Dynamic Noisy Signaling in Discrete Time*

Francesc Dilmé†

October 30, 2014

Abstract

We analyze a dynamic noisy-signaling model in discrete time. We fully characterize it by constructing the equilibrium (sender's) payoffs' set for each prior about her type. It exhibits a self-replicating step structure that resembles a devil's staircase. As a consequence, the equilibrium posterior about the type of the sender is highly discontinuous (discontinuous on a dense set) in the initial continuation value of the sender, and the effort put on signaling is highly non-monotone (with infinite peaks and valleys) in such posterior. We argue that similar undesirable properties are likely to be present in other dynamic models, which hints some limitations of the otherwise desirable discrete-time modeling choice. By mapping our model into a reputations model, we show that reputation may be a permanent phenomenon even under imperfect monitoring, and it can be sustained without building-milking reputation phases.

Keywords: Dynamic Noisy Signaling, Endogenous Effort.

JEL Classification: D82, D83, C73, J24.

---

*Work in progress. Early versions of this paper were titled "Slowing Learning Down". I thank George Mailath, Andrew Postlewaite and Qingmin Liu for their comments and suggestions, as well as the participants in the UPenn Micro Lunch, Columbia's Theory Colloquium and BGSE Microworkshop.

†University of Bonn. fdilme@uni-bonn.de

1 Introduction

This paper provides a full characterization of the set of perfect Bayesian equilibria of a fully-dynamic noisy-signaling model in discrete time. The equilibrium set is found to have complex and "undesirable" mathematical properties. We argue that these properties are likely to be present in other dynamic models with asymmetric information, highlighting the limitations of the otherwise desirable discrete-time modeling choice.

We focus our analysis on a simple dynamic trade model. We consider a seller (entrepreneur) who wants to sell an asset (firm), which can have a low or a high underlying quality (underlying value). Only the seller observes the quality of her asset, and she can exert unobservable managerial effort that generates observable noisy returns.[1] Potential (short-lived) buyers, who stochastically arrive over time, observe the history of returns and make offers to the seller. If she accepts an offer, the asset is sold, and the game ends. Otherwise, the seller continues managing her asset until the arrival of the next buyer.

Even though the setting studied in this paper is simple, its equilibrium structure is quite complex. We explicitly construct the "equilibrium continuation set," which contains the pairs formed by the posteriors about the quality of the asset being high and their continuation payoffs. Since buyers only make acceptable offers when the posterior is above a given threshold, the continuation payoff of the seller is discontinuous at this threshold. This discontinuity replicates itself due to the recursive structure of the continuation values, giving the equilibrium continuation set a step structure. In particular, it may take the form of a devil's staircase, that is, a non-constant continuous function that is flat almost everywhere. As a result, when the composition of the market is endogenized by introducing an entry fee, the equilibrium market composition is discontinuous in a dense set with respect to the entry fee. Also, the expected equilibrium managerial effort features an infinite number of peaks and valleys.

Other discrete-time dynamic models with asymmetric information are likely to feature equilibrium sets with properties similar to ours. The reason is that discrete actions like whether to make acceptable offers or whether to accept a given offer tend to follow cutoff strategies in the belief about some asymmetric information. As a result, a discontinuity in the equilibrium continuation values is likely to appear at this threshold. Due to the recursive nature of the continuation values, the discontinuity will replicate itself and not only complicate the characterization of the equilibrium set, but also give it some undesirable properties. Hence, even though the discrete-time modeling choice

[1] In many economically-relevant situations, such as in markets of heterogeneous assets or in education, the sender has not only inside information about the value of her good, but also information on actions cannot be observed by outsiders, and which stochastically affect some signal.

allows a characterization of behavior without the need of restricting strategies or the equilibrium concept,[2] its use is likely to be limited due to the technical complexity of the equilibrium objects that generates.

In our model, buyers use their local monopoly power to extract all surplus from the high-quality seller in all equilibria. Hence, even in the presence of signaling motives, the high-quality seller has strict incentives to manage her firm optimally. Due to our assumption of no gains from trade for the low-quality asset, buyers never make offers only acceptable by the low-quality seller. Therefore, separation in our model is only driven by different effort choices across the types of the seller.

In most of the periods, our low-quality seller randomizes between exerting her efficient (low) managerial effort versus masquerading her type by exerting a suboptimal (high) effort. Intuitively, if the low-quality seller is supposed to efficiently manage her asset, the signal becomes very informative, so she has incentives to undertake cost-inefficient (but revenue-generating) effort. In this case, high returns would convince future buyers that the quality of the asset is high, so the seller would increase the expected revenue from selling the asset. The reverse is true if she is supposed to put high effort into signaling. The managerial effort of the low-quality seller is small when buyers' beliefs about the quality are close to being degenerated.

Our model is equivalent to a reputations model where a seller repeatedly sells goods to short-lived customers, who only observe noisy signals about the quality of the goods sold in the previous periods. In this case, efficient and inefficient management correspond, respectively, to "reputation milking" and "reputation building." We find that under some parametric assumptions reputation can be a permanent phenomenon, even when monitoring is imperfect, and it is sustained without reputation building-milking cycles.

In the next section we review the literature related to our paper. In Section 2, we set our base model and the main results of the paper. Section 3 concludes. The Appendix provides the proofs omitted in the previous sections.

1.1 Literature Review

There has been some recent interest in dynamic noisy-signaling models. Examples of them are Daley and Green (2012), Dilmé (2013) and Heinsalu (2014), who analyze continuous-time models where an (endogenously or exogenously) informative signal pro-

[2] Differently from most of the previous literature (see the literature review below), we do not need impose any restriction on players' behavior (for example, measurability or continuity conditions on strategies, etc.) and we do not further restrict our focus to a particular equilibrium subset (Markov strategies, refinements, etc.).

gressively reveals information over time about the type of the sender.[3] These models construct dynamic versions of the classic Spence (1973) model and try tanalyze how the sender optimally (and incentive-compatibly) decides when she accepts an offer and/or the effort she puts on separating/pooling with the other type. The technical complexity of these models requires several (technical) assumptions on the strategies available to the sender, and the analysis has to be restricted to a restrictive subset of equilibria. We address similar questions in a discretized model. This allows us to have a full-characterization of the equilibrium set, and therefore a better understanding its structure and its implied behavior.

Our model is also related to the reputations literature, inaugurated by the seminal works of Kreps and Wilson (1982) and Milgrom and Roberts (1982), since it can be reinterpreted as a reputations model. In particular, we show that the vanishing reputations result in Cripps, Mailath and Samuelson (2004) may not hold when short-lived players' payoff is a function of the type, not the action. So, we show that reputation may be sustained without building-milking cycles (in the previous literature, among other mechanisms, building-milking cycles were generated through bounded recall (Liu, 2011), through replenishing types (Mailath and Samuelson, 2001, Board and Meyer-ter-Vehn, 2013) and through adjustment costs (Dilmé, 2014)).

2 Basic Model

2.1 Setting

Time is discrete, $t = 0, 1, ...$ There is an entrepreneur (seller) who wants to sell a firm (or asset). The asset is either of low quality ($\theta = L$) or high quality ($\theta = H$). The seller discounts future payoffs at a discount factor $\delta \in (0, 1)$.

There is a pool of homogeneous short-lived buyers. At every period, there is a probability $1 - \lambda$ that no buyer arrives, and a probability λ that (exactly) one buyer arrives.[4] Buyers value an asset of quality $\theta \in \{L, H\}$ at U_θ, with $U_L < U_H$. If a buyer arrives at t, he makes a take-it-or-leavie-it offer to the seller $P_t \in \mathbb{R}$. If the seller accepts the offer, the asset is sold and the game ends. Otherwise, the game continues.

At every period t, the entrepreneur decides on the effort $e_t \in \{0, 1\}$ put into man-

[3] Kremer and Skrzypacz (2007) and Kaya and Kim (2013) introduce the possibility that the buyer receives one signal about the seller (at the end of the signaling time, in the first case, or upon his arrival, in the second), but not repeated signaling.

[4] The random arrival of buyers has recently been introduced in bargaining models (for example Fuchs and Skrzypacz (2007)) and dynamic lemons markets (for example, Hörner and Vieille (2009) and Kim (2013)), where it may be interpreted as a search friction.

aging the firm. If the effort exerted at t is e_t, the asset generates returns at t equal to $\pi_G \equiv \pi > 0$ with probability $\nu\, e_t$ and $\pi_B \equiv 0$ with probability $1 - \nu\, e_t$, for some $\nu \in (0, 1)$. We use $\xi_t \in \{B, G\}$ to denote the realization of the returns at time t, where B denotes that returns were low at t ("bad" signal), while G denotes that returns are were high at t ("good" signal). The timing of the game is schematically displayed in Figure 1.

The cost of providing effort e is type-dependent and normalized to $c_\theta\, \nu\, e$ for $\theta \in \{L, H\}$. Note that c_θ can be interpreted as "the cost per unit of the probability of generating high returns." We assume that $c_H < \pi < c_L$. This implies that, in autarchy, high effort is optimal for the H-seller, but not for the L-seller. We define $\overline{V}_L \equiv 0$ and $\overline{V}_H \equiv \nu\, \frac{\pi - c_H}{1-\delta}$ as the autarchy values of the L-seller and the H-seller, respectively. Also, we assume $U_H > \overline{V}_H$ (gains from trade for the H-asset) and $U_L < \overline{V}_L$ (no gains from trade for the L-asset).[5]

A (unterminated) *public history* is an element of $\mathcal{H} \equiv \cup_{t=0}^{\infty} \{B, G\}^t$ and it encodes the returns realized in the past. A (unterminated) *private history* is an element of $\tilde{\mathcal{H}} \equiv \cup_{t=0}^{\infty} \big(\{B, G\} \times \{0, 1\} \times (\{-\infty\} \cup \mathbb{R})\big)^t$, that is, a public history plus the effort choices by the seller and the offers made by the buyers, where an offer equal to $-\infty$ at time t corresponds to no buyer arriving in this period. A *terminated private history* $(\tilde{h}^t, P) \in \tilde{\mathcal{H}} \times \mathbb{R}$ is composed of a private history and the offer accepted after it (at time $t + 1$).

A strategy of a buyer who arrives at time t with public history h^t is given by a distribution over the price offers $\tilde{P}(h^t) \in \Delta(\mathbb{R})$. A strategy by the θ-seller, for $\theta \in \{L, H\}$, is an acceptance decision rule $\beta_\theta : \tilde{\mathcal{H}} \times \mathbb{R} \to [0, 1]$, where $\beta_\theta(\tilde{h}^t, P)$ is the probability of accepting an offer P at history $\tilde{h}^t$, and an effort choice $\alpha_\theta : \tilde{\mathcal{H}} \times (\mathbb{R} \cup \{-\infty\}) \to [0, \nu]$, where α/ν is the probability of choosing effort equal to 1.[6] The strategy of the θ-seller is public if only conditions the acceptance offer on the public part of the history and the current offer, and the effort on the public part of the history. Given a terminated private history $((\xi^t, e^t, P^t), P_t)$, the payoff of the corresponding

[5]Since $\overline{V}_L = 0$ we have $U_L < 0$, which may seem counter-intuitive, especially given the usual assumption of free disposal of the asset. Nevertheless, notice that $\pi_B = 0$ is just a normalization, so in general $\overline{V}_L = \frac{\pi_B}{1-\delta}$. Furthermore, transaction (legal/taxes) costs may reduce the buyers' valuation of the asset.

[6]It is notationally convenient to use the probability of generating high returns as the choice of the seller, instead of the effort. It is clear that if the probability of issuing a dividend is restricted to belong to $[0, \nu]$, the two modeling choices are equivalent. Also, since the action at period t is taken after the rejection of an offer (or if there has been no offer at all; see Figure 1), α is a function of $\tilde{\mathcal{H}} \times (\mathbb{R} \cup \{-\infty\})$.

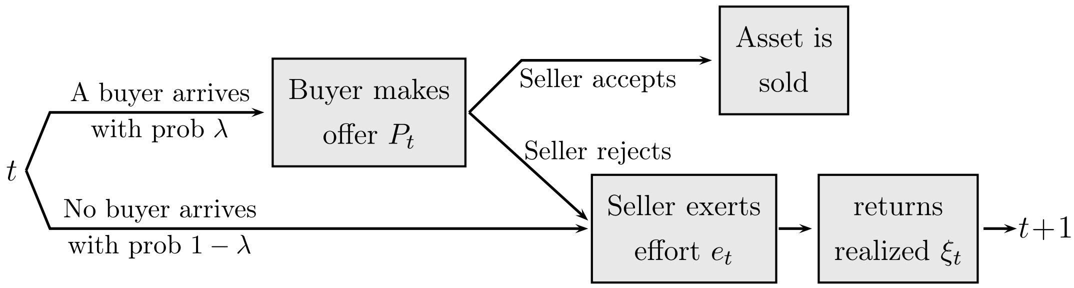

Figure 1: Timing of the model.

payoff for the θ-seller at time 0 is given by

$$V_{\theta,0}(\tilde{h}^t, P_t) = \sum_{s=0}^{t-1} \delta^s \left(\pi_{\xi_s} - \nu\, c_\theta\, e_s\right) + \delta^t\, P_t\ . \tag{2.1}$$

The payoff of a never-terminated path of play (i.e., no offer is accepted by the seller) is defined as usual. Given a strategy profile, the expected payoff at 0 is defined as the expected payoff over the terminated and the never-terminated paths of play.

