Only Time Will Tell: Credible Dynamic Signaling*

Egor Starkov†

July 19, 2020

Abstract

This paper explores a model of dynamic signaling without commitment. It is known that separating equilibria do not exist if the sender cannot commit to future costly actions, since no single action can have enough weight to be an effective signal. This paper, however, shows that informative and payoff-relevant signaling can occur even without commitment and without resorting to unreasonable off-path beliefs. Such signaling can only happen through attrition, when the weakest type mixes between revealing own type and pooling with the stronger types. The possibility of full information revelation in the limit hence depends crucially on the assumptions about the state space. We illustrate the results by exploring a model of dynamic price signaling and show that prices may be informative of product quality even if the seller cannot commit to future prices, with both high and low prices being able to signal high quality.

Keywords: dynamic signaling, repeated signaling, reputation, attrition

JEL Codes: C73, D82, D83, L15

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. In the context of Spence's story of signaling ability through education, this problem was formulated by 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

*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 and University of Copenhagen for valuable feedback and helpful comments.

†Department of Economics, University of Copenhagen, Øster Farimagsgade 5, bygning 26, 1353 København K, Denmark; e-mail: egor@starkov.email.

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

would offer wages appropriate to a high ability worker *before* enough time has elapsed to present an effective screen" (p.363). In other words, if neither workers can commit to completing their education in the future, nor firms can commit to not hire undereducated workers, then education cannot serve as an effective signal.[2] To give a particular example: when low-ability workers are not supposed to pursue a college degree, a single day spent in that pursuit would imply that a worker's ability is not low, and market wages for college dropouts would be correspondingly high – too high to actually deter low-ability workers from enrolling in college only to drop out soon thereafter.

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]), evolving sender's type to create the need for maintaining reputation as opposed to establishing it once (Roddie [2012a,b]), or receivers observing noisy outcomes instead of the sender's actions (Dilmé [2017], Heinsalu [2018]).[3] 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 undertakes the mission to characterize all informative outcomes that can arise in a general model of dynamic (a.k.a. repeated) signaling without commitment. Our main result shows that the scope for signaling, while limited, does in fact exist in dynamic settings, contrary to the intuition of Admati and Perry [1987]. In particular, payoff-relevant signaling is possible via what is effectively a war of attrition, in which 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. The contribution of this paper is both in showing the existence of a wedge between signaling and cheap talk in the setting under consideration – a wedge presumed nonexistent by the existing literature, – and characterizing this wedge explicitly.

The conclusion regarding the uniqueness of attrition as an informative equilibrium outcome leads into another message of the paper, which is methodological. The mechanism of attrition of the lowest type can yield full separation asymptotically if there are only two types in the game but not if there are more (but finitely many). Further, one can show that in an analog of our model with a continuous type space, full asymptotic revelation is possible again.[4] This aims to show that one's modelling assumptions may crucially affect the result even when they are about an object that is as seemingly abstract and arbitrary as type space.

While the attrition structure is restrictive, it nonetheless allows for a nontrivial equilibrium multiplicity. In addition to various possible combinations of informative and uninformative periods,

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

[3] See Whitmeyer [2019] for a discussion regarding the sender-optimal amount of noise in signaling.

[4] An example of such outcome is presented by Fuchs and Skrzypacz [2010].

multiple attrition outcomes can, ceteris paribus, be sustained in equilibrium in any given period. These outcomes would be different in their informativeness, i.e., the probability of separation of the lowest type. This dimension of multiplicity (as well as the range of informative outcomes sustainable in equilibrium in a given period more generally) depends on the richness of the action set.

Our paper can be seen as formalizing the folk wisdom, since the attrition equilibrium structure outlined above ("the high type chooses action a', the low type mixes between a' and a'' with positive probabilities") is encountered in many applied dynamic models.[5] However, in spite of the overwhelming presence in the applied theory literature, the issue has never received a rigorous treatment in signaling literature. Our paper amends that, demonstrating that the uniqueness of attrition as an equilibrium structure arises in general settings well beyond the specific models explored previously.

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 Bond and Zhong [2016], *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.[6]

To illustrate, consider again the job market signaling example. In the story of Spence [1973], high-ability students commit (during the college applications period) to obtaining a college degree. This commitment is too costly for low-ability students, who then forego college altogether and accept lower wages. In the absence of commitment – if a student must every day decide whether to stay in college or drop out and pursue a job at a competitive wage – this equilibrium can still be sustained by unreasonable employers' beliefs off path. Such beliefs would treat any high school graduate as high-ability *as long as* they are on track for a college degree – but any college dropout is immediately downgraded to low-ability in employers' eyes. This belief system incentivizes high-ability students to endure all four years of college to obtain higher wages and disincentivizes low-ability students from pursuing higher education. However, this belief system is internally inconsistent: only high-ability students are believed to go to college in the first place, so dropouts cannot be of low ability! NDOC rules out exactly this kind of inconsistency.

Furthermore, it feels tongue-in-cheek to even call equilibria like the one above "separating", since the sender of some type can never properly "separate" from other types in that scenario. While the

[5] Some examples include 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 bargaining and bilateral trade; Strebulaev, Zhu, and Zryumov [2016] in corporate finance; Vettas [1997], Aköz, Arbatli, and Çelik [2017], Gryglewicz and Kolb [2019], Smirnov and Starkov [2020] in industrial organization/marketing; Smirnov and Starkov [2019] in cheap talk; De Angelis, Ekström, and Glover [2018] in Dynkin games.

[6] C.f. Banks and Sobel [1987] and Cho and Kreps [1987] for signaling and chapter 4 in Myerson [1997] for extensive-form games respectively.

receiver's belief may assign probability zero to those other types as long as the sender plays by the rules, these types are still perceived as *possible*, and the sender's deviation from the equilibrium path would lead the receiver to recognize these types as *probable*. 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* signaling – that which is not reversed by future deviations. NDOC has been widely used in applied models.[7] On a separate note, NDOC has been criticized as leading to possible equilibrium nonexistence (see Madrigal, Tan, and Werlang [1987] and Nöldeke and van Damme [1990b]). We thus characterize the equilibria conditional on existence, without making any existence claims. Section 5, however, provides some examples of such equilibria.

In order to illustrate our results, we explore a simple model of price signaling. In this model a firm is privately informed about the quality of its product, and sets the price of the product in every period in an attempt to signal this quality and increase the consumers' willingness to pay. We construct a family of informative equilibria in this setting and show that both inefficiently low and inefficiently high prices are equally fit to signal high quality in equilibrium, while the literature (some mentioned in Section 5) typically focused on one of the two as a signal. As a consequence, the price path in informative equilibria is highly indeterminate (even beyond the dichotomy outlined above, as we show), which makes it difficult to test empirically whether price signaling is taking place in a given industry.[8]

The remainder of this paper is organized as follows. Section 2 describes the model. 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 is then explored in Section 4. Section 5 considers an application to price signaling and illustrates how our results can be used to construct informative equilibria. Section 6 concludes. All proofs and a number of summplementary results are contained in Appendix A. Further applications of the general model are briefly outlined in Appendix B.

2 Model

2.1 Primitives

We will be looking at a continuous limit of a discrete-time infinite-horizon game. Time is indexed by $t \in \mathcal{T} \equiv \{0, dt, 2dt, ...\}$; period length dt is assumed to be arbitrarily small. There is a long-lived *agent* (sender) who has some persistent *type* $\theta \in \Theta$, where Θ is a finite ordered set. Alternatively, θ can be the state of the world that the agent is privately informed of.

In every period t the agent has to choose an *action* $a_t \in A$, where A is some compact set. Agent's action choices affect public *outcome* $x_t \in X$, which is a random process, and its distribution at time t depends on θ and a_t (more generally, it can depend on the whole past history).[9] We assume that outcomes never allow to perfectly identify θ: the support of x_t conditional on a_t does not depend

[7] Away from Bond and Zhong [2016], one can also find analogs of NDOC in Grossman and Perry [1986], LeBlanc [1992], Vettas [1997], Kraus, Wilkenfeld, and Zlotkin [1995], Sen [2000], Feinberg and Skrzypacz [2005], Lai [2014], Gryglewicz and Kolb [2019], Smirnov and Starkov [2019, 2020].

[8] A similar point was made by Kaya [2013] in relation to signaling product quality via advertising expenditures.

[9] 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.

on θ. Let $h_t \equiv \{a_s, x_s\}_{s\in\mathcal{T},s<t}$ denote the *history* of agent's past action choices and outcomes up to (but not including) time $t \in \mathcal{T}$. Let $\mathcal{H}_t$ denote the set of all such time-t histories and $\mathcal{H} \equiv \cup_{t\in\mathcal{T}}\mathcal{H}_t$ the set of all such histories.

