Belief profile p'' will then satisfy (NDOC-P) and be consistent with the strategy profile α''. The receiver's strategy $\mathbf{b}''$ is, by construction, optimal at all histories. The agent's strategies α''_θ will be optimal at off-path histories by Lemma 2. Optimality of α''_θ for type θ at any on-path history h_t can be verified by observing that value $V(a|h_t, \mathbf{b}'', \theta)$ is the same as in the original equilibrium for all on-path actions $a \in A(h_t)$ and weakly smaller for off-path actions $a \in A \backslash A(h_t)$ (since any such action generates a pointwise lower path of future reputation and (MON) holds). I.e., the choice between any pair of on-path actions is unaffected by the off-path modifications, while deviations to off-path actions are less appealing in the new equilibrium. We conclude that $(\alpha'', \mathbf{b}'', p'')$ is an equilibrium. □

A.2. Proofs: Two types

Proof of Theorem 1. Statement 1. Suppose first, by way of contradiction, that there exist $h_t \in \mathcal{H}$ and $a' \in A$ such that $\alpha_H(a'|h_t) > 0$ but $\alpha_L(a'|h_t) = 0$. Then $p(h_t, a', x_t) = \delta_H$ for any $x_t \in X$ by (NDOC-P). By playing a' at h_t the low type receives the highest possible continuation utility after t (since by Lemma 2 he can play the myopically optimal action thereafter), while by following the equilibrium path he receives strictly less due to (MON). The utility is bounded, hence for dt small enough deviating to a' at h_t is optimal for L — a contradiction.[25]

Payoff-equivalence is shown as follows: for any two $a', a'' \in A$ such that $\alpha_H(a'|h_t) > 0$ and $\alpha_H(a''|h_t) > 0$ it must be that $V(a'|h_t, \mathbf{b}, H) = V(a''|h_t, \mathbf{b}, H)$, otherwise the high type would only play one of the actions and not the other. The first part of the argument showed that $\alpha_L(a'|h_t) > 0$ and $\alpha_L(a''|h_t) > 0$, hence $V(a'|h_t, \mathbf{b}, L) = V(a''|h_t, \mathbf{b}, L)$ by the same logic.

Statement 2. Begin with the first part (that $a' \in A^*(h_t, \mathbf{b}, L)$). For any such a' that $\alpha_H(a'|h_t) = 0$ and $\alpha_L(a'|h_t) > 0$ and any outcome x_t, we have $p(h_{t+dt}) = \delta_L$, where $h_{t+dt} = (h_t, a', x_t)$. By Lemma 2, at all $h_s \succ h_t$, only bliss actions are played: $a_s \in A^*(h_t, \mathbf{b}, L)$. If $a' \notin A^*(h_t, \mathbf{b}, L)$ then playing a bliss action $a'' \in A^*(h_t, \mathbf{b}, L)$ at h_t instead – and continuing with a_s at all subsequent histories – yields a strictly higher stage payoff at h_t and the same continuation payoff. Hence playing a' at h_t was not optimal.

The second part of the second statement follows from the same argument as did payoff equivalence for L in the first statement. □

[25] More formally, by belief consistency and rationality there must exist $a'' \in A$ s.t. $\alpha_L(a''|h_t) > 0$ and $p(h_t, a'', x_t)(H) < p(h_t, a', x_t)(H)$ for all x_t, hence $V((h_t, a', x_t), \mathbf{b}, L)$ is bounded away from $V((h_t, a'', x_t), \mathbf{b}, L)$ for all $h_t, x_t, \mathbf{b}$. Therefore, there exists $\Delta(h_t, \mathbf{b})$ s.t. if $dt < \Delta(h_t, \mathbf{b})$ then period-t gains for L from playing a'' compared to a' cannot outweigh the losses from $t+dt$ onwards.

A.3. Proofs: Finite types

Before proceeding to the proofs of Proposition 1 and Theorem 2, it is convenient to spin parts of them off into supplementary lemmas. We begin by arguing in Lemma 3 that at no history can actions lead to separation of types into disjoint sets that can be compared by a strong set order — unless one of these sets is a singleton coinciding with the lower bound of the other set. In particular, we show that sets of types in the support of two different actions have to necessarily overlap (not in the sense of having common elements, but in the sense of upper and lower bounds).

Lemma 3. *Fix any limit equilibrium that satisfies (MON) and any history $h_t \in \mathcal{H}$. Then for any $a', a'' \in A(h_t)$ we have $\bar{S}(h_t, a') \geq \underline{S}(h_t, a'')$, with equality only if $S(h_t, a')$ is a singleton.*[26]

Proof. Assume by contradiction that $\bar{S}(h_t, a') < \underline{S}(h_t, a'')$ for some $a', a'' \in A(h_t)$. Pick any type $\theta \in S(h_t, a')$ and any strategy $\mathbf{a}'$ on path for θ at h_t. Construct strategy $\mathbf{a}''$ as $\mathbf{a}''(h_t) = a''$ and $\mathbf{a}''(h_s) = \mathbf{a}'(h_s)$ for all $h_s \succ h_t$. This strategy constitutes a profitable deviation for θ at h_t. To see this, observe that the agent's lifetime utility can be written as

$$U(\mathbf{a}|h_t, \mathbf{b}, \theta) \equiv \mathbb{E}\left[\tilde{u}\Big(\mathbf{a}(h_t), p(h_t, \mathbf{a}(h_t), x_t), \theta\Big)dt + \sum_{s \in \mathcal{T}, s>t} e^{-r(s-t)}\tilde{u}\Big(\mathbf{a}(h_s), p(h_s, \mathbf{a}(h_s), x_s), \theta\Big)dt \mid h_t, \theta\right].$$

Since $p(h_s, \mathbf{a}''(h_s), x_s) \geq \delta_{\underline{S}(h_t, a'')} > \delta_{\bar{S}(h_t, a')} \geq p(h_s, \mathbf{a}'(h_s), x_s)$ for any x_s and all $h_s \succ h_t$, with the comparisons being with respect to the FOSD order, (MON) implies that

$$U(\mathbf{a}''|h_t, \mathbf{b}, \theta) - U(\mathbf{a}'|h_t, \mathbf{b}, \theta) \geq \mathbb{E}\left[\tilde{u}\big(a'', p(h_t, a'', x_t), \theta\big) - \tilde{u}\big(a', p(h_t, a', x_t), \theta\big)\right]dt + \Delta U, \tag{A.1}$$

where

$$\Delta U \equiv \mathbb{E}\left[\sum_{s \in \mathcal{T}, s>t} e^{-r(s-t)}\tilde{u}\Big(\mathbf{a}''(h_s), \delta_{\underline{S}(h_t, a'')}, \theta\Big)dt - \sum_{s \in \mathcal{T}, s>t} e^{-r(s-t)}\tilde{u}\Big(\mathbf{a}'(h_s), \delta_{\bar{S}(h_t, a')}, \theta\Big)dt \mid h_t, \theta\right].$$

