

The Implications of Pricing on Social Learning

Itai Arieli¹, Moran Koren², and Rann Smorodinsky¹

¹Faculty of Industrial Engineering, Technion—Israel Institute of Technology.

²School of Mathematical Sciences, Tel Aviv University.

January 25, 2022

Abstract

Two firms produce substitute goods of unknown quality. At each stage the firms set prices and a consumer with private information and unit demand buys from one of the firms. Both firms and consumers see the entire history of prices and purchases. Will such markets aggregate information? Will the firm with the superior product necessarily prevail? We adapt the classic social-learning model by introducing strategic dynamic pricing. We provide necessary and sufficient conditions for asymptotic learning. In contrast to previous results, we show that asymptotic learning can occur when signals are bounded, namely, happens when the density of the consumers at the boundaries of the posterior belief distribution goes to zero. We refer to this property of the signal structure as the “vanishing margins” property.

JEL classification: D43, D83, L13.

1 Introduction

In many markets of substitute products, the value of the various alternatives may depend on some unknown variables. These may take the form of a future change in regulation, a technological shock, an environmental development, or prices in related upstream markets, etc. Although this information is unknown, individual consumers may receive some private information about these fundamentals. We ask whether markets aggregate information correctly and the ex-post superior product eventually dominates the market in such an environment.

Research supported by GIF research grant no. I-1419-118.4/2017, Ministry of Science and Technology grant 19400214, ISF grants 2030524, 2018889, the Bernard M. Gordon Center for Systems Engineering at the Technion, and the TASP Center at the Technion.

Moran Koren acknowledges the financial support of the Fulbright Visiting Postdoc Fellowship, the Management Science and Engineering Department of Stanford University, the Center of Mathematical Sciences and Applications of Harvard University, and the Ministry of Science and Technology of Israel (Yitzhak Shamir Fellowship). me@mkoren.org

For example, consider two competing pharmaceutical companies that produce alternative treatments (i.e., drugs) for a particular medical condition. One firm’s product is established while the other’s treatment is new. The clinical trials performed during the new drug’s FDA approval process induce a common prior over whichever product is superior. Before the new product is commercially launched, doctors receive a sample to be used within their patient community. Therefore, the doctors obtain some private information. Note that these signals are likely to be bounded as the number of free samples given to each doctor is often small. Additionally, as communities differ (e.g., in genetics and demography), the realized success rates of each treatment may differ from one doctor to another. As a result, doctors observe different signals. We ask whether society will correctly aggregate these signals and whether the better drug will necessarily prevail.

Whenever prices are fixed, classic results from the social learning theory tell us that doctors will herd on one of the drugs (possibly the inferior one). Our results, however, argue that when the drug firms adjust their prices dynamically, the aggregation of information depends only on the distribution of the *idiosyncratic* communities, i.e., those are the communities that drive the significant results. We capture the exact condition by the newly introduced notion of “vanishing margins”.¹

We study whether the learning process mentioned above guarantees an efficient outcome. We isolate the role of learning by introducing a simple duopoly model of common value. In our model, consumers, with a unit demand, choose between two substitute products, each with zero marginal cost of production. The timing of the interaction is as follows. Nature randomly chooses one of two states and thus determines the identity of the firm with the superior product. At each stage, both firms observe the entire history of the market – past prices and consumption decisions – and simultaneously set prices. After that, a single consumer arrives and receives a private signal regarding the state of nature. The consumer, based on his signal, the pair of product prices, and the market’s history, decides which product to buy (if any). Our main goal is to identify conditions under which the information in the market fully aggregates asymptotically, i.e., *asymptotic learning* holds.

When prices are set exogenously and are fixed throughout, the above model is precisely the standard herding model (Banerjee, 1992; Bikhchandani et al., 1992). In that model, as shown by Smith and Sørensen (2000), the characterization of asymptotic learning crucially depends upon the quality of agents’ private signals. In particular, one must distinguish between two families of signals: bounded versus unbounded. In the unbounded case, the agent’s private beliefs can, with positive probability, be arbitrarily close to zero and one. Therefore, no matter how many people herd on one alternative, the probability that the next agent will choose the other alternative is always positive. This property entails asymptotic learning.

The learning results in our model diverge from those of the canonical model when

¹For a recent example, consider the pricing of treatments for spinal muscular atrophy (SMA). Until recently, the only treatment for SMA was Biogen’s Spinraza treatment. In April 2019, Novartis received FDA approval for a competing treatment called Zolgensma. The research that led to the FDA approval was performed on 150 patients and thus contained little information about the treatment’s effect on the general population. Biogen responded to the threat by offering discounts to several large healthcare providers (see Gatlin (2019); Reuters (2019)).

signals are bounded.² In the herding model, there is always a positive probability that all agents will eventually choose the suboptimal alternative. However, intuition suggests that when prices are endogenized, they serve to prevent such a herding phenomenon. Hypothetically, once a herd develops on one firm’s product, the other firm will lower its product price to attract new consumers, and learning will not cease. It turns out that this intuition, although not entirely correct, does have some merit. In order for the intuitive argument to hold, signals must exhibit a property that we shall term *vanishing margins*.

We say that signals exhibit *vanishing margins* if the *density* of consumers at the posterior belief’s boundaries is zero. These consumers, i.e., consumers who receive signals that induce the most extreme posterior beliefs, are those who are likely to go against a herd and purchase the less popular product. From the market leader’s perspective, they comprise the *tail* of the distribution. The property of vanishing or nonvanishing margins serves as a measure of the tail’s thickness. Therefore, thin-tailed distributions are those that exhibit the vanishing margins property.

When society herds, each agent follows in the footsteps of his predecessors. Therefore, intuitively, one expects that a thick tail, i.e., a case in which there is a positive probability of seeing a consumer with an extreme signal, will induce learning. Our main result shows that the opposite occurs. When firms are myopic, signals are bounded, and prices are strategically determined, asymptotic learning holds if and only if signals have the vanishing margins property. We extend this result to forward-looking firms; however, to rule out collusive behavior that prevents learning, we need the assumptions that signals are informative enough and that firms use Markovian strategies that depend only on the public belief and not on the calendar time.

The intuition behind our main result is as follows. Consider a setting where the public belief is sufficiently extreme, and a clear market leader emerges. This leader faces the following dilemma. It can either capture the entire market by setting a low price or forego the “tail” consumers by setting a high price. Whenever signals exhibit thin margins, the latter option turns out to be optimal for the leader. Consequently, when “tail” consumers do arrive, the market is completely turned. When signals exhibit nonvanishing margins, aggressive pricing eventually prevails, thus halting any further information aggregation.

1.1 Related Literature

Our work primarily contributes to the herding literature initiated by Bikhchandani et al. (1992) and Banerjee (1992), who introduced models of social learning with agents who act sequentially. Their main contribution was to point out the possibility of rational herds that induce market failure. Smith and Sørensen (2000) noticed that such market failure happens only when signals are bounded. The lion’s share of follow-up studies focused on examining the robustness of the aforementioned

²When signals are unbounded, the rationale underlying Smith and Sørensen’s learning result applies to our model and so learning prevails.

condition in more elaborate settings.³

The first to incorporate dynamic pricing into herding models were Avery and Zemsky (1998). They considered a single firm whose product value is associated with an (unknown) state of nature. Instead of having the product offered at a fixed price, as in the earlier papers (e.g., Welch (1992)), they assumed that the price is set dynamically. In their model a market maker computes, at each stage, the expected value of the product and sets the price accordingly. By contrast, in our model prices are set endogenously by the profit-maximizing firms. Moreover, Avery and Zemsky (1998) showed that the presence of a market maker and dynamic pricing result in learning. In our setting such learning requires the addition of an extra condition, vanishing margins, to the information structure.

A model that is reminiscent of our model is that of Bose et al. (2006, 2008) who studied a herding model with a forward-looking monopolist that sells a good of uncertain quality to consumers. Consumers arrive sequentially and decide whether to purchase the product of the monopolist based on their predecessors' decisions, past prices, and an additional private signal. Bose et al. (2006) restricted attention to information structures with finitely many signals and Bose et al. (2008) to symmetric binary signals. In both models, it was shown that herding is inevitable. Additionally, they showed that if the public belief is sufficiently in favor of the monopoly, then the monopolist will price low enough to attract all consumers, regardless of their realized signal. As we show, their results rely on finite signals where the vanishing margins condition is never satisfied. The methodology and techniques discussed in the present paper may be used to show that in the monopolistic setting, i.e., when there is a single forward-looking firm that competes against an outside option, an information structure that exhibits vanishing margins guarantees asymptotic learning and one that doesn't exhibit vanishing margins (at both ends) guarantees that asymptotic learning fails.

Moscarini and Ottaviani (1997) studied the duopoly case in a static setting with a single-stage interaction between two firms and a single knowledgeable consumer. In fact, their model is a special case of our stage game $(\Gamma(\mu))$, which we study in Section 3. Similar to Bose et al. (2006, 2008), they restricted attention to finite, in fact binary and symmetric signal space. They showed that whenever the prior belief is above (or below) some threshold, all equilibria in their model are *deterrence equilibria* (see Definition 6). That is, in all equilibria, one firm prices out the other firm. Clearly, the emergence of a deterrence equilibrium implies that learning stops in the repeated model. In addition, the authors provided comparative statics over the threshold public belief for which learning stops as a function of the informativeness of the signal (and this is where the restricted signal space is leveraged). As signals become more informative the thresholds move to the extremes. Our result for the stage game, Theorem 2, argues that deterrence occurs, and hence learning stops, whenever the *vanishing margins* condition does not hold. As this condition

³For example, Lee (1993) presented a model with more than two states. Goeree et al. (2006) extended the model from a pure common value to include a private value ingredient. Eyster et al. (2014) studied a model where agents' utility is affected by congestion, and Acemoglu et al. (2010), Mossel et al. (2015), and Arieli and Mueller-Frank (2019) studied a model where agents observe only a partial set of their predecessors.

can never hold for a finite signal space (see Section 5.1), the result in Moscarini and Ottaviani (1997) follows as a corollary.

Mueller-Frank introduced a pair of models with dynamic pricing of a monopoly Mueller-Frank (2016) and a duopoly Mueller-Frank (2012). The models are very similar to ours with the distinction that for Mueller-Frank the firms have the informational advantage and know the true state of the world.⁴ Mueller-Frank asked whether social learning is sufficient to drive consumers to the optimal choice in the long run (“asymptotic efficiency”). Counterintuitively, he demonstrated equilibria in which this is not the case. By contrast, learning entails asymptotic efficiency in our setting (see Corollary 1).

While our major contribution is to the literature on social learning, the vanishing margins property and its effect on firms’ strategic behavior has interesting implications for market behavior and, in particular, for market entry and the adoption of new technologies. Previous studies on such questions assumed that incumbents have either an informational advantage (Bagwell (2007)) or a “first move” advantage, and that they can preempt entry by increasing capacity, investing in R&D (Acemoglu and Cao, 2015; Barrachina et al., 2014), or both (see Milgrom and Roberts, 1982a,b). Our stage game is an example of predatory pricing behavior, in which both incumbent and entrant act simultaneously, and no firm has an informational advantage.

The paper is organized as follows. Section 2 presents the model and the main theorem for the case where firms are myopic. In Section 3 we provide an equilibrium analysis of a stage game, which is central to the analysis of the learning model. We then leverage the analysis of the stage game to prove the aforementioned theorem for the myopic case. Section 4 is an extension of our model and results to the case where firms are farsighted. Section 5 informally discusses related issues.

2 Social Learning and Myopic Pricing

Our model comprises a countably infinite number of consumers, indexed by $t \in \mathbb{N}$, and two firms: Firm 0 and Firm 1. There are two states of nature $\Omega = \{0, 1\}$. In state ω , firm $\omega \in \{0, 1\}$ produces the superior product. We normalize the value of the superior product V_H to 1 and the value of the inferior product V_L to 0. In every time period t the two firms first set (nonnegative) prices $(\tau_0^t, \tau_1^t) \in [0, 1]^2$ for their products.⁵ Then consumer t receives a private signal and must decide whether to buy product 0, product 1, or neither product. Formally, the action set of every consumer is $A = \{0, 1, e\}$, where the action $a = i \in \{0, 1\}$ corresponds to the decision to buy from firm i and the action $a = e$ corresponds to the decision to exit and not to buy.⁶ The payoff of every consumer t , given the price vector (τ_0, τ_1)

⁴Mueller-Frank pointed out that when firms have an informational advantage, the equilibrium analysis crucially hinges on consumers’ off-equilibrium beliefs. This is not the case in our model, which consequently allows for robust observations.

⁵We conjecture that extending our model to allow for negative prices would have little effect on the asymptotic analysis and leave this question for future research.

⁶Our proofs go through with almost no changes in the case where $V_H > V_L \geq V_e \geq 0$ where V_e is the value from existing. In fact, we believe that the only required conditions for our analysis are $V_H > 0$

as a function of the realized state ω , is

$$u(a, \tau_0, \tau_1, \omega) = \begin{cases} 0 & \text{if } a = e \\ 1 - \tau_a & \text{if } a = \omega \\ -\tau_a & \text{otherwise.} \end{cases} \quad (1)$$

For simplicity, we assume that both firms have no marginal cost of production. Hence, firm i 's stage payoff, given a price vector (τ_0, τ_1) , can be described as a function of the consumer's decision as follows:

$$\pi_i(a, \tau_0, \tau_1, \omega) = \begin{cases} \tau_i & \text{if } a = i \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

We assume that the state ω is drawn at stage $t = 0$ according to a commonly known prior distribution, such that $P(\omega = 0) = \mu_0 = 1 - P(\omega = 1)$. The state ω is unknown to both the firms and the consumers. Each consumer $t \in \mathbb{N}$ forms a belief on the state using two sources of information: the history of prices and actions, $h_t \in H_t = ([0, 1]^2 \times \{0, 1, e\})^{t-1}$, and a private signal $s_t \in S$ (where S is some abstract measurable signal space). The firms observe only the realized history $h_t \in H_t$ at every time t and receive no private information. Conditional on the state ω , signals are independently drawn according to a probability measure F_ω . We refer to the tuple (F_0, F_1, S) as an *information structure*. We assume throughout that F_0 and F_1 are mutually absolutely continuous with respect to each other.⁷ The prior μ_0 and the functions F_0 and F_1 are common knowledge among consumers and firms.

Let $H = \cup_{t \geq 1} H_t$ be the set of all finite histories and let $H_\infty = ([0, 1]^2 \times \{0, 1, e\})^\infty$ be the set of all infinite histories. We let $\mathcal{A} \subset \Delta(\{0, 1, e\})^{[0, 1]^2 \times S}$ be the set of decision rules for the consumer; i.e., \mathcal{A} is the set of all measurable functions that map pairs consisting of a price vector and a signal to a (random) consumption decision. A *strategy for consumer t* is a measurable function $\sigma_t : H_t \rightarrow \mathcal{A}$ that maps every history $h_t \in H_t$ to a decision rule. We denote by $\bar{\sigma} = (\sigma_t)_{t \geq 1}$ a pure strategy profile for the consumers. We can view $\bar{\sigma}$ as a function $\bar{\sigma} : H \rightarrow \mathcal{A}$. A (behavioral) strategy for firm i is a (measurable) mapping $\bar{\phi}_i : H \rightarrow \Delta([0, 1])$. We note that the strategy profile $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ together with the prior μ_0 and the information structure (F_0, F_1, S) induces a probability distribution $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$ over $\Omega \times H_\infty \times S^\infty$.

Let $\mu_t = \mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\omega = 0 | h_t)$ be the probability that the state is $\omega = 0$ conditional on the realized history $h_t \in H$. We call μ_t *the public belief at time t* . We note that $\{\mu_t\}_{t=1}^\infty$ is a martingale and therefore, by the martingale convergence theorem, it must converge almost surely to a limit random variable $\mu_\infty \in [0, 1]$.

A strategy profile $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ and a history h_t induce both an expected payoff $\Pi_i^t(\tau_0, \tau_1, \bar{\sigma} | h_t)$ for every firm i and an expected consumer utility $U_t(\tau_0, \tau_1, \bar{\sigma} | h_t)$. We can now define the notion of a Bayesian Nash equilibrium for myopic firms.

Definition 1. A strategy profile $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ constitutes a *myopic Bayesian Nash equilibrium* if for every time t the following conditions hold for almost every history $h_t \in H_t$ that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$:

and $V_H > \max\{V_L, V_e\}$.

⁷ F_0 and F_1 are mutually absolutely continuous whenever $F_0(\hat{S}) > 0 \iff F_1(\hat{S}) > 0$ for any measurable set $\hat{S} \subset S$. Note that under this assumption the probability of a fully revealing signal, for which the posterior probability is either 0 or 1, is zero.

- For every $\tau \in [0, 1]$ and $i = 1, 2$,

$$\Pi_i^t(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t) \geq \Pi_i^t(\tau, \bar{\phi}_{-i}, \bar{\sigma}_t|h_t).$$

- For every price vector $(\tau_0, \tau_1) \in [0, 1]^2$, and every decision rule $\sigma \in \mathcal{A}$,

$$U_t(\tau_0, \tau_1, \bar{\sigma}(h_t)|h_t) \geq U_t(\tau_0, \tau_1, \sigma|h_t).$$

In words, a strategy profile $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ constitutes a myopic Bayesian Nash equilibrium if, for every time t and for almost every history $h_t \in H_t$ that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$, it holds that $\bar{\phi}_i(h_t)$ maximizes the conditional expected stage payoff to every firm i and $\bar{\sigma}(h_t)$ maximizes the conditional expected payoff to consumer t with respect to *every* price vector (τ_0, τ_1) .

Note that our notion of equilibrium is weaker than the notion of a subgame perfect equilibrium (henceforth SPE); however, it still eliminates equilibria with non-credible threats by consumers. One such equilibrium with non-credible threats is the following equilibrium: both firms ask for a price of 0 in every time period. Every consumer t never buys a product (i.e., plays e) unless both firms ask for a price of 0, in which case he buys Product 0 when $\mu_t \geq \frac{1}{2}$ and Product 1 when $\mu_t < \frac{1}{2}$. Note that this equilibrium is sustained by non-credible threats made by the consumer. Such threats are eliminated by the second condition, which requires that, conditional on the realized history h_t , the decision rule $\bar{\sigma}(h_t)$ be optimal with respect to *every* price vector (τ_0, τ_1) , and not just with respect to (τ_0^t, τ_1^t) .