A passive player (receiver), who is also long-lived, starts the game with some prior *belief* $p_0 \in \Delta(\Theta)$ about the agent's type. The receiver updates this belief p_t every period upon observing the agent's action choice a_t and outcome x_t. In shaping this belief the receiver employs Bayes' rule whenever possible. Off-path beliefs are described further. For specificity, we use $p(h_t)$ to denote belief after history h_t. Hereinafter, belief p_t is also referred to as the agent's *reputation.* 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 the agent receives flow payoff $u^\theta(a_t, p_t)$. For simplicity, we assume that payoff u^θ does not depend on outcome x_t except through reputation p_t. This assumption can be relaxed as long as the payoff effect of outcomes is weak relative to reputation and the assumption of payoff monotonicity as given in 2.3 is adjusted accordingly. Define a *bliss* (myopically optimal) action set for type θ given reputation p as

$$B(p|\theta) \equiv \arg\max_{a\in A} \left\{ u^\theta(a,p) \right\}.$$

All in all, the intra-period timing is as follows: (1) action a_t is chosen; (2) outcome x_t is realized; (3) belief p_t is updated; (4) payoff $u^\theta(a_t, p_t)$ is awarded.

The agent maximizes his expected discounted sum of utilities. A *pure strategy* for the agent of type θ is $\mathbf{a}^\theta : \mathcal{H} \to A$. Given some belief system p, let $U^\theta(\mathbf{a}^\theta|h_t)$ denote the expected discounted continuation utility of type θ from following strategy $\mathbf{a}^\theta$ starting from $h_t \in \mathcal{H}$:

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

where r is the agent's discount rate. The expectation is taken over future histories or, equivalently, over the future outcomes. Strategy $\mathbf{a}^\theta$ is *optimal* for the agent of type θ given belief system $p(h_t)$ if it maximizes his continuation payoff at every history $h_t \in \mathcal{H}$:

$$U^\theta(\mathbf{a}^\theta|h_t) = U^\theta(h_t) \equiv \max_{\mathbf{a}} \left\{ U^\theta(\mathbf{a}|h_t) \right\}. \tag{1}$$

With a slight abuse of notation we let $U^\theta(h_t \cup a) \equiv \mathbb{E}_x \left[\max_{\mathbf{a}} \left\{ U^\theta(\mathbf{a}|h_t \cup (a,x)) \right\} \right]$ denote the highest expected continuation utility that type θ can achieve conditional on taking action a at history h_t.

Finally, a *(behavioral) strategy* for the agent of type θ is $\alpha^\theta : \mathcal{H} \to \Delta(A)$. We let $\alpha^\theta(h_t)(a)$ 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 lowest common denominator among the solution concepts used for this class of games 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 is given by the agent's strategy profile $\alpha = \{\alpha^\theta\}_{\theta\in\Theta}$ with $\alpha^\theta : \mathcal{H} \to \Delta(A)$ and the receiver's belief system $p : \mathcal{H} \to \Delta(\Theta)$ such that:*

1. *the agent's strategy profile α is optimal for all types θ;*
2. *the observer's belief p is updated using Bayes' rule whenever possible.*

PBE is a maximally permissive solution concept. Our main results characterize signaling in all PBE that satisfy NDOC (as defined in the following subsection), hence they will also apply if one imposes additional restrictions or equilibrium refinements on top of PBE with NDOC.

2.3 Assumptions

The two sections above define the primitives of the model but impose only very minimal restrictions on them. Throughout the paper, we will also impose the following assumptions:

(MON) Flow payoff function $u^\theta(a_t, p_t)$ is weakly increasing in p_t w.r.t. FOSD mass shifts. I.e., for any $p', p'' \in \Delta(\Theta)$ such that $p'(\theta') > p''(\theta')$, $p'(\theta'') < p''(\theta'')$ for some $\theta' > \theta''$, and $p'(\theta) = p''(\theta)$ for all $\theta \in \Theta\backslash\{\theta', \theta''\}$, it should be that $u^\theta(a_t, p') \geq u^\theta(a_t, p'')$. Further, for any θ and a_t, if $p_t >_{FOSD} \delta_\theta$ then $u^\theta(a_t, p_t) > u^\theta(a_t, \delta_\theta)$.

(FIN) Equilibrium strategy α has finite support for all $h_t \in \mathcal{H}$.

(NDOC) Process p_t is progressively absolutely continuous. I.e., for any $h_s \supset h_t$, $p(h_s)$ is absolutely continuous w.r.t. $p(h_t)$.

In the above, as well as everything that follows, δ_θ is the Dirac delta: "$p(h_t) = \delta_\theta$ for some $\theta \in \Theta$" is equivalent to saying that $p(\theta|h_t) = 1$ and $p(\theta'|h_t) = 0$ for all $\theta' \neq \theta$.

Below are some conditions equivalent or related to the above, with the relations between the respective pairs of conditions described in the subsequent text and summarized in Lemma 1.

(MON-2) $\Theta = \{H, L\}$ and $u^\theta(a_t, p_t)$ is weakly increasing in $p_t(H)$ and $u^L(a_t, p_t) > u^L(a_t, \delta_L)$ for all a_t and $p_t \neq \delta_L$.

(FIN-M) Action set A is finite.

(NDOC-P) After any action a that is not on path at $h_t \in \mathcal{H}$: $p(h_t \cup (a, x_t)) = \delta_{\min S(h_t)}$ for any $x_t \in X$.[10]

The first assumption, (MON), requires the sender's flow payoff function to be monotone w.r.t. reputation p_t. It is sufficient to have weak monotonicity (w.r.t. FOSD order on beliefs) with the exception that it must always be strictly beneficial to pool with the higher types.[11] This assumption

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

[11] The strict part of (MON) simplifies the analysis but rules out some relevant cases. E.g., it disallows the payoff to be a step function, which is the case when the agent only cares about his reputation being above some cutoff.

will be imposed throughout, and in the two-type model ($|\Theta| = 2$) it will also be the only restriction on payoffs needed for the result. Further, in case of two types it can be written more simply as (MON-2). In other words, if $|\Theta| = 2$ then (MON) and (MON-2) are equivalent.

The next pair of assumptions is (FIN) and (FIN-M). The former is an equilibrium refinement that demands that the sender's equilibrium strategy has finite support at any history. The latter is a restriction on the model in that the action set of the sender is finite. The analysis in the remainder of the paper relies on one of these two conditions to hold in order to avoid the problem of Bayesian inference from zero probability events. This problem is illustrated by the following example.

Example 1. *Let $\Theta = \{L, H\}$, $A = [0, 1]$, and suppose that outcomes are uninformative: $x_t \equiv 0$. Fix some equilibrium and history $h_t \in \mathcal{H}$ therein. Suppose that the high type's strategy $\alpha^H(h_t)$ assigns weight one to action $a = 1$, while the low type's strategy $\alpha^L(h_t)$ mixes uniformly over all actions $a \in [0, 1]$. By Bayes' rule, the receiver's belief at history $h_{t+dt} = h_t \cup (1, 0)$ must assign probability one to type H and probability zero to type L. However, it is not immediate whether the receiver* must *in such cases rule out type L at all histories following h_{t+dt} (which he does if (NDOC) holds).*

It is immediate that (FIN-M) is sufficient for (FIN) to hold, therefore in the remainder of this paper we use (FIN). However, (FIN-M) is a useful reminder that in finite games no further equilibrium refinements are required, apart from (NDOC).

Finally, (NDOC) is the assumption on 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 $h_s \supset h_t$ in $\mathcal{H}$. Note that this applies both on and off the equilibrium path. For the discussion of (NDOC), refer to Secton 1. To simplify the analysis, we strengthen (NDOC) to (NDOC-P), which requires that off the equilibrium path, the receiver's belief $p(h_t)$ must be pessimistic – it must put all weight on the lowest type among those not yet ruled out by the receiver. Given (MON), this condition imposes the strongest possible punishment on the sender for any deviation, among those punishments that satisfy (NDOC). Therefore, we argue that for any equilibrium that satisfies (NDOC), there exists an equivalent one that satisfies (NDOC-P), despite the latter being a stronger condition.

The claims made in this section are summarized by the following lemma.

Lemma 1. *Model assumptions are connected through the following relations.*

1. *If $|\Theta| = 2$ then (MON) and (MON-2) are equivalent.*
2. *(FIN-M) implies (FIN).*
3. *If (MON) holds then for any equilibrium that satisfies (FIN) and (NDOC), there exists a payoff-equivalent and on-path strategy-equivalent equilibrium that satisfies (FIN) and (NDOC-P).*

2.4 Discussion of Assumptions

This section discusses the assumptions that are implicit in the model set-up so that the reader can get a clearer picture of which aspects of the model are important for the results, and which modelling assumptions were made purely for expositional simplicity.