By construction of $\mathbf{a}''$, $\mathbf{a}''(h_s) = \mathbf{a}'(h_s)$ for all $h_s \succ h_t$, hence $\tilde{u}\Big(\mathbf{a}''(h_s), \delta_{\underline{S}(h_t, a'')}, \theta\Big) = \tilde{u}\Big(\mathbf{a}'(h_s), \delta_{\underline{S}(h_t, a'')}, \theta\Big)$. By assumption, $\underline{S}(h_t, a'') > \bar{S}(h_t, a')$, hence (MON) implies $\tilde{u}\Big(\mathbf{a}''(h_s), \delta_{\underline{S}(h_t, a'')}, \theta\Big) > \tilde{u}\Big(\mathbf{a}'(h_s), \delta_{\bar{S}(h_t, a')}, \theta\Big)$ for all $h_s \succ h_t$. We conclude that $\Delta U > 0$. Further, value of ΔU is strictly positive regardless of dt.[27] For dt small enough, it dominates the first term on the RHS of

[26] This Lemma and the remainder of the Appendix uses $S(h_t, a)$ to denote "$S(h_t, a, x_t)$ for all $x_t \in X$ in the support". This object is well defined in equilibrium for on-path histories and actions because the support of x_t is type-independent and equilibrium beliefs must be consistent. We are adopting the simplifying assumption that the same holds off the equilibrium path, but this is not necessary for the arguments to go through as long as (NDOC-P) holds.

[27] This is in the sense that for any limit strategy $\mathbf{a}'$ and any outcome path $\mathbf{x}$, it holds by (MON) that $\int_t^\infty e^{-r(s-t)}\tilde{u}\Big(\mathbf{a}'(h_s), \delta_{\underline{S}(h_t, a'')}, \theta\Big)dt > \int_t^\infty e^{-r(s-t)}\tilde{u}\Big(\mathbf{a}'(h_s), \delta_{\bar{S}(h_t, a')}, \theta\Big)dt$, hence the expectation is strictly positive as well.

(A.1) because u is bounded, implying that for small enough dt, $U(\mathbf{a}''|h_t, \mathbf{b}, \theta) > U(\mathbf{a}'|h_t, \mathbf{b}, \theta)$ which contradicts $\mathbf{a}'$ being optimal for θ at h_t.

Now suppose $\bar{S}(h_t, a') = \underline{S}(h_t, a'')$. Suppose by way of contradiction that $|S(h_t, a')| > 1$, meaning $\underline{S}(h_t, a') < \underline{S}(h_t, a'')$. Let $p' \equiv \mathbb{E}[p(h_t, a', x_t) \mid h_t, a']$ be the belief induced by action a' alone (note $p' < \delta_{\bar{S}(h_t,a')} = \delta_{\underline{S}(h_t,a'')}$). Consider type $\underline{\theta} \equiv \underline{S}(h_t, a')$. This type $\underline{\theta}$ must have an on-path strategy $\mathbf{a}'$ that yields reputation $p \not> p'$ at all histories $h_s \succ (h_t, a')$ with positive probability (with certainty if outcomes x are uninformative).[28] Construct strategy $\mathbf{a}''$ as above. Then inequality (A.1) holds with

$$\Delta U \equiv \mathbb{E}\left[\sum_{s\in\mathcal{T}, s>t} e^{-r(s-t)}\left[\tilde{u}\Big(\mathbf{a}''(h_s), \delta_{\underline{S}(h_t,a'')}, \theta\Big) - \tilde{u}\Big(\mathbf{a}'(h_s), p(h_s, \mathbf{a}'(h_s), x_s), \theta\Big)\right] dt \mid h_t, \theta\right] > 0,$$

which is positive since $\tilde{u}\Big(\mathbf{a}''(h_s), \delta_{\underline{S}(h_t,a'')}, \theta\Big) = \tilde{u}\Big(\mathbf{a}'(h_s), \delta_{\underline{S}(h_t,a'')}, \theta\Big) > \tilde{u}\Big(\mathbf{a}'(h_s), p(h_s, \mathbf{a}'(h_s), x_s), \theta\Big)$ for all $h_s \succ h_t$ by the same argument as above. Unlike in the previous case, ΔU is not bounded away from zero for all dt. However, for the agent to prefer $\mathbf{a}'$, i.e., for $U(\mathbf{a}'|h_t, \mathbf{b}, \theta) \geq U(\mathbf{a}''|h_t, \mathbf{b}, \theta)$ to hold, we must have that $\Delta U \to 0$ as $dt \to 0$. This would require that $p' \to \delta_{\bar{S}(h_t,a')}$ as $dt \to 0$, which yields a contradiction in the limit, since $S(p') = S(h_t, a') \supsetneq \{\bar{S}(h_t, a')\}$. □

Next, Lemma 4 puts the (SC) property to use, establishing a form of monotonicity of optimal strategies with respect to type ("higher types play higher strategies"). The main problem in the dynamic setting is the lack of any nice complete order over strategies $\mathbf{a}$, so given two arbitrary strategies, we generally cannot say which one of them is "higher". Therefore, we rephrase this monotonicity result to say instead that if a given strategy (or its equivalent) is optimal for two agent types, then it must also be optimal for all types in between. We cannot say with certainty that the given strategy is chosen on equilibrium path by any of these types in between, but we can claim that any strategy they play must be payoff-equivalent to the one under consideration.

Lemma 4. *Fix any $dt > 0$, an equilibrium $(\alpha, \mathbf{b}, \theta)$ that satisfies (SC) and history $h_t \in \mathcal{H}$. If there exists a pair of strategies $\underline{\mathbf{a}}, \bar{\mathbf{a}} \curlywedge h_t$ that are payoff-equivalent at h_t and are on path at h_t for some types $\underline{\theta}$ and $\bar{\theta} > \underline{\theta}$ respectively, then any strategy $\hat{\mathbf{a}} \curlywedge h_t$ on path at h_t for any $\hat{\theta} \in (\underline{\theta}, \bar{\theta})$ must be payoff-equivalent at h_t to $\bar{\mathbf{a}}, \underline{\mathbf{a}}$.*

Proof. Fix any such $\hat{\mathbf{a}}$. Strategy $\bar{\mathbf{a}}$ has to be optimal for type $\bar{\theta}$. In particular, when evaluated at h_t, it has to be better than $\hat{\mathbf{a}}$:

$$U(\bar{\mathbf{a}}|h_t, \mathbf{b}, \bar{\theta}) \geq U(\hat{\mathbf{a}}|h_t, \mathbf{b}, \bar{\theta}).$$

The same holds for type $\underline{\theta}$, since $\bar{\mathbf{a}}$ and $\underline{\mathbf{a}}$ are payoff-equivalent:

$$U(\bar{\mathbf{a}}|h_t, \mathbf{b}, \underline{\theta}) = U(\underline{\mathbf{a}}|h_t, \mathbf{b}, \underline{\theta}) \geq U(\hat{\mathbf{a}}|h_t, \mathbf{b}, \underline{\theta}).$$

At the same time, $\hat{\theta}$ at least weakly prefers $\hat{\mathbf{a}}$ to $\bar{\mathbf{a}}$, meaning that the converse holds for $\hat{\theta}$:

$$U(\bar{\mathbf{a}}|h_t, \mathbf{b}, \hat{\theta}) \leq U(\hat{\mathbf{a}}|h_t, \mathbf{b}, \hat{\theta}).$$

If this inequality is strict, then this is a direct contradiction with (SC), which requires that $U(\bar{\mathbf{a}}|h_t, \mathbf{b}, \theta) - U(\hat{\mathbf{a}}|h_t, \mathbf{b}, \theta)$ as a function of θ either crosses zero at most once, or is exactly zero. □

Lemma 5 below is the final step before we can move on to the proofs of main results. It can be seen as a weaker version of Proposition 1, claiming that the highest and lowest types at any history have a strategy in common.