The reason we focus on this set of equilibria instead of its more natural subset of SPE is the following. Our results apply either to all equilibria or to none. Therefore, our results hold for the subset of SPE. In addition, resorting to myopic Bayesian equilibria allows us to circumvent the nontrivial requirement of specifying off-equilibrium beliefs.

As is common in the literature, we define *asymptotic learning* as follows.

Definition 2. Fix an information structure (F_0, F_1, S) . Let $\mu_0 \in (0, 1)$ be the prior and let $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ be a strategy profile of the corresponding game. We say that *asymptotic learning holds* for μ_0 and $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ if the belief martingale sequence converges almost surely to a point belief assigning probability 1 to the realized state.

Let f_ω denote the Radon–Nikodym derivative of F_ω with respect to the probability measure $\frac{F_0 + F_1}{2}$. We consider the random variable $p(s) \equiv \frac{f_0(s)}{f_0(s) + f_1(s)}$, which is the posterior probability that $\omega = 0$, conditional on the signal s , when the prior over Ω is $(0.5, 0.5)$. Let $G_\omega(x) = F_\omega(\{s \in S | p(s) < x\})$, $\omega = 0, 1$, be the two cumulative distribution correspondences of the random variable $p(s)$ induced by the two probability distributions, F_ω , $\omega = 0, 1$, over S . As is standard in the literature, let $\text{co}(\text{supp}(p)) = [\bar{\alpha}, \underline{\alpha}]$ be the convex hull of the support of p .

The main goal of our paper is to provide a characterization of asymptotic learning under strategic pricing in terms of the information structure (F_0, F_1, S) . Such a characterization is provided by Smith and Sørensen (2000) for the standard herding model where prices are set exogenously. We start by presenting the formal definition of *bounded and unbounded signals* due to Smith and Sørensen (2000).

Definition 3. The information structure (F_0, F_1, S) is called *unbounded* if $\underline{\alpha} = 0$ and $\bar{\alpha} = 1$. The information structure (F_0, F_1, S) is *bounded* if $\underline{\alpha} > 0$ and $\bar{\alpha} < 1$.

In words, an information structure is unbounded if for every $\beta \in (0, 1)$ the two sets $\{s : p(s) > \beta\}$ and $\{s : p(s) < \beta\}$ have positive probability under $(F_\omega)_{\omega=0,1}$. Smith and Sørensen's characterization shows that in the standard herding model asymptotic learning holds under an unbounded information structure and fails under a bounded information structure.

2.1 Characterization of Asymptotic Learning

For ease of exposition we make the following assumption on $(G_\omega(x))_{\omega=0,1}$. We refer the reader to Section 5 for the general case.

Assumption 1. We assume that the functions $\{G_\omega(x)\}_{\omega=0,1}$ are differentiable on $(\underline{\alpha}, \bar{\alpha})$ with continuous derivatives $(g_\omega(x))_{\omega=0,1} : [\underline{\alpha}, \bar{\alpha}] \rightarrow \mathbb{R}_+$.

Definition 4. An information structure (F_0, F_1, S) exhibits *vanishing margins* if $g_1(\underline{\alpha}) = g_0(\bar{\alpha}) = 0$.

We next show how information aggregation depends on the vanishing margins property. The following theorem provides a full characterization of asymptotic learning in our model.

Theorem 1. *If signals are unbounded or if signals are bounded and exhibit vanishing margins, then asymptotic learning holds for every prior and every equilibrium. If signals are bounded and do not exhibit vanishing margins then asymptotic learning fails for every prior and every equilibrium.*

The rationale underlying the statement of Theorem 1 is as follows. The public belief gravitates toward one of the firms, say Firm 0, providing it with an opportunity to set a positive *deterrence price* that will drive the other firm out of the market, in which case learning stops. Raising the price above the deterrence price will drive the ultra-marginal consumers away from Firm 0 but will increase its profit from the rest of the consumers. The condition of vanishing margins captures the case where such an increase is always profitable, and therefore, learning continues.

Note that whenever asymptotic learning fails, only one firm, possibly the inferior one, prevails. This implies that consumers may consistently buy the inferior product with positive probability. However, when asymptotic learning holds, consumers and firms eventually learn the superior product. Does this imply that they will eventually buy from this firm or will the other firm be able to attract consumers periodically by offering low prices? In Corollary 1 we show that the former outcome holds and the probability of buying from the superior firm converges to one when asymptotic learning occurs.⁸

Corollary 1. *Let $(\bar{\sigma}, \bar{\tau}_0, \bar{\tau}_1)$ be a myopic Bayesian Nash equilibrium. If asymptotic learning holds, then conditional on state $\omega \in \Omega$,*

$$\lim_{t \rightarrow \infty} \mathbf{P}_{(\bar{\sigma}, \bar{\tau}_0, \bar{\tau}_1)}(\{\bar{\sigma}(h_t)(s, (\tau_0, \tau_1)) = \omega\} | \omega) = 1.$$

⁸In some variants of the herding model, such as those studied in Mueller-Frank (2012, 2016); Koren and Mueller-Frank (2022), asymptotic learning does not entail asymptotic efficiency.

Corollary 1 follows from the proof of our main theorem and its proof is relegated to Appendix D.

In Theorem 1 we distinguish between vanishing and nonvanishing margins. One may ask whether this condition is robust in the sense that will consumers learn the identity of the superior firm with high probability when the proportion of the consumers in the tail is small enough, but not zero. Our next result shows that this transition is continuous. Namely, as margins become thinner the associated thresholds approach zero and one. This, in turn, implies that the probability of herding on the optimal firm approaches one. To capture the notion of “thinner margins” we consider the density at $\underline{\alpha}$. A similar definition and result can be obtained for $\bar{\alpha}$.

Definition 5. An information structure exhibits the δ -margins property, for $\delta > 0$, if $g_0(\underline{\alpha}) \leq \delta$.

Let $\bar{\mu}$ be the upper *deterrence threshold* of an information structure with nonvanishing margins. This is the smallest prior such that, for every $\mu \geq \bar{\mu}$, deterrence occurs in $\Gamma(\mu)$.⁹

Proposition 1. For any $\varepsilon > 0$ and $\underline{\alpha}, \bar{\alpha} > 0$, there exists $\delta = \delta(\varepsilon, \underline{\alpha}) > 0$ such that if the information structure exhibits the δ -margins property, then the deterrence threshold satisfies $\bar{\mu} > 1 - \varepsilon$.

The proof of Proposition 1 is relegated to Appendix E.

3 The Proof of Theorem 1

In the proof of Theorem 1 we rely on the analysis of the following three-player stage game $\Gamma(\mu)$. The game comprises two firms and a single consumer and is derived from our sequential game by restricting the game to a single period. That is, in $\Gamma(\mu)$ the state is realized according to the prior μ (state 0 is realized with probability μ and state 1 with probability $1 - \mu$). The two firms post a price simultaneously (possibly at random) and a single consumer receives a private signal in accordance with (F_0, F_1, S) . Based on his private signal and the realized vector of prices, the consumer takes an action $a \in \{0, 1, e\}$. The utility for the consumer is determined by equation (1) and the utility for the firms is determined by equation (2).¹⁰

To guarantee the existence of an equilibrium, we allow firms to use mixed strategies. A mixed strategy for firm i is denoted by $\phi_i \in \Delta[0, 1]$. For a strategy profile $\phi = (\phi_0, \phi_1, \sigma)$, let $Pr_{\phi, \mu}$ be the probability over $\Omega \times [0, 1]^2 \times S$; the state, the price vector, and the signal set S are induced by ϕ , μ , and F_0, F_1 , respectively.

⁹Theorem 2, below, shows that whenever the information structure exhibits nonvanishing margins, $\bar{\mu} < 1$.

¹⁰This auxiliary stage-game model is reminiscent of a few models from the IO literature such as the duopolistic competition model with horizontal differentiation due to Hotelling (1929) (see chapter 7 in Tirole, 1988).

The Consumer's Best Reply

Given a prior μ and a pair of prices (τ_0, τ_1) , we let $v_\mu(\tau_0, \tau_1) \in [\underline{\alpha}, \bar{\alpha}]$ be the threshold in terms of the private belief above which Firm 0 is the unique best reply for the consumer. That is, choosing Firm 0 is uniquely optimal for the consumer if and only if $p(s) > v_\mu(\tau_0, \tau_1)$. A precise functional form of $v_\mu(\tau_0, \tau_1)$ is derived in equation (6) in Appendix A. We can therefore suppress the behavior of the consumer, which, under Assumption 1, is determined uniquely for every price vector (τ_0, τ_1) and almost every signal realization $s \in S$. Thus, we henceforth suppress the reference to the strategy of the consumer when we describe equilibrium strategies.

The Firms' Best Reply

We can write the expected profit of Firm 0 in the game $\Gamma(\mu)$ for the price vector $\tau = (\tau_0, \tau_1)$ as follows:

$$\begin{aligned} \Pi_0(\tau_0, \tau_1, \mu) = & \\ & (\mu(1 - G_0(v_\mu(\tau_0, \tau_1))) + (1 - \mu)(1 - G_1(v_\mu(\tau_0, \tau_1)))) \tau_0, \end{aligned} \quad (3)$$

where $(\mu(1 - G_0(v_\mu(\tau_0, \tau_1))) + (1 - \mu)(1 - G_1(v_\mu(\tau_0, \tau_1))))$ is the probability that the consumer buys from Firm 0 given the price vector (τ_0, τ_1) . A similar equation can be derived for $\Pi_1(\tau_0, \tau_1, \mu)$, the profit of Firm 1.

We make a distinction between two forms of perfect Bayesian Nash equilibria of the game $\Gamma(\mu)$: a *deterrence equilibrium*, where one of the firms is deterred and sells its product with probability zero, and a *non-deterrence equilibrium*, where both firms sell with positive probability. That is,

Definition 6. Let (ϕ_0, ϕ_1) be a SPE of $\Gamma(\mu)$. Say that firm j is *deterred* if

$$Pr_{\phi, \mu}(\sigma(\mu, s, \tau) = j) = 0.$$

In case one of the firms is deterred we refer to (ϕ_0, ϕ_1, σ) as a *deterrence equilibrium*.¹¹

We next study the properties of a deterrence equilibrium in the game $\Gamma(\mu)$. We denote the consumer's posterior belief after observing the signal s by $p_\mu(s)$. It follows readily from Bayes' rule that

$$p_\mu(s) = \frac{\mu p(s)}{\mu p(s) + (1 - \mu)(1 - p(s))}.$$

Since $p(s) \in [\underline{\alpha}, \bar{\alpha}]$ the above equation implies that $p_\mu(s) \in [\underline{\alpha}_\mu, \bar{\alpha}_\mu]$, where

$$\underline{\alpha}_\mu = \frac{\mu \underline{\alpha}}{\mu \underline{\alpha} + (1 - \mu)(1 - \underline{\alpha})} \text{ and } \bar{\alpha}_\mu = \frac{\mu \bar{\alpha}}{\mu \bar{\alpha} + (1 - \mu)(1 - \bar{\alpha})}.$$

Thus $\underline{\alpha}_\mu$ represents a tight lower bound on the posterior probability that the consumer assigns to Firm 0 being the superior firm. Similarly, $\bar{\alpha}_\mu$ represents a tight upper bound on the posterior probability that the consumer assigns to Firm 0 being

¹¹Note that in any SPE at most one firm is not deterred. Otherwise, both firms' expected profit would be zero. This is impossible since the a priori preferred firm can guarantee a positive expected profit, regardless of the other firm's strategy, by setting a sufficiently low positive price.

the superior firm. Assume that $\mu \geq \frac{1}{2}$ and consider a price vector (τ_0, τ_1) where $\tau_0 = 2\alpha_\mu - 1$. In this case the expected profit of a consumer who buys from Firm 0 is at least $\alpha_\mu - \tau_0 = \alpha_\mu - (2\alpha_\mu - 1) = 1 - \alpha_\mu$. By contrast, the expected profit of a consumer who buys from Firm 1 is at most $(1 - \alpha_\mu) - \tau_1 \leq 1 - \alpha_\mu$. Therefore, by setting a price $\tau_0^d := 2\alpha_\mu - 1$ Firm 0 guarantees that the consumer buys its product even if Firm 1 gives away its product for free (namely, even when $\tau_1 = 0$). The following proposition shows that, indeed, in a deterrence equilibrium where Firm 1 is deterred the price of Firm 0 will be $2\alpha_\mu - 1$. A symmetric claim holds for Firm 1.

Hereafter we abuse notation and write $\phi_i = \tau_i$ to denote a pure strategy of firm i that assigns probability one to the price τ_i .

Proposition 2. *Assume that (ϕ_0, ϕ_1) is a deterrence equilibrium in $\Gamma(\mu)$; then either*

$$\begin{aligned} & \text{Firm 1 is deterred, } \phi_0 = 2\alpha_\mu - 1, \Pi_0(\phi_0, \phi_1) = 2\alpha_\mu - 1, \text{ and } \alpha_\mu \geq \frac{1}{2}; \text{ or} \\ & \text{Firm 0 is deterred, } \phi_1 = 1 - 2\bar{\alpha}_\mu, \Pi_1(\phi_0, \phi_1) = 1 - 2\bar{\alpha}_\mu, \text{ and } \bar{\alpha}_\mu \leq \frac{1}{2}. \end{aligned}$$

As a corollary of Proposition 2 we have

Corollary 2. *If (ϕ_0, ϕ_1) is a deterrence equilibrium where firm i is deterred, then for $j \neq i$ it holds that $Pr_{\phi, \mu}(\sigma(\mu, s, \tau) = j) = 1$.*

Thus, whenever one firm is deterred, the other firm takes full control of the market and sells its product with probability one.

The following theorem summarizes the main characteristics of equilibria in the stage game $\Gamma(\mu)$. This characterization is the driving force behind the proof of Theorem 1.

Theorem 2. *Let $\mu \in (0, 1)$ and let (ϕ_0, ϕ_1, σ) be a Bayesian Nash subgame perfect equilibrium of the game $\Gamma(\mu)$:*

1. *If signals are unbounded, then no firm is deterred.*
2. *If signals are bounded and exhibit the vanishing margins property, then no firm is deterred.*
3. *If signals are bounded and do not exhibit the vanishing margins property, then*
 - (a) *If $g_1(\underline{\alpha}) > 0$, then for some sufficiently high prior $\bar{\mu} \in (0, 1)$, Firm 1 is deterred when $\mu > \bar{\mu}$.*
 - (b) *If $g_0(\bar{\alpha}) > 0$, then for some sufficiently low prior $\underline{\mu} \in (0, 1)$, Firm 0 is deterred when $\mu < \underline{\mu}$.*

To see why Theorem 2 is correct, assume without loss of generality that Firm 0 is the a priori preferred firm and offers the deterrence price, τ_0^d . The only possible profitable deviation is a price increase. Such a deviation will have two offsetting effects. On the one hand, it will increase the profit per sale but on the other hand it will result in a loss of market share, in particular a loss of consumers whose signal is least favorable toward Firm 0. When the vanishing margins condition is satisfied the

loss of market share is insignificant and is compensated by the profit per sale and so a price increase is profitable. By contrast, when the vanishing margins condition is not satisfied and the public belief is sufficiently skewed toward Firm 0, the market share loss becomes the dominant effect and the deviation is not profitable.

We relegate the proof of Theorem 2 and the complete analysis of the above stage game to Appendices A and B, respectively.

The Limit Arguments

In the following lemma, which is a direct implication of Definition 1, we connect the stage game with the sequential model.

Lemma 1. *A strategy profile $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ constitutes a myopic Bayesian Nash equilibrium if and only if for every time t , and for almost every history $h_t \in H_t$ that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$, the tuple $(\bar{\phi}_0(h_t), \bar{\phi}_1(h_t), \bar{\sigma}(h_t))$ is a subgame perfect equilibrium (SPE) of $\Gamma(\mu_t)$.*

The strong connection of $\Gamma(\mu)$ to our sequential game allows us to derive some insight into information aggregation from the subgame perfect equilibrium properties of $\Gamma(\mu)$, which we analyze next. We note that, under Assumption 1, in every perfect Bayesian Nash equilibrium of the game $\Gamma(\mu)$, the strategy σ prescribes a unique action for the consumer almost everywhere.

The proof of Theorem 1 leverages the results of Theorem 2 and connects the possibility of deterrence to asymptotic learning. Consider the case where signals exhibit vanishing margins. Theorem 2 tells us that at every stage t , no firm is deterred in $\Gamma(\mu_t)$ and so the consumer's actions are actually informative of his signal. This, in turn, implies that the public belief keeps evolving, which, by similar arguments to those underlying the results of Smith and Sørensen (2000), leads to asymptotic learning. On the other hand, assume learning is possible when signals do not exhibit vanishing margins. This means that eventually the public belief will be sufficiently in favor of the superior firm. This, by Theorem 2, the superior firm will price sufficiently aggressively to deter the other firm. In other words, all consumers in the stage game will necessarily buy from the aggressive firm, actions will then be uninformative, and learning will stall. Note that when prices are set exogenously this cannot happen.

We now turn to the formal proof of Theorem 1. We start with the following corollary of Lemma 6 (the proof can be found in Appendix B).

Corollary 3. *If signals exhibit vanishing margins or if signals are unbounded, then for every $\varepsilon > 0$ there exist some $r > \underline{\alpha}$ and $\delta' > 0$ such that if $\mu \in [\varepsilon, 1 - \varepsilon]$ and $\phi = (\phi_0, \phi_1)$ is a SPE of $\Gamma(\mu)$, then*

$$\mathbf{P}_{\mu, \phi}(v_\mu(\tau_0, \tau_1) \geq r) > \delta'.$$

A similar condition holds for Firm 1.

In words, by Theorem 2, if signals exhibit vanishing margins or if signals are unbounded, then the probability of a consumer going against the herd is positive. Corollary 3 argues that this probability cannot be arbitrarily close to zero if the prior is bounded away from the edges.