To start with, the model setup includes a number of assumptions that impede with instantaneous separation of types, namely: persistent sender's type θ, vanishing period length, compact action set and finite action costs. The former is required for (NDOC) to have any bite: if type could change over time then the sender's reputation would need constant re-verification, meaning that credible signaling is impossible by design. The other three assumptions are meant to remove any implicit commitment power the sender may have (since in discrete time he can effectively commit to not revise his action until the next period) and to remove the potential of any given single action to be informative. All of these assumptions restrict us to the world in which, according to Weiss [1983] and Admati and Perry [1987], perfect separation is impossible, since this is the world we are interested in exploring.

This paper's message is not about arguing that this is the only plausible set of assumptions in dynamic signaling. Indeed, there are many settings in which a single action has the weight to be informative enough by itself, or the sender has at least some commitment power, or the sender's type is, in fact, volatile. This paper argues instead that there exist real-world settings, to which the aforementioned set of assumptions applies, and we as economists care about characterizing them. For example, one such setting is price signaling by the firm – be it signaling of the firm's product quality to consumers or signaling of its production costs to existing and potential competitors. Prices are typically perfectly observable and can be changed frequently at no cost.

The above raises the question of why we limit ourselves to discrete time, thus giving the sender limited commitment power, rather than exploring a proper continuous time model. The answer is simplicity. While the essence of our results carries over to the continuous time case, their statements become less clear-cut, and the analysis of such model becomes encumbered by the specifics of the continuous-time analysis. Furthermore, one can argue that discrete-time model is more general, since continuous time is its limit (special) case.

On the other hand, finiteness of the type space is a crucial assumption. For example, if type θ is distributed on an interval, then attrition takes a very different form from what is stated in Theorems 1 and 2. Instead of the lowest type separating with positive probability in every period (in which payoff-relevant signaling happens), we could have a positive mass of types at the lower end of the support separating every period. In continuous time, the lower bound of the support of types would increase smoothly over time along the pooling path, shrinking the support; an example of such equilibrium is constructed by Fuchs and Skrzypacz [2010]. Importantly, such attrition could lead to full separation in the limit as $t \to \infty$, unlike in the case with finitely many (but more than two) types. Furthermore, we can no longer guarantee that with a continuum of types, attrition is the only way in which payoff-relevant signaling can proceed.

Finally, we assume that the receiver is passive, and the sender receives utility from reputation p. One may see this as a reduced form of a repeated Stackelberg game in which in every period t the sender first chooses action a_t, which together with the respective outcome x_t determines his current reputation p_t, and the receiver then responds with some action b_t, after which both players $i \in \{S, R\}$ receive utilities $u_i^\theta(a_t, b_t)$.[12] This is a standard reduction used both in signaling literature (Kaya [2009], Roddie [2012a,b]) and other literatures (e.g., Bayesian Persuasion – see Kamenica and Gentzkow [2011]). Given that this is by now a standard technique, we do not describe the full game in order to economize on notation. However, our results can be easily extended to both repeated Stackelberg games in which the sender and the receiver act in sequence in every period, and (with slightly more effort) to repeated games in which both act simultaneously.

[12]The receiver's utility may depend on true θ as long as the receiver does not observe his own utility flow.

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 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 new insight 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. All this is summarized by the second part of the theorem. The statement does not claim existence of any such separating actions, since they, of course, need not exist in any given case. However, Section 5 presents an example of a setting in which such informative equilibrium exists.

Theorem 1. *Suppose that $\Theta = \{L, H\}$ and (MON-2) holds. In any equilibrium such that (FIN) and (NDOC) 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 $U^\theta(h_t \cup a)$ is the same across such a for all θ.*
2. *if $\alpha^H(a|h_t) = 0$ and $\alpha^L(a|h_t) > 0$ then $a \in B(\delta_L|L)$ and $U^L(h_t \cup a) = U^L(h_t \cup a')$ 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 equilibrium. Firstly, any separating action perfectly identifies the low type, meaning it cannot reflect any signaling motives, and must hence be myopically optimal for the low type. Secondly, if the low type mixes between pooling and separating, then he must be indifferent between the two – meaning that pooling with the high type must yield exactly the same expected payoff for the low type as separation. Gains from pooling in this scenario (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. In other words, if your game fits the following framework:

- dynamic game with continuous time or short time intervals or patient players,
- binary state of the world known by one player but not other(s),
- the informed player has a preference over other(s)' beliefs, and the direction of this preference does not depend on the state,
- the informed player chooses an action every period but cannot verifiably reveal the state (action set is independent of type),

then *the only* informative equilibrium structure that can arise in this game (unless you are willing to allow for NDOC-nonconformant beliefs off the equilibrium path) is attrition. Under attrition, the high type is playing some pooling action, while the low type mixes between that and a separating action.

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 assumption on payoffs:

(SC) $U^\theta(\mathbf{a}|h_t)$ satisfies single-crossing in $(\theta, \mathbf{a})$ at all $h_t \in \mathcal{H}$. I.e., for any $\mathbf{a}', \mathbf{a}'' \in \cup_{h_t \in \mathcal{H}} \cup_{\theta \in S(h_t)} \arg\max_{\mathbf{a}} U^\theta(\mathbf{a}|h_t)$ and all $h_t \in \mathcal{H}$, function $\mathcal{U}(\theta) \equiv U^\theta(\mathbf{a}''|h_t) - U^\theta(\mathbf{a}'|h_t)$ either crosses zero at most once, or is identically zero.

This assumption belongs to a family of single-crossing conditions widely encountered in the literature on signaling, monotone comparative statics, and mechanism design.[13] The purpose of our condition is standard: to ensure that the agent's preferences over strategies satisfy a kind of monotonicity w.r.t. his type. Our condition, however, has three distinctive features which differentiate it slightly from other single-crossing conditions in the literature.

Feature 1. This is an assumption about an equilibrium object, since belief system $p(h_t)$ – which enters $U^\theta(\mathbf{a}|h_t)$ – is endogenous to equilibrium. The simplest way to justify the assumption in this respect is strengthening the assumption to all combinations of actions and beliefs $(\mathbf{a}, p)$, i.e., to assume that

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

satisfies single-crossing for all pairs $(\mathbf{a}', p'), (\mathbf{a}'', p'')$.

Feature 2. Unlike most standard single-crossing assumptions, (SC) does not require that $\mathbf{a}'' > \mathbf{a}'$ (and that single-crossing happens from below).[14] In contrast, we also require the statement to hold for any pair of unordered strategies and respective belief profiles. The reason for that is while the grand set of choices $(A \times \Delta(\Theta))^{\mathcal{H}}$ is a lattice w.r.t. the product order for some given order on A, the subset of these choices available to the sender at any given history is not necessarily its sublattice. The standard monotone comparative statics results are then of limited use. On the upside, however, we only require that single-crossing is satisfied for those strategies that may be optimal for any type. I.e., if one can rule out some actions as certainly suboptimal at certain histories, these actions may be safely ignored.

[13] 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].

[14] The order implied here is the product order on the set $(A \times \Delta(\Theta))^{\mathcal{H}}$ of all collections $(a(h_t), p(h_t))$, composed of some order on A (although we have not imposed any) and FOSD order on $\Delta(\Theta)$.

Feature 3. (SC) is a condition on the expectation of a discounted sum $\mathbb{E}\sum_t e^{-rt}u^\theta(a_t, p_t)$ rather than on the flow utility $u^\theta(a, p)$. While the latter would be more preferable, aggegating 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 due to feature 2 above.

All of the above means that (SC) is quite a non-trivial condition and may be difficult to verify in many models. If anything, verifying (SC) might as well be the main impediment to exploiting our results in applied models. However, this task is not impossible. Some examples of applied models, in which payoff functions can be easily verified to satisfy (SC) are presented in Section 5 and Appendix B.