Lemma 5. *Fix a limit equilibrium $(\alpha, \mathbf{b}, p)$ that satisfies (MON) and (SC), and fix any $h_t \in \mathcal{H}$. There exist h_t-payoff-equivalent strategies $\bar{\mathbf{a}}, \underline{\mathbf{a}} \curlywedge h_t$, on path at h_t for types $\bar{S}(h_t)$ and $\underline{S}(h_t)$ respectively.*

Proof. We will proceed by induction on the support size $|S(h_t)|$. The claim of the lemma holds trivially for $|S(h_t)| = 1$, and by Theorem 1 it also holds for $|S(h_t)| = 2$. The remainder of the proof shows that if the claim holds when $|S(h_t)| = k - 1$ for $k \geq 3$, then it also holds when $|S(h_t)| = k$. Let $\mathbf{x} : \mathcal{H} \to X$ denote an outcome profile which prescribes some outcome for every history. Fix some $\mathbf{x}$. Coupled with some pure strategy of the agent, the receiver's equilibrium strategy $\mathbf{b}$ and the equilibrium belief system p, it fully determines the path of play and the agent's payoffs.

Begin the second layer of induction, iterating forwards on time periods from t. At h_t and any subsequent history $h_s \succ h_t$, one of the following must apply:

1. There is an action a on path for both types $\bar{S}(h_t)$ and $\underline{S}(h_t)$ at h_s. If this is the case, call h_s a *non-splitting* history and continue to $h_{s+dt} = (h_s, a, \mathbf{x}(h_s))$.
2. There is no action a on path for both $\bar{S}(h_t)$ and $\underline{S}(h_t)$ at h_s. If this is the case, call h_s a *splitting* history.

Proceed along the non-splitting path (according to the chosen $\mathbf{x}$) until the first splitting history h_s. Pick arbitrary actions $\bar{a}$ and $\underline{a}$ that are on path for $\bar{S}(h_t)$ and $\underline{S}(h_t)$ at h_s respectively, and consider two continuation histories $\bar{h}_{s+dt} \equiv (h_s, \bar{a}, \mathbf{x}(h_s))$ and $\underline{h}_{s+dt} \equiv (h_s, \underline{a}, \mathbf{x}(h_s))$. Then we have that $|S(h_{s+dt})| < |S(h_s)| = k$ for both continuation histories, because $S(\bar{h}_{s+dt}) \subseteq S(h_s)\backslash\underline{S}(h_s)$ and $S(\underline{h}_{s+dt}) \subseteq S(h_s)\backslash\bar{S}(h_s)$. Therefore, by the induction assumption, the statement of the lemma holds at both $\bar{h}_{s+dt}$ and $\underline{h}_{s+dt}$.

In particular, the statement of the lemma for $\bar{h}_{s+dt}$ states that there exist two $\bar{h}_{s+dt}$-payoff-equivalent strategies on path at $\bar{h}_{s+dt}$ for $\bar{S}(h_s)$ and $\underline{S}(\bar{h}_{s+dt})$ respectively. Playing $\bar{a}$ at h_s is on path for both of these types, hence there also exists a pair of strategies on path at h_s for the two types respectively, which grant the same payoff *conditional on* $\mathbf{x}$.[29] However, the argument above applies to any outcome profile $\mathbf{x}$ and, in particular, to any outcome $\mathbf{x}(h_s)$, hence there also exists a pair of strategies $\bar{\mathbf{a}}', \bar{\mathbf{a}}''$ on path at h_s for the two types $\bar{S}(h_s)$ and $\underline{S}(\bar{h}_{s+dt})$ respectively, which are payoff-equivalent at h_s (unconditionally).

By a mirror argument, there also exists a pair of h_s-payoff-equivalent strategies $\underline{\mathbf{a}}', \underline{\mathbf{a}}''$ on path at h_s for $\underline{S}(h_s)$ and $\bar{S}(\underline{h}_{s+dt})$ respectively. Note further that by Lemma 3 we have that $\bar{S}(\underline{h}_{s+dt}) > \underline{S}(\bar{h}_{s+dt})$. Lemma 4 hence applies: $\underline{\mathbf{a}}''$ must be payoff-equivalent to $\bar{\mathbf{a}}', \bar{\mathbf{a}}''$, thus so is $\underline{\mathbf{a}}'$. We have shown that the statement of the lemma holds at h_t if $|S(h_t)| = k$ and h_t is a splitting history.

We are left to cover non-splitting histories. Suppose h_t is non-splitting. Fix $\mathbf{x}$. Then we know that the statement of the lemma holds at the first splitting history h_s following h_t along the path of pooling actions and fixed outcomes $\mathbf{x}$. Therefore, there exists a pair of strategies on path at h_t for $\bar{S}(h_t)$ and $\underline{S}(h_t)$, which grant the same payoff at h_t conditional on $\mathbf{x}$. This applies to any outcome profile $\mathbf{x}$, hence there exists a pair of strategies $\bar{\mathbf{a}}, \underline{\mathbf{a}}$ on path at h_t for $\bar{S}(h_t)$ and $\underline{S}(h_t)$, which are payoff-equivalent at h_t. This concludes the induction argument and the proof of the lemma. □

---

[28] This follows from the observation that beliefs are correct in equilibrium, hence at any h_s there must exist an action a_s on path for $\underline{\theta}$ at h_s that weakly increases the probability that the receiver assigns to type $\underline{\theta}$:

$$\mathbb{E}[p(h_s, a_s, x_s)(\underline{\theta}) \mid h_s, a_s] \geq p(h_s)(\underline{\theta}).$$

[29] It does not matter for our argument if all types assign probability zero to outcome $\mathbf{x}(h_s)$ conditional on $\bar{a}$.

Proof of Proposition 1. Let $\bar{\theta} \equiv \bar{S}(h_t)$. Note that the statement of the proposition holds trivially if $|S(h_t)| = 1$, so for the remainder of this proof we assume that this is not the case (i.e., $\bar{\theta} \neq \underline{\theta}$). From Lemma 5 we know there exist h_t-payoff-equivalent $\bar{\mathbf{a}}, \underline{\mathbf{a}} \curlywedge h_t$ on path at h_t for $\bar{\theta}$ and $\underline{\theta}$ respectively. Then by Lemma 4, any pure strategy $\mathbf{a} \curlywedge h_t$ on path at h_t for any $\theta \in S(h_t)\backslash\{\bar{\theta}, \underline{\theta}\}$ is payoff-equivalent at h_t to $\bar{\mathbf{a}}, \underline{\mathbf{a}}$.