Proof of Theorem 1. We start the proof of Theorem 1 by showing that if the information structure (F_0, F_1, S) does not exhibit vanishing margins, then the martingale of the public belief must converge to an interior point. Indeed, let us assume without loss of generality that $g_1(\underline{\alpha}) > 0$. Let $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ be a myopic equilibrium. By Lemma 1, for almost every history $h_t \in H_t$ that is realized in accordance with $P_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$, the profile $(\bar{\sigma}(h_t), \bar{\phi}_0(h_t), \bar{\phi}_1(h_t))$ is a SPE of $\Gamma(\mu_t)$. By Theorem 2, there exists $\bar{\mu}$ such that, for all $\mu \in (\bar{\mu}, 1)$, there is a unique Bayesian Nash subgame perfect equilibrium of $\Gamma(\mu)$ in which the consumer chooses Firm 0 almost everywhere (i.e., Firm 1 is deterred by Firm 0). This implies that if $\mu_t \in (\bar{\mu}, 1]$, then $\mu_{t+1} = \mu_t$ almost everywhere. We note that since signals are never fully informative it must be the case that $\mu_t < 1$ for all t (almost everywhere). Therefore, if the vanishing margins property does not hold then asymptotic learning fails.

Next we show that if the vanishing margins property holds, then the public belief martingale converges to a limit belief in which the true state is assigned probability one. By Lemma 1, $(\bar{\phi}_0(h_t), \bar{\phi}_1(h_t), \bar{\sigma}(h_t))$ is a SPE of $\Gamma(\mu_t)$ for $P_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$ in almost every history $h_t \in H_t$. Corollary 3 implies that if $\mu_t \in [\epsilon, 1 - \epsilon]$ then for some $\delta' > 0$ and $r > \underline{\alpha}$, the realized price vector (τ_0, τ_1) satisfies $v_{\mu_t}(\tau_0, \tau_1) \geq r$ with probability at least δ' .

Since the distribution $G_0(\cdot)$ first-order stochastically dominates $G_1(\cdot)$ (see Lemma 14 in Appendix D), under any such price vector (τ_0, τ_1) there exists a probability at least $G_0(r) > 0$ that the consumer will not buy from Firm 0. Note that (again by Lemma 14)

$$\frac{G_0(v_{\mu_t}(\tau_0, \tau_1))}{G_1(v_{\mu_t}(\tau_0, \tau_1))} \leq \frac{G_0(r)}{G_1(r)} = \beta < 1.$$

Therefore, it follows from Bayes' rule that with probability at least $G_0(r)\delta'$ the public belief μ_{t+1} satisfies

$$\frac{\mu_{t+1}}{1 - \mu_{t+1}} = \frac{\mu_t}{1 - \mu_t} \frac{G_0(v_{\mu_t}(\tau_0, \tau_1))}{G_1(v_{\mu_t}(\tau_0, \tau_1))} \leq \frac{\mu_t}{1 - \mu_t} \beta. \quad (4)$$

Hence, in particular, if $\mu_t \in [\epsilon, 1 - \epsilon]$ then there exists a positive constant $\eta > 0$ such that $|\mu_{t+1} - \mu_t| > \eta$, with probability at least $G_0(r)\delta$.

By the martingale convergence theorem, the limit $\mu_\infty = \lim_{t \rightarrow \infty} \mu_t$ exists. By the above argument $\mu_\infty \in \{0, 1\}$ almost everywhere. This shows that asymptotic learning holds. \square

4 Social Learning and Farsighted Firms

In this section we show that by and large our result carries through to a setting where the firms are farsighted and maximize a discounted expected revenue stream. We extend our sequential model to the non-myopic case by defining the *non-myopic sequential consumption game*. In this model, as in the myopic case, each firm sets a price in every time period, except that now each firm tries to maximize its discounted sum of the stream of payoffs. We still retain the perfection assumption with respect to consumers.

Let $\Pi_i^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ denote the repeated game and $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ be a strategy profile in it. The expected payoff to firm i when the discount factor is $\delta > 0$ is

$$\Pi_i^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}) = E_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})} \left((1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \Pi_i^t(\bar{\phi}_0(h_t), \bar{\phi}_1(h_t), \bar{\sigma}(h_t) | h_t) \right).$$

We define a *Bayesian Nash equilibrium* as follows.

Definition 7. A strategy profile $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ constitutes a *Bayesian Nash equilibrium* if:

- For every $i = 1, 2$ and every strategy $\bar{\psi}_i$ of firm i ,

$$\Pi_i^\delta(\bar{\phi}_i, \bar{\phi}_{-i}, \bar{\sigma}) \geq \Pi_i^\delta(\bar{\psi}_i, \bar{\phi}_{-i}, \bar{\sigma}).$$

- For every time t , almost every history $h_t \in H_t$ that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$, every price vector $(\tau_0, \tau_1) \in [0, 1]^2$, and every decision rule $\sigma \in \mathcal{A}$, the following condition holds:

$$U_t(\tau_0, \tau_1, \bar{\sigma}(h_t) | h_t) \geq U_t(\tau_0, \tau_1, \sigma | h_t).$$

In the repeated interaction case the key impediment to asymptotic learning is that firms collude and “split” the market. This can happen when the discount factor is close enough to one and, say, in even-numbered time periods Firm 1 asks for a very high price and Firm 0 takes the full market (by playing α_{μ_t}) and in odd-numbered time periods Firm 0 asks for a very high price and Firm 1 takes the full market. Indeed, under this strategy profile, learning stops. This sort of equilibrium, however, is ruled out by the Markovian property that will be defined next.

Definition 8. A strategy $\bar{\phi}_i$ of firm i is called *Markovian* if there exists a measurable function $\psi_i : [0, 1] \rightarrow \Delta([0, 1])$ such that for every strategy of the other firm j and the consumer $(\bar{\phi}_j, \bar{\sigma})$ it holds, for every time t , that $\bar{\phi}_i(h_t) = \psi_i(\mu_t)$ for almost every history $h_t \in H_t$ that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$. A *Markovian equilibrium* is a Bayesian Nash equilibrium $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ such that $\bar{\phi}_0$ and $\bar{\phi}_1$ are Markovian strategies.

Here we focus our analysis on Markovian equilibria. Let $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ be a strategy profile, let $h_t \in H_t$, and denote by $\Pi_i^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma} | h_t)$ the continuation payoff to firm i in the subgame starting in history $h_t \in H_t$. Note that in a Markovian equilibrium there exists a measurable function $V_i : [0, 1] \rightarrow [0, 1]$ such that one can write $\Pi_i^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma} | h_t) = V_i(\mu_t)$ for almost every history h_t that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$. Thus the continuation payoff of firm i at time t depends only on the public belief μ_t . By Definition 7, if $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ constitutes a Bayesian Nash equilibrium, then $\bar{\phi}_i$ maximizes the continuation payoff $\Pi_i^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma} | h_t)$ of firm i for almost every history $h_t \in H_t$ that is realized in accordance with $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}$.

The use of Markovian strategies is common in economics. In our case Markovian strategies rule out tit-for-tat strategies and hence reflect the idea that firms are competing rather than colluding.¹² Indeed, under Markovian strategies prices are

¹²In a different setting Bhaskar et al. (2013) demonstrate that all equilibria that are robust to payoff perturbations are Markovian.

functions of the public belief only and so firms cannot rely on the calendar time. We note, however, that even when restricting attention to Markovian strategies public randomization may serve as a coordination device that the firms can use to split the market. In fact, public randomization drives folk theorem results such as those obtained in Yamamoto (2019). Our model does not include public randomization.

We next discuss another possibility that may lead to failure of learning. Consider signals with $\underline{\alpha}$ and $\bar{\alpha}$ that are close enough to $\frac{1}{2}$ and thus the signals are sufficiently uninformative. In such a case there exist strategy profiles such that in every time period both firms achieve a positive payoff but for every pair of realized prices (τ_0, τ_1) in the support of the two strategies only a single firm takes the full market. Thus, while there exists a positive probability for every firm to extract a positive payoff, for every realized pair of prices (τ_0, τ_1) consumers buy from a unique firm i with probability one, and therefore learning stops.

Such strategies cannot constitute a Markovian equilibrium since one can show that, for one of the firms, one price in the support yields a zero payoff in the stage game and so this firm could deviate by assigning probability zero to this price. However, what firms could potentially do is to play a strategy that approximates the aforementioned strategies so that for every pair of realized prices (τ_0, τ_1) in the support of the two strategies only a single firm takes the full market with very high probability that is less than one. In such strategies the public belief may keep slightly and vanishingly changing. This change in the public belief could serve as a calendar time that will be used by the firm to collude. To rule out this possibility we have added the requirement that signals satisfy $\underline{\alpha} < \frac{1}{3}$ and $\bar{\alpha} > \frac{2}{3}$.

We can now state our first result for farsighted firms.

Theorem 3. *Consider a bounded information structure (F_0, F_1, S) that exhibits the vanishing margins property. Asymptotic learning holds for any discount factor $\delta < 1$ in every pure Markovian equilibrium. If, in addition, $\underline{\alpha} < \frac{1}{3}$ and $\bar{\alpha} > \frac{2}{3}$, then asymptotic learning holds for any discount factor $\delta < 1$ in every Markovian equilibrium.*

We next outline the proof of Theorem 3. We first show that if $\underline{\alpha} < \frac{1}{3}$ and $\bar{\alpha} > \frac{2}{3}$ holds, then in every Markovian equilibrium where $\mu_t \geq \frac{1}{2}$, the identity of the dominant firm in any strategy profile as above is determined only by the realized price τ_0 of Firm 0 (see Lemma 10). That is, for any realized price τ_0 it is the case that either almost all consumers buy from Firm 0 or almost all consumers buy from Firm 1. This property implies that in the case where $\lim_{t \rightarrow \infty} \mu_t \geq \frac{1}{2}$ the profit for Firm 1 must approach zero as t goes to infinity. Otherwise, Firm 0 can eventually make a profitable deviation (see Lemma 11). Thus the only way that learning can stop is for one firm to take full control of the market from some time t onwards.

In order to rule out this possibility, we need to show that the public belief $\mu_t \in (0, 1)$ never reaches a point where one of the firms takes over the market and plays the deterrence price from time t onwards. This case is ruled out when signals have vanishing margins. The reason for this is that the leading firm can slightly increase the price for the product above the deterrence price and then play again the corresponding deterrence price from the next period onward. This is a profitable deviation since the firm increases its current period profit since the deterrence price

is not optimal in the one-stage game and also increases its continuation payoff since the deterrence price profit is convex (see Lemma 9).

For the converse direction we establish the following weaker result.

Theorem 4. *If signals are bounded and do not exhibit vanishing margins, then asymptotic learning fails for any discount factor $\delta > 0$ in every pure Markovian equilibrium.*

We prove Theorem 4 by way of contradiction. We show that if asymptotic learning holds then with positive probability the belief martingale μ_t reaches a point that is arbitrarily close to either zero or one such that the expected continuation payoff of the dominating firm i is bounded by $V_i(\mu_t) + CE|\mu_t - \mu_{t+1}|$ for some constant $C > 0$. We then show, as in the myopic case, that the dominating firm makes a profitable deviation to the deterrence price where all consumers buy its product with probability one.

The proofs for Theorem 3 and Theorem 4 are relegated to Appendix C.

5 Discussion

We now turn to discuss three natural questions that arise from our model and analysis:¹³

- Do our conclusions hold when the differentiability assumption on the signal distribution (Assumption 1) is relaxed?
- Is Blackwell order consistent with asymptotic learning?
- What if, instead of profit maximizing firms, the price pairs are set by a social planner who wishes to maximize welfare?

5.1 General Signals

Throughout the analysis we have restricted our attention to information structures (F_0, F_1, S) that satisfy Assumption 1. In many applications this assumption fails to hold. In particular, Assumption 1 does not hold when the set of signals is countable or finite. It is therefore important to understand whether our condition can be stated more generally to capture all signal distributions.

Fortunately, it turns out that such a general condition does exist. Let (F_0, F_1, S) be a general signal distribution and let G_ω be the CDFs of the posterior beliefs, as defined in Section 2. Define $g_0, g_1 \in [0, \infty]$ as follows:

$$g_0 = \liminf_{x \rightarrow \alpha^+} \frac{G_0(x)}{x - \alpha} \text{ and } g_1 = \liminf_{x \rightarrow \bar{\alpha}^-} \frac{1 - G_1(x)}{\bar{\alpha} - x}.$$

Obviously, g_0, g_1 are both well defined. Note that g_0 is defined using the limit from the left ($x \rightarrow \alpha^+$) whereas g_1 uses the limit from the right ($x \rightarrow \bar{\alpha}^-$).

We can now state the more general condition for vanishing margins as follows.

Definition 9. The information structure (F_0, F_1, S) satisfies vanishing margins if $g_0 = g_1 = 0$.

¹³We thank the anonymous reviewers for raising these questions.

Note that if (F_0, F_1, S) satisfies Assumption 1, then the condition in Definition 9 coincides with the condition in Definition 4. Moreover, note that for finite signal distribution we have $g_0 = g_1 = \infty$, and thus vanishing margins fails. Our results for myopic firms hold verbatim under the more general definition of vanishing margins.¹⁴

We omit the proofs for the general setting but note that the underlying ideas for the proofs are similar whereas their exposition becomes more cumbersome.¹⁵ The primary reason for this is that with an arbitrary information structure the consumer can be indifferent between two options (e.g., indifferent between the two products or between a product and exiting) with positive probability. Therefore, given a price pair, the consumer may have more than one best reply. In addition, it is not necessarily the case that any such best reply induces a two-player game between the firms that admits an equilibrium. The underlying reason is that the consumer strategy may lead to discontinuity in firms' payoffs as a function of prices. Under Assumption 1 the consumer has a unique best reply almost everywhere and such discontinuity can be ignored.

An additional challenge posed by the aforementioned discontinuity pertains to the mere existence of an equilibrium in $\Gamma(\mu)$. Absent this equilibrium, our results become vacuous. Fortunately, we can use the result of Reny (1999) to overcome this.

Consider the following specific best-reply consumer strategy: whenever a consumer is indifferent between buying from one firm and the outside option he always chooses to buy from the firm. Whenever a consumer is indifferent between buying from Firm 0 and Firm 1, and his expected utility from purchasing a product is at least zero, he chooses the firm that is a priori preferred. That is, in this case he chooses Firm 0 whenever $\mu \geq \frac{1}{2}$ and Firm 1 whenever $\mu < \frac{1}{2}$. In all other cases he strictly prefers one alternative and therefore chooses this alternative.

Under this consumer strategy, game $\Gamma(\mu)$ satisfies Reny's (1999) *better-reply secure* condition for any $\mu \in [0, 1]$. Theorem 3.1 in Reny (1999) thus guarantees the existence of a mixed subgame perfect equilibrium in $\Gamma(\mu)$.

5.2 Blackwell Ordering and Social Learning

In a classic paper, Blackwell (1953) defines a partial order over information structures. Roughly speaking, one information structure Blackwell dominates another if the superior information structure can be derived from the inferior information structure by virtue of having an additional signal. Blackwell shows that this order is consistent with the Bayesian decision maker's utility for all decision problems over the underlying state space. That is, one information structure Blackwell dominates another if and only if for any decision problem the Bayesian decision maker is (weakly) better off with the dominant one. A natural question is whether the Blackwell order is consistent with asymptotic learning.

¹⁴Although we have not written a rigorous proof for the case where firms are farsighted, we believe that the results carry through.

¹⁵A previous version of the paper with general signals, including all proofs, is available online at <http://bit.ly/SLPricingMK>.

Surprisingly, the answer is no. To see this, it is enough to show that the vanishing margins condition is inconsistent with Blackwell ordering. This is demonstrated next.¹⁶

Consider the following two information structures. G^1 induces a posterior distribution that is equal to the uniform distribution on $[\frac{1}{4}, \frac{3}{4}]$. Thus $[\underline{\alpha}, \bar{\alpha}] = [\frac{1}{4}, \frac{3}{4}]$ and G^1 is obtained using a constant density of $g^1 = 2$ over $[\frac{1}{4}, \frac{3}{4}]$. G^2 induces a posterior distribution that is obtained from a triangular density g^2 that is equal to $16x - 4$ on $[\frac{1}{4}, \frac{1}{2}]$ and $-16x + 12$ on $[\frac{1}{2}, \frac{3}{4}]$.

We note that $g_1^1(\frac{1}{4}), g_0^1(\frac{3}{4}) > 0$ whereas $g_1^2(\frac{1}{4}) = g_0^2(\frac{3}{4}) = 0$. Thus G^1 does not satisfy the vanishing margins condition and G^2 does satisfy the vanishing margins condition. By contrast, G^1 Blackwell dominates G^2 . This holds true since G^1 is a *mean-preserving spread* of G^2 (this can be easily verified through a simple calculation).

Paradoxically, when society does not asymptotically learn (as in G^1) the concealment of some information from the consumers (as in G^2) may lead to an improvement for society since it will now asymptotically learn.¹⁷

5.3 The Planner's Problem

In our model the price pair, at each stage, is driven by the two profit-maximizing firms. Alternatively, one could study a model where the price pair is set, at each stage, by a social planner. The planner has access to the same information as the firms do but wants to maximize social welfare, as measured by the discounted sum of utilities (see Smith et al., 2017). One could then study necessary and sufficient conditions on the information structures for the planner to asymptotically learn the realized state. In particular, one could ask whether, the vanishing margins condition is instrumental for learning in this case as well.

Apparently, the vanishing margins condition does not characterize asymptotic learning in the planner's problem. To get some intuition, recall the observation in Section 5.2 that the vanishing margins condition does not respect Blackwell ordering. Thus asymptotic learning may fail for a certain information structure but carry through for another information structure that is inferior in the Blackwell-ordering sense. We argue that this does not hold in the planner's problem. This follows from the fact that if the planner stopped experimenting at a certain prior $\mu > \frac{1}{2}$ for the superior information structure, then he cannot guarantee a social welfare that is higher than μ starting at the prior μ .¹⁸ Since the planner is always better off under the superior information structure, he cannot guarantee more than μ with respect to the inferior information structure. Thus he must stop experimenting also in the problem where the consumer's private information is obtained from the inferior information structure.

¹⁶We thank the coeditor for pointing out this inconsistency and suggesting the example.