We use $\psi_t \equiv \psi_t(h^t) \equiv \Pr(\theta = H|h^t)$ to denote the posterior of the buyers about the type of the seller being H at t. Also, we use $V_{\theta,t} \equiv V_\theta(h^t)$ to denote the highest continuation value of the seller at time t. Note that it only depends on the public history and the type, even if the seller follows a non-public strategy. The reason is that previous previous effort choices or offers received are not observed by future buyers, so the distribution of future offers depends only on the current (and future) public history. Also, past effort choices do not change the payoffs of a continuation play. So, as usual, we focus on public strategies by the seller.[7]

Definition 2.1. An *(public Bayesian-perfect) equilibrium* is a strategy profile for the seller $(\beta^*_\theta, \alpha^*_\theta)_{\theta\in\{L,H\}}$, a strategy for the buyers $\tilde{P}^*$ and a beliefs process ψ that satisfy:

1. For any $h^t \in \mathcal{H}$, $\tilde{P}^*(h^t)$ solves the following problem:
$$\tilde{P}^* \in \arg\max_{\tilde{P}'} \mathbb{E}_t\left[\psi_t\, \beta_{H,t}(\tilde{P}')\, U_H + (1-\psi_t)\, \beta_{L,t}(\tilde{P}')\, U_L - \tilde{P}'|\tilde{P}'\right]\ . \tag{2.2}$$

2. For any $h^t \in \mathcal{H}$, $(\beta^*_\theta, \alpha^*_\theta)$ are optimal policy functions of the θ-seller's problem, for $\theta \in \{L, H\}$:
$$V_\theta(h^t) = \lambda\, \mathbb{E}_t\left[\max_{\beta_\theta\in[0,1]} (\beta_\theta\, \tilde{P}_t + (1-\beta_\theta)\, W_\theta(h^t)\right] + (1-\lambda)\, W_\theta(h^t)\ , \tag{2.3}$$
$$W_\theta(h^t) = \max_{\alpha\in[0,\nu]} \left(\alpha\,(\pi - c_\theta) + \delta\left(\alpha\, V_\theta(h^t, G) + (1-\alpha)\, V_\theta(h^t, B)\right)\right)\ . \tag{2.4}$$

[7] As usual, a strategy of our game is an equilibrium if and only if it is outcome-equivalent to an equilibrium in public strategies, so characterizing the public equilibria is enough to characterize all equilibria in the game.

3. Whenever possible, ψ is updated using the Bayes' rule.

Remark 2.1. Note that in our model signaling is productive in the sense that high managerial effort generates high returns that are valuable for the seller. Nevertheless, this is very different from the usual productive-signaling models of education, where the signal increases the productivity of the student, and therefore the value that she has for the uninformed part of the market.

2.2 Equilibrium Characterization

Existence

We begin with a result stating the existence of equilibria.

Proposition 2.1. *For all ψ_0, an equilibrium exists.*

Buyers' Behavior

Let $\psi^\star$ be the lowest posterior such that if a buyer offers $\overline{V}_H$ and it is accepted by the seller for sure (independently of her type), the buyer makes non-negative profits. Formally:

$$\psi^\star U_H + (1-\psi^\star)\, U_L - \overline{V}_H = 0 \quad \Rightarrow \quad \psi^\star \equiv \frac{\overline{V}_H - U_L}{U_H - U_L} \in (0,1)\ .$$

The next proposition establishes the equilibrium behavior of the buyers. In this proposition, as in those in the rest of the paper, the phrase "In any equilibrium" will be omitted.

Proposition 2.2. *Fix a public history $h^t \in \mathcal{H}$.*

1. *If $\psi(h^t) < \psi^\star$ then no equilibrium offer at h^t is accepted.*
2. *If $\psi(h^t) = \psi^\star$ then an equilibrium offer $P(h^t)$ is accepted only if $P(h^t) = \overline{V}_H$.*
3. *If $\psi(h^t) > \psi^\star$ then $\Pr\left(P(h^t) = \overline{V}_H\right) = 1$.*

Part 1 derives from the assumption $U_L < \overline{V}_L$. Indeed, an offer that is accepted only by the L-seller is clearly suboptimal, since there are no gains from trade for the L-quality asset. Also, since $\psi(h^t) < \psi^\star$, attracting the H-seller (by offering at least $\overline{V}_H$) generates losses. The rationale behind part 3 is very similar to "Diamond's Paradox" and can be explained as follows. Assume that the maximum continuation value for the H-seller is $V_H^* > \overline{V}_H$. This is only possible if there is an equilibrium offer higher than V_H^*. Nevertheless, a buyer can always offer slightly less than V_H^* and, given that the seller discounts the future, she accepts the offer for sure, which generates a profitable deviation.

High-Quality Seller

Let's now determine the behavior of the high-quality buyer.

Proposition 2.3. *The H-seller always exerts high effort and accepts all offers equal to $\overline{V}_H$.*

The first part of Proposition 2.3 is a consequence of Proposition 2.2. Indeed, given that equilibrium offers do not exceed $\overline{V}_H$, exerting low effort lessens the continuation payoff below $\overline{V}_H$. Since the H-seller can guarantee herself a payoff of $\overline{V}_H$, exerting low effort is strictly suboptimal. The second part comes from the fact that if the H-seller accepts an equilibrium offer with probability less than one, the buyer can increase it an arbitrarily small amount so that the seller accepts it for sure. This would be a profitable deviation.

Pooling on Accepting Offers, Separation on Effort Choices

Given that the behavior of the H-seller is fully determined by Proposition 2.3, we now focus on the behavior of the L-seller. In order to save notation we will drop the subindexes in the continuation values and strategies of the L-seller, using V_t to denote $V_{L,t}$ and α_t to denote $\alpha_{L,t}$.

A trivial corollary of Propositions 2.2 and 2.3 is that the L-seller accepts with probability one equilibrium offers equal to $\overline{V}_H$ and rejects all other equilibrium offers. Indeed, given that $c_L > c_H$, it is clear that $V(h^t) < \overline{V}_H$ for all histories $h^t \in \mathcal{H}$. Therefore, at every history, the posterior is updated using only the returns ξ_{t-1} and the expected equilibrium effort of the L-seller $\alpha(h^{t-1})$. In particular, $\psi(h^t) = \psi_{\xi_{t-1}}(\psi(h^{t-1}), \alpha(h^{t-1}))$ for all $h^t \in \mathcal{H}$, where[8]

$$\psi_G(\psi,\alpha) \equiv \frac{\nu\,\psi}{\nu\,\psi + \alpha\,(1-\psi)} \quad \text{and} \quad \psi_B(\psi,\alpha) \equiv \frac{(1-\nu)\,\psi}{(1-\nu)\,\psi + (1-\alpha)\,(1-\psi)}\ , \tag{2.5}$$

If $\alpha = \nu$ (i.e. th L-seller pools with the H-seller) the signal is totally uninformative, so $\psi_\xi(\psi,\nu) = \psi$ for all $\xi \in \{B, G\}$. This is likely to happen when the cost of mimicking is low. The following proposition establishes a sufficient condition in order to have separation in the effort decision:

Proposition 2.4. *If $c_L - \pi > \delta\,\lambda\,\overline{V}_H$, the L-seller never fully mimics the H-seller, i.e., $\alpha(h^t) < \nu$ for all $h^t \in \mathcal{H}$.*

[8]Note that, since $\nu \in (0,1)$, Bayes' rule is well defined except for the case $\psi = 0$ and $\alpha \in \{0,1\}$. Nevertheless, note that buyers are never perfectly convinced that the type of the seller is L. Indeed, there is no public history that has 0 probability when the type of the seller is H.

Proposition 2.4 is very helpful in simplifying the arguments and intuitions. The reason is that, at each history, one could interpret the strategy of the L-seller "as if" she could choose the realization of the returns (generating high returns at cost c_L,) that is, as if she had a perfect control over the realized public history. In order to provide a neat characterization of the equilibrium set in our model and clear intuitions about our results, in the remainder of the paper we will assume that $c_L - \pi > \delta\,\lambda\,\overline{V}_H$, except for Section 2.4, where we will relax it.

Continuation Payoffs Dynamics

Let $\underline{v} \equiv \lambda\,\overline{V}_H$ be the expected revenue from selling the asset at a given period, provided that if a buyer arrives he offers $\overline{V}_H$ to the seller. Note that, by Proposition 2.2, if $\psi(h^t) > \psi^\star$ then $V(h^t) \geq \underline{v}$. Also, let $\bar{v} \equiv \frac{\lambda\,\overline{V}_H}{1-(1-\lambda)\,\delta}$ be the upper bound on the payoff that can be achieved by the L-seller. This consists of the continuation payoff for the L-seller if she does not exert any effort in the continuation play, and whenever a buyer arrives, he offers $\overline{V}_H$ to the seller.

For any $v \in \mathbb{R}$, let's define

$$V_B(v) \equiv \frac{v - \underline{v}}{\delta} \quad \text{and} \quad V_G(v) \equiv \frac{v + (1-\lambda)\,(c_L - \pi) - \underline{v}}{\delta}\,. \tag{2.6}$$

It is easy to show that if $v \in [0, \bar{v}]$ then $V_L(v) < v < V_H(v)$. These functions characterize the dynamics of the continuation values of the L-seller.

Proposition 2.5. *If $\psi(h^t) > \psi^\star$ then $V(h^t, \xi) = \min\{\bar{v}, V_\xi(V(h^t))\}$ for all $\xi \in \{B, G\}$ and $h^t \in \mathcal{H}$.*

Intuitively, when the continuation value of the L-seller is not too high (not too low), then she has to be indifferent about generating high (low) returns. The reason is that, otherwise, high (low) returns would convince future buyers that the type of the seller is H, effectively providing her with a continuation value of $\bar{v}$ in the next period. The indifference conditions of the L-seller impose the restriction that the next period's continuation payoff is obtained using the functions $V_\xi(\cdot)$ for $\xi \in \{B, G\}$. When the continuation value of the seller is high (i.e. $V(h^t) \geq V_G^{-1}(\bar{v})$), the L-seller can convince future buyers that her type is H by generating high returns, but she does not do this because the increment on her continuation value does not compensate for the cost of generating high returns.

Proposition 2.5 is very useful because it isolates the dynamics of the posterior from the dynamics of the continuation payoff. Indeed, the continuation payoff in the next period is only a function of the current continuation payoff and the dividend issuance on the current period, independently of the previous history or equilibrium played.

The Set of Equilibrium Payoffs

Let's now characterize the equilibrium continuation payoffs. As we will see, this will be useful in order to characterize the equilibrium strategies. The *equilibrium payoffs correspondence* $\hat{V} : [0,1] \to [0, \bar{v}]$ is given by

$$v \in \hat{V}(\psi) \quad \Leftrightarrow \quad \text{when } \psi_0 = \psi \text{ there exists a NE with } V(\emptyset) = v \ . \tag{2.7}$$

The following result establishes some basic properties of $\hat{V}(\cdot)$:

Proposition 2.6. *$\hat{V}(\psi) = \{0\}$ if $\psi < \psi^\star$, $\hat{V}(\psi^\star) = [0, \underline{v}]$ and $\hat{V}(\psi) \subset [\underline{v}, \bar{v}]$ if $\psi > \psi^\star$.*

Using Proposition 2.6 we can schematically depict the equilibrium payoffs correspondence as in Figure 2. The first statement of Proposition 2.6 is a consequence of the fact that, by Proposition 2.4, the L-seller never has strict incentives to mimic the H-seller. In particular, the L-seller is indifferent about exerting 0 effort forever, which implies that returns are low thereafter. In this outcome, if $\psi(h^t) < \psi^\star$, the posterior remains below $\psi^\star$ forever. Therefore, no future buyer makes a positive offer to the seller, so the L-seller gets its autarchy value $\overline{V}_L = 0$.