Suppose now there exists a pure strategy $\bar{\mathbf{a}}' \curlywedge h_t$ on path at h_t for $\bar{\theta}$, which is payoff-distinct at h_t from $\bar{\mathbf{a}}$. By (SC), all types $\theta \in S(h_t)\backslash\bar{\theta}$ must have a strict preference at h_t between $\bar{\mathbf{a}}$ and $\bar{\mathbf{a}}'$. The former is optimal for these types, hence $\bar{\mathbf{a}}'$ is only on path for $\bar{\theta}$. The two strategies cannot prescribe different actions $\bar{\mathbf{a}}(h_t) \neq \bar{\mathbf{a}}'(h_t)$ at h_t, since this is in violation of Lemma 3. The same, however, applies to any subsequent history h_s such that $|S(h_s)| > 1$. At all $h_s \succ h_t$ s.t. $|S(h_s)| = 1$, Lemma 2 implies that all pure strategies on path at h_s are h_s-payoff-equivalent. Both facts together imply that $\bar{\mathbf{a}}(h_s) = \bar{\mathbf{a}}'(h_s)$ for all $h_s \succ h_t$. This contradicts $\bar{\mathbf{a}}$ and $\bar{\mathbf{a}}'$ being payoff-distinct, hence such $\bar{\mathbf{a}}'$ does not exist. Therefore, any pure strategy $\mathbf{a}$ on path at h_t that is h_t-payoff-distinct from $\bar{\mathbf{a}}$ is only on path for $\underline{\theta}$. This concludes the proof. □

Proof of Theorem 2. From Proposition 1, all actions $a \in A(h_t)$ are on path for $\underline{\theta}$, which proves the first statement of the theorem.

From the fact that payoff-relevant signaling happens at h_t we know that there exist two pure strategies $\underline{\mathbf{a}}, \bar{\mathbf{a}}$ that are payoff-distinct at h_t and prescribe different actions at h_t: $\underline{a} \equiv \underline{\mathbf{a}}_t \neq \bar{a} \equiv \bar{\mathbf{a}}_t$. From Proposition 1 we know at least one of these strategies – suppose $\underline{\mathbf{a}}$ – is on path for $\underline{\theta}$ but not for any other $\theta \in S(h_t)\backslash\underline{\theta}$ at h_t. Furthermore, it follows from the definition of payoff-relevant signaling that there is no $\bar{\mathbf{a}}'$ s.t. $\bar{\mathbf{a}}' \curlywedge h_t$ and $\bar{\mathbf{a}}'_t = \underline{a}$, and which is payoff-equivalent to $\bar{\mathbf{a}}$ at h_t. Therefore, $\underline{a}$ is only on path for $\underline{\theta}$, while $\bar{a}$ is optimal for all $\theta \in S(h_t)$ at h_t.

We now show that $\underline{a} \in A^*\left(h_t, \mathbf{b}, \underline{\theta}\right)$. Suppose not. Then type $\underline{\theta}$ can play some $a_s \in A^*\left(h_s, \mathbf{b}, \underline{\theta}\right)$ at every history $h_s \succeq h_t$. Compared to following $\underline{\mathbf{a}}$, this strategy would yield the same payoff at all times $s > t$ and a strictly higher payoff at t (same as in the proof of Theorem 1), hence $\underline{\mathbf{a}}$ is not optimal for $\underline{\theta}$ at h_t — a contradiction.

To complete the proof of statements 2 and 3 of the theorem, we need to show that $\bar{a} \notin A^*\left(h_t, \mathbf{b}, \underline{\theta}\right)$. Assume not. Consider the strategy of playing $\bar{a}$ at h_t and all subsequent histories. Compared to following $\underline{\mathbf{a}}$, this strategy would yield $\underline{\theta}$ a weakly higher payoff at all times $s > t$ and a strictly higher payoff at t (due to $p(h_t, (\bar{a}, x_t)) > \delta_{\underline{\theta}}$ for all x_t and to the strict part of (MON)), hence it is a profitable deviation from $\underline{\mathbf{a}}$ for $\underline{\theta}$ at h_t — a contradiction.

This completes the proof of Theorem 2. □

Proof of Corollary 1. The low type must be indifferent between taking a separating action $\underline{a}$ at h_t and pooling on $\bar{a}$ at h_t and separating at h_{t+dt}. This indifference dictates that one period of pooling must be exactly as attractive as one period of being revealed as $\underline{\theta}$, i.e., $\mathbb{E}_x\left[\tilde{u}\left(\bar{a}, p(h_t, (\bar{a}, x)), \underline{\theta}\right) | \underline{\theta}\right] = \tilde{u}\left(\underline{a}, \delta_{\underline{\theta}}, \underline{\theta}\right)$. □

Proof of Corollary 2. Proposition 1 states that all pure strategies $\mathbf{a}'$ on path at h_t for any $\theta \in S(h_t)\backslash\underline{S}(h_t)$ are payoff-equivalent at h_t. Since there is no payoff-relevant signaling in equilibrium, the set of such strategies is a singleton: if there is more than one then there exists $h_s \succ h_t$ at which the two prescribe different actions, but that constitutes payoff-irrelevant signaling at h_s (the two strategies coincide on $[t, s)$, hence they are payoff-equivalent at h_s).

Therefore, at any h_t there exists some $\bar{a} \in A$ such that $\alpha_\theta(\bar{a}|h_t) = 1$ for all $\theta \in S(h_t)\backslash\underline{S}(h_t)$. Together with part 1 of Theorem 2, this means that $S(h_t, \bar{a}) = S(h_t)$. By part 3 of the theorem, $\bar{a}$ is the unique element of $A(h_t)\backslash A^*\left(h_t, \mathbf{b}, \underline{\theta}\right)$. By part 2 of the theorem, for any $\underline{a} \in A(h_t) \cap A^*\left(h_t, \mathbf{b}, \underline{\theta}\right)$ we have $S(h_t, \underline{a}) = \underline{S}(h_t)$. Since all on-path histories h_{t+dt} can be written as $h_{t+dt} = (h_t, a, x_t)$ for some $a \in A(h_t)$ and $x_t \in X$, and outcome x_t does not change support S, we obtain that for any pair of on-path histories h_t, h_{t+dt}, it must be that either $S(h_{t+dt}) = S(h_t)$, or $S(h_{t+dt}) = \{\underline{S}(h_t)\}$. Applying this observation iteratively from h_0 (for which $S(h_0) = \Theta$) completes the proof. □

A.4. Proofs: Separable settings

Proof of Proposition 2. It is immediate that if $\phi_1(a, p)$ is weakly increasing in p and $\psi(\theta) \geq 0$ then $\phi_1(a, p)\psi(\theta)$ is weakly increasing in p. Together with the assumption that $\phi_0(a, p)$ is weakly increasing in p, this implies trivially that (4) is weakly increasing in p, which is what is required by (MON). □