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 equilibrium at any history. We begin, however, by stating a weaker result which 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 equilibrium of the game can be split into two classes. The first class consists of pooling strategies played by all types. While there may be many such strategies nominally, 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 differ 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 following results we need to introduce some additional notation and definitions. Firstly, denote the two boundaries of the belief support as $\bar{S}(h_t) \equiv \max S(h_t)$ and $\underline{S}(h_t) \equiv \min S(h_t)$ respectively. Furthermore, in a manner similar to type support S, given an equilibrium strategy profile let us define action support as

$$A^\theta(h_t) \equiv \left\{a \in A | \alpha^\theta(h_t)(a) > 0\right\},$$
$$A(h_t) \equiv \cup_{\theta \in S(h_t)} A^\theta(h_t).$$

We say that a pure strategy $\mathbf{a}$ *arrives at* $h_t = \{a_s(h_t), x_s(h_t)\}_{s\in\mathcal{T}, s<t}$ – and denote it as $\mathbf{a} \supset h_t$ – if $\mathbf{a}(h_{t'}) = a_{t'}(h_t)$ for any $h_{t'} = \{a_s(h_t), x_s(h_t)\}_{s\in\mathcal{T}, s<t'}$ with $t' < t$.[15] Further, say that $\mathbf{a}$ is *on path for* θ *at* h_t if $\mathbf{a} \supset h_t$ and $\mathbf{a}$ is on path according to type θ's equilibrium strategy α starting from h_t, and $\mathbf{a}$ is *on path at* h_t if it is on path at h_t for some $\theta \in S(h_t)$. Note that action supports and "on-pathness" only work properly if (FIN) is satisfied. To give a counterexample, if some $\alpha^\theta(h_t)$ prescribed mixing uniformly over an interval of actions $a \in [0, 1]$ in violation of (FIN), then none of these actions would be included in the support (since they are all played with zero probability) and none of them would be on path according to the definitions above.

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

Definition 2. *Fix an equilibrium and history* $h_t \in \mathcal{H}$. *Any two pure strategies* $\mathbf{a}', \mathbf{a}'' \supset h_t$ *are:*

- payoff-distinct at h_t *if there exists* $\theta \in S(h_t)$ *such that* $U^\theta(\mathbf{a}'|h_t) \neq U^\theta(\mathbf{a}''|h_t)$*;*

[15] We also use notation $\mathbf{a} \supset h_t \cup a$ to state "$\mathbf{a} \supset h_t$ and $\mathbf{a}_t = a$".

- payoff-equivalent at h_t *if they are not payoff-distinct at* h_t.

Note also that while using the notation for full pure strategies, throughout the whole analysis we actually work with continuation strategies from some history $h_t \in \mathcal{H}$. When discussing strategies conditional on some history h_t we ignore all game paths that are ruled out by h_t. In particular, two pure strategies $\mathbf{a}', \mathbf{a}'' \supset h_t$ that prescribe the same actions at all $h_s \supseteq h_t$ but differ at some $h_s \not\supseteq h_t$ are treated as the same strategy for all means and purposes. We avoid introducing the continuation strategies explicitly in order to economize on notation, which is quite heavy as is.

The result can now be stated as follows.

Proposition 1. *Suppose the payoff function* u^θ *satisfies (MON) and (SC). Fix an equilibrium such that (FIN) and (NDOC) hold. Fix some history* $h_t \in \mathcal{H}$. *Then, defining* $\underline{\theta} \equiv \underline{S}(h_t)$, *the following hold:*

1. *all pure strategies* $\mathbf{a}'$ *on path at* h_t *for any* $\theta \in S(h_t)\backslash\underline{\theta}$ *are payoff-equivalent and optimal for all* $\theta \in S(h_t)$ *at* h_t, *and at least one of these strategies is on path for* $\underline{\theta}$ *at* h_t;
2. *any pure strategy* $\mathbf{a}''$ *that is on path at* h_t *and payoff-distinct at* h_t *from any such* $\mathbf{a}'$ *is only on path for* $\underline{\theta}$.

The proposition implies, in particular, that any pure strategy $\mathbf{a}$ that is on path for some type $\theta \in S(h_t)$ is also on path for 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.

The proposition above is stated in terms of *strategies* rather than actions, and so provides only limited insight into how equilibrium actions look in any given period. That is, if $\mathbf{a}'$ as defined in the proposition is unique then it is relatively straightforward that in every period there will be some single pooling action as prescribed by $\mathbf{a}'$ that all types above $\underline{\theta}$ will play for sure, while the lowest type will somehow mix between this pooling action and some number (between zero and infinity) of separating actions. The challenge, however, comes from possible non-uniqueness of $\mathbf{a}'$. If there are many pooling actions, then they may be informative in that different pooling actions convey different information – even though (or exactly because) all types are indifferent between them. The following sections explore this issue in more detail and works around it to characterize the within-period signaling outcomes.

4.3 Payoff-Relevant Signaling

To talk about signaling in relation to individual actions (rather than whole strategies), we need to define more precisely what "signaling" means in a dynamic context with many types. It is clear that if all types pool on the same action in a given period then no information is revealed, while if every type plays an action different from all others then full separation occurs, which is the most informative signaling outcome. The two grey zones are partitioning – when, for example, some types play action a' and some others play a'' – and mixing, – when one type plays action a' for sure and another type mixes between two actions a' and a''. Both of the aforementioned outcomes are usually dubbed as "semi-separation" in static settings and considered informative outcomes in signaling models. In dynamics, however, there are further complications.

In a dynamic setting, it matters – for both sender's payoff and reputation – not only what action the sender plays in a given period, but also his past and future actions. In particular, a single costly action is inconsequential by itself, having only infinitesimal effect on payoff and, as a result, reputation – unless, that is, it is backed up by costly actions at future periods. Symmetrically, future actions also have the power to negate the payoff consequences of past actions. This issue is illustrated by the following example, and its consequences are discussed further.

Example 2. *Suppose* $\Theta = \{1, 2, 3\}$*, types are ex ante equiprobable,* $A = \mathbb{R}_+$*, and* $u^\theta(a, p) = \mathbb{E}_p(\theta) - a$*. Then the following would be an equilibrium: type* $\theta = 2$ *plays some* a'' *at* $t = 0$ *and* $a = 0$ *at all* $t \geq dt$*, while types* $\theta = 1, 3$ *play* $a' = a''(1 - e^{-rdt})$ *at all* $t \geq 0$*. This is a PBE of the game as long as* $a'' \leq 1$*. In this PBE some information about type is conveyed in period zero – namely, type* $\theta = 2$ *separates from* $\theta = 1, 3$*, – but this signaling is not relevant to the sender's payoff.*

The same is not necessarily true for the receiver. In particular, we can think of this example as a game between a worker (sender) and a firm (receiver), where a *is the worker's effort and* θ *is his ability. Suppose the receiver's flow payoff is given by* $v(a, p, \theta) = \theta a^2 - \mathbb{E}_p(\theta)$*, with the first term being the worker's output, and* $\mathbb{E}_p(\theta)$ *in both players' payoffs is the worker's wage, dictated by the market. In this case the firm's expected discounted profit from hiring a worker of type* $\theta = 2$ *at* $t = 0$ *equals* $2(a'')^2 - \frac{2}{1-e^{-rdt}}$*, while that from hiring a worker of type* $\theta \in \{1, 3\}$ *is* $\frac{2\left(a''(1-e^{-rdt})\right)^2 - 2}{1-e^{-rdt}} = 2(a'')^2(1 - e^{-rdt}) - \frac{2}{1-e^{-rdt}}$*, which is strictly less. This is reversed for hiring decisions made at* $t \geq dt$*.*

The example above illustrates that arbitrary information can in principle be conveyed via *payoff-irrelevant signaling* – when different types play different actions, but nonetheless all types are indifferent between all actions and respective continuations. This can be seen as a manifestation of "cheap talk": for small dt, actions in a single period are effectively costless, and so can be used as a pure communication device with no regard for action costs. Situations in which informative communication arises through cheap talk have been studied extensively, see Sobel [2013] for a recent survey. Further, one can claim that there *will* often be scope for such cheap/payoff-irrelevant communication in our model. This is due to multi-dimensionality of reputation p: with more than two types there are bound to be situations in which an agent is indifferent between two kinds of average reputation – one that says the agent is of average type for sure and the one that says the agent is of either high or low type with comparable probabilities, just like in the example above.[16] Consequently, in this paper we focus on *payoff-relevant signaling* – communication that relies on heterogeneity across the agent's types of costs or benefits from different actions, in the spirit of Spence [1973]. The information revealed through payoff-relevant signaling is that which cannot be communicated via plain cheap talk. Our contribution to characterizing the informative outcomes in dynamic signaling games without commitment should then be seen as complementary to that of the cheap talk literature.

[16] Battaglini [2002] and Chakraborty and Harbaugh [2007, 2010] explore cheap talk with multidimensional information (which our setting is an instance of in case $|\Theta| > 2$) and show that there generally exist informative equilibria in which the receiver can perfectly learn about $N - 1$ out of N dimensions of uncertainty from a single sender.

We now provide the formal definitions of payoff-relevant and irrelevant signaling in our setting.