The intuition behind the second statement is that if $\psi(h^t) = \psi^\star$ and a buyer arrives, he is indifferent about offering $\overline{V}_H$ or an unacceptable offer, so he may potentially randomize. Therefore, the flow expected revenue from selling the asset in this period can be any value in $[0, \underline{v}]$. Also, by Proposition 2.4 the L-seller does not fully mimic the H-seller; therefore, low returns lower the posterior below $\psi^\star$, which provides a continuation value equal to 0, which implies that the continuation value is no higher than $\underline{v}$.[9]

The third statement is a direct consequence of Proposition 2.2. If $\psi(h^t) > \psi^\star$ and a buyer arrives, he is going to offer $\overline{V}_H$ for sure. Since a buyer arrives with probability λ, the expected payoff is at least $\underline{v} = \lambda \, \overline{V}_H$.

Beliefs Dynamics

For a given v, $\underline{\psi}(v)$ (resp. $\overline{\psi}(v)$) provides us with the lowest (resp. highest) initial prior where an equilibrium with an initial continuation payoff equal to v exists. Formally:

$$\underline{\psi}(v) \equiv \inf \left\{ \psi \in [0,1] \,\middle|\, v \in \hat{V}(\psi) \right\} \quad \text{and} \tag{2.8}$$

$$\overline{\psi}(v) \equiv \sup \left\{ \psi \in [0,1] \,\middle|\, v \in \hat{V}(\psi) \right\} \ . \tag{2.9}$$

In Figure 2 we have a graphical depiction of both concepts. For a given continuation value v, $\underline{\psi}(v)$ corresponds to the horizontal infimum at height v of the graph of $\hat{V}$, while

[9]While the first and the third parts of Proposition 2.6 still hold when $c_L - \pi < \tilde{\delta} \, \lambda \, \overline{V}_H$, the second part only holds under this assumption. Section 2.4 discusses the $c_L - \pi < \tilde{\delta} \, \lambda \, \overline{V}_H$ case.

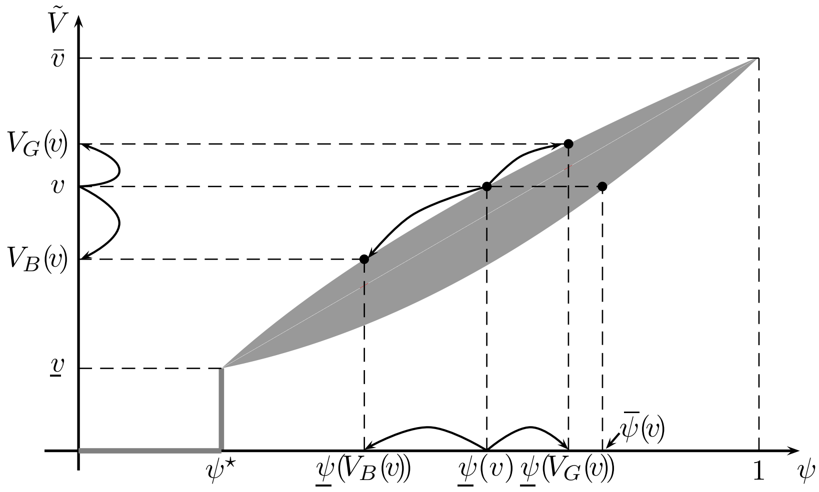

Figure 2: Schematic depiction of the equilibrium payoffs correspondence. By Proposition 2.6 we have that $\hat{V}(\psi) = \{0\}$ when $\psi < \psi^\star$, and $\hat{V}(\psi^\star) = [0, \underline{v}]$. The arrows indicate the equilibrium change in the posterior and continuation payoff at $(\underline{\psi}(v), v)$. By Lemma 2.1, in all equilibria, if $\psi(h^t) = \underline{\psi}(v)$ and $V(h^t) = v$, then $\psi(h_t, \xi_t) = \underline{\psi}(V_{\xi_t}(v))$ and $V(h^t, \xi_t) = V_{\xi_t}(v)$.

$\overline{\psi}(v)$ corresponds to the horizontal supremum. The following result establishes that if (ψ_t, V_t) is part of the boundary of the graph of $\hat{V}$, with $V_t > \underline{v}$, then (ψ_{t+1}, V_{t+1}) also is.

Lemma 2.1. *If $\psi(h^t) = \underline{\psi}(V(h^t))$ for some h^t and $V(h^t) > \underline{v}$, then $\psi(h^t, \xi) = \underline{\psi}(V(h^t, \xi)))$ for both $\xi \in \{B, G\}$. The same holds for $\overline{\psi}(\cdot)$.*

Let's provide some intuition as to why Lemma 2.1 is true. Assume, by contradiction, that there is some equilibrium and h^t (normalize it to $h^t \equiv \emptyset$) such that $\psi(\emptyset) = \underline{\psi}(V(\emptyset))$ and $\psi(G) > \tilde{\psi}_+$ for some $\tilde{\psi}_+ \in \hat{V}^{-1}(V(\emptyset))$ (the case with B is analogous). It is then easy to find $\tilde{\psi} < \psi(\emptyset)$ and $\tilde{\alpha} \in (0,1)$ such that $\psi_B(\tilde{\psi}, \tilde{\alpha}) = \psi(B)$ and $\psi_H(\tilde{\psi}, \tilde{\alpha}) = \tilde{\psi}_+$. This implies that when $\psi_0 = \tilde{\psi}$, there exists an equilibrium providing a continuation value equal to $V(\emptyset)$ to the L-seller, so $\tilde{\psi} \in \hat{V}^{-1}(V(\emptyset))$. This contradicts our initial assumption that $\psi(\emptyset) = \underline{\psi}(V(\emptyset))$.

Step Structure

Let $\mathcal{V} \subset [\underline{v}, \bar{v}]$ be the smallest set that contains $\underline{v}$ and

$$v \in \mathcal{V} \cup [0, \underline{v}) \quad \Leftarrow \quad V_\xi(v) \in \mathcal{V} \text{ for some } \ \xi \in \{B, G\} \ .$$

Note that $\mathcal{V}$ is countable, and therefore $[0, \bar{v}] \backslash \mathcal{V}$ is dense in $[0, \bar{v}]$. Intuitively, $\mathcal{V}$ is composed of the continuation values v such that there exists some ψ_0 and history h^t such that $V(\emptyset) = v$ and $V(h^t) = \underline{v}$, satisfying $V(h^s) \geq \underline{v}$ for all $s < t$.

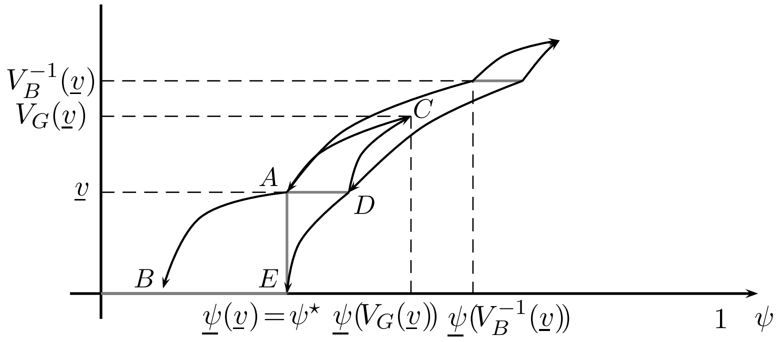

Figure 3: Intuition of the self-replication of the steps in the equilibrium payoffs correspondence. Since $\underline{v} \in \hat{V}(\psi^\star)$ (point A in the picture), we can find two continuation equilibria with posteriors and continuation values at points B and C. Then, there exists an equilibrium with posterior and continuation value D that uses $E \equiv (\psi^\star, 0)$ and C as continuation equilibria. This generates the first step (all points between A and D). The next step is constructed similarly using the first step.

Proposition 2.7. *For all $v \in (\underline{v}, \bar{v})$ we have $\overline{\psi}(v) = \inf\{\underline{\psi}(v') | v' > v \ \&\ v' \in \mathcal{V}\}$ and $\underline{\psi}(v) = \sup\{\overline{\psi}(v') | v' < v \ \&\ v' \in \mathcal{V}\}$. Furthermore, $\underline{\psi}(v) < \overline{\psi}(v)$ if and only if $v \in \mathcal{V} \cup \{0\}$.*

A corollary of Proposition 2.7 is that $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ are increasing. Indeed, if there existed $\underline{v} \leq v < v' \leq \bar{v}$ such that $\underline{\psi}(v) > \underline{\psi}(v')$ then for any $v'' \in (v, v')$ we would have $\overline{\psi}(v'') \leq \underline{\psi}(v') < \underline{\psi}(v) \leq \overline{\psi}(v) \leq \underline{\psi}(v'')$, which is a contradiction. Intuitively, the existence of an equilibrium where the initial continuation payoff is high requires the prior about the type being H to be also high. Also, since by Proposition 2.1 $\hat{V}(\psi_0)$ is non-empty for all $\psi_0 \in [0, 1]$, we have that $\hat{V}^{-1}(v) = [\underline{\psi}(v), \overline{\psi}(v)]$, which implies that the graph of $\hat{V}(\cdot)$ has a step structure.

The fact that, by Proposition 2.7, if $v \in \mathcal{V} \cup \{0\}$ then $\underline{\psi}(v) < \overline{\psi}(v)$, highlights one of the main difficulties of considering discrete (time/types/signals/effort choices) models: the jumps in the posterior belief. Since at every period, after the realization of the signals, there is a jump in the posterior, the continuation value jumps accordingly. In our model, due to the discontinuity in the behavior of the buyers (established in Proposition 2.2), the set of continuation payoffs suddenly increases at $\psi^\star$ (see Proposition 2.6). The effect of this discontinuity is replicated for higher posteriors and payoffs and makes $\hat{V}$ have non-standard properties, as described below.

The intuition behind Proposition 2.7 is depicted in Figure 3. Assume $\psi_0 = \psi^\star$ and $V(\emptyset) = \underline{v}$. By Proposition 2.4 we have that in such an equilibrium $\psi(B) < \psi^\star$, and therefore, by Proposition 2.6, $V(B) = 0$. It is then easy to show that if we increase ψ_0 above but close to $\psi^\star$, there exists an equilibrium (denoted using tildes) where $\tilde{\psi}(B) < \psi^\star$ and $\tilde{\psi}(G) = \psi(G)$ (constructed using the same continuation plays as in our original equilibrium). The same argument can be applied to $V_B^{-1}(\underline{v})$, and by induction to all elements in $\mathcal{V}$.

Given that, by Proposition 2.7, the set $\mathcal{V}$ determines the heights of the "steps" in the graph of $\hat{V}$, its density plays a crucial role in determining its properties. The following proposition establishes the properties of $\mathcal{V}$.

Proposition 2.8. *1. If $V_G^{-1}(\bar{v}) \leq \underline{v}$ then $\bar{v}$ is the only accumulation point of $\mathcal{V}$.*

2. If $V_G^{-1}(\bar{v}) \in (\underline{v}, V_B^{-1}(\underline{v})]$ then $\mathcal{V}$ is not dense in $(\underline{v}, \bar{v}]$ but has countable many accumulation points.

3. If $V_G^{-1}(\bar{v}) > V_B^{-1}(\underline{v})$ then $\mathcal{V}$ is dense in $(\underline{v}, \bar{v}]$.

Note that when $\mathcal{V}$ is not dense (i.e., in the first two cases of Proposition 2.8) there exist continuation values in the interior of $[\underline{v}, \bar{v}]\backslash\mathcal{V}$. This implies that $\hat{V}$ is not a function on $(\psi^\star, 1]$, i.e., $\hat{V}(\psi)$ is multivalued for some $\psi > \psi^\star$. In order to see this, assume that $(v, v') \cap \mathcal{V} = \emptyset$ for some $v, v' \in (\underline{v}, \bar{v})$ with $v < v'$. Then, Proposition 2.7 establishes that $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ are constant in (v, v'), i.e., $(v, v') \subset \hat{V}(\underline{\psi}(v))$.

In Figure 4 we depict the graph of $\hat{V}(\cdot)$ when $\mathcal{V}$ is not dense. The reason for the existence of vertical segments in the graph (other than the one situated at $\psi^\star$) is, as it was for the existence of horizontal steps, the discrete nature of our model. For the sake of clarity, focus on the case of $\underline{v} \geq V_G^{-1}(\bar{v})$. By Proposition 2.6, continuation values in $(0, \underline{v})$ can only be achieved if the posterior is equal to $\psi^\star$. Since when $\underline{v} \geq V_G^{-1}(\bar{v})$ the only effort exerted in equilibrium by the L-seller is 0,[10] $\psi(B^t) = \psi^\star$ for some t only if ψ_0 belongs to $\underline{\psi}(\mathcal{V})$, which is a discrete, non-dense set.