Proof of Proposition 3. For some fixed $\mathbf{a}', \mathbf{a}'', h_t$, the function $\mathcal{U}(\theta)$ under representation (4) is given by

$$\begin{aligned}\mathcal{U}(\theta) = \mathbb{E}\Bigg[&\sum_{s\in\mathcal{T}, s\geq t} \left[\phi_0\left(\mathbf{a}''(h_s), p(h_s)\right) - \phi_0\left(\mathbf{a}'(h_s), p(h_s)\right)\right] dt \\ &+ \sum_{s\in\mathcal{T}, s\geq t} \phi_1\left(\mathbf{a}''(h_s), p(h_s)\right)\psi(\theta)dt \\ &- \sum_{s\in\mathcal{T}, s\geq t} \phi_1\left(\mathbf{a}'(h_s), p(h_s)\right)\psi(\theta)dt \mid h_t, \mathbf{b}, \theta \Bigg] \\ = \;&\Phi_0 + \Phi_1\psi(\theta)\end{aligned}$$

for some Φ_0, Φ_1 that do not depend on θ. In the above, the expectation is over outcomes, hence by assumption that outcomes are uninformative, the expectation is resolved trivially, since strategies $\mathbf{a}, \mathbf{b}$ then fully determine the path of play. It then follows immediately that if $\psi(\theta)$ is strictly monotone, then $\mathcal{U}(\theta)$ is either strictly monotone, or equivalently zero (if $\Phi_1 = 0$), hence (SC) holds. □

A.5. Proofs: Existence

Proof of Theorem 3. Since outcomes x are uninformative, they are suppressed throughout the proof.

Part 1. Fix first $u \in W_0$ and let $a^p \equiv a(u, 0)$ denote the pooling action that corresponds to u. It is immediate that playing a^p at all $t \in \mathcal{T}$ yields expected discounted utility $\frac{dt}{1-e^{-rdt}}u(\theta)$ to type θ, and that there exists a Δ_ϵ such that for $dt < \Delta_\epsilon$, this payoff is within ϵ of $\frac{u(\theta)}{r}$. Consider a candidate equilibrium that consists of all types of the sender following a^p at all on-path histories and playing some myopically optimal $a \in A^*(h_t, \mathbf{b}, \theta)$ at all off-path histories h_t, and the receiver playing some optimal strategy $\mathbf{b}$. The receiver's belief remains at $p(h_t) = p_0$ at all on-path histories h_t (i.e., those at which the sender chose a^p in all periods $\{0, \ldots, t - dt\}$), and drops to $p(h_t) = \delta_{\underline{\theta}}$ otherwise. It is immediate that beliefs are consistent and satisfy (NDOC-P) and hence (NDOC). By definition, optimality holds for the sender at off-path histories and for the receiver at all histories. We only need to show that sender's on-path optimality holds. This amounts to the requirement that the following holds for all θ and all on-path h_t:

$$\begin{aligned}\sum_{s\in\mathcal{T}, s\geq t} e^{-r(s-t)}u(\theta)dt &\geq \max_{a\in A} u\left(a, \mathbf{b}(h_t, a, x_t), \theta\right) dt \\ &\quad + \sum_{s\in\mathcal{T}, s>t} e^{-r(s-t)}\underline{u}(\theta)dt \\ \iff \frac{e^{-rdt}}{1-e^{-rdt}}\left(u(\theta) - \underline{u}(\theta)\right) &\geq \max_{a\in A} u\left(a, \mathbf{b}(h_t, a, x_t), \theta\right) - u(\theta). \end{aligned} \tag{A.2}$$

Recall that all utility functions are bounded, and $u(\theta) > \underline{u}(\theta)$ for all θ by definition of W_0. The fraction on the left-hand side of (A.2) is decreasing in $dt > 0$, with $\lim_{dt\to 0+} \frac{e^{-rdt}}{1-e^{-rdt}} = +\infty$, there exists $\Delta_\theta > 0$ s.t. for all $dt < \Delta_\theta$ the inequality holds for a given θ. Taking $\Delta \equiv \min\left\{\Delta_\epsilon, \{\Delta_\theta\}_{\theta\in\Theta}\right\}$ proves part 1 for $u \in W_0$.

The statement of part 1 for $u \in co(W_0)\backslash W_0$ can be proved as follows. Carathéodory's Convexity Theorem (Aliprantis and Border, 2006, Theorem 5.32) implies that u can then be expressed as a convex combination of a *finite* number of elements of W_0, denote them $\{u^1, \dots, u^K\}$ (the theorem also implies in our case that $K \le N+2$). Lemma 3.7.1 in Mailath and Samuelson (2006) (originally due to Sorin, 1986) argues that in this case, there exists Δ_c s.t. for any $dt < \Delta_c$ there exists a sequence $\{u_t\}_{t=0}^\infty$ with $u_t \in \{u^1, \dots, u^K\}$ for all t such that $\sum_{t\in\mathcal{T}} e^{-rt} u_t(\theta) dt = \frac{u(\theta)dt}{1-e^{-rdt}}$. As above, there exists Δ_ϵ s.t. for all $dt < \Delta_\epsilon$, $\frac{u(\theta)dt}{1-e^{-rdt}}$ is within ϵ of $\frac{u(\theta)}{r}$. One can then consider a candidate equilibrium as above, in which the sender follows the sequence of actions corresponding to $\{u_t\}_{t=0}^\infty$ on the equilibrium path, and everything else is as before. Condition (A.2) for given θ, h_t would then transform to

$$\sum_{s\in\mathcal{T}, s\ge t} e^{-r(s-t)} u_t(\theta) dt \ge \max_{a\in A} u\left(a, \mathbf{b}(h_t, a, x_t), \theta\right) dt + \sum_{s\in\mathcal{T}, s> t} e^{-r(s-t)} \underline{u}(\theta) dt.$$

Since $u_t(\theta) \ge \min_k \left\{u^k(\theta)\right\} > \underline{u}(\theta)$, the condition above surely holds if

$$\frac{e^{-rdt}}{1-e^{-rdt}} \left(\min_k \left\{u^k(\theta)\right\} - \underline{u}(\theta)\right) \ge \max_{a\in A} u\left(a, \mathbf{b}(h_t, a, x_t), \theta\right) - \min_k \left\{u^k(\theta)\right\},$$

which again holds for all $dt < \Delta_\theta$ for some Δ_θ. Setting $\Delta \equiv \min\left\{\Delta_c, \Delta_\epsilon, \{\Delta_\theta\}_{\theta\in\Theta}\right\}$, the result in part 1 follows.