Definition 3. *Fix an equilibrium 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 $U^\theta(h_t \cup a') \neq U^\theta(h_t \cup a'')$.*
- Payoff-irrelevant signaling *happens at h_t if there exist $a', a'' \in A(h_t)$ such that $p(h_t \cup a') \neq p(h_t \cup a'')$ but $U^\theta(h_t \cup a') = U^\theta(h_t \cup a')$ for all $\theta \in S(h_t)$.*

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.

4.4 From Action Profiles to Actions

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, albeit at the cost of restricting model scope to payoff functions that satisfy (SC).

Theorem 2. *Suppose the payoff function u^θ satisfies (MON) and (SC). Fix an equilibrium such that (FIN) and (NDOC) 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. *$A(h_t) \cap B\left(\delta_{\underline{\theta}} | \underline{\theta}\right)$ is nonempty, and any $\underline{a}$ in the intersection is on path only for $\underline{\theta}$ at h_t;*
3. *any action $\bar{a} \in A(h_t) \backslash B\left(\delta_{\underline{\theta}} | \underline{\theta}\right)$ is optimal at h_t for all $\theta \in S(h_t)$.*

What the theorem says is that in any 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 $\underline{\theta}$, higher types may select different actions despite all types being indifferent between all of these actions (and the continuations they induce).

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 *flow payoffs* the low type gets from the separating and pooling actions must be the same. The low type must be indifferent between separating at t and $t + dt$, so one period of pooling must 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* $h_{t+dt} \equiv h_t \cup (\bar{a}, x)$ *for some* $\bar{a}$ *and all* x *in the support. Then such* $\bar{a}$ *must satisfy* $\mathbb{E}_x \left[u^{\underline{\theta}} \left(\bar{a}, p(h_t \cup (\bar{a}, x)) \right) | \underline{\theta} \right] = u^{\underline{\theta}} \left(\underline{a}, \delta_{\underline{\theta}} \right)$.

Finally, the theorem applies to all histories, including those off the equilibrium path. Applying it inductively starting from the root history, we get the following corollary, which states that in the absence of payoff-irrelevant signaling, only the lowest type $\min \Theta$ can ever separate from the rest, while the remaining ones can never separate from one another. It is worth noting that there may be histories h_t at which $p(h_t)$ assigns arbitrarily small weight to the lowest type. What is important is that this type can never be ruled out completely along the pooling path.

Corollary 2. *In any 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. $S(h_t) = \{\min \Theta\}$.

The corollary above together with Theorem 2 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 $\underline{a} \in B \left(\delta_{\underline{\theta}} | \underline{\theta} \right)$ only taken by the lowest type $L \equiv \min \Theta$ and a pooling action $\bar{a}$ that satisfies the condition in the last part of the theorem – the latter action will be played by L 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 L separates in that given period. Hence by changing these probabilities we will be able to sustain different pooling actions $\bar{a}$ in equilibrium. Finally, we need to verify that from time T onwards, the pooling strategy is such that L is exactly indifferent at T (or the last period before T) between separating and following this pooling path.

The following section takes this cookbook and uses it to construct an informative equilibrium in a concrete setting.

5 Application to Price Signaling

5.1 Background

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 in an informative equilibrium.

5.2 Model Setup

There is a long-lived firm i that faces a continuum of consumers $j \in [0,1]$ every period $t \in \mathcal{T} \equiv \{0, dt, 2dt, ...\}$, where period length dt is "small". The firm offers for sale a single product of privately known quality $\theta \in \Theta = \{H, L\}$ (in Section 5.4 we discuss the case $|\Theta| > 2$). 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. A consumer'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 $\mathbb{E}[\theta] \geq \frac{a_t}{v_j}$.

The population of consumers is renewed every period. The newly arriving consumers base their belief $p(h_t)$ about the product quality on the prior $p_0 \in \mathtt{int}\Delta(\Theta)$, the whole price path $\{a_s\}_{s \leq t}$, and product reviews as described below.[17] With probability $1 - e^{-\phi dt} \approx \phi dt$ for $\phi \in [0,1)$ the population of consumers in a given period 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 flow profit of a firm of type θ receives in period t after setting price a_t is then given by

$$u^\theta(a_t, p_t) \equiv a_t \left(1 - \frac{a_t}{\mathbb{E}[\theta|p_t]}\right)_+ . \tag{2}$$

According to Theorem 1, (payoff-relevant) signaling must necessarily take the form of attrition. Under such attrition signaling, type $\theta = L$ mixes between some pooling price a_t^p and separating to bliss price $a^{L|L} = \frac{L}{2}$, whereas type H sets the pooling price. We then show that Theorem 2 applies as well, meaning that the result translates to the case $|\Theta| > 2$.

5.3 Equilibrium

We will be looking for an equilibrium of the game that satisfies (NDOC-P) and (FIN). 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., $p(h_t) \neq \delta_\theta$). 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 ($p(h_t) = \delta_L$) then there are no signaling motives – it sets price $a_t^{\theta|L} \equiv \frac{L}{2}$ and earns profit $u^\theta\left(a_t^{\theta|L}, \delta_L\right) = \frac{L}{4}$. Similarly, if $p(h_t) = \delta_H$ then $a_t^{\theta|L} \equiv \frac{H}{2}$ and $u^\theta\left(a_t^{\theta|H}, \delta_H\right) = \frac{H}{4}$.

[17] Here $\mathtt{int}\Delta(\Theta)$ denotes the interior of $\Delta(\Theta)$, i.e., we assume that the prior p_0 does not rule out any states in Θ. This assumption is not strictly necessary, but all results are trivial without it.

[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.

The latter is due to (NDOC) – without this assumption we could enforce a wide spectrum of prices using the threat of losing reputation in case the firm deviates from some prescribed price level. With (NDOC) such threats are ruled out, hence the firm's behavior at histories with $p(h_t) = \delta_H$ must be myopically optimal.

We begin by deriving the pooling price a_t^p that renders the low type indifferent between separating and pooling. The former yields the continuation value equal to

$$U^L(a^{L|L}|h_t) = \frac{dt}{1-e^{-rdt}}\frac{L}{4} \approx \frac{L}{4r},$$

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

$$U^L(a_t^p|h_t) = a_t^p\left(1 - \frac{a_t^p}{\mathbb{E}[\theta|p_t]}\right)dt + e^{-rdt}\mathbb{E}_{x_t}U^L(h_t \cup (a_t^p, x_t)), \tag{3}$$

where $p_t \equiv p(h_t \cup (a_t^p, \varnothing))$ and the continuation value can be written as

$$\mathbb{E}_{x_t}U^L(h_t \cup (a_t^p, x_t)) = (1 - e^{-\phi dt})U^L(h_t \cup (a_t^p, L)) + e^{-\phi dt}U^L(h_t \cup (a_t^p, \varnothing)).$$

After a bad review $x_t = L$ the consumers are sure that the firm is bad, i.e., $p(h_t \cup (a_t^p, L)) = \delta_L$, so we have $U^L(h_t \cup (a_t^p, L)) = \frac{L}{4r}$. After no review $x_t = \varnothing$ the consumers' belief is inconclusive, hence in our construction signaling should continue. This means L must be indifferent between pooling and separating once again, so $U^L(h_t \cup (a_t^p, \varnothing)) = \frac{L}{4r}$ as well. Finally, L's indifference at h_t yields $U^L(a_t^p|h_t) = U^L(a^{L|L}|h_t) = \frac{L}{4r}$. Plugging all of these into (3), we obtain that flow payoff from pooling must coincide with that from separating:

$$\frac{L}{4} = a_t^p\left(1 - \frac{a_t^p}{\mathbb{E}[\theta|p_t]}\right).$$

The solution to the above is given by

$$a_t^p = \frac{\mathbb{E}[\theta|p_t]}{2}\left[1 \pm \sqrt{1 - \frac{L}{\mathbb{E}[\theta|p_t]}}\right]. \tag{4}$$

Therefore, for a fixed $\mathbb{E}[\theta|p_t]$ we have two equivalid candidates for the pooling price 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 a_t^p (whichever root to (4) we choose) is a function of the seller's reputation

$$p(h_t \cup (a_t^p, \varnothing)) = \frac{p(h_t)}{1 - \lambda_t(1 - p(h_t))}, \tag{5}$$

where λ_t is the probability with which type L separates at h_t. In particular, we can choose these probabilities freely and construct an equilibrium for arbitrary λ_t. Conversely, if we are restricted in our choice of a_t^p – e.g., if this pooling price must amount to an integer number of dollars, – this constrains the set of $p(h_t \cup a_t^p)$ and, consequently, λ_t that we can implement in equilibrium at any given history h_t.[19]

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 setting the pooling price a_t^p is optimal for H. His continuation value from doing so is

$$U^H(a_t^p|h_t) = \frac{L}{4}dt + e^{-rdt}\left[(1-e^{-\phi dt})\frac{H}{4r} + e^{-\phi dt}U^H(h_t \cup (a_t^p, \varnothing))\right].$$

Indeed, his flow payoff is the same as for L (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 $\frac{H}{4}$ in every future period. The same, however, applies to any history with $p(h_t) \in \mathtt{int}\Delta(\Theta)$, hence $U^H(h_t \cup (a_t^p, \varnothing)) = U^H(a_t^p|h_t)$, which allows us to conclude that $U^H(a_t^p|h_t) = \frac{1}{4r}\frac{rL+\phi H}{r+\phi}$. Setting any price other than a_t^p results in $p(h_{t+dt}) = \delta_L$ (by (NDOC-P)), and hence yields value of at most $\frac{L}{4r} < \frac{1}{4r}\frac{rL+\phi H}{r+\phi}$. Therefore, pooling is indeed optimal for H. All of the above proves the following proposition.