In the third case of Proposition 2.8, instead, $\hat{V}(\psi)$ is univaluated when $\psi \in (\psi^\star, 1)$. In this case, the region of the continuation payoffs space where the seller is indifferent about the effort choice (given by $(\underline{v}, V_G^{-1}(\bar{v}))$) is wide enough so that all its elements can be approached by iteratively applying $V_B^{-1}(\cdot)$ and $V_B^{-1}(\cdot)$ to $\underline{v}$. Since learning is endogenous and mixing probabilities cannot be discretized, the size of the jumps in the posterior is endogenous and non-discretizable. Therefore, mixing probabilities are adjusted in such a way that even though $\underline{\psi}(v) = \psi^\star$ for all $v \in (0, \underline{v}]$, $\underline{\psi}(V_B^{-1}(v))$ is increasing in v when $v \in (0, \underline{v}]$.

Effort Choice

Let's now establish the (generic) uniqueness of the public outcomes:

[10] Indeed, $V(v) > \bar{v}$ for all $v > \underline{v}$, so high effort is never exerted in equilibrium.

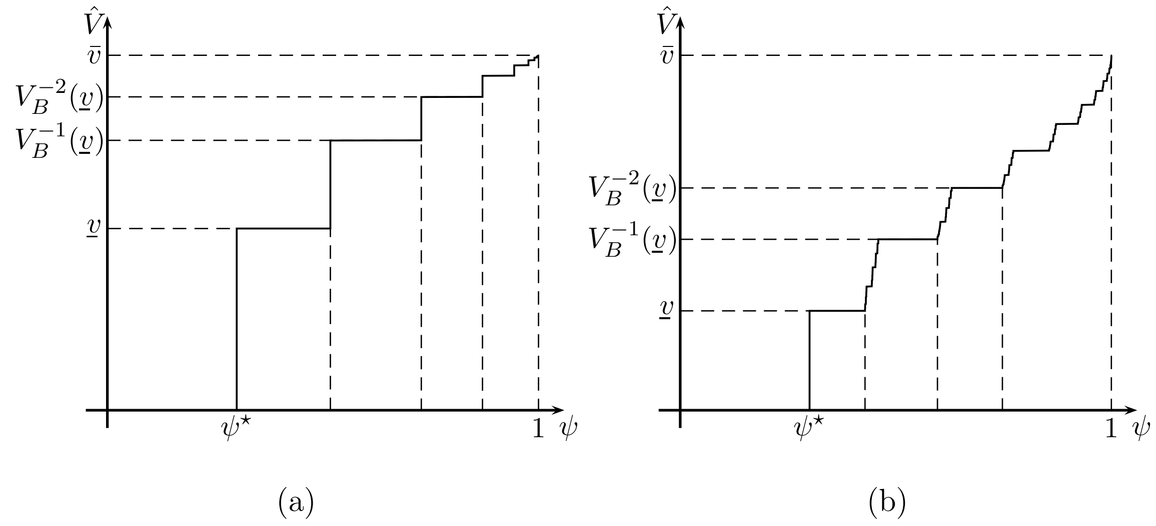

Figure 4: Graph of $\hat{V}(\cdot)$ when $\mathcal{V}$ is not dense. (a) corresponds to the first case in Proposition 2.8, while (b) corresponds to the second case. While in (a) we have a "regular" staircase, in (b) we see that each step has a substructure, provided by the existence of multiple accumulation points in $\mathcal{V}$.

Proposition 2.9. *The equilibrium distribution over public outcomes is generically*[11] *unique.*

When $\underline{v} \geq V_G^{-1}(\bar{v})$ (which happens if, for example, $c_L - \pi$ is large, δ is small or λ is large) the equilibrium choice is always 0, that is, the L-seller does not put any effort into masquerading. The reason is that either the seller has a high cost of signaling, current costs of masquerading are highly valued compared with the future reward or the seller is very confident that a buyer will arrive soon, so there is no need to keep the posterior high.

Let's finally state a result about the monotonicity of the generically unique equilibrium effort $\alpha_t \equiv \alpha(\psi_t)$. It implies that when $V_G^{-1}(\bar{v}) > V_B^{-1}(\underline{v})$ the equilibrium effort has an infinite number of peaks and valleys. Figure 5 illustrates the result.

Lemma 2.2. *Let $\psi \in (\underline{\psi}(v), \overline{\psi}(v))$ for some $v \in \mathcal{V}$, and let $\xi \in \{B, G\}$ be the generically unique signal such that $V_\xi(v) \in \mathcal{V}$.*[12] *Then, $\alpha(\cdot)$ is decreasing at ψ when $\xi = B$ and increasing at ψ when $\xi = G$.*

---

[11] By generically we mean that it is unique except for a set of measure zero in the parameter space. For example, for any choice of parameters $(\psi_0, \lambda, \delta, c_H, \pi, U_L, U_H)$, there is a unique equilibrium for all choices of c_L satisfying our assumptions except maybe for a countable (i.e., measure-zero) set.

[12] Note that given the previous properties of our model, the measure of ψ satisfying $\psi \in (\underline{\psi}(v), \overline{\psi}(v))$ for some $v \in \mathcal{V}$ is $1 - \psi^\star$, that is, this is generically the case when $\psi > \psi_0$.

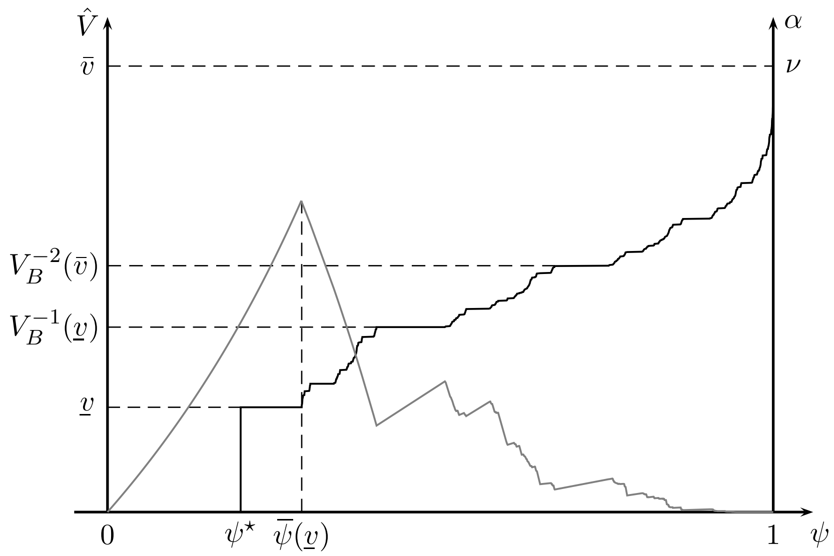

Figure 5: In black, graph of $\hat{V}$ when $V_B^{-1}(\underline{v}) < V_G^{-1}(\bar{v})$. For $\psi > \psi^\star$ this is the graph of a devil's staircase, that is, a continuous, non-constant and increasing function that is flat almost everywhere. In gray (right vertical axis), equilibrium effort choice in the (generically) unique equilibrium. We plot them together in order to see the relationship between the steps and the peaks of both curves.

Devil's Staircase and Ubiquitously Discontinuous Market Composition

Let's assume $V_B^{-1}(\underline{v}) < V_G^{-1}(\bar{v})$. In this case, by Proposition 2.8, the set of heights of the "steps" in $\mathcal{V}$ defined previously is dense. Therefore, the graph of $\hat{V}$ looks like an upward going "staircase" with an infinite number of steps continuously put one after the other. The next lemma states this formally:

Lemma 2.3. *If $\mathcal{V}$ is dense, the $\hat{V}(\psi)$ is a singleton for all $\psi \in (\psi^\star, 1]$. In this case, $\hat{V} : (\psi^\star, 1] \to \mathbb{R}$ is a devil's staircase, i.e., it is continuous, differentiable with derivative equal to 0 almost everywhere and globally increasing.*

Figure 5 (a) shows the graph of $\hat{V}$ when $\mathcal{V}$ is dense (i.e., $V_B^{-1}(\underline{v}) < V_G^{-1}(\bar{v})$). As we see, the steps in the staircase make the continuation payoffs correspondence flat almost everywhere. This implies, as the following Lemma 2.4 below states, that its inverse is not well behaved.

Figure 5 (b) depicts the mixing probability in the effort choice in the generically unique equilibrium when $\mathcal{V}$ is dense. We see that it is low in the extremes (i.e. when learning is slow). The effort choice peaks around $\psi^\star$, since if a B signal is observed, the posterior falls below $\psi^\star$; so in the following period no high offer will be made. Intuitively, the L-seller gets "scared," so the equilibrium expected effort increases.

Lemma 2.4. *If $V_B^{-1}(\underline{v}) < V_G^{-1}(\bar{v})$ then any function $f : (\psi^\star, 1] \to [\underline{v}, \bar{v}]$ that satisfies $f(\psi_0) \in \hat{V}^{-1}(\psi_0)$ for all $\psi_0 \in (\psi^\star, 1]$ is discontinuous in a dense set.*

The following remark explains why the previous result may generate highly discontinuous functions if we endogenize the initial prior about the quality of the asset.

Remark 2.2. Consider an extension of our model where we endogenize the initial composition of the qualities of the asset (i.e., ψ_0) in the following way. We assume that there is a big pool of entrepreneurs with low-value ideas and a much smaller pool of entrepreneurs with high-value ideas.[13] Setting up a firm requires paying a fixed cost $\overline{V}_0 \in (\underline{v}, \bar{v})$. Since $\overline{V}_H > \bar{v}$, all entrepreneurs with highly valuable ideas set up the firm for sure. Heuristically, free entry imposes the restriction that the mass of entrepreneurs with low-value ideas that set up the firm is such that $V(\emptyset) = \overline{V}_0$. In this case, from Lemma 2.3, it is easy to show that any $\psi_0 : (\underline{v}, \bar{v}) \to (\psi^\star, 1]$ is strictly increasing and discontinuous in a dense set. This makes the composition of the market extremely sensitive to changes in the entry costs by the firms.

2.3 Reputations Interpretation

In order to interpret our model as a reputations model, let's first establish the following result:

Proposition 2.10. *Fix any equilibrium $((\alpha_\theta, \beta_\theta)_{\theta \in \{L,H\}}, \tilde{P})$ of our model, and let $\mu(h^t) \equiv \Pr(\tilde{P}(h^t) = \overline{V}_H)$ for all $h^t \in \mathcal{H}$. Then the continuation value of the L-seller solves*

$$V(h^t) = \max_{\alpha \in [0,\nu]} \left(\mu(h^t)\, \underline{v} - \alpha\, c + \tilde{\delta} \left(\alpha\, V(h^t, G) + (1-\alpha)\, V(h^t, B) \right) \right) \tag{2.10}$$

where $\tilde{\delta} \equiv \delta\,(1-\lambda)$ and $c \equiv (1-\lambda)\,(c_L - \pi)$. Conversely, if (α, μ, ψ) solve (2.10), with ψ updated following Bayes' rule (2.5) and $\mu(h^t) = \mathbb{I}_{\psi(h^t) > \psi^\star}$ whenever $\psi(h^t) \neq \psi^\star$, then they are equilibrium strategies in our model.

In the previous proposition, $\underline{v}$ is interpreted as the flow payoff provided by the fact that when $\psi(h^t) > \psi^\star$, there is a positive probability (equal to λ) that an offer $\overline{V}_H$ will be received. c is the net cost of choosing high effort per unit of probability, adjusted for the fact that it is incurred only when no offer is accepted. Finally, $\tilde{\delta}$ is the effective discount factor, which incorporates the fact that the L-seller discounts the future, but also that buyers arrive randomly, so it takes some time for offers to arrive.

This renormalization has an interpretation of the model as a reputations model (see Remark 2.3). It is particularly useful because, even though the probability that the

[13] We interpret the quality of the firm as the value of the patents it owns. So, entrepreneurs first get (good or bad) ideas, then set up the firm, and finally sell (or not) the firm.

game ends by accepting an offer changes depending on the posterior about the type of the seller, the effective discount rate ($\tilde{\delta}$) is independent of the probability of receiving a high offer.

Remark 2.3. Our model can be reinterpreted as a reputations model in which a seller (for example, a restaurant) repeatedly sells a good to short-lived customers at price $\underline{v}$. Customers buy the good only if they are reasonably convinced that the quality of the good is high (i.e., $\psi_t \geq \psi^\star$). In this case, the returns in our original model can be reinterpreted as a noisy signal (online reviews, etc.) about the quality of the signal. While a high-quality product "looks good" (i.e., good signals appear frequently), making a low-quality product "look good" (i.e. send good signals as if it were high-quality) is costly. In this case, the continuation value of the L-seller would follow equation (2.10).