Part 2. Again, we begin by showing the statement for $u \in \underline{W}$, as opposed to $co(\underline{W})$. Fix some $(u, \nu) \in \underline{W}$ and let $a^p \equiv a(u, \nu)$ denote the action that corresponds to u. It follows from part 1 of this theorem that if the receiver's belief at some history h_t is $p(h_t) = \tilde{p}(\nu)$, then there exists a continuation equilibrium from that history onwards that yields expected discounted utility within ϵ of $\frac{u}{r}$ for all types of the sender. Consider now a candidate equilibrium that is as described in part 1 from period $t = dt$ onwards, whereas in period $t = 0$ (at history $h_0 = \varnothing$) types $\theta \ne \underline{\theta}$ of the sender play a^p with probability 1 (and the receiver's updates her belief to $\tilde{p}(\nu)$ in this case), type $\underline{\theta}$ plays a^p with probability $1-\nu$ and some myopically optimal $\underline{a} \in A^*(h_0, \mathbf{b}, \underline{\theta})\backslash a^p$ with probability ν (leading to the receiver's belief $p(h_1) = \delta_{\underline{\theta}}$; condition (AS) guarantees that the separating action exists). It is immediate that beliefs are consistent and satisfy (NDOC-P) and hence (NDOC). As in part 1, there exists Δ_ϵ s.t. for $dt < \Delta_\epsilon$, types $\theta \ne \underline{\theta}$ get payoffs within ϵ of $\frac{u}{r}$ in period 0, hence the argument from part 1 confirms that it is optimal for type θ to abide by this equilibrium strategy if $dt < \Delta_\theta$ for some Δ_θ. Type $\underline{\theta}$ must be indifferent between a^p and $\underline{a}$. Since $u(\underline{\theta}) = \underline{u}(\underline{\theta})$, this indifference holds, both in period $t = 0$ and for all future on-path histories. Hence the behavior prescribed in the candidate equilibrium is optimal for type $\underline{\theta}$ as well, and for $dt < \Delta_\epsilon$ yields payoff within ϵ of $\frac{u(\theta)}{r}$. Therefore, the candidate equilibrium is indeed an equilibrium for dt small enough. Setting $\Delta \equiv \min\left\{\Delta_\epsilon, \{\Delta_\theta\}_{\theta\in\Theta\backslash\underline{\theta}}\right\}$ proves part 2 of the theorem for $(u, \nu) \in \underline{W}$.

Finally, to complete the proof we need to show that the statement of part 2 holds also for $(u, \nu) \in co(\underline{W})\backslash\underline{W}$. Invoking again the Carathéodory's Convexity Theorem, there exists a finite collection $((u^1, \nu^1), \dots, (u^K, \nu^K))$, with $(u^k, \nu^k) \in \underline{W}$ for all k, and weights $(\xi_1, \dots, \xi_K)$ s.t. $\sum_{k=1}^K \xi_k = 1$ and $\sum_{k=1}^K \xi_k \cdot u^k = u$. W.l.o.g. order the points so that $\nu^1 < \dots < \nu^K$. Further, let $a_k \equiv a(u^k, \nu_k)$ denote the action corresponding to the utility profile u^k.

Consider now a candidate equilibrium that off the equilibrium path looks the same as the equilibria constructed above (i.e., at any off-path h_t, $p(h_t) = \delta_\theta$, type-θ sender plays some myopically optimal $a \in A^*(h_t, \mathbf{b}, \theta)$ distinct from the pooling action described below, and the receiver follows some optimal strategy $\mathbf{b}$). On the equilibrium path:

- the receiver follows some optimal strategy $\mathbf{b}$;
- types $\theta \ne \underline{\theta}$ play pooling action a_k in all periods $t \in \mathcal{T}$ s.t. $\tau_k \le t < \tau_{k+1}$, where the switching times are defined given by $\tau_k \equiv \max\left\{t \in \mathcal{T} \mid \sum_{s\in\mathcal{T}, s<t} e^{-rs} dt \le \sum_{k'=1}^{k-1} \xi_{k'}\right\}$, with $\tau_1 \equiv 0$ and $\tau_{K+1} \equiv +\infty$;
- type $\underline{\theta}$ follows the pooling strategy above in all periods $t \notin \{\tau_1, \tau_2, \dots, \tau_K\}$, while in period τ_k he plays pooling action a_k with probability $\frac{\nu_k - \nu_{k-1}}{1-\nu_{k-1}}$ (with $\nu_0 \equiv 0$) and plays a myopically optimal action $\underline{a} \in A^*\left(h_{\tau_k}, \mathbf{b}, \underline{\theta}\right)$ with probability $\frac{1-\nu_k}{1-\nu_{k-1}}$;
- the receiver's belief is given by $p(h_t) = \tilde{p}(\nu_k)$ for all on-path histories h_t s.t. $\tau_k \le t < \tau_{k+1}$, and $p(h_t) = \delta_{\underline{\theta}}$ for all other histories.

To verify that this is indeed an equilibrium, note the following:

1. The receiver's beliefs are consistent on path. To see this, let h_k denote the period-τ_k on-path history, $k \in \{0, 1, \dots, K\}$, where $h_{\tau_0} \equiv \varnothing$. Then Bayes' rule implies that the probability the receiver assigns to, e.g., type $\underline{\theta}$ at h_k is
$$\frac{p(\underline{\theta}|h_k)}{1-p(\underline{\theta}|h_k)} = \frac{p(\underline{\theta}|h_{k-1})}{1-p(\underline{\theta}|h_{k-1})} \cdot \frac{1-\nu_k}{1-\nu_{k-1}} = \frac{p(\underline{\theta}|\varnothing)\cdot(1-\nu_k)}{1-p(\underline{\theta}|\varnothing)} \equiv \frac{\tilde{p}(\underline{\theta}|\nu_k)}{1-\tilde{p}(\underline{\theta}|\nu_k)},$$
where the second equality is obtained by iterative substitution of $(p(\underline{\theta}|h_{k-1}), \dots, p(\underline{\theta}|h_0), p(\underline{\theta}|\varnothing))$. Similar logic shows that $p(\theta|h_k) = \tilde{p}(\theta|\nu_k)$ for all $\theta \ne \underline{\theta}$, meaning $p(h_k) = \tilde{p}(\nu_k)$. It is further immediate that the beliefs are not updated at on-path histories t with $t \notin \{\tau_1, \dots, \tau_K\}$.
2. The receiver's beliefs on path trivially satisfy (NDOC-P).
3. The sender of type $\underline{\theta}$ receives payoff $\underline{u}(\underline{\theta})$ at every history (on- and off-path), hence has no profitable deviations.
4. By the same argument as in the proof of part 1 for $u \in co(W_0)\backslash W_0$, there exists Δ_θ such that it is optimal for the sender of type $\theta \ne \underline{\theta}$ to follow the equilibrium strategy if $dt < \Delta_\theta$.

The final thing we need to verify is that the sender of type θ receives expected discounted utility within ϵ of $\frac{u(\theta)}{r}$. Type $\underline{\theta}$ gets payoff $\underline{u}(\underline{\theta})$ at all on-path histories, hence same as in part 1, there exists Δ_ϵ s.t. for all $dt < \Delta_\epsilon$ the discounted utility of type $\underline{\theta}$ is within ϵ of $\frac{\underline{u}(\theta)}{r} = \frac{u(\theta)}{r}$. For type $\theta \ne \underline{\theta}$, the discounted utility at $t = 0$ is given by

$$\sum_{k=1}^K \sum_{t\in\mathcal{T}, \tau_k \le s < \tau_{k+1}} e^{-rt} u^k dt,$$

and recall that $\sum_{k=1}^K \xi_k \cdot u^k = u$ By definition of τ_k, for all $k \in \{1, \dots, K\}$:

$$\left| \sum_{s\in\mathcal{T}, \tau_k \le s < \tau_{k+1}} e^{-rs} dt - \xi_k \right| \le 2e^{-rdt} dt,$$

and all utilities are bounded, which means that

$$\Delta'_u \equiv \max_{p\in P_s, \theta\in\Theta} \left\{\max_{a\in A} \tilde{u}(a, p, \theta) - \min_{a\in A} \tilde{u}(a, p, \theta)\right\} < \infty.$$

Hence

$$\left|\sum_{k=1}^{K}\sum_{t\in\mathcal{T},\tau_{k-1}\leq s<\tau_k} e^{-rt}u^k dt-\epsilon\right|\leq \Delta_u'\cdot 2e^{-rdt}dt, \tag{A.3}$$

meaning there exists Δ_u s.t. for all $dt < \Delta_u$ the RHS of (A.3) is smaller than ϵ. Setting $\Delta \equiv \min\left\{\Delta_u, \Delta_\epsilon, \{\Delta_\theta\}_{\theta\in\Theta\setminus\underline{\theta}}\right\}$, this concludes the proof of part 2 of the theorem. □

Proof of Corollary 3. The corollary follows immediately from the construction for $(u, v) \in co(\underline{W})$ in the proof of part 2 of Theorem 3. □

References

Abreu, D., Gul, F., 2000. Bargaining and reputation. Econometrica 68 (1), 85–117. http://dx.doi.org/10.1111/1468-0262.00094.