¹⁷By contrast, unbounded signals do respect the Blackwell ordering.

¹⁸If the planner stops experimenting at μ , then all the consumers will buy the product of Firm 0, which generates a welfare of μ .

References

- Daron Acemoglu and Dan Cao. Innovation by entrants and incumbents. *Journal of Economic Theory*, 157:255–294, 2015. ISSN 10957235. doi: 10.1016/j.jet.2015.01.001. URL <http://dx.doi.org/10.1016/j.jet.2015.01.001>.
- Daron Acemoglu, Asuman Ozdaglar, and Ali ParandehGheibi. Spread of (mis)information in social networks. *Games and Economic Behavior*, 70(2): 194–227, 2010. ISSN 10902473. doi: 10.1016/j.geb.2010.01.005. URL <http://dx.doi.org/10.1016/j.geb.2010.01.005>.
- Daron Acemoglu, Munther A. Dahleh, Ilan Lobel, and Ozdaglar. Bayesian learning in social networks. *Review of Economic Studies*, 78:1201–1236, 2011. ISSN 0034-6527. doi: 10.1093/restud/rdr004. URL <http://restud.oxfordjournals.org/content/78/4/1201.abstract><http://www.sciencedirect.com/science/article/pii/S0899825603001441>.
- Itai Arieli and Manuel Mueller-Frank. Multidimensional social learning. *Review of Economic Studies*, 86(3):913–940, 2019. ISSN 0034-6527. doi: 10.1093/restud/rdy029. URL <https://academic.oup.com/restud/article/86/3/913/5034182>.
- Christopher Avery and Peter Zemsky. Multidimensional uncertainty and herd behavior in financial markets. *The American Economic Review*, 88(4):724–748, 1998.
- Kyle Bagwell. Signalling and entry deterrence: A multidimensional analysis. *RAND Journal of Economics*, 38(3):670–697, 2007. ISSN 0741-6261. doi: 10.1111/j.0741-6261.2007.00106.x. URL <http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=28548028&site=ehost-live&scope=site>.
- Abhijit V. Banerjee. A Simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, 1992. ISSN 00335533. doi: 10.2307/2118364. URL <http://www.jstor.org/stable/2118364>.
- Alex Barrachina, Yair Tauman, and Amparo Urbano. Entry and espionage with noisy signals. *Games and Economic Behavior*, 83:127–146, 2014. ISSN 08998256. doi: 10.1016/j.geb.2013.10.005. URL <http://dx.doi.org/10.1016/j.geb.2013.10.005>.
- Venkataraman Bhaskar, George J Mailath, and Stephen Morris. A foundation for markov equilibria in sequential games with finite social memory. *Review of Economic Studies*, 80(3):925–948, 2013.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. ISSN 0022-3808. doi: 10.1086/261849.
- David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24:265–272, 1953.

- Subir Bose, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. Dynamic monopoly pricing and herding. *The RAND Journal of Economics*, 37(4):910–928, 2006. ISSN 07416261. doi: 10.1111/j.1756-2171.2006.tb00063.x. URL <http://hdl.handle.net/2381/3394>.
- Subir Bose, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. Monopoly pricing in the binary herding model. *Economic Theory*, 37(2):203–241, 2008. ISSN 0938-2259. doi: 10.1007/s00199-007-0313-9. URL <http://link.springer.com/10.1007/s00199-007-0313-9>.
- Darrell Duffie, Semyon Malamud, and Gustavo Manso. Information percolation in segmented markets. *Journal of Economic Theory*, 153(1):1–32, 2014. ISSN 00220531. doi: 10.1016/j.jet.2014.05.006. URL <http://linkinghub.elsevier.com/retrieve/pii/S0022053114000726>.
- Erik Eyster, Andrea Galeotti, Navin Kartik, and Matthew Rabin. Congested observational learning. *Games and Economic Behavior*, 87:519–538, 2014. ISSN 10902473. doi: 10.1016/j.geb.2014.06.006.
- Allison Gatlin. Gene therapy Zolgensma gains approval in Spinal Muscular Atrophy treatment; Novartis stock pops. *Investor’s Business Daily*, 2019.
- Jacob K. Goeree, Thomas R. Palfrey, and Brian W. Rogers. Social learning with private and common values. *Economic Theory*, 28(2):245–264, 2006. ISSN 09382259. doi: 10.1007/s00199-005-0642-5.
- Helios Herrera and Johannes Hörner. Biased social learning. *Games and Economic Behavior*, 80:131–146, 2013. ISSN 08998256. doi: 10.1016/j.geb.2012.12.006. URL <http://dx.doi.org/10.1016/j.geb.2012.12.006>.
- Harold Hotelling. Stability in Competition. *The Economic Journal*, 39(153):41, mar 1929. ISSN 00130133. doi: 10.2307/2224214.
- Moran Koren and Manuel Mueller-Frank. The welfare costs of informationally efficient prices. *Games and Economic Behavior*, 131:186–196, 2022. ISSN 0899-8256. doi: <https://doi.org/10.1016/j.geb.2021.11.014>. URL <https://www.sciencedirect.com/science/article/pii/S0899825621001603>.
- In Ho Lee. On the convergence of informaional cascades. *Journal of Economic Theory*, 61:395–411, 1993.
- Paul Milgrom and John Roberts. Limit pricing and entry under incomplete information : An equilibrium Analysis. *Econometrica*, 50(2):443–459, 1982a.
- Paul Milgrom and John Roberts. Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2):280–312, 1982b. ISSN 00220531. doi: 10.1016/0022-0531(82)90031-X.
- Giuseppe Moscarini and Marco Ottaviani. Economic Models of social learning. pages 265–298, 1997.

- Elchanan Mossel, Allan Sly, and Omer Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015. ISSN 0012-9682. doi: 10.3982/ECTA12058. URL <http://arxiv.org/abs/1209.5527>.
- Manuel Mueller-Frank. Market power, fully revealing prices and welfare. 2012.
- Manuel Mueller-Frank. Price efficiency and welfare. 2016.
- Philip J. Reny. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67(5):1029–1056, 1999. doi: 10.2307/2999512. URL <http://www.jstor.org/stable/2999512>.
- Reuters. NHS England, Biogen reach deal on pricey drug for deadly disorder. *Reuters*, 2019.
- Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- Lones Smith, Peter Norman Sørensen, and Jianrong Tian. Informational herding, optimal experimentation, and contrarianism. *Working paper*, 2017.
- Jean Tirole. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press, 1988.
- Ivo Welch. Sequential sales, learning, and cascades. *Journal of Finance*, 47(2): 695–732, 1992.
- Yuichi Yamamoto. Stochastic games with hidden states. *Theoretical Economics*, 14(3):1115–1167, 2019.

A Proofs of the Stage Game

A.1 Equilibrium Analysis of $\Gamma(\mu)$

We begin by studying the consumer's best-reply strategy in $\Gamma(\mu)$. Recall that the consumer's posterior belief after observing the signal s is $p_\mu(s) = \frac{\mu p(s)}{\mu p(s) + (1-\mu)(1-p(s))}$.

Fix a price vector $\tau = (\tau_0, \tau_1)$ and note that the consumer optimizes his expected utility against τ if he follows the following strategy:

$$\sigma(\mu, s, \tau) = \begin{cases} a = 0 & \text{if } p_\mu(s) - \tau_0 \geq \max\{(1 - p_\mu(s)) - \tau_1, 0\} \\ a = 1 & \text{if } (1 - p_\mu(s)) - \tau_1 \geq \max\{p_\mu(s) - \tau_0, 0\} \\ a = e & \text{otherwise.} \end{cases} \quad (5)$$

Every realized price vector (τ_0, τ_1) induces two possible market scenarios. One is *fully covered market scenario*, where, under σ , the consumer never uses the outside option e and always buys from one of the firms for almost all signal realizations. The other is a *partially covered market scenario*, where $\sigma(\mu, s, \tau) = e$ holds with positive probability. We can infer from (5) that when the market is fully covered, the consumer buys from Firm 0 whenever $p_\mu(s) - \tau_0 \geq (1 - p_\mu(s)) - \tau_1$ and when the market is not fully covered, the consumer buys from Firm 0 whenever $p_\mu(s) - \tau_0 \geq 0$.

Given a prior μ and a pair of prices (τ_0, τ_1) , recall that $v_\mu(\tau_0, \tau_1)$ is the threshold in terms of the private belief above which buying from Firm 0 is the unique best-reply of the consumer. That is, choosing Firm 0 is uniquely optimal for the consumer if and only if $p(s) > v_\mu(\tau_0, \tau_1)$. One can easily see from the above equations that $v_\mu(\tau_0, \tau_1)$ has the following form:

$$v_\mu(\tau_0, \tau_1) = \begin{cases} \frac{(1-\mu)(1+\tau_0-\tau_1)}{2\mu-(2\mu-1)(1+\tau_0-\tau_1)} & \text{if the market is fully covered,} \\ \frac{(1-\mu)\tau_0}{\mu-(2\mu-1)\tau_0} & \text{otherwise.} \end{cases} \quad (6)$$

Note that $v_\mu(\tau_0, \tau_1)$ is a continuous function of (μ, τ_0, τ_1) .

We start with some preliminary results concerning equilibrium behavior in the game $\Gamma(\mu)$.

For $\mu \in [0, 1]$ we use the following shorthand: $G_\mu(x) = \mu G_0(x) + (1 - \mu)G_1(x)$. It follows by equation (3) that whenever the consumer's strategy σ obeys equation (5), the expected utility of Firm 0 in the game $\Gamma(\mu)$, $\Pi_0(\tau_0, \tau_1, \sigma)$, can be written as follows:

$$\Pi_0(\tau_0, \tau_1, \sigma) = (1 - G_\mu(v_\mu(\tau_0, \tau_1)))\tau_0. \quad (7)$$

For a mixed strategy profile (ϕ_0, ϕ_1) , let $\phi \in \Delta([0, 1] \times [0, 1])$ be the price probability distribution (ϕ_0, ϕ_1) induced over $[0, 1] \times [0, 1]$. By equation (7), Firm 0's payoff from the mixed strategy profile (ϕ_0, ϕ_1) can be written as follows:

$$\begin{aligned} \Pi_0(\phi_0, \phi_1, \sigma) &= \Pi_0(\phi_0, \phi_1) = \\ &= \int (\mu(1 - G_0(v_\mu(\tau_0, \tau_1)) + (1 - \mu)(1 - G_1(v_\mu(\tau_0, \tau_1))))\tau_0 d\phi(\tau_0, \tau_1), \end{aligned} \quad (8)$$

where $v_\mu(\cdot, \cdot)$ is defined as equation (6).

The next lemma provides an alternative way to write $v_\mu(\tau_0, \tau_1)$ and its derivative. This will turn out to be useful in the sequel. Consider the following function $\bar{v}_\mu : [0, 1]^2 \rightarrow \mathbb{R}$:

$$\bar{v}_\mu(\tau_0, \tau_1) \equiv \begin{cases} \log\left(\frac{1+\tau_0-\tau_1}{1-\frac{1+\tau_0-\tau_1}{2}}\right) - \log\left(\frac{\mu}{1-\mu}\right) & \text{if the market is fully covered,} \\ \log\left(\frac{\tau_0}{1-\tau_0}\right) - \log\left(\frac{\mu}{1-\mu}\right) & \text{if the market is not full.} \end{cases} \quad (9)$$

Lemma 2. *It holds that*

$$\frac{\partial \bar{v}_\mu(\tau)}{\partial \tau_0} \Big|_{\tau^*(\mu)} = \begin{cases} \frac{2}{1-(\tau_0-\tau_1)^2} & \text{if the market is fully covered,} \\ \frac{1}{\tau_0(1-\tau_0)} & \text{if the market is not full} \end{cases} \quad (10)$$

and

$$\frac{\partial v_\mu(\tau)}{\partial \tau_0} = \frac{e^{\bar{v}_\mu(\tau_0, \tau_1)}}{(1 + e^{\bar{v}_\mu(\tau)})^2} \frac{\partial \bar{v}_\mu(\tau_0, \tau_1)}{\partial \tau_0}. \quad (11)$$

Proof. The proof makes standard use of the log-likelihood ratio transformation (see, e.g., Smith and Sørensen (Smith and Sørensen, 2000), Herrera and Hörner (Herrera and Hörner, 2013), and Duffie et al. (Duffie et al., 2014)). The log-likelihood ratio of a belief $p \in [0, 1]$ is given by $\log(\frac{p}{1-p})$. In particular, the log likelihood ratio of the posterior belief is

$$\log\left(\frac{p_\mu(s)}{1-p_\mu(s)}\right) = \log\left(\frac{\mu}{1-\mu}\right) + \log\left(\frac{p(s)}{1-p(s)}\right). \quad (12)$$

It follows from equation (6) that a consumer with private belief $p_\mu(s)$ prefers Firm 0 if and only if

$$\log\left(\frac{p_\mu(s)}{1-p_\mu(s)}\right) \geq \log\left(\frac{v_\mu(\tau)}{1-v_\mu(\tau_0, \tau_1)}\right) = \bar{v}_\mu(\tau_0, \tau_1).$$

Equation (11) then follows directly from the fact that $v_\mu(\tau) = \frac{e^{\bar{v}_\mu(\tau)}}{1+e^{\bar{v}_\mu(\tau)}}$. \square

The following simple observation will be useful in our analysis.

Observation 1. *Let $\mu \in [0, 1]$ and let (ϕ_0, ϕ_1) be a SPE of $\Gamma(\mu)$. The following properties hold:*

$$\phi_0([2\alpha_\mu - 1, 1]) = 1 \text{ and } \phi_1([1 - 2\bar{\alpha}_\mu, 1]) = 1.$$

Proof. We prove the observation for Firm 0. Note that if $\alpha_\mu \leq \frac{1}{2}$ we have nothing to prove. Assume that $\alpha_\mu > \frac{1}{2}$; then, if $\tau_0 = 2\alpha_\mu - 1$, the consumer will buy from Firm 0 almost everywhere for almost every signal realization s and every price $\tau_1 \geq 0$ of Firm 1. To see this, note that $p_\mu(s) > \alpha_\mu$ for almost every signal $s \in S$. Therefore,

$$p_\mu(s) - (2\alpha_\mu - 1) > 1 - p_\mu(s).$$

This shows that for a price $\tau_0 = 2\alpha_\mu - 1$ the consumer buys from Firm 0 almost everywhere even for $\tau_1 = 0$. In particular, under any price $\tau_0 \leq 2\alpha_\mu - 1$ the expected profit of Firm 0 is τ_0 . Therefore, if $\alpha_\mu > \frac{1}{2}$, the price $2\alpha_\mu - 1$ strictly dominates all prices $\tau_0 < 2\alpha_\mu - 1$ for Firm 0. \square

A.2 Properties of Deterrence Equilibria

A key property of a deterrence equilibrium is given in the following lemma.

Lemma 3. *Let (ϕ_0, ϕ_1) be a deterrence equilibrium in the game $\Gamma(\mu)$. If Firm 1 is deterred, then $\underline{\alpha}_\mu \geq \frac{1}{2}$. Symmetrically, if Firm 0 is deterred, then $\bar{\alpha}_\mu \leq \frac{1}{2}$.*

In words, if firm i is driven out of the market (in the sense that the consumer surely does not buy from it), it must be the case that the consumer's posterior belief assigns a probability of at most $\frac{1}{2}$ that firm i is the superior firm.

Proof. Assume to the contrary that $\underline{\alpha}_\mu < \frac{1}{2}$ and that (ϕ_0, ϕ_1) is a deterrence equilibrium in which Firm 1 is deterred. In this case $\Pi_1(\phi_0, \phi_1) = 0$. Consider a deviation of Firm 1 to the pure strategy $\tau_1 = \frac{1-2\underline{\alpha}_\mu}{2} > 0$. By equation (5) we can conclude that any consumer whose signal falls in the set $\{s \in S | p_\mu(s) \in [\underline{\alpha}_\mu, \underline{\alpha}_\mu + \frac{\underline{\alpha}_\mu}{2} + \frac{1}{4}]\}$ will choose Firm 1 almost everywhere for any equilibrium strategy ϕ_0 for Firm 0.

Note that the set $\{s \in S | p_\mu(s) \in [\underline{\alpha}_\mu, \underline{\alpha}_\mu + \varepsilon]\}$ has positive probability for every $\varepsilon > 0$ and in particular for $\varepsilon = \underline{\alpha}_\mu + \frac{\underline{\alpha}_\mu}{2} + \frac{1}{4}$. Therefore this deviation entails a positive expected utility for Firm 1 and hence a profitable deviation, thus contradicting the equilibrium assumption. \square

We next turn to prove Proposition 2.

Proof of Proposition 2. Let us assume without loss of generality that Firm 1 is deterred and so $\underline{\alpha}_\mu \geq \frac{1}{2}$ (by Lemma 3). It follows from Observation 1 that $\phi_0([2\underline{\alpha}_\mu - 1, 1]) = 1$. Assume by way of contradiction that $\phi_0[2\underline{\alpha}_\mu - 1 + \delta, 1] > 0$ for some positive $\delta > 0$ and consider the price $\tilde{\tau}_1 = \frac{\delta}{2}$ for Firm 1 (the deterred firm). In this case, for any realized $\tau_0 \in [2\underline{\alpha}_\mu - 1 + \delta, 1]$, any consumer with a private signal s such that $p_\mu(s) \in [\underline{\alpha}_\mu, \underline{\alpha}_\mu + \frac{\delta}{4}]$, an event whose probability is positive, will buy from Firm 1, which, in turn, will have a positive utility. In the deterrence equilibrium Firm 1's utility is obviously zero and hence the price $\tilde{\tau}_1 = \frac{\delta}{2}$ constitutes a profitable deviation, thus contradicting the equilibrium assumption. Therefore $\phi_0[2\underline{\alpha}_\mu - 1 + \delta, 1] = 0$ for any $\delta > 0$. Hence Firm 0 plays $\tau_0 = 2\underline{\alpha}_\mu - 1$ almost everywhere, as claimed. \square

By Lemma 3, the condition $\underline{\alpha}_\mu \geq \frac{1}{2}$ is necessary in order for a deterrence equilibrium (in which Firm 1 is deterred) to exist. We now turn to study the implications of this condition.