2.4 Mimicking Case

Let's now relax the condition made in Proposition 2.4 (assumed in the subsequent results in Section 2.2). Assume, therefore, that $c_L - \pi \leq \delta\,\lambda\,\overline{V}_H$.

In this case, we are not going to repeat the previous analysis, although most results still apply.[14] Instead, we will provide a result that exposes the main qualitative difference when we assume $c_L - \pi \leq \delta\,\lambda\,\overline{V}_H$ instead of $c_L - \pi > \delta\,\lambda\,\overline{V}_H$

Proposition 2.11. *Assume $c_L - \pi \leq \delta\,\lambda\,\overline{V}_H$. Then there exist equilibria where $\psi(h^t) \geq \psi^\star$ for all $h^t \in \mathcal{H}$. In those equilibria,* $\lim_{t\to\infty} \Pr(\psi_t = \psi^\star) = 1$.

Proposition 2.11 states that there are equilibria where the amount of information released is limited. Even though it is not efficient for the L-seller to exert effort (we still assume $\pi - c_L < 0$), the prospect of obtaining a high offer makes her willing to incur payoff losses for arbitrarily long periods of time. As the second part of the proposition hints, incentives to exert high effort are provided when the posterior is $\psi^\star$ by buyers mixing over the offer. When good signals are observed, (indifferent) buyers make a high offer with high probability, while if bad signals are observed, the probability of offering a high price decreases.

This result is in sharp contrast to the result in Cripps, Mailath and Samuelson (2004), especially given the reputations interpretation of our model. Indeed, they show that under imperfect monitoring, reputation asymptotically disappears. Hence, a continuation strategy of the game with perfect information is played with probability 1 in the long run. Note that in our model with perfect information the L-asset is not sold in any equilibrium, since there are no gains from trade for this asset.

[14]For example, it is not difficult to show that the equilibria that we found are still equilibria under the assumption $c_L - \pi \leq \delta\,\lambda\,\overline{V}_H$.

The reason why our model features equilibria with permanent reputation effects is the following. In classical reputation models the uninformed (sequence of short-lived) players' payoff depends on the actions taken by the informed players and not directly on the types. In our model, instead, their payoff is a function of the type, not the action. In the classical reputations model, a customer who believes that the seller exerts high effort would buy the good independently of the type of the seller. In our model, instead, if a customer is convinced enough that the type of the seller is L, he does not buy, independently of the effort exerted by the seller.

3 Discussion and Conclusions

3.1 Discussion

We have focussed our analysis on (arguably) one of the simplest fully-dynamic learning models with non-trivial equilibrium dynamics. This has allowed us to provide a full characterization of the equilibrium dynamics under the only assumption of sequential rationality,[15] and obtain implications for the equilibrium behavior of the players.

Our approach contrasts with that in most of the previous literature (see the literature review above), which mainly uses continuous-time models (with several restrictions on the strategies of the players) and focusses in small subsets of the equilibrium set (continuous, Markov, Pareto, etc) in order to characterize some equilibria of the model. In these models it remains an open question how approachable are their results as a limit of discrete-time versions of the model and, more importantly, what equilibrium predictions has the model when some of the restrictions on behavior (or on the subset of analyzed equilibria) are relaxed.

Even though our setting is simple, the techniques required to solve it are quite involving, and the obtained equilibrium objects are relatively complex. The main difficulty in solving dynamic games with hidden types and actions like ours lies in solving the fixed-point problem between, in one side, the of the buyers' actions and beliefs about the type of the seller and the equilibrium effort of each type and, in the other side, the incentive-compatible effort choices of each type of the seller given the buyers' actions and beliefs. The solution is determined globally (local conditions are not enough to characterize strategies), which complicates the comparative statics and perturbation analysis. This complexity and the need to solve the model globally make difficult to know how robust are the results we have obtained. Still, we argue that most of the features of our model are likely to be obtained in other dynamic games.

[15] Even though for expositional clarity we focus on perfect Bayesian equilibria, it is easy to show that Propositions 2.2 and 2.3 hold in general, and therefore the rest of the arguments are also valid.

In our model, as in any discrete-time model with learning, the posterior about hidden information makes (potentially endogenous) jumps in the periods where there some information is revealed. As a consequence, if the behavior of some of the players is discontinuous at some given threshold in the beliefs (which is likely to happen if there are discrete actions such as sell a good or not), the discontinuity replicates itself. Indeed, since the continuation value at a given history is constructed recursively using the continuation values in the next period's possible histories, the effect of the discontinuity depends on the number of "jumps" that it takes the posterior to pass the threshold. This happens even when the size of the jump is endogenous, as in our model. As a result, highly discontinuous equilibrium objects will naturally appear in discrete versions of some dynamic trade models.

So, even though understanding the generality of our techniques and findings is beyond the scope of this paper, they can be useful as a guide to characterize the equilibrium structure of other economic models. Our model is a first example of a full characterization of the equilibrium set of a fully-dynamic stochastic model in discrete time, and therefore establishes a step forward towards a deeper understanding the structure of dynamic noisy trade models.

3.2 Conclusions

Our analysis highlights the technical challenges that dynamic learning models pose. Even though we consider arguably one of the simplest fully-dynamic stochastic models of repeated noisy signaling, the characterization of its equilibrium set is technically demanding. The main difficult in our analysis arises from the discontinuity of some behavior in the posterior, which generates a self-replicating step structure in the continuation payoff set. Similar structures are likely to be present in discrete time versions of other models in the literature, since some actions are discrete in nature.

Endogenizing the information that a signal conveys implies that the equilibrium speed of learning cannot be extreme. Fast learning induces the low type to masquerade, which slows learning down. Slow learning, instead, reduces her incentives to mimic, which speeds learning up. As a result, some information is revealed in each period, specially when the posterior is close to being degenerated, so beliefs updating is slow and each type chooses her myopically-optimal action. Also, when we reinterpret our model as a reputations model, we find that reputation may be sustained in the long run without the need of building-milking reputation phases.

Overall, repeated signaling models provide us with valuable insights into the dynamic incentives of traders of heterogeneous assets. It is then important to understand the structure of their equilibrium set and their implied behavior. We provide a new step in this direction by providing a neat example of a simple model that can be fully characterized. The generalization of our insights and techniques to more general models is left to future research.

References

Board, Simon and Moritz Meyer ter Vehn (2013) "Reputation for Quality," *Econometrica*, Vol. 81, pp. 2381–2462.

Cripps, Martin W., George J. Mailath, and Larry Samuelson (2004) "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, Vol. 72, pp. pp. 407–432.

Daley, Brendan and Brett Green (2012) "Waiting for News in the Market for Lemons," *Econometrica*, Vol. 80, pp. 1433–1504.

Diamond, Peter A. (1971) "A model of price adjustment," *Journal of Economic Theory*, Vol. 3, pp. 156–168.

Dilmé, Francesc (2013) "Dynamic Quality Signaling with Hidden Actions," working paper.

——— (2014) "Reputation Building through Costly Adjustment," working paper.

Fuchs, William and Andrzej Skrzypacz (2007) "Bargaining with Arrival of New Traders," 2007 Meeting Papers 186, Society for Economic Dynamics.

Heinsalu, Sander (2014) "Noisy signalling over time," working paper.

Hörner, Johannes and Nicolas Vieille (2009) "Public vs. Private Offers in the Market for Lemons," *Econometrica*, Vol. 77, pp. 29–69.

Kaya, Ayça and Kyungmin Kim (2013) "Trading Dynamics in the Market for Lemons," working paper.

Kim, Kyungmin (2013) "Information about Sellers' Past Behavior in the Market for Lemons," working paper.

Kremer, Ilan and Andrzej Skrzypacz (2007) "Dynamic Signaling and Market Breakdown," *Journal of Economic Theory*, Vol. 133, pp. 58–82.

Kreps, David M. and Robert Wilson (1982) "Reputation and imperfect information," *Journal of Economic Theory*, Vol. 27, pp. 253–279.

Liu, Qingmin (2011) "Information Acquisition and Reputation Dynamics," *Review of Economic Studies*, Vol. 78, pp. 1400–1425.

Mailath, George J and Larry Samuelson (2001) "Who Wants a Good Reputation?" *Review of Economic Studies*, Vol. 68, pp. 415–41.

Milgrom, Paul and John Roberts (1982) "Predation, reputation, and entry deterrence," *Journal of Economic Theory*, Vol. 27, pp. 280–312.

Nöldeke, Georg and Eric van Damme (1990) "Signalling in a Dynamic Labour Market," *Review of Economic Studies*, Vol. 57, pp. 1–23.

Spence, A Michael (1973) "Job Market Signaling," *Quarterly Journal of Economics*, Vol. 87, pp. 355–74.

Swinkels, Jeroen M (1999) "Education Signalling with Preemptive Offers," *Review of Economic Studies*, Vol. 66, pp. 949–70.

A Omitted Proofs

Proof of Proposition 2.1 (page 7)

Existence follows from the explicit construction of equilibria of the other results in the paper. In particular, it is easy to show that the structure established in Lemma A.1 for $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ guarantees that if $\psi_0 \in [\underline{\psi}(v), \overline{\psi}(v)]$ then there exists some α such that $\psi_\xi(\psi_0, \alpha) \in [\underline{\psi}(V_\xi(v)), \overline{\psi}(V_\xi(v))]$. Given the properties of $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ established in Proposition 2.7, it follows that we can construct a continuation play consistent with equilibrium for any $\psi_0 \in [0, 1]$.

Proof of Proposition 2.2 (page 7)

Let's first show that no offer exceeds $\overline{V}_H$. Let's define $v_H^{\max} \equiv \sup_{h^t \in \mathcal{H}} V_H(h^t)$. Note that $v_H^{\max} \leq U_H$, since an equilibrium offer never exceeds U_H. Assume that $v_H^{\max} > \overline{V}_H = \frac{\nu\,(\pi - c_H)}{1-\delta}$ and let's define $\varepsilon \equiv (1-\delta)\, v_H^{\max} - \nu\,(\pi - c_H) > 0$. Let h^t be such that $V_H(h^t) \in (v_H^{\max} - \varepsilon, v_H^{\max}]$. Then, there is some $\alpha_H \in [0, \nu]$ such that

$$
\begin{aligned}
W_H(h^t) &= \overbrace{\alpha_H\,(\pi - c_H)}^{\leq \nu\,(\pi - c_H)} + \delta \left(\overbrace{\alpha_H\, V_H(h^t, G) + (1 - \alpha_H(h^t))\, V_H(h^t, B)}^{\leq v_H^{\max}} \right) \\
&< v_H^{\max} - \varepsilon \ .
\end{aligned}
$$

Therefore, it must be the case that, in equilibrium, if a buyer arrives at h^t the equilibrium distribution of his offer, $\tilde{P}(h^t)$, is such that

$$
v_H^{\max} - \varepsilon < V_H(h^t) \leq \lambda\, \mathbb{E}[\max\{\tilde{P}(h^t), v_H^{\max} - \varepsilon\}] + (1-\lambda)\,(v_H^{\max} - \varepsilon) \ .
$$

In particular, this requires that $\Pr[\tilde{P}_t > v_H^{\max} - \varepsilon] > 0$. Nevertheless, an offer $P_t > v_H^{\max} - \varepsilon$ is dominated by offering $v_H^{\max} - \varepsilon$,[16] which is accepted by the H-seller for sure. So, $v_H^{\max} \leq \overline{V}_t$.

[16]It is strictly dominated because since $V_H(h^t) \geq V_L(h^t)$, in both cases, the offer is going to be accepted by all types of seller.

Assume first $\psi_t > \psi^\star$. Assume that a buyer arrives at t. If the H-seller accepts $\overline{V}_H$ for sure then it is optimal to offer $\overline{V}_H$. Otherwise, assume that the H-seller accepts an offer $\overline{V}_H$ with probability $\gamma \in [0,1)$. Then, if the buyer offers $\overline{V}_H + \varepsilon$, with $\varepsilon > 0$, the offer is accepted by the H-seller for sure (since her continuation payoff is $\overline{V}_H$), so this guarantees a payoff to the buyer equal to

$$\psi_t \, \overline{V}_H + (1-\psi_t)\, \overline{V}_L - (\overline{V}_H + \varepsilon) \ .$$

Note that since $\psi_t > \psi^\star$, the sum of the first two terms is strictly higher than $\overline{V}_H$. Therefore, if $\varepsilon > 0$ is small enough, the whole expression is positive, and higher than $\psi_t \, \gamma \, \overline{V}_H + (1-\psi_t)\, \overline{V}_L - \overline{V}_H$. Since the payoff is decreasing in ε, all such offers are strictly dominated. Therefore, in equilibrium, $\overline{V}_H$ is offered and accepted with probability one.