Admati, A.R., Perry, M., 1987. Strategic delay in bargaining. Rev. Econom. Stud. 54 (3), 345–364. http://dx.doi.org/10.2307/2297563.

Aköz, K.K., Arbatli, C.E., Celik, L., 2020. Manipulation through biased product reviews. J. Ind. Econ. 68 (4), 591–639. http://dx.doi.org/10.1111/joie.12240.

Aliprantis, C.D., Border, K.C., 2006. Infinite Dimensional Analysis: A Hitchhiker's Guide, third ed. Springer, http://dx.doi.org/10.1007/3-540-29587-9.

Athey, S., 2002. Monotone comparative statics under uncertainty. Q. J. Econ. 117 (1), 187–223. http://dx.doi.org/10.1162/003355302753399481.

Aumann, R.J., 1964. Mixed and behavior strategies in infinite extensive games. In: Advances in Game Theory, vol. 52, Princeton University Press, pp. 627–650. http://dx.doi.org/10.1515/9781400882014-029.

Ausubel, L.M., Cramton, P., Deneckere, R.J., 2002. Bargaining with incomplete information. In: Handbook of Game Theory, vol. 3, pp. 1897–1945. http://dx.doi.org/10.1016/S1574-0005(02)03013-8.

Ausubel, L.M., Deneckere, R.J., 1989. A direct mechanism characterization of sequential bargaining with one-sided incomplete information. J. Econom. Theory 48 (1), 18–46. http://dx.doi.org/10.1016/0022-0531(89)90118-X.

Ausubel, L.M., Deneckere, R.J., 1998. Bargaining and forward induction. mimeo.

Bagwell, K., Riordan, M.H., 1991. High and declining prices signal product quality. Am. Econ. Rev. 224–239.

Banks, J.S., Sobel, J., 1987. Equilibrium selection in signaling games. Econometrica 55 (3), 647–661. http://dx.doi.org/10.2307/1913604.

Beaudry, P., Poitevin, M., 1993. Signalling and renegotiation in contractual relationships. Econometrica 61 (4), 745–782. http://dx.doi.org/10.2307/2951762.

Bond, P., Zhong, H., 2016. Buying high and selling low: Stock repurchases and persistent asymmetric information. Rev. Financ. Stud. 29 (6), 1409–1452. http://dx.doi.org/10.1093/rfs/hhw005.

Chen, Y., 2011. Perturbed communication games with honest senders and naive receivers. J. Econom. Theory 146 (2), 401–424. http://dx.doi.org/10.1016/j.jet.2010.08.001.

Cho, I.-K., Kreps, D.M., 1987. Signaling games and stable equilibria. Q. J. Econ. 102 (2), 179–221. http://dx.doi.org/10.2307/1885060.

Daley, B., Green, B., 2012. Waiting for news in the market for lemons. Econometrica 80 (4), 1433–1504. http://dx.doi.org/10.3982/ECTA9278.

Daley, B., Green, B., 2020. Bargaining and news. Amer. Econ. Rev. 110 (2), 428–474. http://dx.doi.org/10.1257/aer.20181316.

De Angelis, T., Ekström, E., Glover, K., 2022. Dynkin games with incomplete and asymmetric information. Math. Oper. Res. 47 (1), 560–586. http://dx.doi.org/10.1287/moor.2021.1141.

Deneckere, R., Liang, M.-Y., 2006. Bargaining with interdependent values. Econometrica 74 (5), 1309–1364. http://dx.doi.org/10.1111/j.1468-0262.2006.00706.x.

Dilmé, F., 2017. Noisy signaling in discrete time. J. Math. Econom. 68, http://dx.doi.org/10.1016/j.jmateco.2016.10.002.

Dilmé, F., Li, F., 2016. Dynamic signaling with dropout risk. Am. Econ. J. Microecon. 8 (1), 57–82. http://dx.doi.org/10.1257/mic.20120112.

Ely, J.C., Välimäki, J., 2003. Bad reputation. Q. J. Econ. 118 (3), 785–814. http://dx.doi.org/10.1162/00335530360698423.

Feinberg, Y., Skrzypacz, A., 2005. Uncertainty about uncertainty and delay in bargaining. Econometrica 73 (1), 69–91. http://dx.doi.org/10.1111/j.1468-0262.2005.00565.x.

Fuchs, W., Skrzypacz, A., 2010. Bargaining with arrival of new traders. Amer. Econ. Rev. 100 (3), 802–836. http://dx.doi.org/10.1257/aer.100.3.802.

Gabaix, X., 2019. Behavioral inattention. In: Handbook of Behavioral Economics: Applications and Foundations 1, vol. 2. Elsevier, pp. 261–343. http://dx.doi.org/10.1016/bs.hesbe.2018.11.001.

Grossman, S.J., Perry, M., 1986. Sequential bargaining under asymmetric information. J. Econom. Theory 39 (1), 120–154. http://dx.doi.org/10.1016/0022-0531(86)90023-2.

Gryglewicz, S., Kolb, A., 2022. Dynamic signaling with stochastic stakes. Theor. Econ. 17 (2), 539–559. http://dx.doi.org/10.3982/TE3710.

Gul, F., Pesendorfer, W., 2012. The war of information. Rev. Econom. Stud. 79 (2), 707–734. http://dx.doi.org/10.1093/restud/rds017.

Gul, F., Sonnenschein, H., Wilson, R., 1986. Foundations of dynamic monopoly and the coase conjecture. J. Econom. Theory 39 (1), 155–190. http://dx.doi.org/10.1016/0022-0531(86)90024-4.

Handel, B., Schwartzstein, J., 2018. Frictions or mental gaps: What's behind the information we (don't) use and when do we care? J. Econ. Perspect. 32 (1), 155–178. http://dx.doi.org/10.1257/jep.32.1.155.

Heinsalu, S., 2018. Dynamic noisy signaling. Am. Econ. J. Microecon. 10 (2), 225–249. http://dx.doi.org/10.1257/mic.20160336.

Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. Amer. Econ. Rev. 101 (6), 2590–2615. http://dx.doi.org/10.1257/aer.101.6.2590.

Kaya, A., 2009. Repeated signaling games. Games Econom. Behav. 66 (2), 841–854. http://dx.doi.org/10.1016/j.geb.2008.09.030.