Proposition 2. *The following constitutes an equilibrium of the price signaling game for any profile of λ_t. At any history $h_t \in \mathcal{H}$:*

1. *if $p(h_t) = \delta_\theta$ then all types play $\frac{\theta}{2}$ and $p(h_{t+dt}) = p(h_t)$ for all $h_{t+dt} \supset h_t$;*
2. *if $p(h_t) \in \mathtt{int}\Delta(\Theta)$ then:*
   - *(a) type $\theta = H$ plays a_t^p as given by (4) with probability one;*
   - *(b) type $\theta = L$ plays a_t^p w.p. $1-\lambda_t$ and $a^{L|L} = \frac{L}{2}$ w.p. λ_t;*
   - *(c) belief $p(h_t \cup (a_t^p, \varnothing))$ is computed according to (5); $p(h_t \cup (a_t^p, H)) = \delta_H$, and $p(h_{t+dt}) = \delta_L$ for all other $h_{t+dt} \supset h_t$.*

5.4 Many Types

The equilibrium described in the Proposition 2 traslates immediately to the case with $|\Theta| > 2$. In that case all types $\theta \in \Theta \backslash L$ would behave as the high type above. Therefore, signaling through attrition is still possible. We now show that Theorem 2 can be applied in this problem as well to verify that payoff-relevant signaling is possible *only* through attrition, even despite the presence of informative reviews.

To do so, we need to verify that (SC) holds for the seller's payoff function. 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 sets a myopically optimal price, thus obtaining payoff $\frac{\theta}{4}$ per period. Therefore, we only need to consider strategies that are non-trivial at histories with non-degenerate beliefs $p(h_t)$. The agent's value from following a given strategy $\mathbf{a}$ starting from any such history h_t in the reduced game is given by

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

[19] Such an integer constaint will also apply to the bliss action $a^{L|L}$, but this would not affect the overall argument.

where $\mathbb{E}_s[\theta|\mathbf{a}] \equiv \mathbb{E}\left[\theta|h_t \cup \mathbf{a}_{(t,s]}\right]$. It is easy to see that this value function satisfies (SC): for any $\mathbf{a}', \mathbf{a}''$ we have

$$\begin{aligned}\mathcal{U}(\theta) &\equiv U^\theta(\mathbf{a}''|h_0) - U^\theta(\mathbf{a}'|h_0) \\ &= \sum_{s\in\mathcal{T}, s\geq t} e^{-(r+\phi)(s-t)-\phi dt}\left[\mathbf{a}''_t\left(1-\frac{\mathbf{a}''_t}{\mathbb{E}_s[\theta|\mathbf{a}'']}\right)_+ - \mathbf{a}'_t\left(1-\frac{\mathbf{a}'_t}{\mathbb{E}_s[\theta|\mathbf{a}']}\right)_+\right]dt,\end{aligned}$$

which is independent of θ. Therefore, the conclusions of Theorem 2 apply, and payoff-relevant signaling is only possible through attrition.

5.5 Takeaways

The price signaling model presented in this section does, despite being highly stylized, demonstrate that:

1. informative price signaling is possible without commitment, cost advantages, and with or without consumers learning from experiences;
2. 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;
3. multiple informative equilibria exist that differ in the speed of separation;
4. due to the above, the empirical identification of price signaling in a given market is a daunting task.

6 Conclusion

This paper explores a model of dynamic signaling without commitment. 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 assumed signaling to be impossible in such setting, unless strong assumptions about off-equilibrium-path beliefs are adopted. This paper overturns this view, demonstrating that signaling is, in fact, possible even under reasonable off-path beliefs.

Contrary to the literature, we allow for *suggestive* signaling rather than requiring *conclusive* separation. We show that such signaling must necessarily happen through the attrition of the lowest type of the agent. In this attrition scenario, all types pool on the same action (or split across some payoff-equivalent actions), while the lowest type also plays some separating action with positive intensity.

The paper also contains a methodological contribution. In particular, we demonstrate the importance of seemingly innocuous assumptions regarding the type space for the conclusions one obtains. In particular, our results imply that perfect learning is possible in the limit as $t \to \infty$ in the model with two sender types but not possible with finitely many types, while the literature demonstrates that perfect learning is possible with a continuum of sender types.

Finally, we explore an application of our results to a model of dynamic price signaling. We construct an informative equilibrium in which prices set by the firm contain information about the quality of its product. We show that price signaling can happen through both inefficiently low and inefficiently high prices, thus reconciling some of the disparate conclusions in the literature and arguing that empirical identification of price signaling in the data is a complicated venture.

References

A. R. Admati and M. Perry. Strategic delay in bargaining. *Review of Economic Studies*, 54(3): 345–364, 1987.

K. K. Aköz, C. E. Arbatli, and L. Çelik. Manipulation through biased product reviews. *SSRN Electronic Journal*, 2017. doi: 10.2139/ssrn.3068345.

S. Athey. Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1):187–223, 2002.

L. M. Ausubel, P. Cramton, and R. J. Deneckere. Bargaining with incomplete information. *Handbook of game theory*, 3:1897–1945, 2002.

K. Bagwell and M. H. Riordan. High and declining prices signal product quality. *The American Economic Review*, pages 224–239, 1991.

J. S. Banks and J. Sobel. Equilibrium selection in signaling games. *Econometrica: Journal of the Econometric Society*, 55(3):647–661, 1987.

M. Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, 2002.

P. Beaudry and M. Poitevin. Signalling and renegotiation in contractual relationships. *Econometrica*, 61(4):745–782, 1993.

P. Bond and H. Zhong. Buying high and selling low: Stock repurchases and persistent asymmetric information. *The Review of Financial Studies*, 29(6):1409–1452, June 2016. doi: 10.1093/rfs/hhw005.

A. Chakraborty and R. Harbaugh. Comparative cheap talk. *Journal of Economic Theory*, 132(1): 70–94, 2007.

A. Chakraborty and R. Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5): 2361–82, 2010.

I.-K. Cho and D. M. Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221, 1987. doi: 10.2307/1885060.

B. Daley and B. Green. Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504, 2012.

T. De Angelis, E. Ekström, and K. Glover. Dynkin games with incomplete and asymmetric information. *arXiv preprint arXiv:1810.07674*, 2018.

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

F. Dilmé. Noisy signaling in discrete time. *Journal of Mathematical Economics*, 68, 2017.

F. Dilmé and F. Li. Dynamic signaling with dropout risk. *American Economic Journal: Microeconomics*, 8(1):57–82, Feb. 2016. doi: 10.1257/mic.20120112.

Y. Feinberg and A. Skrzypacz. Uncertainty about uncertainty and delay in bargaining. *Econometrica*, 73(1):69–91, 2005.

W. Fuchs and A. Skrzypacz. Bargaining with arrival of new traders. *American Economic Review*, 100(3):802–36, 2010.

S. J. Grossman and M. Perry. Sequential bargaining under asymmetric information. *Journal of Economic Theory*, 39(1):120–154, 1986.

S. Gryglewicz and A. Kolb. Strategic pricing in volatile markets. *SSRN Electronic Journal*, 2019. doi: 10.2139/ssrn.3154372.

S. Heinsalu. Dynamic noisy signaling. *American Economic Journal: Microeconomics*, 10(2):225–249, May 2018. doi: 10.1257/mic.20160336.

E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

A. Kaya. Repeated signaling games. *Games and Economic Behavior*, 66(2):841–854, 2009. doi: 10.1016/j.geb.2008.09.030.

A. Kaya. Dynamics of price and advertising as quality signals: anything goes. *Economics Bulletin*, 2(1):1556–1564, 2013.

A. Kaya and K. Kim. Trading dynamics with private buyer signals in the market for lemons. *The Review of Economic Studies*, 85(4):2318–2352, 2018.