Lemma 4. *If (ϕ_0, ϕ_1) is a non-deterrence Bayesian Nash SPE of $\Gamma(\mu)$, then the following conditions hold: $\phi_0((2\underline{\alpha}_\mu - 1, 1)) > 0$, $\Pi_0(\phi_0, \phi_1, \sigma) \geq 2\underline{\alpha}_\mu - 1$, and $\Pi_1(\phi_0, \phi_1) > 0$. Symmetrically for Firm 1, $\phi_1((1 - 2\bar{\alpha}_\mu, 1)) > 0$, $\Pi_1(\phi_0, \phi_1, \sigma) \geq 1 - 2\bar{\alpha}_\mu$, and $\Pi_0(\phi_0, \phi_1) > 0$.*

Proof. We prove the first part of the lemma. Lemma 1 implies that $\phi_0([2\underline{\alpha}_\mu - 1, 1]) = 1$. We further note that if (ϕ_0, ϕ_1) is a SPE profile for which ϕ_0 is the Dirac measure on $2\underline{\alpha}_\mu - 1$, then $\Pi_0(\phi_0, \phi_1) = 2\underline{\alpha}_\mu - 1$, which means that the consumer buys from Firm 0 almost everywhere. Hence such an equilibrium must be a deterrence equilibrium. Therefore, it must hold that $\phi_0((2\underline{\alpha}_\mu - 1, 1)) > 0$.

The fact that $\Pi_1(\phi_0, \phi_1) > 0$ follows since, as in the proof of Proposition 2, if $\phi_0((2\underline{\alpha}_\mu - 1, 1)) > 0$, then Firm 1 can guarantee a positive payoff against ϕ_0 . \square

B Proof of Theorem 2

Unbounded signals

We begin the proof of Theorem 2, by studying the case of unbounded signals, i.e., where $\underline{\alpha} = 0$ and $\bar{\alpha} = 1$. The following corollary shows that whenever signals are unbounded there cannot be a deterrence equilibrium. In fact all equilibria are non-deterrence equilibria.

Corollary 4. *If signals are unbounded then there are no deterrence equilibria in $\Gamma(\mu)$.*

Proof. Since $\bar{\alpha} = 0$ and $\underline{\alpha} = 1$ it follows that $\underline{\alpha}_\mu = 0$ and $\bar{\alpha}_\mu = 1$. The proof now follows from Lemma 3. \square

Bounded signals with vanishing margins

We now consider the case where signals are bounded, i.e., $\underline{\alpha}, \bar{\alpha} \in (0, 1)$, and signals exhibit the vanishing margins property, i.e., $g_1(\underline{\alpha}) = 0$. In the following lemma we show that the vanishing margins property also yields that $g_0(\underline{\alpha}) = 0$.

Lemma 5. *If the information structure (F_0, F_1, S) exhibits vanishing margins then $g_0(\underline{\alpha}) = 0$.*

Proof. Since the vanishing margins condition holds we have that $g_1(\underline{\alpha}) = 0$. Assume to the contrary that $g_1(\underline{\alpha}) < g_0(\underline{\alpha})$. Since F_0, F_1 are absolutely mutually continuous, there exists ε such that $\int_{\underline{\alpha}}^{\underline{\alpha}+\varepsilon} g_0(s)ds > \int_{\underline{\alpha}}^{\underline{\alpha}+\varepsilon} g_1(s)ds \Rightarrow G_0(\underline{\alpha} + \varepsilon) > G_1(\underline{\alpha} + \varepsilon)$. This stands in contradiction to Lemma 14 in Appendix D, which shows that G_0 first order stochastically dominates G_1 . \square

The second part of Theorem 2 is proved in the following proposition.

Proposition 3. *If the information structure (F_0, F_1, S) exhibits vanishing margins, then for every $\mu \in (0, 1)$ there is no deterrence equilibrium in $\Gamma(\mu)$.*

Proof. Without loss of generality assume that $\mu \in (\frac{1}{2}, 1)$ and assume to the contrary that there exists a deterrence equilibrium in $\Gamma(\mu)$. By Lemma 3, the only possible deterrence equilibrium is one in which Firm 1 is deterred and by Proposition 2 it must take the form of $(2\underline{\alpha}_\mu - 1, \phi_1)$. Therefore, $\Pi_0(2\underline{\alpha}_\mu - 1, \phi_1) = 2\underline{\alpha}_\mu - 1$. We first claim that it is sufficient to show that

$$\Pi_0(2\underline{\alpha}_\mu - 1 + \varepsilon, 0) - \Pi_0(2\underline{\alpha}_\mu - 1, 0) > 0 \quad (13)$$

for some $\varepsilon > 0$. To see this, note that $\Pi_0(2\underline{\alpha}_\mu - 1, \phi_1) = 2\underline{\alpha}_\mu - 1$ for any mixed strategy ϕ_1 of Firm 1. In addition, for any fixed price τ_0 of Firm 0 the payoff $\Pi_0(\tau_0, \tau_1)$ is (weakly) decreasing in τ_1 . Therefore, the inequality in (13) implies that $\Pi_0(2\underline{\alpha}_\mu - 1 + \varepsilon, \phi_1) > 2\underline{\alpha}_\mu - 1$ for any mixed strategy ϕ_1 of Firm 1. Therefore, if $2\underline{\alpha}_\mu - 1 + \varepsilon$ yields a profitable deviation to Firm 0 against price $\tau_1 = 0$ it also yields a profitable deviation with respect to any strategy ϕ_1 of Firm 1.

To establish equation (13) note that

$$\begin{aligned} & \frac{\partial \Pi_0(\tau_0, 0)}{\partial \tau_0} \Big|_{\tau_0=2\alpha_\mu-1} = \\ & 1 - (2\alpha_\mu - 1) \left(\frac{\partial v_\mu(\tau_0, 0)}{\partial \tau_0} \Big|_{2\alpha_\mu-1} \right) (\mu g_0(\alpha) + (1 - \mu)g_1(\alpha)). \end{aligned} \quad (14)$$

Since vanishing margins holds, by Lemma 5 we have that $g_1(\alpha) = g_0(\alpha) = 0$. Therefore equation (14) implies that

$$\frac{\partial \Pi_0(\tau_0, 0)}{\partial \tau_0} \Big|_{\tau_0=2\alpha_\mu-1} = 1.$$

Hence $\Pi_0(2\alpha_\mu - 1 + \varepsilon, 0) - \Pi_0(2\alpha_\mu - 1, 0) > 0$ for all sufficiently small $\varepsilon > 0$. \square

If a deterrence equilibrium does not exist then at each stage of our sequential setting the actual action of the consumer will give us additional information and the public belief will shift. Intuitively, this drives the learning result. However, it turns out that this is not enough. Herrera and Hörner (Herrera and Hörner, 2013) show that in the herding model the fact that $\mu_t \neq \mu_{t+1}$ almost everywhere does not imply that asymptotic learning holds. In order to establish asymptotic learning the following stronger result is required.

Lemma 6. *If the information structure (F_0, F_1, S) exhibits the vanishing margins condition or signals are unbounded, then for every $\varepsilon > 0$ there exists $\delta > 0$ such that if $\mu \in [\varepsilon, 1 - \varepsilon]$ and $\phi = (\phi_0, \phi_1, \sigma)$ is a SPE of $\Gamma(\mu)$, then $P_{\mu, \phi}(\sigma(\tau_0, \tau_1, s) = a) \leq 1 - \delta$ for any Firm $i = 0, 1$.*

Note that Lemma 6 subsumes Proposition 3 in that under vanishing margins, a deterrence equilibrium does not exist. This applies that there exists an upper bound on the probability that the consumer buys from any given firm. This in turn implies that if μ is bounded away from zero and one, then the distance between μ and the posterior probability, conditional on the action of the current consumer, is bounded away from zero.

Proof. We prove the lemma under the assumption that the information structure (F_0, F_1, S) exhibits vanishing margins. The proof for the unbounded case is similar and therefore omitted.

Assume by way of contradiction that there exists an $\varepsilon > 0$ and a sequence of SPE $\phi^k = (\phi_0^k, \phi_1^k, \sigma_k)$ of $\Gamma(\mu_k)$ such that $\mu_k \leq 1 - \varepsilon$ and

$$\lim_{k \rightarrow \infty} P_{\mu_k, \phi^k}(\sigma_k(\tau_0, \tau_1, s) = 0) = 1.$$

In words, as k increases, the prior approaches 1. We show that this entails that the probability of the consumer buying from Firm 0 approaches 1 as well.

We can clearly assume (possibly by considering subsequences) that the sequence $\{\Pi_\omega(\phi_0^k, \phi_1^k, \sigma_k)\}_{k=1}^\infty$ converges. Similarly, we can assume that $\{(\phi_0^k, \phi_1^k, \mu_k)\}_{k=1}^\infty$ converges to some¹⁹ (ϕ_0, ϕ_1, μ) . As $\lim_{k \rightarrow \infty} P_{\mu_k, \phi^k}(\sigma_k(a = 0)) = 1$, the limit profit of Firm 1 shrinks to zero:

$$\lim_{k \rightarrow \infty} \Pi_1(\phi_0^k, \phi_1^k, \sigma_k) = 0.$$

¹⁹The convergence of ϕ_ω^k is assumed with respect to the weak topology.

It follows that the limit price of Firm 0, $\lim_{k \rightarrow \infty} \phi_0^k$, is the pure deterrence price $2\alpha_\mu - 1$. To see this assume by way of contradiction that $\phi_0((2\alpha_\mu - 1 + \eta, 1)) > 0$ for some $\eta > 0$. We claim that Firm 1 can guarantee a positive profit against ϕ_0 by playing $\tau_1' = \frac{\eta}{2}$ in $\Gamma(\mu)$. In this case, the consumers for whom $p_\mu(s) \in [\alpha_\mu, \alpha_\mu + \frac{\delta}{4})$ will strictly prefer to buy from Firm 1. This yields a positive expected payoff that is bounded away from zero, for all sufficiently large k . A contradiction to the fact that $(\phi_0^k, \phi_1^k, \mu_k)$ is a SPE for every k .

Consider the game $\Gamma(\mu_k)$ and the strategy profile $(\phi_0^k, \phi_1^k) = (2\alpha_{\mu_k} - 1, \phi_1^k)$. A standard continuity consideration implies that since $\phi^k = (\phi_0^k, \phi_1^k)$ is a SPE of $\Gamma(\mu_k)$ and $\lim_{k \rightarrow \infty} (\phi_0^k, \phi_1^k, \mu_k) = (\phi_0, \phi_1, \mu)$, it holds that (ϕ_0, ϕ_1) is a SPE of $\Gamma(\mu)$. Therefore $(\phi_0, \phi_1) = (2\alpha_\mu - 1, \phi_1)$. Under the price $2\alpha_\mu - 1$, the consumer buys from Firm 0 almost everywhere.

This yields that (ϕ_0, ϕ_1) is a deterrence equilibrium of $\Gamma(\mu)$, which stands in contradiction to Proposition 2. \square

We get the following corollary of Lemma 6.

Corollary 3. If signals exhibit vanishing margins or if signals are unbounded, then for every $\varepsilon > 0$ there exists some $r > \underline{\alpha}$ and $\delta' > 0$ such that if $\mu \in [\varepsilon, 1 - \varepsilon]$ and $\phi = (\phi_0, \phi_1)$ is a SPE of $\Gamma(\mu)$, then

$$P_{\mu, \phi}(v_\mu(\tau_0, \tau_1) \geq r) > \delta'.$$

A similar condition holds for Firm 1.

Bounded signals without vanishing margins

The following lemma shows that the consumer's threshold signal approaches the lower bound $\underline{\alpha}$ as μ approaches 1 in every SPE.

Lemma 7. Let $\{\mu_k\}_{k=1}^\infty \subseteq (0, 1)$ be a sequence of priors such that $\lim_{k \rightarrow \infty} \mu_k = 1$. Let $\phi^k = (\phi_0^k, \phi_1^k, \sigma_k)$ be a SPE for the game $\Gamma(\mu_k)$. Then the following holds for every $\epsilon > 0$:

$$\lim_{k \rightarrow \infty} P_{\mu_k, \phi^k}(v_{\mu_k}(\tau_0, \tau_1) \in [\underline{\alpha}, \underline{\alpha} + \epsilon]) = 1.$$

Proof. Assume by way of contradiction that there exists some $\epsilon_0 > 0$ and $\delta > 0$ for which the following holds (possibly considering a subsequence):

$$\lim_{k \rightarrow \infty} P_{\mu_k, \phi^k}(v_{\mu_k}(\tau_0, \tau_1) \in [\underline{\alpha}, \underline{\alpha} + \epsilon]) < 1 - \delta.$$

This implies that the payoff to Firm 0 is at most $1 - \delta G_0(\underline{\alpha} + \epsilon_0) < 1$. To see this note that with a probability of at least $\delta > 0$ it holds for sufficiently large k that $v_{\mu_k}(\tau_0, \tau_1) > \underline{\alpha} + \epsilon$. Therefore, with probability at least δ the profit of Firm 0 is bounded by $1 - G_0(\underline{\alpha} + \epsilon_0)$. Therefore, the expected profit of Firm 0 is bounded by

$$\delta(1 - G_0(\underline{\alpha} + \epsilon_0)) + (1 - \delta) = 1 - \delta G_0(\underline{\alpha} + \epsilon_0).$$

Since signals are bounded and $\lim_{k \rightarrow \infty} \mu_k = 1$ it must hold, for sufficiently large k , that

$$2\alpha_{\mu_k} - 1 > 1 - \delta G_0(\underline{\alpha} + \epsilon).$$

In the game $\Gamma(\mu_k)$, consider a deviation by Firm 0 to the pure price $\tau_0 = 2\alpha_{\mu_k} - 1$. Firm 0 then guarantees an expected revenue of

$$2\alpha_{\mu_k} - 1 > 1 - \delta G_0(\alpha + \epsilon),$$

which implies a contradiction. \square

The following corollary shows that as μ approaches 1, it holds that for any SPE of $\Gamma(\mu)$, the equilibrium price of Firm 0 approaches 1.

Corollary 5. *Let $\{\mu_k\}_{k=0}^{\infty} \subset (0, 1)$ be a sequence of priors that converges to 1, and let $(\phi_0^k, \phi_1^k, \sigma_k)$ be a SPE of $\Gamma(\mu_k)$ for any k . Then,*

$$\lim_{k \rightarrow \infty} \phi_0^k = 1.$$

Corollary 5 follows from Proposition 2 and Lemma 4.

The following lemma provides an upper limit to the support of Firm 1 in a non-deterrence equilibrium.

Lemma 8. *If (ϕ_0, ϕ_1) is a non-deterrence equilibrium, then $\phi_0([2\alpha_{\mu} - 1, \bar{\alpha}_{\mu}]) = 1$ and $\phi_1([1 - 2\bar{\alpha}_{\mu}, 1 - \alpha_{\mu}]) = 1$.*

Proof. It follows from Proposition 2 that $\Pi_0(\phi_0, \phi_1) > 0$. Note further that for any price $\tau_0 > \bar{\alpha}_{\mu}$ the consumer would be strictly better off choosing e than buying from Firm 0. Therefore, we must have that $\phi_0((\bar{\alpha}_{\mu}, 1]) = 0$ for otherwise a profitable deviation could have been constructed for Firm 0. \square

Finally, we present a proof of the third part of Theorem 2, which considers the case of nonvanishing margins. In such a case, whenever the prior is sufficiently biased in favor of one firm, there is a unique equilibrium in which the a priori unfavorable firm is deterred.

Proposition 4. *If $g_0(\alpha) > 0$, then $\exists \bar{\mu} \in (0, 1)$ such that any SPE of $\Gamma(\mu)$ is a deterrence equilibrium for all $\mu > \bar{\mu}$. Symmetrically, if $g_1(\bar{\alpha}) > 0$, then $\exists \underline{\mu} \in (0, 1)$ such that any SPE of $\Gamma(\mu)$ is a deterrence equilibrium for all $\mu < \underline{\mu}$.*

Proof. We prove the first part of the proposition. The proof of the second part follows from symmetric considerations.

Assume by way of contradiction that there exists a sequence of priors $\{\mu_k\}$ such that $\lim_{k \rightarrow \infty} \mu_k = 1$ and a corresponding sequence of SPEs, $\{(\phi_0^k, \phi_1^k)\}_{k=1}^{\infty}$, such that $\phi^k = (\phi_0^k, \phi_1^k)$ is a non-deterrence equilibrium of $\Gamma(\mu_k)$ for all values of k .

Note that it must be the case that for almost every realized price τ_0 (with respect to ϕ_0^k) of Firm 0,

$$\Pi_0(\tau_0, \phi_1^k) = \Pi_0(\phi_0^k, \phi_1^k)$$

(otherwise Firm 0 would have a profitable deviation).

Let τ_0^k be the highest price in the support of ϕ_0^k . It follows from the above that

$$\Pi_0(\tau_0^k, \phi_1^k) = \Pi_0(\phi_0^k, \phi_1^k). \quad (15)$$

Since (ϕ_0^k, ϕ_1^k) is a non-deterrence equilibrium, Lemma 4 implies that $\phi_0((2\alpha_{\mu_k} - 1, 1]) > 0$ for all $k \geq 1$.

We next show that for all sufficiently large k there exists $\varepsilon > 0$ such that $\Pi_0(\tau_0^k - \varepsilon, \phi_1^k) - \Pi_0(\tau_0^k, \phi_1^k) > 0$.

We claim first that $v_{\mu_k}(\tau_0^k, \tau_1) > \underline{\alpha}$ for almost every realized τ_1 (with respect to ϕ_1^k). Assume that there exists a measurable subset $T \subset [0, 1]$ with $\phi_1^k(T) > 0$ such that $v_{\mu_k}(\tau_0^k, \tau_1) = \underline{\alpha}$ for all τ_0^k . Since $v_{\mu_k}(\tau_0, \tau_1)$ is increasing in τ_0 for every fixed τ_1 , it follows from the definition of τ_0^k that $v_{\mu_k}(\tau_0^k, \tau_1) = \underline{\alpha}$ for ϕ_0^k almost all realized prices τ_0^k of Firm 0. Therefore, we must have that the profit of Firm 1, conditional on $\tau_1 \in T$, is zero. By Lemma 4, Firm 1's expected payoff under ϕ^k is strictly positive, and hence we must have a profitable deviation for Firm 1.