Kaya, A., 2013. Dynamics of price and advertising as quality signals: Anything goes. Econ. Bull. 2 (1), 1556–1564.

Kaya, A., Kim, K., 2018. Trading dynamics with private buyer signals in the market for lemons. Rev. Econom. Stud. 85 (4), 2318–2352. http://dx.doi.org/10.1093/restud/rdy007.

Kirmani, A., Rao, A.R., 2000. No pain, no gain: A critical review of the literature on signaling unobservable product quality. J. Mark. 64 (2), 66–79. http://dx.doi.org/10.1509/jmkg.64.2.66.18000.

Kraus, S., Wilkenfeld, J., Zlotkin, G., 1995. Multiagent negotiation under time constraints. Artif. Intell. 75 (2), 297–345. http://dx.doi.org/10.1016/0004-3702(94)00021-R.

Laffont, J.-J., Martimort, D., 2002. The Theory of Incentives: The Principal-Agent Model. Princeton University Press, http://dx.doi.org/10.1515/9781400829453.

Lai, E.K., 2014. Expert advice for amateurs. J. Econ. Behav. Organ. 103, 1–16. http://dx.doi.org/10.1016/j.jebo.2014.03.023.

LeBlanc, G., 1992. Signalling strength: Limit pricing and predatory pricing. Rand J. Econ. 493–506. http://dx.doi.org/10.2307/2555901.

Lee, J., Liu, Q., 2013. Gambling reputation: Repeated bargaining with outside options. Econometrica 81 (4), 1601–1672. http://dx.doi.org/10.3982/ECTA9200.

Leland, H.E., Pyle, D.H., 1977. Informational asymmetries, financial structure, and financial intermediation. J. Finance 32 (2), 371–387. http://dx.doi.org/10.2307/2326770.

Madrigal, V., Tan, T.C.C., Werlang, S.R.d.C., 1987. Support restrictions and sequential equilibria. J. Econom. Theory 43 (2), 329–334. http://dx.doi.org/10.1016/0022-0531(87)90063-9.

Mailath, G.J., Samuelson, L., 2006. Repeated Games and Reputations: Long-Run Relationships. Oxford University Press, http://dx.doi.org/10.1093/acprof:oso/9780195300796.001.0001.

Milgrom, P., Roberts, J., 1982a. Limit pricing and entry under incomplete information: An equilibrium analysis. Econometrica 50 (2), 443–459. http://dx.doi.org/10.2307/1912637.

Milgrom, P., Roberts, J., 1982b. Predation, reputation, and entry deterrence. J. Econ. Theory 27 (2), 280–312. http://dx.doi.org/10.1016/0022-0531(82)90031-X.

Milgrom, P., Roberts, J., 1986. Price and advertising signals of product quality. J. Political Econ. 94 (4), 796–821. http://dx.doi.org/10.1086/261408.

Milgrom, P., Shannon, C., 1994. Monotone comparative statics. Econometrica 62 (1), 157–180. http://dx.doi.org/10.2307/2951479.

Myerson, R.B., 1997. Game Theory: Analysis of Conflict. Harvard University Press, p. 568. http://dx.doi.org/10.2307/j.ctvjsf522.

Nöldeke, G., van Damme, E., 1990a. Signalling in a dynamic labour market. Rev. Econom. Stud. 57 (1), 1–23. http://dx.doi.org/10.2307/2297540.

Nöldeke, G., van Damme, E., 1990b. Switching away from probability one beliefs. mimeo URL https://www.researchgate.net/profile/Georg-Noeldeke/publication/4923610_Switching_Away_From_Probability_One_Beliefs/links/00b7d526a6ff6db029000000/Switching-Away-From-Probability-One-Beliefs.pdf.

Osborne, M.J., Rubinstein, A., 1990. Bargaining and Markets. Academic Press Limited, ISBN: 0-12-528632-5.

Pei, H., 2021. Trust and betrayals: Reputational payoffs and behaviors without commitment. Theor. Econ. 16, 449–475. http://dx.doi.org/10.3982/TE4182.

Quah, J.K.-H., Strulovici, B., 2012. Aggregating the single crossing property. Econometrica 80 (5), 2333–2348. http://dx.doi.org/10.3982/ecta9869.

Riley, J.G., 2001. Silver signals: Twenty-five years of screening and signaling. J. Econ. Lit. 39 (2), 432–478. http://dx.doi.org/10.1257/jel.39.2.432.

Roddie, C., 2012a. Signaling and reputation in repeated games, I: Finite games. SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.1994378.

Roddie, C., 2012b. Signaling and reputation in repeated games, II: Stackelberg limit properties. SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.2011835.

Rubinstein, A., 1985. A bargaining model with incomplete information about time preferences. Econometrica 53 (5), 1151–1172. http://dx.doi.org/10.2307/1911016.

Sen, A., 2000. Multidimensional bargaining under asymmetric information. Internat. Econom. Rev. 41 (2), 425–450. http://dx.doi.org/10.1111/1468-2354.00070.

Smirnov, A., Starkov, E., 2019. Timing of Predictions in Dynamic Cheap Talk: Experts Vs. Quacks, vol. 334, University of Zurich, Department of Economics Working Papers, http://dx.doi.org/10.2139/ssrn.3485707.

Smirnov, A., Starkov, E., 2022. Bad news turned good: Reversal under censorship. Am. Econ. J. Microecon. 14 (2), 506–560. http://dx.doi.org/10.1257/mic.20190379.

Sorin, S., 1986. On repeated games with complete information. Math. Oper. Res. 11 (1), 147–160. http://dx.doi.org/10.1287/moor.11.1.147.

Spence, M., 1973. Job market signaling. Q. J. Econ. 87 (3), 355–374. http://dx.doi.org/10.2307/1882010.

Strebulaev, I.A., Zhu, H., Zryumov, P., 2016. Optimal Issuance Under Information Asymmetry and Accumulation of Cash Flows, vol. 164, Rock Center for Corporate Governance at Stanford University Working Paper, http://dx.doi.org/10.2139/ssrn.2356144.

Swinkels, J.M., 1999. Education signalling with preemptive offers. Rev. Econom. Stud. 66 (4), 949–970. http://dx.doi.org/10.1111/1467-937x.00115.

Vettas, N., 1997. On the informational role of quantities: Durable goods and consumers' word-of-mouth communication. Internat. Econom. Rev. 915–944. http://dx.doi.org/10.2307/2527222.

Vincent, D.R., 1990. Dynamic auctions. Rev. Econom. Stud. 57 (1), 49–61. http://dx.doi.org/10.2307/2297542.

Vincent, D.R., 1998. Repeated signalling games and dynamic trading relationships. Internat. Econom. Rev. 39 (2), 275–294. http://dx.doi.org/10.2307/2527293.

Vong, A., 2021. The crisis of expertise. SSRN Sch. Pap. 3768185, http://dx.doi.org/10.2139/ssrn.3768185.

Weiss, A., 1983. A sorting-cum-learning model of education. J. Polit. Econ. 91 (3), 420–442. http://dx.doi.org/10.1086/261156.

Whitmeyer, M., 2021. Opacity design in signaling games. arXiv preprint arXiv:1902.00976.