A. Kirmani and A. R. Rao. No pain, no gain: A critical review of the literature on signaling unobservable product quality. *Journal of marketing*, 64(2):66–79, 2000.

S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial intelligence*, 75(2):297–345, 1995.

J.-J. Laffont and D. Martimort. *The theory of incentives: the principal-agent model.* Princeton university press, 2002. ISBN 9780691091846.

E. K. Lai. Expert advice for amateurs. *Journal of Economic Behavior & Organization*, 103:1–16, 2014.

G. LeBlanc. Signalling strength: limit pricing and predatory pricing. *The RAND Journal of Economics*, pages 493–506, 1992.

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

H. E. Leland and D. H. Pyle. Informational asymmetries, financial structure, and financial intermediation. *The journal of Finance*, 32(2):371–387, 1977.

V. Madrigal, T. C. C. Tan, and S. R. d. C. Werlang. Support restrictions and sequential equilibria. *Journal of Economic Theory*, 43(2):329–334, Dec. 1987. doi: 10.1016/0022-0531(87)90063-9.

P. Milgrom and J. Roberts. Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, pages 443–459, 1982a.

P. Milgrom and J. Roberts. Predation, reputation, and entry deterrence. *Journal of economic theory*, 27(2):280–312, 1982b.

P. Milgrom and J. Roberts. Price and advertising signals of product quality. *Journal of political economy*, 94(4):796–821, 1986.

P. Milgrom and C. Shannon. Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, pages 157–180, 1994.

R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997. ISBN 9780674341166.

G. Nöldeke and E. van Damme. Signalling in a dynamic labour market. *Review of Economic Studies*, 57(1):1–23, 1990a. doi: 10.2307/2297540.

G. Nöldeke and E. van Damme. Switching away from probability one beliefs. mimeo, 1990b.

J. K.-H. Quah and B. Strulovici. Aggregating the single crossing property. *Econometrica*, 80(5): 2333–2348, 2012.

J. G. Riley. Silver signals: Twenty-five years of screening and signaling. *Journal of Economic literature*, 39(2):432–478, 2001.

C. Roddie. Signaling and reputation in repeated games, I: Finite games. *SSRN Electronic Journal*, 2012a. doi: 10.2139/ssrn.1994378.

C. Roddie. Signaling and reputation in repeated games, II: Stackelberg limit properties. *SSRN Electronic Journal*, 2012b. doi: 10.2139/ssrn.2011835.

A. Sen. Multidimensional bargaining under asymmetric information. *International Economic Review*, 41(2):425–450, 2000.

A. Smirnov and E. Starkov. Timing of predictions in dynamic cheap talk: experts vs. quacks. *University of Zurich, Department of Economics Working Papers*, 334, 2019.

A. Smirnov and E. Starkov. Bad news turned good: Reversal under censorship. *American Economic Journal: Microeconomics*, forthcoming, 2020. doi: 10.1257/mic.20190379.

J. Sobel. Giving and receiving advice. *Advances in economics and econometrics*, 1:305–341, 2013.

M. Spence. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374, Aug. 1973. doi: 10.2307/1882010.

I. A. Strebulaev, H. Zhu, and P. Zryumov. Optimal issuance under information asymmetry and accumulation of cash flows. *Rock Center for Corporate Governance at Stanford University Working Paper*, 164, 2016.

J. M. Swinkels. Education signalling with preemptive offers. *Review of Economic Studies*, 66(4): 949–970, 1999.

N. Vettas. On the informational role of quantities: Durable goods and consumers' word-of-mouth communication. *International Economic Review*, pages 915–944, 1997. doi: 10.2307/2527222.

D. R. Vincent. Dynamic auctions. *Review of Economic Studies*, 57(1):49–61, 1990. doi: 10.2307/2297542.

A. Weiss. A sorting-cum-learning model of education. *Journal of Political Economy*, 91(3):420–442, 1983.

M. Whitmeyer. On optimal transparency in signaling. *arXiv preprint arXiv:1902.00976*, 2019.

Appendix A. Proofs and Supplementary Results

A.1 Proofs: Preliminaries

Our 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.

Lemma 2. *In any equilibrium that satisfies (NDOC), at any $h_t \in \mathcal{H}$, if $|S(h_t)| = 1$ then for all θ and all $h_s \supseteq h_t$:* $\alpha^\theta \left(B(\delta_{S(h_t)}|\theta) \mid h_s\right) = 1$.

Proof. By (NDOC), for all $h_s \supseteq h_t$: $p(h_s) = \delta_{S(h_t)}$. In particular, $p(h_s)$ is independent of all actions and outcomes during $[t, s)$. Therefore, the solution to (1) is given by pointwise maximization of the flow utility. □

Lemma 2 above is the direct consequence of (NDOC): actions cannot affect 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. Parts 1 and 2 are trivial. For part 3, denote the original equilibrium strategy and belief profile as α_1 and p_1 respectively. Construct the new equilibrium (α_2, p_2) by setting $\alpha_2(h_t) = \alpha_1(h_t)$ and $p_2(h_t) = p_1(h_t)$ for all on-path histories h_t. Then for all off-path histories h_t set $p_2(h_t) = \delta_{\underline{S}(h)}$ where $h \subset h_t$ is the last on-path history preceding h_t. The strategy $\alpha_2(h_t)$ for off-path histories h_t is set in conformance with Lemma 2.

Belief profile p_2 will then satisfy (NDOC-P) and be consistent with the strategy α_2. The strategy itself will be optimal at off-path histories by Lemma 2. Optimality of α_2 at any on-path history h_t can be verified by observing that $U^\theta(h_t \cup a)$ is the same in both equilibria for all $a \in A(h_t)$ and weakly smaller in the newly constructed equilibrium for $a \in A \backslash A(h_t)$. I.e., the choice between any pair of off-path actions is unaffected by the off-path modifications, while deviations to off-path actions are less appealing in the new equilibrium. Therefore, (α_2, p_2) is an equilibrium. $\square$

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 $p(h_t \cup (a, x)) = \delta_H$ for any $x \in X$ by (NDOC). 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. The utility is bounded, hence for dt small enough deviating to a at h_t is optimal for L – a contradiction.

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 $U^H(h_t \cup a) = U^H(h_t \cup a')$, 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 $U^L(h_t \cup a) = U^L(h_t \cup a')$ by the same logic.

Statement 2. Begin with the first part (that $a \in B(\delta_L|L)$). For any such a that $\alpha^H(a|h_t) = 0$ and $\alpha^L(a|h_t) > 0$ and any outcome x, we have $p(h_{t+dt}) = \delta_L$, where $h_{t+dt} = h_t \cup (a, x)$. By Lemma 2, $a' \in B(\delta_L|L)$ must be played at all histories beginning with h_{t+dt}. If $a \notin B(\delta_L|L)$ then playing a' at h_t instead – and continuing with a' at all subsequent histories – yields a strictly higher flow 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. $\square$

A.3 Proofs: Finite Types

Before stating the proof of Theorem 2, we need some 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. *Suppose (MON) holds and dt is small enough. Fix any equilibrium and any history $h_t \in \mathcal{H}$. Then for any $a', a'' \in A(h_t)$ we have $\bar{S}(h_t \cup a') \geq \underline{S}(h_t \cup a'')$, with equality only if $S(h_t \cup a')$ is a singleton.*[20]

Proof. Assume by contradiction that $\bar{S}(h_t \cup a') < \underline{S}(h_t \cup a'')$ for some $a', a'' \in A(h_t)$. Pick any type $\theta \in S(h_t \cup a')$ and any strategy $\mathbf{a}'$ on path for θ at h_t. Then deviating to a'' at h_t and following $\mathbf{a}'$ after t is strictly better for θ than following $\mathbf{a}'$ throughout. To see this, recall that u^θ is increasing in p_t by (MON) – and reputation $p_s \geq \delta_{\underline{S}(h_t \cup a'')}$ generated by the deviation for all $s > t$ is strictly higher than any reputation on equilibrium path, since $\underline{S}(h_t \cup a'') > \bar{S}(h_t \cup a')$. This contradicts $\mathbf{a}'$ being optimal for θ as long as dt (period length and, hence, utility weight on the current period) is small enough.

[20]This Lemma and the remainder of the Appendix uses the notation $S(h_t \cup a) \equiv S(h_t \cup (a, x))$ for all $x \in X$ in the support. This object is well defined in equilibrium for on-path histories and actions because the support of x 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) holds.