Assume now $\psi_t \leq \psi^\star$. Note that from the first part of the proof we have $W_L(h^t) < W_H(h^t) = \overline{V}_H$ for all histories. Therefore, an offer strictly lower than $\overline{V}_H$ would attract (at most) only the L-seller. This offer is never profitable (given that $U_L < \overline{V}_L$). Also, if a buyer offers $\overline{V}_H$ the L-seller accepts for sure. Therefore, the profit for the buyer from offering $\overline{V}_H$ is no larger than $\psi_t \, U_H + (1-\psi_t)\, U_L - \overline{V}_H$. If $\psi_t < \psi^\star$, this is negative, so it is dominated by an unacceptable offer. Only if $\psi_t = \psi^\star$ and the H-seller accepts an offer equal to $\overline{V}_H$ with probability one, the buyer is indifferent about offering it or not, so he could potentially randomize.

Proof of Proposition 2.3 (page 8)

Note that Proposition 2.2 implies that $V_H(h^t) = \overline{V}_H$ for all histories. Assume that there is a history where $\alpha_H(h^t)/\nu \in [0,1)$. Then,

$$V_H(h^t) = \overbrace{\alpha_H(h^t)\,(\pi - c_H)}^{<\nu\,(\pi - c_H)} + \delta \left(\overbrace{\alpha_H(h^t)\, V_H(h^t, G) + (1-\alpha_H(h^t))\, V_H(h^t, B)}^{=\overline{V}_H} \right) \ .$$

Trivially, $V_H(h^t) < \overline{V}_H$, which is a contradiction. Finally, in the proof of Proposition 2.2 we see that an offer equal to $\overline{V}_H$ is made only if it is accepted with probability one by the H-seller.

Proof of Proposition 2.4 (page 8)

We prove this result by first solving a simpler version of our model, referred as to the "perfect-monitoring model," where we allow the L-seller to choose G for sure (so high returns are generated for sure) at cost c_L (we assume this in all proofs of Propositions 2.5-2.10 [note that none of the proofs of these propositions uses Proposition 2.4]). We then verify (in this proof) that for all equilibria and histories of the perfect-monitoring model $\alpha(h^t) \leq \nu$, so all equilibria found under the perfect-monitoring model are also equilibria of the original model. We finally show (in this proof) that the set of equilibrium payoffs in both models is the same, so our result holds. We recommend the reader first go over the proofs of Propositions 2.5-2.10 before reading this proof.

We first prove that $\alpha(h^t) < \nu$ for all h^t. Note that if $V(h^t) = v \in \mathcal{V}$ then $\overline{\psi}(V_B(v)) \leq \underline{\psi}(v) < \overline{\psi}(v) \leq \underline{\psi}(V_G(v))$ (using Proposition 2.7,) so $\alpha(h^t) < \nu$. If, instead, $v \notin \mathcal{V}$ then note that there exists $v' \in \mathcal{V}$ such that $V_B(v) < v' < v$. The reason is that $V_L^{-n}(\underline{v}) \in \mathcal{V}$ for all $n \in \mathbb{N}$ and there exists some n_* such that $V_B^{n_*}(v) \in (0, \underline{v})$, so $V_B^{-(n_*-1)}(\underline{v}) \in (V_B(v), v)$. This implies (by Proposition 2.7) that $\overline{\psi}(V_B(v)) \leq \underline{\psi}(v') < \overline{\psi}(v') \leq \underline{\psi}(v)$. Therefore, the expected effort choice is strictly lower than ν.

In order to compare the equilibrium payoff correspondences of the two models, for each $v \in [0, \bar{v}]$, let $\underline{\psi}_1(v)$ and $\underline{\psi}_\nu(v)$ be defined as in (2.8) for the perfect-monitoring and original models, respectively. Our goal is to prove that $\underline{\psi}_1(v) = \underline{\psi}_\nu(v)$ for all $v \in [0, \bar{v}]$ (we can proceed similarly to prove the same for $\overline{\psi}(\cdot)$ defined in (2.9)). Note that $\underline{\psi}_1(v) \geq \underline{\psi}_\nu(v)$ for all $v \in [0, \bar{v}]$, since as is shown previously in this proof, $\alpha(h^t) \leq \nu$ for all equilibria in our reduced model. Let's define for each $v \in [0, \bar{v}]$, $d_q(v) \equiv 1/\underline{\psi}_\nu(v) - 1/\underline{\psi}_1(v) \geq 0$, and let $d_q^{\max} \equiv \sup_v(d_q(v))$. Assume, by contradiction, that $d_q^{\max} > 0$.

The main difficulty in proving $\underline{\psi}_1(v) = \underline{\psi}_\nu(v)$ for all v is that there may be histories in our original model where the L-seller has strict incentives to exert high effort. It is easy to show that also in the perfect-monitoring model the L-seller is indifferent about exerting high or low effort at some history h^t only if $V(h^t, \xi) = V_\xi(V(h^t))$ for all $\xi \in \{B, G\}$), as in the perfect-monitoring model. She is strictly willing to exert effort only if $V(h^t, G) > V_G(V(h^t))$ and $V(h^t, B) < V_B(V(h^t))$, in which case $\psi(h^t, \xi) = \psi(h^t)$ for all $\xi \in \{B, G\}$. Also, the seller is willing to exert low effort if $V(h^t, B) = V_B(V(h^t))$ and $V(h^t) \leq V_G^{-1}(\bar{v})$.

Note that for any h^t and $V(h^t)$ such that $\psi(h^t) = \underline{\psi}_1(V(h^t))$, either $\psi(h^t, \xi) = \underline{\psi}_1(V_\xi(V(h^t)))$ for $\xi \in \{B, G\}$ (if $\alpha(h^t) < \nu$, using the same argument as in Lemma 2.1) or $\underline{\psi}_1(V(h^t, G)) \leq \underline{\psi}_1(V(h^t))$ (if $\alpha(h^t) = \nu$). Note that since $\underline{\psi}_1(V_G(v)) > \underline{\psi}_1(v) + m$ for some $m > 0$, and $\underline{\psi}(V(h^t, G)) \geq V_G(V(h^t))$, we have that in the second case $d_q(V(h^t, G)) > d_q(V(h^t, G))$. Therefore, if $\psi(h^t, \xi) \neq \underline{\psi}(V_\xi(V(h^t)))$, then $d_q(V(h^t, H)) > d_q(V(h^t)) + m'$ for some $m' > 0$.

Consider a strictly decreasing sequence $(\varepsilon_n)_n$ such that $\lim_{n\to\infty} \varepsilon_n = 0$. Also, let's define

$$v^{\inf} \equiv \lim_{n \to 0} \underbrace{\inf \left\{ v > 0 \,\middle|\, d_q(v) > d_q^{\max} - \varepsilon_n \right\}}_{\equiv v_n} .$$

Now, we can use an argument similar to the proof of Lemma A.1 to complete the proof. Note that it can be used since equation (A.4) still holds, given that otherwise there would exist $v > v^{\inf}$ with $d_q(v) > d_q(v) + m'$, which is a contradiction. Therefore, as in the proof of Lemma A.1, $d_q^{\max} = 0$, so $\underline{\psi}_1(\psi) = \underline{\psi}_\nu(\psi)$ for all $\psi \in [0, 1]$.

Proof of Proposition 2.5 (page 9)

To prove the result, it is important to notice that, in equilibrium, if at some public history h^t we have $\Pr(\tilde{P}(h^t) = \overline{V}_H | \tilde{P}(h^t) \neq -\infty) = 1$ (i.e., conditional on a buyer arriving at h^t he offers $\overline{V}_H$ for sure), then the L-seller is indifferent between exerting effort or not if and only if $V(h^t,\xi) = V_\xi(V(h^t))$ for all $\xi \in \{B,G\}$ (it follows from the indifference conditions).[17] This implies that the seller is willing to generate returns equal to π_ξ if and only if $V(h^t,\xi) = V_\xi(V(h^t))$, for all $\xi \in \{B,G\}$.

We prove each part separately:

1. Assume first $V_B(V(h^t)) \geq 0$ (so $V(h^t) \geq \underline{v}$). This implies, by Proposition 2.6 that if a buyer arrives at t he offers $\overline{V}_H$ for sure (recall that we do not use Proposition 2.5 to prove it). By contradiction, assume $V(h^t,B) \neq V_B(V(h^t))$. If $V(h^t,B) < V_B(V(h^t))$, then the L-seller has strict incentives to generate high returns. Therefore, low returns would convince future buyers that the type of the seller is H, providing the L-seller with a continuation payoff $\bar{v}$, which clearly dominates exerting high effort. If $V(h^t,B) > V_B(V(h^t))$, then generating low returns reports a value of $\underline{v} + \delta\,(1-\lambda)\,V(h^t,B)$, which, since $V(h^t) < \bar{v}$, is strictly higher than $V(h^t)$, so we also have a contradiction.

2. Assume now $V_G(V(h^t)) \leq \bar{v}$. If $V(h^t,G) > V_G(V(h^t))$ then the L-seller has strict incentives to generate high returns. Nevertheless, if low returns are generated, she convinces future buyers that her type is H, and this clearly dominates generating high returns. If, instead, $V(h^t,G) < V_G(V(h^t))$, then the L-seller strictly prefers not to generate high returns. If she instead generates high returns, she convinces the buyers that her type is H. Using the expression of $V_G(\cdot)$ and the fact that $V(h^t) < V_G^{-1}(\bar{v})$, we see that this is a profitable deviation. Therefore, $V(h^t,G) = V_G(V(h^t))$.

Proof of Proposition 2.6 (page 10)

Let's prove the propsition by parts.

1. Assume $\psi(h^t) < \psi^\star$ for some $h^t \in \mathcal{H}$. By Proposition 2.2 if a buyer arrives at h^t, he makes an unacceptable offer. Assume $\alpha(h^t) \leq \nu$ (the other case is analogous). Since $\psi(h^t,B) \leq \psi(h^t)$ the L-seller is weakly willing to proceed to a history where, if a buyer arrives, he offers an unacceptable offer. This argument can be iteratively used, so we have that the L-seller is willing to choose a continuation play where an acceptable offer is not made at any point in the future. This proves our result.

2. The fact that $\hat{V}(\psi^\star) \supset [0,\underline{v}]$ comes from the explicit construction of equilibria that we provide in Section 2.2.[18] In order to prove that $\hat{V}(\psi^\star) \subset [0,\underline{v}]$, note first that if there is some h^t such that $V(h^t) > \underline{v}$ and $\psi(h^t) = \underline{\psi}$ it needs to be the case that $\alpha(h^t) = \nu$. Otherwise, $\psi(h^t,\xi) < \underline{\psi}$ for some $\xi \in \{B,G\}$, so $V(h^t,\xi) = 0$. If the L-seller is willing to choose ξ, then $V(h^t) = V_\xi^{-1}(0) \leq 0$, which is a contradiction. So, assume that $\alpha(h^t) = \nu$ and $\psi(h^t) = \psi^\star$. Let $v^{\max} \equiv \sup \hat{V}(\psi^\star)$, and assume $v^{\max} > \underline{v}$.

[17]As mentioned in the proof of Proposition 2.4, the proofs of Propositions 2.5-2.10 are done allowing the L-seller to choose G for sure (so she generates high returns for sure) at cost c_L.

[18]In this section we assume that an equilibrium exists and we then construct it explicitly.

Consider $\varepsilon > 0$ small and an equilibrium and history such that $V(h^t) \geq v^{\max} - \varepsilon$. By the observation made at the beginning of the proof of Proposition 2.5, we have that $V(h^t, G) = V_G(V(h^t)) \geq V_G(v^{\max} - \varepsilon)$. Therefore, since $V_G(v^{\max}) > v^{\max}$, if ε is close enough to 0, we have $V(h^t, G) > v^{\max}$, which is a contradiction.

3. Proving $\hat{V}(\psi) \subset [\underline{v}, \bar{v}]$ if $\psi > \psi^\star$ is trivial, since $\bar{v}$ is clearly the maximum payoff achievable by a L-seller, while at least she gets the probability that a buyer arrives multiplied by a high offer $\overline{V}_H$, that is, $\lambda \overline{V}_H$, that is exactly $\underline{v}$.