Using equation (7) we can write

$$\begin{aligned} & \frac{\partial \Pi_0(\tau_0, \phi_1^k)}{\partial \tau_0} \Big|_{\tau_0 = \tau_0^k} = \\ & \int (\mu_k(1 - G_0(v_{\mu_k}(\tau_0, \tau_1)) + (1 - \mu_k)(1 - G_1(v_{\mu_k}(\tau_0, \tau_1)))) d\phi_1^k(\tau_1) \\ & - \tau_0^k \left(\frac{\partial v_{\mu_k}(\tau_0, \tau_1)}{\partial \tau_0} \Big|_{\tau_0^k} \right) (\mu_k g_0(v_{\mu_k}(\tau_0^k, \tau_1)) + (1 - \mu_k) g_1(v_{\mu_k}(\tau_0^k, \tau_1))) d\phi_1^k(\tau_1). \end{aligned} \quad (16)$$

Since $\lim_{k \rightarrow \infty} \mu_k = 1$, it follows from Lemma 7 that

$$\lim_{k \rightarrow \infty} \mathbb{P}_{\mu_k, \phi_1^k}(v_{\mu_k}(\tau_k, \tau_1^k) - \underline{\alpha} > \delta) = 0,$$

for any $\delta > 0$.

Since the information structure (F_0, F_1, S) does not exhibit the vanishing margins property, it follows that $g_1(\underline{\alpha}) > 0$ and, by Lemma 5, that $g_0(\underline{\alpha}) > 0$. Therefore, for some $\beta > 0$,

$$\lim_{k \rightarrow \infty} \mathbb{P}_{\mu_k, \phi_1^k}(\mu_k g_0(v_{\mu_k}(\tau_0^k, \tau_1)) + (1 - \mu_k) g_1(v_{\mu_k}(\tau_0^k, \tau_1)) > \beta) = 1.$$

We further note that $\phi_1^k([0, 1 - \alpha_{\mu_k}]) = 1$ by Lemma 8.

Since $\lim_{k \rightarrow \infty} \mu_k = 1$ we have that $\lim_{k \rightarrow \infty} \phi_1^k = 0$. Moreover, Corollary 5 implies that $\lim_{k \rightarrow \infty} \tau_0^k = \lim_{k \rightarrow \infty} \phi_0^k = 1$. Therefore, $\lim_{k \rightarrow \infty} (\tau_0^k - \tau_1^k)^2 = 1$. Hence equation (9) and equation (10) of Lemma 2 imply that

$$\lim_{k \rightarrow \infty} \left(\frac{\partial v_{\mu_k}(\tau_0, \tau_1^k)}{\partial \tau_0} \Big|_{\tau_0^k} \right) = \infty, \quad (17)$$

for any choice of τ_1^k in the support of ϕ_1^k . Therefore, equation (16) implies that $\frac{\partial \Pi_0(\tau_0, \tau_1)}{\partial \tau_0} \Big|_{\tau_0 = \tau_0^k} < 0$ for all sufficiently large values of k . Hence, in particular, for all sufficiently large values of k there exists a sufficiently small $\varepsilon > 0$ such that

$$\Pi_0(\tau_0^k - \varepsilon, \phi_1^k) - \Pi_0(\tau_0^k, \phi_1^k) > 0.$$

Therefore equation (15) implies that for all sufficiently large k , Firm 0 has a profitable deviation from ϕ_0^k . This stands in contradiction to the assumption that (ϕ_0^k, ϕ_1^k) is an equilibrium strategy. \square

Theorem 2 consolidates Corollary 4, Proposition 3, and Proposition 4.

C Proofs for the Farsighted Firms

We state a lemma that will prove useful for obtaining the results for farsighted firms.

Lemma 9. $2\alpha\mu - 1$ is a strictly convex and strictly increasing function of μ on $[0, 1]$ with a derivative that is bounded by $\frac{2(1-\alpha)}{(1-2\alpha)^2}$.

Proof. Let

$$h(\mu) = 2\alpha\mu - 1 = 2 \frac{\mu\alpha}{\mu\alpha + (1-\mu)(1-\alpha)} - 1. \quad (18)$$

The first derivative of $h(\mu)$ is $h'(\mu) = \frac{2\alpha(1-\alpha)}{[\mu(2\alpha-1) + (1-\mu)]^2}$, which is positive and bounded by $\frac{2(1-\alpha)}{(1-2\alpha)^2}$. This establishes that $2\alpha\mu - 1$ is strictly increasing with a bounded derivative.

The second derivative of $h(\mu)$ is

$$\frac{d^2h(\mu)}{d\mu^2} = \frac{4(1-\alpha)\alpha(2\alpha-1)}{(\mu(1-2\alpha) - (1-\alpha))^3}. \quad (19)$$

Recall that $\alpha < 0.5$ and so the numerator in equation (19) is negative. In addition, as $\mu \leq 1$ we conclude that $1 - \alpha > 1 - 2\alpha \geq \mu(1 - 2\alpha)$ and so the denominator of (19) is also negative. Thus, $\frac{d^2h(\mu)}{d\mu^2} > 0$ and so $h(\mu)$ must be strictly convex. \square

C.1 Proof of Theorem 3

Theorem 3 Consider a bounded information structure (F_0, F_1, S) that exhibits the vanishing margins property. Asymptotic learning holds for any discount factor $\delta > 0$ in every pure Markovian equilibrium. If, in addition, $\alpha < \frac{1}{3}$ and $\bar{\alpha} > \frac{2}{3}$, then asymptotic learning holds for any discount factor $\delta > 0$ in every Markovian equilibrium.

We recall that for every Markovian equilibrium $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ there exists functions $V_i : [0, 1] \rightarrow \mathbb{R}$ for $i = 0, 1$ such that the continuation payoff of firm i $\Pi_i^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma} | h_t) = V_i(\mu_t)$ is a function of the public belief μ_t only.

We call a Markovian equilibrium *dominant* if

$\lim_{t \rightarrow \infty} \max_{i \in \{0, 1\}} \mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\bar{\sigma}(h_t)(\tau_0, \tau_1) = i | h_t) = 1$ with probability one. Thus, in a dominant equilibrium, when time goes to infinity, the conditional probability that there exists a unique firm i that dominates the market approaches one. In contrast with the myopic case where all equilibria are dominant, in the general case an equilibrium may be non-dominant. However, one can easily show that if the discount factor δ is sufficiently small, then all equilibria are dominant. In general, we have

Lemma 10. Consider a bounded information structure (F_0, F_1, S) and let $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ be a Markovian equilibrium. If $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is a pure equilibrium or if $\alpha < \frac{1}{3}$ and $\bar{\alpha} > \frac{2}{3}$, then $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is a dominant equilibrium.

It follows from Lemma 10 that in order to prove Theorem 3 it is sufficient to show that if the vanishing margins condition holds, then asymptotic learning holds in any dominant equilibrium. We first show the following claim.

Claim 5. *If (F_0, F_1, S) is a bounded information structure that exhibits the vanishing margins property and $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is a Markovian dominant equilibrium, then $\min_{i=0,1} \liminf_{t \rightarrow \infty} V_i(\mu_t) = 0$ holds with probability one.*

Proof of the claim. Assume by way contradiction that the claim doesn't hold for some Markovian dominant equilibrium $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$. Therefore, there exists a positive-measure subset of histories for which $\liminf_{t \rightarrow \infty} V_i(\mu_t) \geq c > 0$ for both firms. This implies that with positive probability $\mu_\infty \in [\eta, 1 - \eta]$ for some constant $0 < \eta < \frac{1}{2}$. To see this, note that otherwise we would have that $\lim_{t \rightarrow \infty} \mu_t \in \{0, 1\}$ with probability one, which stands in contradiction to $\liminf_{t \rightarrow \infty} V_i(\mu_t) \geq c > 0$ for both firms. Let $\tilde{H} \subseteq H_\infty$ be the set of histories for which $\lim_{t \rightarrow \infty} \mu_t = \mu_\infty \in [\eta, 1 - \eta]$ and $\liminf_{t \rightarrow \infty} V_i(\mu_t) \geq c > 0$ for both Firms i . By our assumption it holds that $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\tilde{H}) = r > 0$.

Recall that one can write the continuation payoff of firm i given a history h_t as follows:

$$V_i(\mu_t) = (1 - \delta)\Pi_i(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t) + \delta E_{\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}}^-(V_i(\mu_{t+1})|h_t). \quad (20)$$

We claim that with positive probability it holds that

$$E_{\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}}^-(V_0(\mu_{t+1})|h_t) + \theta < \Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t) \quad (21)$$

for some $\theta > 0$ and infinitely many times t . That is, with positive probability the stage-game payoff at time t is larger by θ than the expected continuation payoff at time $t + 1$. To see this, note that equation (20) implies that if $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\bar{\sigma}(h_t)(\tau_0, \tau_1) = j|h_t) \geq 1 - \epsilon$ for some firm j , then for the other firm i

$$V_i(\mu_t) \leq \epsilon + \delta E_{\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}}^-(V_0(\mu_{t+1})|h_t). \quad (22)$$

Therefore if the condition in (21) does not hold with positive probability we have that for every $h \in \tilde{H}$ and $\epsilon > 0$ it holds for all sufficiently large t that $E_{\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}}^-(V_0(\mu_{t+1})|h_t) \geq \Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t) - \epsilon$. In addition, equation (22) implies that $\frac{1}{\delta}(V_0(\mu_t) - \epsilon) \leq E_{\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}}^-(V_0(\mu_{t+1})|h_t)$ infinitely often with positive probability. Since for every $h \in \tilde{H}$ we have that $V_0(\mu_t) \geq c/2$ from some time period onward, we must have that $V_0(\mu_t) > 1$ with positive probability. This stands in contradiction to the fact that $V_0(\mu) \leq 1$ for every $\mu \in [0, 1]$.

Note next that if $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\bar{\sigma}(h_t)(\tau_0, \tau_1) = i|h_t) \geq 1 - \epsilon$, then $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(v_{\mu_t}(\tau_0, \tau_1) \leq \alpha + \eta(\epsilon)|h_t) \geq 1 - \eta(\epsilon)$ for some $\eta(\epsilon)$ that goes to zero when ϵ goes to zero.

Let $\beta > 0$. Consider a deviation of Firm 0 at time t that is obtained by reducing every realized price above β in the support of $\bar{\phi}_0(h_t)$ by β . Such a deviation $\bar{\phi}'_0$ when applied at time t guarantees a stage-game payoff of $\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t) - \beta$. In addition, for every history h_t when (21) holds we have $\mathbf{P}_{(\bar{\phi}'_0, \bar{\phi}_1, \bar{\sigma})}(v_{\mu_t}(\tau_0, \tau_1) = \alpha|h_t)$ with a probability that approaches one as t goes to infinity. Therefore, such a deviation guarantees a stage-game payoff that is arbitrarily close to $\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t)$ and guarantees that $\mu_{t+1} = \mu_t$ with a probability that approaches one as t goes to infinity.

Applying the deviation repeatedly in all subsequent time periods implies that Firm 0 can guarantee a continuation payoff that is arbitrarily close to $\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t)$ as time goes to infinity. Equation (21) implies that for sufficiently large t it holds

with positive probability that there exists a history h_t at which Firm 0 has a profitable deviation. This stands in contradiction to the fact that $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is an equilibrium. Therefore we have that for almost every infinite history $h \in H'$ it holds that the liminf continuation payoff of Firm 1 $\liminf_{t \rightarrow \infty} V_1(\mu_t) = 0$. This concludes the proof of the claim. \square

To conclude the proof of Theorem 3 we prove the following claim.

Claim 6. *If (F_0, F_1, S) is a bounded information structure that exhibits the vanishing margins and $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is a Markovian equilibrium for which $\min_{i=0,1} \liminf_{t \rightarrow \infty} V_i(\mu_t) = 0$ holds with probability one, then asymptotic learning holds.*

Proof. Assume by way of contradiction that the claim does not hold. Then there exists a positive measure subset of histories for which $\lim_{t \rightarrow \infty} \mu_t \notin \{0, 1\}$. Assume without loss of generality that the subset of histories $h \in H$ for which $\lim_{t \rightarrow \infty} \mu_t \in [\frac{1}{2}, 1 - \eta]$ has a positive probability for some $\eta > 0$. We denote this subset by H' . We can now use similar arguments to those invoked in the proof of Proposition 2 to conclude that for almost every history $h \in H'$ it holds that $\liminf_{t \rightarrow \infty} |\Pi_0^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma} | h_t) - 2\alpha_{\mu_t} - 1| = 0$.

For every $\epsilon > 0$ let us denote by H_ϵ the set of all *finite* histories h_t for which

$$|\Pi_0^\delta(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma} | h_t) - 2\alpha_{\mu_t} - 1| \leq \epsilon.$$

It follows from the above that there exists $r > 0$ such that for every $\epsilon > 0$,

$$\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\exists t \text{ s.t. } h_t \in H_\epsilon \text{ and } \frac{1}{2} \leq \mu_t \leq 1 - \frac{\eta}{2}) \geq r.$$

As we assume vanishing margins, we can invoke Proposition 3 and conclude that there exists sufficiently small $\epsilon_0 > 0$ such that whenever $\mu \leq 1 - \frac{\eta}{2}$, Firm 0 has some price $\tau_0 \in [0, 1]$ that guarantees the following stage payoff:

$$\Pi_0(\tau_0, \phi_1, \mu) > 2\alpha_\mu - 1 + \epsilon_0, \quad (23)$$

for every strategy $\phi_1 \in \Delta([0, 1])$ of Firm 1.

We define the strategy $\hat{\phi}_0$ for Firm 0 as follows. Let h be some finite history. If $h \in H_{\frac{\epsilon_0}{2}}$ then set $\hat{\phi}_0(h_t) = \tau_0$, where τ_0 is the price that satisfies the inequality in equation (23). If $h \notin H_{\frac{\epsilon_0}{2}}$ but has some prefix $h_t \in H_{\frac{\epsilon_0}{2}}$ then set the price at $2\alpha_{\mu_{t+1}} - 1$ (where μ_{t+1} is the public belief at stage $t + 1$). Note that this implies that from stage $t + 1$ onward Firm 1 is deterred and the public belief remains fixed. Finally, whenever no prefix of h is in $H_{\frac{\epsilon_0}{2}}$ let the price be that which was chosen according to the original strategy $\bar{\phi}_0$.

The continuation payoff of Firm 0 for any finite history in $H_{\frac{\epsilon_0}{2}}$ is

$$(1 - \delta)[2\alpha_{\mu_t} - 1 + \epsilon_0] + \delta E_{(\hat{\phi}_0, \bar{\phi}_1, \bar{\sigma})}[2\alpha_{\mu_{t+1}} - 1 | h_t].$$

That is, the current period deviation of Firm 0 yields, by equation (23), an expected payoff of at least $2\alpha_{\mu_t} - 1 + \epsilon_0$. Thereafter the value of μ_{t+1} is realized and in all subsequent periods $t' > t$ Firm 0 receives a constant payoff of $2\alpha_{\mu_{t+1}} - 1$. As the

function $2\alpha_\mu - 1$ is convex (by Lemma 9) this continuation payoff is guaranteed to satisfy the following inequality:

$$E_{(\bar{\phi}_0, \bar{\phi}_1), \bar{\sigma}}[2\alpha_{\mu_{t+1}} - 1|h_t] \geq 2\alpha_{\mu_t} - 1.$$

Comparing this with the continuation payoff from the original strategy implies that the deviation yields a profit that is at least $(1 - \delta)\frac{\alpha_0}{2}$. This stands in contradiction to the fact that $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is a Bayesian Nash equilibrium. \square

Before proving Lemma 10 we prove the following auxiliary lemma.

Lemma 11. *Let $\mu \geq \frac{1}{2}$, $(\phi_0, \phi_1) \in \Delta([0, 1]) \times \Delta([0, 1])$ a pair of mixed strategies for the firms, and a strategy σ of consumer in $\Gamma(\mu)$. Assume that for every pair (τ_0, τ_1) of prices in the support of (ϕ_0, ϕ_1) it holds that $v_\mu(\tau_0, \tau_1) \in \{\alpha, \bar{\alpha}\}$. If $\Pi_0(\phi_0, \phi_1, \sigma) > 0$, then there exists a price τ_0^H such that $\Pi_0(\phi_0, \phi_1, \mu) \geq \Pi_0(\tau_0^H, \phi_1, \mu)$ and $v_\mu(\tau_0, \tau_1) = \alpha$ for every price τ_1 in the support of ϕ_1 . Moreover, if $\mathbf{P}_{(\phi_0, \phi_1, \sigma)}(\sigma(s, (\tau_0, \tau_1)) = 1) = r$, then $\Pi_0(\tau_0^H, \phi_1, \mu) \geq \frac{1}{1-r}\Pi_0(\phi_0, \phi_1, \mu)$.*

A symmetric statement holds for Firm 1 and $\mu \leq \frac{1}{2}$.

Proof. Let τ_0^H be the supremum across all prices in the support of ϕ_0 such that $v_\mu(\tau_0, \tau_1) = \alpha$ holds with positive probability. We claim first that $v_\mu(\tau_0^H, \tau_1) = \alpha$ holds for any price in the support of ϕ_1 . To see this, note first that $\tau_0^H \leq \alpha_\mu$. This follows since for $\tau_0 > \alpha_\mu$ it holds that $v_\mu(\tau_0, \tau_1) > \alpha$ for any price τ_1 . This inequality follows from the fact that for $\tau_0 > \alpha_\mu$ some consumers have a negative expected profit from buying Firm 0's product. We next contend that

$$\bar{\alpha}_\mu - \tau_0^H > 1 - \bar{\alpha}_\mu. \quad (24)$$

Inequality (24) follows since for $\alpha = \frac{1}{3}$, $\bar{\alpha} = \frac{2}{3}$, if we let $\tau_0^H = \alpha_\mu$, then (24) becomes $2(\frac{2\mu}{2\mu+(1-\mu)}) - (\frac{\mu}{\mu+2(1-\mu)}) - 1 > 0$. One can easily show that the function $2(\frac{2\mu}{2\mu+(1-\mu)}) - (\frac{\mu}{\mu+2(1-\mu)}) - 1$ is strictly positive on $(\frac{1}{2}, 1)$. Therefore, when $\alpha < \frac{1}{3}$, $\bar{\alpha} > \frac{2}{3}$, and $\tau_0^H \leq \alpha_\mu$ the inequality is strict for $\frac{1}{2} \leq \mu < 1$. Inequality (24) implies that if the price of Firm 0 is $\tau_0^H \leq \alpha_\mu$, then even if Firm 1 gives away its product for free, some consumers (those with a posterior belief that is close to $\bar{\alpha}_\mu$) will choose to buy from Firm 0. This together with the fact that $v_\mu(\tau_0, \tau_1) \in \{\alpha, \bar{\alpha}\}$ implies that $v_\mu(\tau_0^H, \tau_1) = \alpha$ for any price τ_1 in the support of Firm 1, as desired.