Now suppose $\bar{S}(h_t \cup a') = \underline{S}(h_t \cup a'')$. Suppose by way of contradiction that $|S(h_t \cup a')| > 1$, meaning $\underline{S}(h_t \cup a') < \underline{S}(h_t \cup a'')$. Then among all types in $S(h_t \cup a')$ there exists such θ that receives reputation $p_s < \delta_{\bar{S}(h_t \cup a')}$ with positive probability for all $s > t$ (this follows from belief consistency). Such θ would strictly benefit from the deviation described in the first part of this proof, yielding a contradiction. □

Lemma 4 puts the (SC) property to use, establishing a form of monotonicity of optimal strategies w.r.t. 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 monotonicity 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. *Suppose (SC) holds. Fix any equilibrium and history $h_t \in \mathcal{H}$. If there exists a pair of strategies $\underline{\mathbf{a}}, \bar{\mathbf{a}} \supset h_t$ 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 $\hat{\mathbf{a}} \supset h_t$ 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 $\hat{\mathbf{a}}$:

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

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

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

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

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

If this inequality is strict, then this is a direct contradiction with (SC), which requires that $U^{\theta}(\bar{\mathbf{a}}|h_t) - U^{\theta}(\hat{\mathbf{a}}|h_t)$ 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. *Suppose (MON) and (SC) hold and dt is small enough. Fix any equilibrium and any $h_t \in \mathcal{H}$. There exist h_t-payoff-equivalent strategies $\bar{\mathbf{a}}, \underline{\mathbf{a}} \supset h_t$, on path at h_t for $\bar{S}(h_t)$ 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 $|S(h_t)| = k-1$ then it also holds when $|S(h_t)| = k \geq 3$. 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 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 $h_s \supset h_t$, 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 $h_{s+dt} = h_s \cup (a, \mathbf{x}(h_s))$.
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 $\bar{h}_{s+dt} \equiv h_s \cup (\bar{a}, \mathbf{x}(h_s))$ and $\underline{h}_{s+dt} \equiv h_s \cup (\underline{a}, \mathbf{x}(h_s))$. 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, 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* $\mathbf{x}$.[21] 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. □

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 $\bar{\mathbf{a}}, \underline{\mathbf{a}} \supset h_t$ on path at h_t for $\bar{\theta}$ and $\underline{\theta}$ respectively. Then by Lemma 4, any pure strategy $\mathbf{a} \supset h_t$ 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 $\bar{\mathbf{a}}' \supset h_t$ 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 at h_t in equilibrium – $\bar{\mathbf{a}}(h_t) \neq \bar{\mathbf{a}}'(h_t)$ – since this is in violation of Lemma 3. The same, however, applies to any subsequent history, hence $\bar{\mathbf{a}}(h_s) = \bar{\mathbf{a}}'(h_s)$ for all $h_s \supset h_t$.[22] 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

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

[22] To be slightly more precise, the argument applies to any h_s such that $|S(h_s)| > 1$. Otherwise Lemma 2 kicks in and implies that all pure strategies on path at h_s are h_s-payoff-equivalent.

we know at least one of these strategies – suppose $\underline{\mathbf{a}}$ – is on path for $\underline{\theta}$ but not for any other $\theta \in S(h_t)\backslash\bar{\theta}$ at h_t. Furthermore, it follows from the definition of payoff-relevant signaling that there is no $\bar{\mathbf{a}}' \supset h_t \cup \underline{a}$ that 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 $\underline{a} \in B\left(\delta_{\underline{\theta}}|\underline{\theta}\right)$. If this is not true then type $\underline{\theta}$ can play some $a \in B\left(\delta_{\underline{\theta}}|\underline{\theta}\right)$ at h_t and every history after it. 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 $\bar{a} \notin B\left(\delta_{\underline{\theta}}|\underline{\theta}\right)$. 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 $p(h_t \cup \bar{a}) > \delta_{\underline{\theta}}$ and the strict part of (MON)), hence $\underline{\mathbf{a}}$ would not be optimal for $\underline{\theta}$ at h_t – a contradiction.

This completes the proof of Theorem 2. □

Proof of Corollary 1. The statement is proved in the text. 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., $\mathbb{E}_x\left[u^{\underline{\theta}}\left(\bar{a}, p(h_t \cup (\bar{a}, x))\right)|\underline{\theta}\right] = u^{\underline{\theta}}\left(\underline{a}, \delta_{\underline{\theta}}\right)$. □

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 $h_{t'} \supset h_t$ at which the two prescribe different actions, but that constitutes payoff-irrelevant signaling at $h_{t'}$ (the two strategies coincide on $[t, t')$, hence they are payoff-equivalent at $h_{t'}$).

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

Proof of Proposition 2. Contained in the text. □

Appendix B. Application Examples

This appendix presents examples of applied models that fit our framework. These are meant to demonstrate some instances of models yielding additively and/or multiplicatively separable payoff functions that allow (SC) to be verified with little effort, complementing Section 5 in this respect.

B.1 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 $t \in \mathcal{T} \equiv \{0, dt, 2dt, ...\}$ (where period length dt is "small") she chooses whether to acquire education or not, $e \in \{0, 1\}$. Alternatively, in a more general version of the model the candidate could select education intensity $e \in [0, \bar{e}]$. The flow 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. 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 competitive firms will at any history h_t all 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:

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

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

since accepting an offer at h_t is equivalent to receiving a lumpsum payoff of $w(h_t)/r$.

It is easy to see that for any pair of strategies $\mathbf{a}', \mathbf{a}''$, the function $\mathcal{U}(\theta) \equiv U^{\theta}(\mathbf{a}''|h_t) - U^{\theta}(\mathbf{a}'|h_t)$ can be written as

$$\mathcal{U}(\theta) = \sum_{s \in \mathcal{T}, s \geq t} e^{-r(s-t)} \left[(d'(h_s) - d''(h_s)) \frac{w}{r} - \left[(1 - d'(h_s)) \cdot l\left(e'(h_s)\right) - (1 - d''(h_s)) \cdot l\left(e''(h_s)\right) \right] m(\theta) \right],$$

which is linear in $m(\theta)$. 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 $U^{\theta}(\mathbf{a}|h)$ satisfies (SC) and our results apply. The candidate will only be able to signal her ability via attrition: in an informative equilibrium – if it exists – the lowest type will drop out of education and take up a job at a random date, while all other types will go through the whole education process prescribed by the equilibrium and only accept the job afterwards. The latter path will also be taken by some low-ability candidates.

B.2 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. See Ausubel, Cramton, and Deneckere [2002] for a survey of the related literature.

Consider two players, a buyer B and a seller S, who interact repeatedly in every period $t \in \mathcal{T} \equiv \{0, dt, 2dt, ...\}$. There is a unit of indivisible good initially possessed by the seller, who values it at $\theta \in \Theta \subseteq \mathbb{R}_+$, which is their private information. The buyer's valuation is $v(\theta)$. In every period one of the players $b_t \in \{B, S\}$ is chosen as the proposer (the choice rule can be random or deterministic, and/or history-dependent) and can offer a price $p \in \mathbb{R}_+$. The other player then decides $d \in \{0, 1\}$ 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 $p = d = 0$ for both players at all subsequent histories). Otherwise the game continues to the next period. Both players discount the future at rate r.

History $h_t = \{b_s, p_s, d_s\}_{s\in\mathcal{T}, s<t}$ in this game is given by the players' past offers and responses, as well as identities of the proposing player. A player's strategy is given by $\mathbf{a}_i = \{p_i(h), d_i(p, e)\}_{h\in\mathcal{H}, e\in[0,1]}$. Denoting $\beta(h_s) \equiv \mathbb{P}(b(h_s) = B \mid h_s)$, the seller's expected payoff from following some strategy $\mathbf{a}_S$ conditional on the buyer's strategy $\mathbf{a}_B$ and some history $h_t \in \mathcal{H}$, by which no offer had been accepted, is given by

$$U^\theta(\mathbf{a}_S|h_t, \mathbf{a}_B) \equiv \sum_{s\in\mathcal{T}, s\geq t} e^{-r(s-t)} \left[\beta(h_s) d_S(h_s) \left(p_B(h_s) - \theta\right) + \left(1 - \beta(h_s)\right) d_B(h_s) \left(p_S(h_s) - \theta\right)\right].$$

It is easy to see that $U^\theta(\mathbf{a}_S|h_t, \mathbf{a}_B)$ is linear in θ, hence so is $\mathcal{U}(\theta) \equiv U^\theta(\mathbf{a}''_S|h_t, \mathbf{a}_B) - U^\theta(\mathbf{a}'_S|h_t, \mathbf{a}_B)$ for any pair of seller's strategies $\mathbf{a}'_S, \mathbf{a}''_S$ and any given buyer's strategy $\mathbf{a}_B$. Therefore, (SC) holds in this model and our results apply – independently of $v(\theta)$. The exact shape that the equilibria can take does, however, depend on both $v(\theta)$ and $b(h_t)$. See Ausubel et al. [2002] for more details.