Proof of Lemma 2.1 (page 11)

Note that the standard Bayes rule implies

$$\frac{1}{\psi(h^t)} = \frac{1-\rho}{\psi(h^t, G)} + \frac{\rho}{\psi(h^t, B)} \ . \tag{A.1}$$

We will prove the result for $\underline{\psi}(\cdot)$, and the proof for $\overline{\psi}(\cdot)$ is analogous. Assume that the claim of the lemma is not true. Then, there exists an equilibrium and history with $V(h^t) > \underline{v}$ and $\psi(h^t) = \underline{\psi}(V(h^t))$ such that $\psi(h^t, B) > \underline{\psi}(V_B(V(h^t)))$ (the other case, that is, when $\psi(h^t, G) > \underline{\psi}(V_G(V(h^t)))$, is done analogously). Then, let's define

$$\begin{aligned}
\tilde{\psi} &\equiv \frac{1}{\frac{\nu}{\psi(h^t,G)} + \frac{1-\nu}{\underline{\psi}(V_B(V(h^t)))}} < \frac{1}{\frac{\nu}{\psi(h^t,G)} + \frac{1-\nu}{\psi(h^t,B)}} = \psi(h^t) \quad \text{and} \\
\tilde{\alpha} &\equiv \frac{\nu}{\nu + (1-\nu)\,\frac{\psi(h^t,G)\,(1-\underline{\psi}(V_B(V(h^t))))}{(1-\psi(h^t,G))\,\underline{\psi}(V_B(V(h^t)))}} \in (0,1) \ .
\end{aligned}$$

Note that $\tilde{\alpha}$ is such that $\psi_B(\tilde{\psi}, \tilde{\alpha}) = \underline{\psi}(V_B(V(h^t)))$ and $\psi_G(\tilde{\psi}, \tilde{\alpha}) = \psi(h^t, G)$. Since, by assumption, there are equilibrium continuation paths at $\underline{\psi}(V_B(V(h^t)))$ and $\psi(h^t, G)$ providing, respectively, $V_B(V(h^t))$ an $V_G(V(h^t))$, when $\psi_0 = \tilde{\psi}$ there exists an equilibrium providing continuation value $V(h^t)$ to the L-seller. This implies that $\tilde{\psi} \geq \underline{\psi}(V(h^t))$, which is a clear contradiction.

Proof of Proposition 2.7 (page 12)

Let's define

$$\alpha_B(\psi, \psi') \equiv 1 - \frac{\psi\,(1-\psi')}{(1-\psi)\,\psi'}\,(1-\nu) \quad \text{and} \quad \alpha_G(\psi, \psi') \equiv \frac{\psi\,(1-\psi')}{(1-\psi)\,\psi'}\,\nu \ . \tag{A.2}$$

Using the standard Bayes rule, it can be shown that for each history h^t, $\alpha_\xi(\psi, \psi')$ provides the (unique) equilibrium mixing $\alpha(h^t)$ compatible with $\psi = \psi(h^t)$ and $\psi(h^t, \xi) = \psi'$, for all $\xi \in \{B, G\}$.

In order to prove the results, we first introduce a lemma that is very useful to characterize the functions $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$.

Lemma A.1. *$\underline{\psi}(\cdot)$ is the unique function $\underline{q} : (0, \bar{v}] \to [0, 1]$ that satisfies*

1. $\underline{q}(v) = \psi^\star$ for $v \in (0, \underline{v}]$,

2. $\alpha_B\big(\underline{q}(v), \underline{q}(V_B(v))\big) = \alpha_G\big(\underline{q}(v), \underline{q}(V_G(v))\big)$ for $v \in (\underline{v}, V_G^{-1}(\bar{v}))$, and

3. $\alpha_B\big(\underline{q}(v), \underline{q}(V_B(v))\big) = 0$ for $v \in [V_G^{-1}(\bar{v}), \bar{v})$.

Similarly, $\overline{\psi}(\cdot)$ *is the unique function* $\bar{q} : [0, \bar{v}] \to [0, 1]$ *the previous conditions changing* $\bar{q}(v) = \psi^\star$ *for* $v \in [0, \underline{v})$ *in the first part and* $v \in [\underline{v}, V_G^{-1}(\bar{v}))$ *in the second.*

Proof. We do the proof for $\underline{\psi}(\cdot)$ (for $\overline{\psi}(\cdot)$ the proof is analogous). Let's first prove existence. Using Bayes' rule (given in equation (A.1)) it is easy to show that $\underline{q}(\cdot)$ exists if and only if $\underline{\Gamma}(\cdot) \equiv \frac{1-\underline{q}(\cdot)}{\underline{q}(\cdot)}$ exists satisfying:

$$\underline{\Gamma}(v) = \begin{cases} \frac{1-\psi^\star}{\psi^\star} & \text{if } v \in (0, \underline{v}] \\ \nu\, \underline{\Gamma}(V_B(v)) + (1-\nu)\, \underline{\Gamma}(V_G(v)) & \text{if } v \in (\underline{v}, V_G^{-1}(\bar{v})) \\ \nu\, \underline{\Gamma}(V_B(v)) & \text{if } v \in [V_G^{-1}(\bar{v}), \bar{v}] \end{cases}$$

For each $v \in (0, \bar{v})$ let $\underline{\mathcal{H}}(v) \equiv \{h^t | V_{h^t}(v) \le \underline{v} \ \&\ V_{h^s}(v) \in (\underline{v}, \bar{v})\ \forall s < t\}$ be the set of (continuation) histories where, provided that the initial continuation payoff is v, $\underline{\psi}(v)$ reaches $\psi^\star$ for first time.[19] Then, it is easy to show that a solution for $\underline{\Gamma}(\cdot)$ is given by:

$$\underline{\Gamma}(v) \equiv \frac{1-\psi^\star}{\psi^\star} \sum_{h^t \in \underline{\mathcal{H}}(v)} \prod_{s=1}^{t} \nu^{\mathbb{I}(h_s^t = B)}\, (1-\nu)^{\mathbb{I}(h_s^t = G)} \ . \tag{A.3}$$

Note that $\underline{\Gamma}(\cdot)$ is left-continuous.

Using Lemma 2.1 we know that $\underline{\psi}(\cdot)$ satisfies parts 1-3 of Lemma A.1. Consider two functions, $q(\cdot)$ and $\tilde{q}(\cdot)$, both satisfying the conditions of Lemma A.1. Define $d_q(v) \equiv \frac{1}{q(v)} - \frac{1}{\tilde{q}(v)}$ and assume that $d_q^{\max} \equiv \sup_{v \in (0, \bar{v}]} d_q(v) > 0$. Consider a sequence $(\varepsilon_n)_n$, with $\varepsilon_n > 0$ for all n and such that $\lim_{n\to\infty} \varepsilon_n = 0$. Define also

$$v^{\inf} \equiv \liminf_{n \to 0} \underbrace{\inf \left\{\, v > 0 \,\middle|\, d_q(v) > d_q^{\max} - \varepsilon_n \,\right\}}_{\equiv v_n} \ .$$

Note that $(v_n)_n$ is a non-decreasing sequence (indeed, when n increases, the set over which the infimum is taken gets smaller). Also, note that $d_q(v_n) \to d_q^{\max}$. Let's first show that $v^{\inf} \le V_G^{-1}(\bar{v})$. Assume otherwise, so for some sequence $v_n > V_G^{-1}(\bar{v})$ for all n we have that[20]

$$d_q(V_B(v_n)) = \frac{\frac{1}{q(v_n)} - \rho}{1-\rho} - \frac{\frac{1}{\tilde{q}(v_n)} - \rho}{1-\rho} = \frac{d_q(v_n)}{1-\rho} \to \frac{d_q^{\max}}{1-\rho} > d_q^{\max} \ .$$

This is a clear contradiction. Assume then that $v_n \le V_G^{-1}(\bar{v})$ for all n. Note that since $q(v) = \tilde{q}(v) = \psi^\star$ when $v \in (0, \underline{v}]$, we have that $v_n > \underline{v}$ for all n. Using equation (A.1) we have

$$d_q(v_n) = \nu\, d_q(V_G(v_n)) + (1-\nu)\, d_q(V_B(v_n)) \ . \tag{A.4}$$

[19] $V_{h^t}(v)$ is defined recursively as $V_\emptyset(v) = v$ and $V_{h^t}(v) = V_{h_t^t}(V_{h^{t-1}}(v))$.

[20] Note that when $v > V_G^{-1}(\bar{v})$ we have $\alpha = 0$, so the update of beliefs follows $\frac{1}{\psi_B(\psi, 0)} = \frac{1/\psi - \nu}{1-\nu}$.

Note that the LHS is asymptotically (when $n \to \infty$) equal to $d_q^{\max}$. Each of the terms in the RHS is bounded above by $d_q^{\max}$. So, since the LHS is a linear combination of them, their limit must be equal to $d_q^{\max}$. In particular, $\lim_{n\to\infty} d_q(V_B(v_n)) = d_q^{\max}$. So, we have a contradiction, since $\lim_{n\to\infty} V_B(v_n) < v^{\inf}$, but $v^{\inf}$ is an infimum. □

(Continuation of the proof of Proposition 2.7) We first prove that $v \in \mathcal{V} \cup \{0\}$ if and only if $\underline{\psi}(v) < \overline{\psi}(v)$. Let's first prove the "only if" implication. Note first that $\underline{\psi}(\underline{v}) = \psi^\star < \overline{\psi}(\underline{v})$. Indeed, note that the solution for x of

$$\frac{1}{x} = \frac{\nu}{\psi^\star} + \frac{1-\nu}{\underline{\psi}(V_G^{-1}(\underline{v}))}$$

belongs to $(\psi^\star, \underline{\psi}(V_G^{-1}(\underline{v})))$ and, if $\psi_0 = x$, we have continuation payoffs that support an equilibrium with initial continuation value equal to $\underline{v}$. Given that $\underline{\psi}(\underline{v}) < \overline{\psi}(\underline{v})$, it is easy to prove using induction (note that $\mathcal{V}$ can be constructed recursively applying V_B^{-1} and V_G^{-1}) that $\underline{\psi}(v) < \overline{\psi}(v)$ for all $v \in \mathcal{V} \cup \{0\}$.

To prove the "if" implication, we prove that if $v \notin \mathcal{V} \cup [0, \underline{v})$ then $\underline{\psi}(v) = \overline{\psi}(v)$. To do this, consider

$$v^{\inf} \equiv \liminf_{n\to 0} \inf \left\{ v \notin \mathcal{V} \,\middle|\, \tfrac{1}{\underline{\psi}(v)} - \tfrac{1}{\overline{\psi}(v)} > d_q^{\max} - \varepsilon \right\} ,$$

where now $d_q^{\max} \equiv \sup_{v \notin \mathcal{V}} \left(\frac{1}{\underline{\psi}(v)} - \frac{1}{\overline{\psi}(v)}\right)$. Note that $V \notin \mathcal{V} \Rightarrow V_\xi(v) \notin \mathcal{V}\ \forall \xi \in \{B, G\}$. Therefore, applying an argument similar to that in the first part of the proof of Lemma A.1, the result holds. The argument works because when we restrict the domain of q in Lemma A.1 to $(\mathcal{V} \cup [0, \underline{v}))^c$, both $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ are the unique functions satisfying the three conditions of the lemma.[21] Nevertheless, since $\underline{v} \notin (\mathcal{V} \cup [0, \underline{v}))^c$, they satisfy exactly the same conditions, so they are equal. Note that this argument fails when $v \in \mathcal{V}$, since in this case we cannot rule out $\sup_{v \in \mathcal{V} \geq \underline{v}} \left(\frac{1}{\underline{\psi}(v)} - \frac{1}{\overline{\psi}(v)}\right) = \frac{1}{\underline{\psi}(\underline{v})} - \frac{1}{\overline{\psi}(\underline{v})} > 0$.

Finally, we prove part $\underline{\psi}(v) = \sup\{\overline{\psi}(v') | v' < v \ \&\ v' \in \mathcal{V}\}$ (the other case is analogous). First note if $[v, v'] \cap \mathcal{V} = \emptyset$ for some $v < v'$, then $\underline{\psi}(v) = \underline{\psi}(v')$ (and $\overline{\psi}(v) = \overline{\psi}(v')$). Indeed, we can see from equation (A.3) that $\underline{\psi}(v) \neq \underline{\psi}(v')$ only if $\mathcal{H}(v) \neq \mathcal{H}(v')$. It is easy to see that this implies that there exists some $v'' \in (v, v']$ and history h^t such that $V_{h^t}(v') = \underline{v}$. Nevertheless, this implies that $v'' \in \mathcal{V}$, which is a contradiction.