Therefore, for any price $\tau_0 > \tau_0^H$ in the support of ϕ_0 it holds that $v_\mu(\tau_0, \tau_1) = \bar{\alpha}$. Hence $\Pi_0(\tau_0^H, \phi_1, \sigma) = \tau_0^H \geq \Pi_0(\phi_0, \phi_1, \sigma)$.

Finally, note that

$$\Pi_0(\phi_0, \phi_1, \sigma) \leq 0 \times \phi_0(\tau_0 > \tau_0^H) + \phi_0(\tau_0 \leq \tau_0^H)\tau_0^H.$$

The last assertion follows since

$$\phi_0(\tau_0 \leq \tau_0^H) = \mathbf{P}_{(\phi_0, \phi_1, \sigma)}(\sigma(s, (\tau_0, \tau_1)) = 0) = 1 - r. \quad \square$$

We next turn to the proof of Lemma 10.

Proof of Lemma 10. Let $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ be a Markovian equilibrium and assume to the contrary that it is not a dominant equilibrium. As mentioned every equilibrium for which asymptotic learning holds is also a dominant equilibrium. It therefore

follows from Claim 6 that $\liminf_{t \rightarrow \infty} V_i(\mu_t) \geq c$ holds with positive probability for some $c > 0$ and $i = 1, 2$. This implies that the following event: $\liminf_{t \rightarrow \infty} V_i(\mu_t) \geq c$ for some $c > 0$ and $i = 1, 2$ and $\lim_{t \rightarrow \infty} \mu_t = \mu_\infty \in [\eta, 1 - \eta]$ for some $\eta > 0$, holds with positive probability. Let $\tilde{H} \subseteq H_\infty$ be a subset of histories for which $\liminf_{t \rightarrow \infty} V_i(\mu_t) \geq c > 0$ for $i = 1, 2$, and $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\tilde{H}) > 0$. Assume further, without loss of generality, that $\mu_\infty \in [\frac{1}{2}, 1 - \eta]$ for some $\eta > 0$ for any history $h \in \tilde{H}$.

We note that for every $\epsilon > 0$ and a history $h \in H$ there exists a time t' such that for $t \geq t'$ it holds that either $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(v_{\mu_t}(\tau_0, \tau_1) \leq \underline{\alpha} + \epsilon|h_t) \geq 1 - \epsilon$ or $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(v_{\mu_t}(\tau_0, \tau_1) \geq \bar{\alpha} - \epsilon|h_t) \geq 1 - \epsilon$. That is, with probability that approaches one the realized pair of prices (τ_0, τ_1) has the property that $v_{\mu_t}(\tau_0, \tau_1)$ approaches the boundaries of the signal's posterior distribution. To see this note that, as in the proof of Theorem 1, the subset of histories $h \in H_\infty$ for which $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(v_{\mu_t}(\tau_0, \tau_1) \in [\bar{\alpha} + \epsilon, \underline{\alpha} - \epsilon]|h_t) > \epsilon$ holds infinitely often for some $\epsilon > 0$ must lead consumers and firms to learn the identity of the superior firm and hence $\mu_\infty \in \{0, 1\}$. Therefore, by slightly reducing the price of firm 0 that is identified in Lemma 11 we have that for every $\theta > 0$ if $\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t) > \theta$, then one can find a price τ_0^t such that both $\Pi_0(\tau_0^t, \bar{\phi}_1, \bar{\sigma}|h_t)$ approaches $\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t)$ and $v_{\mu_t}(\tau_0, \tau_1) = \underline{\alpha}$ holds with probability that approaches one as t goes to infinity.

We can now consider two cases. If inequality (21) holds for some $\theta > 0$ and infinitely many times t for a positive-measure subset of histories $H' \subseteq \tilde{H}$, then we can use a similar consideration to the one applied in Claim 5 to deduce that Firm 0 has a profitable deviation. Otherwise we must have that $\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t)$ approaches $E_{\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}}(V_0(\mu_{t+1})|h_t)$ as t goes to infinity for almost every history $h \in \tilde{H}$. Since in addition it holds that $\liminf_{t \rightarrow \infty} V_1(\mu_t) \geq c$, we must have that for some $r > 0$ and every $\epsilon > 0$ it holds with probability one that $\mathbf{P}_{(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})}(\bar{\sigma}(h_t)(s, (\tau_0, \tau_1)) = 1|h_t) \geq r$. We can now again use Lemma 11 to deduce that as t goes to infinity Firm 0 can deviate and guarantees a continuation payoff that is arbitrarily close to $\frac{1}{1-r}\Pi_0(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma}|h_t)$. This stands in contradiction to the fact that $(\bar{\phi}_0, \bar{\phi}_1, \bar{\sigma})$ is a Markovian equilibrium. \square

C.2 Proof of Theorem 4

Theorem 4 Consider a bounded information structure (F_0, F_1, S) such that either $g_0(\underline{\alpha}) > 0$ or $g_1(\bar{\alpha}) > 0$. Asymptotic learning fails for any discount factor $\delta > 0$ in every pure Markovian equilibrium.

We now turn to the proof of Theorem 4. We prove the theorem for the case where $g_0(\underline{\alpha}) > 0$ (the case where $g_1(\bar{\alpha}) > 0$ is shown symmetrically). Fix $\delta > 0$. To prove the theorem, assume to the contrary that there exists a pure Markovian equilibrium $(\bar{\phi}_0, \bar{\phi}_1)$ for which asymptotic learning holds. That is, $\bar{\phi}_i : [0, 1] \rightarrow [0, 1]$ is the pure equilibrium strategy of firm i that chooses a price as a function of the public belief. Let $V : [0, 1] \rightarrow [0, 1]$ be the function representing the continuation payoff of Firm 0 as a function of the public belief²⁰ μ .

²⁰Since we analyze the game from the perspective of Firm 0, we suppress the subscript 0.

Given a pair of prices (τ_0, τ_1) and a prior μ , let $\varphi_a(\mu, (\tau_0, \tau_1))$ be the probability that the consumer chooses action $a \in \{0, 1, e\}$. Thus, $\varphi_\omega(\mu, (\tau_0, \tau_1))$ represents the probability that the consumer chooses to buy from Firm ω where $\omega = 0, 1$, and $\varphi_e(\mu, (\tau_0, \tau_1))$ represents the probability that the consumer chooses the outside option e . Let $w_\mu(\tau_0, \tau_1)$ be the supremum over $\alpha \in [\underline{\alpha}, \bar{\alpha}]$ such that an agent with a posterior α_μ will choose Firm 1. We can write

$$\begin{aligned}\varphi_0(\mu, (\tau_0, \tau_1)) &= \mu(1 - G_0(v_\mu(\tau_0, \tau_1))) + (1 - \mu)(1 - G_1(v_\mu(\tau_0, \tau_1))), \\ \varphi_1(\mu, (\tau_0, \tau_1)) &= \mu G_0(w_\mu(\tau_0, \tau_1)) + (1 - \mu)G_1(w_\mu(\tau_0, \tau_1)), \\ \varphi_e(\mu, (\tau_0, \tau_1)) &= \mu(G_0(v_\mu(\tau_0, \tau_1)) - G_0(w_\mu(\tau_0, \tau_1))) + \\ &\quad (1 - \mu)(G_1(v_\mu(\tau_0, \tau_1)) - G_1(w_\mu(\tau_0, \tau_1))).\end{aligned}$$

Furthermore, for $a \in \{0, 1, e\}$, let $\mu_a(\mu, (\tau_0, \tau_1))$ be the posterior probability of $\omega = 0$ conditional on action a of the consumer. This represents the public belief in the next period as a function of the consumer's choice. By Bayes' law we have

$$\begin{aligned}\mu_0(\mu, (\tau_0, \tau_1)) &= \frac{\mu[1 - G_0(v_\mu(\tau_0, \tau_1))]}{\varphi_0(\mu, (\tau_0, \tau_1))}, \quad \mu_1(\mu, (\tau_0, \tau_1)) = \frac{\mu G_0(w_\mu(\tau_0, \tau_1))}{\varphi_1(\mu, (\tau_0, \tau_1))} \\ \text{and } \mu_e(\mu, (\tau_0, \tau_1)) &= \frac{\mu[G_0(v_\mu(\tau_0, \tau_1)) - G_0(w_\mu(\tau_0, \tau_1))]}{\varphi_e(\mu, (\tau_0, \tau_1))}.\end{aligned}$$

Note that when the public belief is $\mu_t = \mu$, the stage t payoff to Firm 0 is $\varphi_0(\mu, (\tau_0, \tau_1))\tau_0$. The continuation payoff in the next stage is $V(\mu_0(\mu, (\tau_0, \tau_1)))$ with probability $\varphi_0(\mu, (\tau_0, \tau_1))$, it is $V(\mu_1(\mu, (\tau_0, \tau_1)))$ with probability $\varphi_1(\mu, (\tau_0, \tau_1))$, and it is $V(\mu_e(\mu, (\tau_0, \tau_1)))$ with probability $\varphi_e(\mu, (\tau_0, \tau_1))$. Overall, we can write Firm 0's expected continuation payoff $\Pi_0^\delta(\tau_0, \tau_1)$ as a function of the pair of prices (τ_0, τ_1) and the prior μ as follows:

$$\begin{aligned}\Pi_0^\delta(\tau_0, \tau_1 | h_t) &= (1 - \delta)\varphi_0(\mu, (\tau_0, \tau_1))\tau_0 + \\ &\quad \delta[\varphi_0(\mu, (\tau_0, \tau_1))V(\mu_0(\mu, (\tau_0, \tau_1))) + \varphi_1(\mu, (\tau_0, \tau_1))V(\mu_1(\mu, (\tau_0, \tau_1)))] + \\ &\quad \delta\varphi_e(\mu, (\tau_0, \tau_1))V(\mu_e(\mu, (\tau_0, \tau_1))).\end{aligned}\tag{25}$$

Note that for $(\tau_0, \tau_1) = (\bar{\varphi}_0(\mu), \bar{\varphi}_1(\mu))$, by the definition of V , we have that $V(\mu) = \Pi_0^\delta(\bar{\varphi}_0(\mu), \bar{\varphi}_1(\mu) | h_t)$.

Let $C > 0$ be a constant and consider an auxiliary payoff function Ψ_μ to Firm 0 that is obtained when one replaces the continuation payoff V in (25) with the function $W_\mu(\hat{\mu}) = V(\mu) + C|\mu - \hat{\mu}|$. That is,

$$\Psi_\mu(\tau_0, \tau_1) = (1 - \delta)\varphi_0(\mu_t, (\tau_0, \tau_1))\tau_0 + \delta E_{(\tau_0, \tau_1)}(W_\mu(\mu_a)).$$

In words, the next-stage continuation payoff to Firm 0 is $W_\mu(\mu_a)$ instead of $V(\mu_a)$ for any realized action $a \in \{0, 1, e\}$.

Let f be a function of τ_0 and possibly other variables. We henceforth use the notation f' to denote its right partial derivative $\frac{\partial f}{\partial \tau_0} = \lim_{\tau \rightarrow \tau_0^+} \frac{f(\tau_0) - f(\tau)}{\tau_0 - \tau}$ with respect to τ_0 . We next show the following lemma.

Lemma 12. *There exists $\beta > 0$, a prior $\hat{\mu} < 1$, and a function $K(\mu, (\tau_0, \tau_1))$ that satisfies the following two conditions: first, $K(\mu, (\tau_0, \tau_1)) \leq -\beta$ for every $\mu > \hat{\mu}$ and any pair of prices (τ_0, τ_1) , and second,*

$$\Psi'_\mu(\tau_0, \tau_1) = (1 - \delta)\varphi_0(\mu, (\tau_0, \tau_1)) + \delta K(\mu, (\tau_0, \tau_1))v'_\mu(\tau_0, \tau_1).\tag{26}$$

Proof of Lemma 12. The continuation payoff in equation (25) comprises three expressions. Differentiating the first expression with respect to τ_0 gives

$$(1 - \delta)[\varphi_0(\mu, (\tau_0, \tau_1)) + \tau_0 \varphi_0'(\mu, (\tau_0, \tau_1))],$$

where

$$\varphi_0'(\mu, (\tau_0, \tau_1)) = (-\mu g_0(v_\mu(\tau_0, \tau_1)) - (1 - \mu)g_1(v_\mu(\tau_0, \tau_1)))v_\mu'(\tau_0, \tau_1). \quad (27)$$

Since $g_0(\underline{\alpha}) > 0$ and $v_\mu(\tau_0, \tau_1)$ approaches $\underline{\alpha}$ as μ goes to one it holds that $-\mu g_0(v_\mu(\tau_0, \tau_1)) - (1 - \mu)g_1(v_\mu(\tau_0, \tau_1))$ approaches -2β for some $\beta > 0$.

In order to complete the proof of the lemma it is sufficient to show that the derivative of the last two expressions of (25) can be written as $v_\mu'(\tau_0, \tau_1)H(\mu, (\tau_0, \tau_1))$, for some function $H(\mu, (\tau_0, \tau_1))$ that goes to zero as μ goes to one.

We show this first for a fully covered market where the outside option e is played with zero probability. Under this assumption the last expression of (25) is zero. The derivative of $\delta\varphi_0(\mu, (\tau_0, \tau_1))W_\mu(\mu_0(\mu, (\tau_0, \tau_1)))$ is

$$\begin{aligned} & \delta[\varphi_0'(\mu, (\tau_0, \tau_1))W_\mu(\mu_0(\mu, (\tau_0, \tau_1)))] + \\ & \varphi_0(\mu, (\tau_0, \tau_1))W_\mu' \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_0(\mu, (\tau_0, \tau_1)))\mu_0'(\mu, (\tau_0, \tau_1)). \end{aligned} \quad (28)$$

Since $\varphi_e(\mu, (\tau_0, \tau_1)) = 0$ it holds that $\varphi_1(\mu, (\tau_0, \tau_1)) = 1 - \varphi_0(\mu, (\tau_0, \tau_1))$. Therefore, the derivative of

$\delta\varphi_1(\mu, (\tau_0, \tau_1))W_\mu(\mu_1(\mu, (\tau_0, \tau_1)))$ is

$$\begin{aligned} & \delta[-\varphi_0'(\mu, (\tau_0, \tau_1))W_\mu(\mu_1(\mu, (\tau_0, \tau_1)))] + \\ & (1 - \varphi_0(\mu, (\tau_0, \tau_1))) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_1(\mu, (\tau_0, \tau_1)))\mu_1'(\mu, (\tau_0, \tau_1)). \end{aligned} \quad (29)$$

We note that the derivative of the second expression in (25) equals the sum of the expressions in equations (28) and (29). Summing the first expression in (28) with the first expression in (29) gives $\delta\varphi_0'(\mu, (\tau_0, \tau_1))[W_\mu(\mu_0(\mu, (\tau_0, \tau_1))) - W_\mu(\mu_1(\mu, (\tau_0, \tau_1)))]$. Since signals are bounded, $|\mu_0(\mu, (\tau_0, \tau_1)) - \mu_1(\mu, (\tau_0, \tau_1))|$ goes to zero as μ goes to one. Hence it also holds that $W_\mu(\mu_0(\mu, (\tau_0, \tau_1))) - W_\mu(\mu_1(\mu, (\tau_0, \tau_1)))$ approaches zero. Therefore, equation (27) implies that the sum can be written as a product $M(\mu, (\tau_0, \tau_1))v_\mu'(\tau_0, \tau_1)$, where $M(\mu, (\tau_0, \tau_1))$ approaches zero with μ .

It remains to show that the sum of the second expression in (28) and the second expression in (29) can be written as $L(\mu, (\tau_0, \tau_1))v_\mu'(\tau_0, \tau_1)$ for some function $L(\mu, (\tau_0, \tau_1))$ that approaches zero with μ . This sum equals

$$\begin{aligned} & \delta\varphi_0(\mu, (\tau_0, \tau_1)) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_0(\mu, (\tau_0, \tau_1)))\mu_0'(\mu, (\tau_0, \tau_1)) + \\ & (1 - \varphi_0(\mu, (\tau_0, \tau_1))) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_1(\mu, (\tau_0, \tau_1)))\mu_1'(\mu, (\tau_0, \tau_1)). \end{aligned}$$

We show this for $(1 - \varphi_0(\mu, (\tau_0, \tau_1))) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_1(\mu, (\tau_0, \tau_1)))\mu_1'(\mu, (\tau_0, \tau_1))$. The fact that it holds also for $\varphi_0(\mu, (\tau_0, \tau_1)) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_0(\mu, (\tau_0, \tau_1)))\mu_0'(\mu, (\tau_0, \tau_1))$ follows similarly.