Note that since $\underline{\psi}(\cdot)$ is increasing and left-continuous (since $\underline{\Gamma}(\cdot)$ defined in Lemma A.1 is increasing and left-continuous), we have that

$$\underline{\psi}(v) = \sup\{\underline{\psi}(v') | v' < v\} \leq \sup\{\overline{\psi}(v') | v' < v\} = \sup\{\overline{\psi}(v') | v' < v \ \&\ v' \in \mathcal{V}'\} ,$$

where the last inequality comes from the fact that, as we showed, $\overline{\psi}(v)$ is constant in the intervals outside $\mathcal{V}$. We apply the same technique as in the proof of Lemma A.1, using

$$v^{\inf} \equiv \liminf_{n\to 0} \inf \left\{ v \,\middle|\, \lim_{w \searrow v} (1/\underline{\psi}(v) - 1/\overline{\psi}(w)) > d_q^{\max} - \varepsilon \right\} ,$$

[21]Indeed, a corollary to Lemma A.1 is given by stating the same result but restricting $\underline{q} : \mathcal{V}^c \cup \to [0, 1]$. In this case, since $0, \underline{v} \notin \mathcal{V}^c$, $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ have exactly the same conditions, so they are equal.

where now $d_q^{\max} \equiv \sup_v \big(\lim_{w \nearrow v} \big(\frac{1}{\underline{\psi}(v)} - \frac{1}{\overline{\psi}(w)} \big)\big)$. Note that the argument works because it is trivially true for $v = \underline{v}$ and $\underline{\psi}(v) = \psi^\star$ for $v \in (\underline{v}, V_L^{-1}(\underline{v}))$, so $v^{\inf} > \underline{v}$.

Proof of Proposition 2.8 (page 13)

Let's first prove that if $V_B^{-1}(\underline{v}) > V_G^{-1}(\bar{v})$ then $\mathcal{V}$ is not dense. Indeed, note that $v \in \mathcal{V}$ only if $V_\xi(v) \in \mathcal{V}$ for some $\xi \in \{B, G\}$. Nevertheless, for any $v \in (V_G^{-1}(\bar{v}), V_B^{-1}(\underline{v}))$ we have $V_G(v) > \bar{v}$ and $V_B(\underline{v}) < \underline{v}$.

If $\underline{v} \geq V_G^{-1}(\bar{v})$ then note that $\alpha(h^t) = 0$ for all equilibria and histories. So, $\mathcal{V} = \{V_B^{-(n-1)}(\underline{v}) | \forall n \in \mathbb{N}\}$, so the only accumulation point of $\mathcal{V}$ is $\bar{v}$ (note that $\lim_{n\to\infty} V_B^{-n}(\underline{v}) = \bar{v}$).

If, instead, $V_G^{-1}(\bar{v}) \in (\underline{v}, V_B^{-1}(\underline{v}))$ then note that $V_G^{-1}(\bar{v})$ is an accumulation point of $\mathcal{V}$. Indeed, $\bar{v}$ is an accumulation point of $\mathcal{V}$ by the same argument as above. Since $V_G(\underline{v}) < \bar{v}$, we have that $V_G^{-1}\big(\{V_B^{-(n-1)}(\underline{v}) | \forall n \in \mathbb{N}\} \cap (V_G(\underline{v}), \bar{v})\big) \subset \mathcal{V}$, so $V_G^{-1}(\bar{v})$ is an accumulation point of $\mathcal{V}$. Many other accumulation points can be found by using a similar procedure.

Finally, let's prove that if $V_B^{-1}(\underline{v}) \leq V_G^{-1}(\bar{v}))$ then $\mathcal{V}$ is dense. Consider otherwise, that is, $\mathcal{V}$ is not dense in $(\underline{v}, \bar{v})$, and let $A \subset (\underline{v}, \bar{v})$ be an interval with maximal length satisfying $A \cap \mathcal{V} = \emptyset$. Note that if $v \notin \mathcal{V}$ then $V_\xi(v) \notin \mathcal{V}$ for all $\xi \in \{B, G\}$. If $\sup(A) < V_G^{-1}(\bar{v})$ then $V_G(A) \cap \mathcal{V} = \emptyset$, but then $V_G(A) \subset (\underline{v}, \bar{v})$, $V_G(A) \cap \mathcal{V} = \emptyset$ and $V_G(A)$ has more length than A,[22] which is a contradiction. If $\inf(A) < V_G^{-1}(\bar{v}) \leq \sup(A)$ then $(V_G(\inf(A)), \bar{v}]$ is not in $\mathcal{V}$. Nevertheless, since $V_B^{-n}(v) \in \mathcal{V}$ for all $v \in \mathcal{V}$ and $n \in \mathcal{N}$ and $\lim_{n\to\infty} V_B^{-n}(v) \to \bar{v}$, this implies that $\mathcal{V}$ is empty, which is a contradiction. If $V_G^{-1}(\bar{v}) \leq \inf(A)$ then $V_B(A) \subset (\underline{v}, \bar{v})$, $V_B(A) \cap \mathcal{V} = \emptyset$ and $V_B(A)$ has more length than A, which again is a contradiction.

Proof of Proposition 2.9 (page 13)

Note that the equilibrium is not unique only if there is some $v \in \mathcal{V}$ such that $V_\xi(v) \in \mathcal{V}$ for all $\xi \in \{B, G\}$. Indeed, otherwise, if $\psi_0 \in \hat{V}^{-1}(v)$ and, for example, $V_G(v) \notin \mathcal{V}$, $\alpha(\emptyset)$ is uniquely given by $\alpha_G(\psi_0, \underline{\psi}(V_G(v)))$, where $\alpha_G(\cdot, \cdot)$ is defined in (A.2).

So, there is multiplicity in equilibrium only if $V_{h^t}(v) = \underline{v} = V_{\tilde{h}^s}(v)$ for two different histories $h^t \neq \tilde{h}^s$ (where V_{h^t} is defined as in footnote 19). It is easy to verify that this does not hold generically. So, since the public outcome distribution is only a function of $\alpha(h^t)$, the result holds.

Proof of Lemma 2.2 (page 14)

Take $v \in \mathcal{V}$. Given the definition of $\mathcal{V}$, it is obvious that, generically, there exists a unique $\xi \in \{B, G\}$ such that $V_\xi(v) \in \mathcal{V}$ and $V_{\bar{\xi}}(v) \notin \mathcal{V}$. This implies that $\underline{\psi}(V_{\bar{\xi}}(v)) = \overline{\psi}(V_{\bar{\xi}}(v))$. If $\xi = B$ then, for all $\psi_0 \in (\underline{\psi}(v), \overline{\psi}(v))$ we have that $\psi(G) = \underline{\psi}(V_{\bar{\xi}}(v))$, so $\alpha(\cdot)$ is clearly increasing. The reverse is true when $\xi = G$.

[22]Indeed, it is easy to verify that $\sup(V_G(A)) - \inf(V_G(A)) = \frac{\sup(A) - \inf(A)}{\delta}$.

Proof of Lemma 2.3 (page 15)

The proof of this lemma is given by the fact that $\mathcal{V}$ is dense, given the properties of $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ stated in Proposition 2.7 (which imply that that $\hat{V}^{-1}(v) = [\underline{\psi}(v), \overline{\psi}(v)]$ for all v and $\underline{\psi}(\cdot)$ and $\overline{\psi}(\cdot)$ are increasing).

Proof of Lemma 2.4 (page 15)

This is trivial given that $\hat{V}(\cdot)$ is a devil's staircase in the domain $(\psi^\star, 1]$. Indeed, since $\underline{\psi}(v) < \overline{\psi}(v)$ for all $v \in \mathcal{V}$ and $\mathcal{V}$ is dense, the result holds.

Proof of Proposition 2.10 (page 16)

Note that, by Proposition 2.2, $\mu(h^t) = 0$ if $\psi(h^t) < \psi^\star$ and $\mu(h^t) = 1$ if $\psi(h^t) > \psi^\star$. Given Propositions 2.2 and 2.3, it is easy to see that $V(\cdot)$ satisfies equations (2.3) and (2.4) if and only if it satisfies the following equation:

$$V(h^t) = \max_{\alpha \in [0,\nu]} \Big(\lambda\, \mu(h^t)\, \overline{V}_H + \alpha\, (1 - \lambda\, \mu(h^t))\, (\pi - c_L) \tag{A.5}$$

$$+ (1 - \lambda\, \mu(h^t))\, \delta\, \big(\alpha\, V(h^t, G) + (1-\alpha)\, V(h^t, B) \big) \Big) \, . \tag{A.6}$$

Note that if $\mu(h^t) = 1$, the previous equation coincides with (2.10).

Let's use tildes to denote a solution to (2.10). We first prove that $\tilde{V}(h^t) = 0$ whenever $\psi(h^t) < \psi^\star$. Assume otherwise; so we have that $v_L^* \equiv \sup \big\{ \tilde{V}(h^t) \,\big|\, h^t \text{ such that } \psi(h^t) < \psi^\star \big\} > 0$. Let h^t be such that $\psi(h^t) < \psi^\star$ and $\tilde{V}(h^t) > v_L^* - \varepsilon$, for some $\varepsilon > 0$. First notice that if $\alpha(h^t) > 0$, then

$$\tilde{V}(h^t, G) \geq \frac{\tilde{V}(h^T) + c_L - \pi}{\delta} \, . \tag{A.7}$$

If $\varepsilon > 0$ is small enough, we have $\tilde{V}(h^t, G) > v_L^*$. Therefore, it must be the case that $\alpha(h^t) < \nu$. This implies that $\tilde{V}(h^t, B) \geq \frac{\tilde{V}(h^t)}{\delta} > v_L^*$, but this is a contradiction, since $\psi(h^t, B) < \psi(h^t) < \psi^\star$. So, $v_L^* = 0$.

Now, let's prove that $\tilde{V}(h^t) = \lambda\, \tilde{\mu}(h^t)\, \overline{V}_H$ when $\psi(h^t) = \psi^\star$. If $\alpha(h^t) < \nu$, then the result trivially holds. In order to finally prove that it also holds for $\alpha(h^t) = \nu$, let's define, similar to before, $v_L^* \equiv \sup\{\tilde{V}(h^t) | h^t \text{ such that } \psi(h^t) = \psi^\star\}$, and let's first prove that $v_L^* \leq \lambda\, \overline{V}_H$. Assume otherwise; i.e., $v_L^* > \lambda\, \overline{V}_H$, and consider an equilibrium and history where $\psi(h^t) = \psi^\star$ and $\tilde{V}(h^t) = \lambda\, \overline{V}_H - \varepsilon$ for some $\varepsilon > 0$. Note that necessarily $\alpha(h^t) = \nu$. Then, it is easy to show that

$$\tilde{V}(h^t, G) - \tilde{V}(h^t) > \frac{c_L - \pi - \lambda\, \overline{V}_H}{\delta} \, .$$

Therefore, if $\varepsilon > 0$ is small enough, $\tilde{V}(h^t, G) > v_L^*$, but since $\alpha(h^t) = \nu$ we have $\psi(h^t, G) = \psi^\star$, which is a contradiction. Therefore, $v_L^* = \lambda\, \overline{V}_H$. Since $\tilde{V}(h^t, B) \geq 0$, equation (A.7) implies that $\tilde{V}(h^t, B) > \lambda\, \overline{V}_H$, so $\alpha(h^t) < \nu$.

As we mentioned, it is trivial to show that equations (2.10) and (A.5) are equivalent when $\psi(h^t) > \psi^\star$. Therefore, the statement holds.

Proof of Proposition 2.11 (page 17)

We will make this proof only for the case $\psi_0 = \psi^\star$, by showing that there exists an equilibrium where $\Pr(\psi_t = \psi^\star) = 1$ for all t. Extending it to the case $\psi_0 > \psi^\star$ only requires "pasting" this continuation play every time $\psi^\star$ is reached. Note that if $c_L - \pi \leq \delta\,\lambda\,\overline{V}_H$ and $\psi_0 = \psi^\star$, then there is an equilibrium where $\psi(h^t) = \psi^\star$ for all h^t. Indeed, consider an equilibrium where

$$\Pr(\tilde{P}(h^t) = \overline{V}_H) = \mu\,\mathbb{I}_{h^t_t = G} \ ,$$

for some $\mu \in [0,1]$ to be determined. In this equilibrium, we want the continuation payoff after a history h^t to depend only on the last signal. To verify that this equilibrium exists, let $V(\xi)$ denote the continuation payoff after signal ξ. Then

$$V(\xi) = \mu\,\mathbb{I}_{\xi=H}\,\lambda\,\overline{V}_H + (1 - \mu\,\mathbb{I}_{\xi=H}\,\lambda)\left(\nu\,(\pi - c_L) + \delta\,(\nu\,V(H) + (1-\nu)\,V(L))\right) \ .$$

The condition $c_L - \pi \leq \delta\,\lambda\,\overline{V}_H$ is necessary to ensure $V(\xi) \geq 0$ for all $\xi \in \{B, G\}$. We can then solve for μ by making the L-seller indifferent about exerting effort, in which case we find

$$\mu = \frac{c_L - \pi}{\lambda\,\overline{V}_H + \lambda\,\nu\,(c_L - \pi)} \in (0,1) \ .$$