Note first that

$$\begin{aligned} & \mu_1'(\mu, (\tau_0, \tau_1)) = \\ & \frac{(1 - \mu)g_1(v_\mu(\tau_0, \tau_1))v_\mu'(\tau_0, \tau_1)(1 - \varphi_0(\mu, (\tau_0, \tau_1))) + \varphi_0'(\mu, (\tau_0, \tau_1))(1 - \mu)G_1(v_\mu(\tau_0, \tau_1))}{[1 - \varphi_0(\mu, (\tau_0, \tau_1))]^2} \end{aligned}$$

Using this and the fact that $\frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_1(\mu, (\tau_0, \tau_1))) = -C$ we have that $(1 - \varphi_0(\mu, (\tau_0, \tau_1))) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_1(\mu, (\tau_0, \tau_1))) \mu'_1(\mu, (\tau_0, \tau_1))$ equals

$$-C \frac{(1 - \mu)g_1(v_\mu(\tau, \tau_1))v'_\mu(\tau, \tau_1)(1 - \varphi_0(\mu, (\tau, \tau_1))) + \varphi'_0(\mu, (\tau, \tau_1))(1 - \mu)G_1(v_\mu(\tau, \tau_1))}{1 - \varphi_0(\mu, (\tau, \tau_1))}.$$

Note that the first expression is $-C(1 - \mu)g_1(v_\mu(\tau, \tau_1))v'_\mu(\tau, \tau_1)$, which satisfies the required condition. The second expression is equal to

$$-C \frac{\varphi'_0(\mu, (\tau, \tau_1))(1 - \mu)G_1(v_\mu(\tau, \tau_1))}{\mu G_0(v_\mu(\tau, \tau_1)) + (1 - \mu)G_1(v_\mu(\tau, \tau_1))}.$$

Since $v_\mu(\tau, \tau_1)$ approaches α as μ goes to one, we can use a standard first-order approximation to deduce that $\frac{G_i(v_\mu(\tau, \tau_1))}{g_i(\alpha)(v_\mu(\tau, \tau_1) - \alpha)}$ approaches one for $i = 0, 1$. This implies that the second expression approaches $-C\varphi'_0(\mu, (\tau, \tau_1)) \frac{(1 - \mu)g_1(\alpha)}{(\mu g_1(\alpha) + (1 - \mu)g_1(\alpha))}$ as μ approaches one. Therefore, since $\frac{g_1(\alpha)}{g_0(\alpha)} = \frac{\alpha}{1 - \alpha}$ it follows that $\frac{(1 - \mu)g_1(\alpha)}{(\mu g_1(\alpha) + (1 - \mu)g_1(\alpha))}$ approaches zero when μ goes to 1. Thus, when the market is fully covered, the lemma follows from equation (27).

Consider the case where the outside option e is played with positive probability. Again, in order to complete the proof of the lemma for this case it is sufficient to show that the derivative of the last two expressions of (25) can be written as $v'_\mu(\tau_0, \tau_1)L(\mu, (\tau_0, \tau_1))$, where $L(\mu, (\tau_0, \tau_1))$ is some function that goes to zero as μ goes to one.

Note that in this case the expression $\delta\varphi_1(\mu, (\tau_0, \tau_1))W_\mu(\mu_1(\mu, (\tau_0, \tau_1)))$ in equation (25) has a derivative of zero with respect to τ_0 . This follows from the fact that when the outside option is played with positive probability the term $w_\mu(\tau_0, \tau_1)$ is constant in some open neighborhood of τ_0 . Similarly, $G_\omega(w_\mu(\tau_0, \tau_1))$ also has a zero derivative for $\omega = 0, 1$. Hence it holds that $\varphi'_e(\mu, (\tau_0, \tau_1)) = -\varphi'_0(\mu, (\tau_0, \tau_1))$. Therefore, the derivative of $\delta\varphi_e(\mu, (\tau_0, \tau_1))W_\mu(\mu_e(\mu, (\tau_0, \tau_1)))$ with respect to τ_0 is

$$\delta[-\varphi'_0(\mu, (\tau_0, \tau_1))W_\mu(\mu_e(\mu, (\tau_0, \tau_1)))] + \varphi_e(\mu, (\tau_0, \tau_1)) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_0(\mu, (\tau_0, \tau_1))) \mu'_e(\mu, (\tau_0, \tau_1)). \quad (30)$$

Summing the first expression in (28) and the first expression in (30) gives $\delta\varphi'_0(\mu, (\tau_0, \tau_1))[W_\mu(\mu_0(\mu, (\tau_0, \tau_1))) - W_\mu(\mu_e(\mu, (\tau_0, \tau_1)))]$, which, as explained above, can be written as $R(\mu, (\tau_0, \tau_1))v'_\mu(\tau_0, \tau_1)$, where $R(\mu, (\tau_0, \tau_1))$ is some function that approaches zero with μ .

Again, it remains to show that

$$\delta\varphi_0(\mu, (\tau_0, \tau_1)) \frac{\partial W_\mu}{\partial \hat{\mu}}(\mu_0(\mu, (\tau_0, \tau_1))) \mu'_0(\mu, (\tau_0, \tau_1)) + \delta\varphi_e(\mu, (\tau_0, \tau_1)) W'_\mu(\mu_e(\mu, (\tau_0, \tau_1))) \mu'_e(\mu, (\tau_0, \tau_1))$$

can be written as $B(\mu, (\tau, \tau_1))v'_\mu(\tau, \tau_1)$, where $B(\mu, (\tau_0, \tau_1))$ is some function that approaches zero with μ . This is shown in a similar way as explained above for the case where the market is fully covered. \square

We next state another auxiliary lemma. By Lemma 9, the function $2\alpha_\mu - 1$ is convex, differentiable, as a function of μ , and has a derivative that is bounded by some positive constant $C > 1$.

Lemma 13. Let $(\bar{\phi}_0, \bar{\phi}_1)$ be the equilibrium under which asymptotic learning holds. For every $\hat{\mu} \in (0, 1)$ it holds with positive probability that there exists a time t such that $\mu_t = \mu \geq \hat{\mu}$ and

$$E_{\bar{\phi}(\mu_t)}(V(\mu_{t+1})) \leq E_{\bar{\phi}(\mu_t)}(W_\mu(\mu_{t+1})) = V(\mu_t) + E_{\bar{\phi}(\mu_t)}(C|\mu_{t+1} - \mu_t|). \quad (31)$$

Proof. For a public belief $\mu = \mu_t$, Firm 0 can guarantee a continuation payoff that is greater than or equal to $2\alpha_\mu - 1$ by repeatedly choosing the price $\tau_0 = \max\{2\alpha_\mu - 1, 0\}$. Therefore, $V(\mu) \geq 2\alpha_\mu - 1$ for every $\mu \in (0, 1)$. Additionally, it is easy to see that $V(\mu) < 1$ for every $\mu \in (0, 1)$.

Consider the function $f_\mu : [0, 1] \rightarrow \mathbb{R}_+$, which is defined as follows: $f_\mu(\tilde{\mu}) = 2\alpha_\mu - 1 + C|\tilde{\mu} - \mu|$. The function f_μ satisfies $f_\mu(\mu) = 2\alpha_\mu - 1$ and $f_\mu(\tilde{\mu}) > 2\alpha_{\mu+|\mu-\tilde{\mu}|} - 1$ for every $\tilde{\mu} \neq \mu$. Therefore, $f_\mu(1) > 2\alpha_1 - 1 = 1$. Hence, for all sufficiently large $\mu < 1$ there exist unique priors μ^1, μ^2 such that $0 < \mu^1 < \mu < \mu^2 < 1$, $|\mu - \mu^1| = |\mu - \mu^2|$, and $f_\mu(\mu^j) = 1$ for $j = 1, 2$. Note that μ^1 and μ^2 are increasing in μ and both approach one as μ approaches one. This follows since $|\mu - \mu^2|$ approaches zero as μ approaches one.

Let μ^* be large enough such that the corresponding μ^{*1} has the property that

$$\alpha_{\mu^{*1}} = \frac{\alpha_{\mu^{*1}}}{\alpha_{\mu^{*1}} + (1 - \alpha)(1 - \mu^{*1})} > \hat{\mu}.$$

This is indeed possible as signals are bounded and $\alpha > 0$.

Since asymptotic learning holds, there exists with positive probability a time \hat{t} such that $\mu^* < \mu_{\hat{t}}$. Assume that from time \hat{t} to time $\hat{t} + k - 1$, inequality (31) does not hold with probability one; then it follows by induction that

$$E_{\bar{\phi}}(V(\mu_{\hat{t}+k})|\mu_{\hat{t}}) > V(\mu_{\hat{t}}) + \sum_{i=0}^{\hat{t}+k-1} E_{\bar{\phi}}\left(C|\mu_{\hat{t}+i+1} - \mu_{\hat{t}+i}| \mid \mu_{\hat{t}}\right) \geq E_{\bar{\phi}}(f_{\mu_{\hat{t}}}(\mu_{\hat{t}+k})|\mu_{\hat{t}}). \quad (32)$$

Let η be the first random time t such that $\mu_\eta \notin [\mu_{\hat{t}}^1, \mu_{\hat{t}}^2]$. Since asymptotic learning holds, η is finite with probability one. In addition, $f_{\mu_{\hat{t}}}(\mu_\eta) > 1$ and $\mu_\eta \geq \alpha_{\mu^{*1}} > \hat{\mu}$ by construction. Therefore, if inequality (31) does not hold from time \hat{t} to time η with probability one, it follows from (32) that $E_{\bar{\phi}}(V(\mu_\eta)|\mu_{\hat{t}}) > 1$. This stands in contradiction to the fact that $V(\tilde{\mu}) < 1$ for every $\tilde{\mu} \in (0, 1)$. Hence with positive probability there exists a time $t > \hat{t}$ such that $\mu_t > \hat{\mu}$ and the inequality (31) holds. \square

We next prove Theorem 4.

Proof of Theorem 4. Let $\mu = \mu_t$ and let $\tau_i = \bar{\phi}_i(\mu_t)$ for $i = 0, 1$ be the corresponding equilibrium prices. Let $\bar{\tau}_0 = \min\{2\alpha_\mu - 1 + \tau_1, \alpha_\mu\}$. Note that $\bar{\tau}_0 > 0$ for all sufficiently large μ . The price $\bar{\tau}_0$ is the maximal price for Firm 0, as a function of τ_1 , for which the consumer chooses Firm 0 with probability one. Since asymptotic learning holds, the probability that the consumer buys from Firm 0 is less than one (for otherwise learning would have stopped), and thus $\bar{\tau}_0 < \tau_0$.

We claim first that there exists $\tilde{\mu} < 1$ such that if $\tilde{\mu} \leq \mu_t = \mu$, then $\Psi_\mu(\tau_0, \tau_1) - \Psi_\mu(\bar{\tau}_0, \tau_1) < 0$. To see this note that Lemma 12 and the mean value theorem imply that

$$\Psi_\mu(\tau_0, \tau_1) - \Psi_\mu(\bar{\tau}_0, \tau_1) = \left(\varphi_0(\mu, (\tau, \tau_1)) + v'_\mu(\tau, \tau_1)K(\mu, (\tau, \tau_1))\right)(\tau_0 - \bar{\tau}_0),$$

for some $\tau \in [\bar{\tau}_0, \tau_0]$. By Lemma 2 and Lemma 4, as μ goes to one τ_1 approaches zero and τ_0 approaches one. Hence, as in the proof of Proposition 4, it follows that $v'_\mu(\tau, \tau_1)$ approaches ∞ as μ goes to one. Since $K(\mu, (\tau, \tau_1)) \leq -\beta$ for all sufficiently large μ we must have that $\Psi_\mu(\tau_0, \tau_1) - \Psi_\mu(\bar{\tau}_0, \tau_1) < 0$.

Let t be a time such that $\mu = \mu_t \geq \tilde{\mu}$ and $E_{\bar{\phi}(\mu_t)}(V(\mu_{t+1})) \leq V(\mu_t) + E_{\bar{\phi}(\mu)}(W_\mu(\mu_{t+1}))$. Such a time t exists with positive probability by Lemma 13. To derive the contradiction, note that

$$\begin{aligned} V(\mu_t) &= (1 - \delta)\varphi_0(\mu_t, (\tau_0, \tau_1))\tau_0 + \delta E_{\bar{\phi}(\mu_t)}(V(\mu_{t+1})) \leq \\ &(1 - \delta)\varphi_0(\mu_t, (\tau_0, \tau_1))\tau_0 + \delta E_{\bar{\phi}(\mu_t)}(W_\mu(\mu_{t+1})) = \Psi_\mu(\tau_0, \tau_1) < \Psi_\mu(\bar{\tau}_0, \tau_1) = \\ &(1 - \delta)\bar{\tau}_0 + \delta V(\mu_t). \end{aligned} \tag{33}$$

Note that the last equality in (33) holds since for the price $\bar{\tau}_0$ the consumer chooses Firm 0 with probability one and the public belief at time $t + 1$ is μ_t . This implies that $V(\mu_t) < (1 - \delta)\bar{\tau}_0 + \delta V(\mu_t)$ and so $V(\mu_t) < \bar{\tau}_0$. If, however, Firm 0 deviates and plays the price $\bar{\tau}_0$ from time t onwards, then, by the Markovian property, this guarantees a continuation payoff of $\bar{\tau}_0$. This yields a profitable deviation to Firm 0 as $V(\mu_t) < \bar{\tau}_0$, in contradiction to the assumption that $(\bar{\phi}_0, \bar{\phi}_1)$ is a Bayesian Nash equilibrium. \square

D Auxiliary Lemmas

Lemma 14. *The ratio $\frac{G_1(r)}{G_0(r)}$ is nonincreasing in r and $\frac{G_1(r)}{G_0(r)} > 1$ for all $r \in (\underline{\alpha}, \bar{\alpha})$. In particular, G_0 first-order stochastically dominates G_1 . Moreover, for any point $x \in [\underline{\alpha}, \bar{\alpha}] \cap (0, 1)$, it holds that $(1 - x)g_0(x) = xg_1(x)$ and $\lim_{x \rightarrow +\underline{\alpha}} \frac{G_1(x)}{G_0(x)} = \frac{\underline{\alpha}}{1 - \underline{\alpha}}$.*

Proof. The fact that $\lim_{x \rightarrow +\underline{\alpha}} \frac{G_0(x)}{G_1(x)} = \frac{\underline{\alpha}}{1 - \underline{\alpha}}$ follows from the relation $(1 - x)g_0(x) = xg_1(x)$ and the fact that $G_\omega(x) = \int_{\underline{\alpha}}^x g_\omega(x)dx$ for $\omega = 0, 1$ when $x \leq \bar{\alpha}$. The proof of the other parts follows from the more general result that appears in Lemma A1 of (Acemoglu et al., 2011). \square

Corollary 1. Let $(\bar{\sigma}, \bar{\tau}_0, \bar{\tau}_1)$ be a myopic Bayesian Nash equilibrium. If asymptotic learning holds, then conditional on state $\omega \in \Omega$,

$$\lim_{t \rightarrow \infty} \mathbf{P}_{(\bar{\sigma}, \bar{\tau}_0, \bar{\tau}_1)}(\{\sigma_t(\mu_t, s, \bar{\tau}(\mu_t)) = \omega\} | \omega) = 1.$$

Proof. Without loss of generality assume that the realized state is $\omega = 0$. Since asymptotic learning holds, we have that $\lim_t \mu_t = 1$ almost everywhere. By Lemma 7 we have that $\lim_{t \rightarrow \infty} v_{\mu_t}(\tau_0^t(\mu_t), \tau_1^t(\mu_t)) = \underline{\alpha}$. Therefore,

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbf{P}_{(\bar{\sigma}, \bar{\tau}_0, \bar{\tau}_1)}(\{\sigma_t(\mu_t, s, \bar{\tau}(\mu_t)) = 0\} | \omega = 0) &= \lim_{t \rightarrow \infty} G_0(v_{\mu_t}(\tau_0^t(\mu_t), \tau_1^t(\mu_t))) = \\ &G_0(\underline{\alpha}) = 1. \end{aligned}$$

\square

E Proof of Proposition 1

Proof. Assume that the proposition is false. Thus, we can find some $\varepsilon > 0$, $\alpha > 0$, and a sequence of information structures $\{G_0^n, G_1^n\}_{n=1}^\infty$, such that:

- (i) The support of G_0^n and G_1^n is $[\alpha, \bar{\alpha}]$ for every n .
- (ii) For every n , $g_1^n(\alpha) > 0$ and $\lim_{n \rightarrow \infty} g_1^n(\alpha) = 0$.
- (iii) The deterrence threshold, $\bar{\mu}^n$, is bounded above by $1 - \varepsilon$.

Let μ be an arbitrary prior in the interval $(1 - \varepsilon, 1)$ and let $\Gamma^n(\mu)$ be the stage game for the prior μ and the information structure (G_0^n, G_1^n) . Recall that (see (14)) for every n ,

$$\frac{\partial \Pi_0(\tau_0, 0)}{\partial \tau_0} \Big|_{\tau_0=2\alpha_\mu-1} = 1 - (2\alpha_\mu - 1) \left(\frac{\partial v_\mu(\tau_0, 0)}{\partial \tau_0} \Big|_{2\alpha_\mu-1} \right) (\mu g_0^n(\alpha) + (1 - \mu)g_1^n(\alpha)).$$

By our assumption, the only equilibrium of the game $\Gamma^n(\mu)$ is a deterrence equilibrium. Thus, for every n ,

$$1 - (2\alpha_\mu - 1) \left(\frac{\partial v_\mu(\tau_0, 0)}{\partial \tau_0} \Big|_{2\alpha_\mu-1} \right) (\mu g_0^n(\alpha) + (1 - \mu)g_1^n(\alpha)) \leq 0. \quad (34)$$

Recall that $v_\mu(\tau)$ is the indifference threshold in the game $\Gamma^n(\mu)$ with prices $\tau = \tau_0, \tau_1$. By equation (6) it is independent of the information structure's shape and is determined solely by the game's prior μ and price vector τ . In addition, G_0^n first-order stochastically dominates G_1^n (Lemma 14). Therefore $g_0^n(\alpha) \leq g_1^n(\alpha)$ and so $(\mu g_0^n(\alpha) + (1 - \mu)g_1^n(\alpha)) \leq g_1^n(\alpha)$. We can now deduce that

$$1 - (2\alpha_\mu - 1) \left(\frac{\partial v_\mu(\tau_0, 0)}{\partial \tau_0} \Big|_{2\alpha_\mu-1} \right) g_1^n(\alpha) \leq 0. \quad (35)$$

Furthermore, since $2\alpha_\mu - 1 < 1$ whenever $\mu < 1$, Lemma 2 implies that $\left(\frac{\partial v_\mu(\tau_0, 0)}{\partial \tau_0} \Big|_{2\alpha_\mu-1} \right) < \infty$. Now note that by (ii) above the limit on the left-hand side of inequality (35) is 1, a contradiction. □