Bayesian learning in markets with common value

Itai Arieli, Moran Koren, and Rann Smorodinsky

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

October 1, 2017

Abstract

Two firms produce substitute goods with 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 superior firm necessarily prevail? We adapt the classical social learning model by introducing strategic dynamic pricing. We provide necessary and sufficient conditions for learning. In contrast to previous results, learning can occur when signals are bounded. This happens when signals exhibit the newly introduced *vanishing likelihood* 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 variable. This may take the form of some future change in regulation, a technological shock, environmental developments, or prices in related upstream markets. Although this information is unknown, individual consumers may receive some private information about these fundamentals. We ask whether, in such an environment, markets aggregate information correctly and the ex-post superior product will eventually dominate the market.

Research supported by GIF research grant no. I-1419-118.4/2017, ISF grant 2018889, Technion VPR grants, the joint Microsoft-Technion e-Commerce Lab, the Bernard M. Gordon Center for Systems Engineering at the Technion, and the TASP Center at the Technion.

To present this in a concrete example, one could think about competing propulsion technologies in the automotive industry, say, electric cars versus traditional fuel-powered cars. The value of a car strongly depends, inter alia, on the future cost of using it, which is primarily driven by its propulsion costs. Which technology will allow for lower costs depends on many unknown factors such as oil prices, regulation, and more. Consumers decide which of these technologies to purchase based on the (limited) information they have about these factors, coupled with the price of each alternative. Consumers who anticipate a possible decrease in oil prices might opt for the fuel option, while others who anticipate "green" subsidies to electric cars would opt for those. Firms, anticipating the possibility of some private information becoming available to the consumers, set prices accordingly.

In this work we focus on the role of social learning in such environments. We study whether the learning process guarantees an efficient outcome. We isolate the role of learning by introducing a simple duopoly model of common value in which consumers, with a unit demand, choose between two substitute products, each with zero marginal cost of production. Conditional on the (unknown) state of nature all consumers agree on the value and identity of the superior product. Consumers arrive sequentially. Each consumer receives some private information on the quality of the products. Thus consumers have an informational advantage over firms (put differently, all the information firms receive is publicly available). In order to distil the learning effects we further assume that the aggregation of all signals fully reveals the state of nature that determines the identity of the superior firm. Our main goal is to identify conditions under which *asymptotic learning* holds, that is, conditions under which information is fully aggregated in the market asymptotically.

The timing of the interaction is as follows. Nature randomly chooses one of two states of nature, and so determines the identity of the firm with the superior product. At each stage the two firms observe the entire history of the market - past prices and consumption decisions - and simultaneously set prices. A single consumer arrives and receives a private signal regarding the state of nature. The consumer makes her consumption decision based on her signal, the pair of prices for each product, and the entire history of prices and consumption decisions. The consumer can also choose to opt out and not to buy any product.[1]

The same model, but with prices set exogenously and fixed throughout, is exactly the standard herding model. In that model, as shown by Smith and Sørensen [28], the characterization of asymptotic learning crucially de-

[1] Our model and results carry through even when past prices are not observed.

pends upon the quality of the private signals of the agents. In particular, one must distinguish between two families of signals: bounded versus unbounded. In the unbounded case the private beliefs of the agents are, with positive probability, arbitrarily close to zero and one. Therefore, no matter how many people herd on one of the alternatives, there is always a positive probability that the next agent will receive a signal that will make him break away from the herd towards the other alternative. Thus, as shown by Smith and Sørensen [28], this property entails asymptotic learning. The same logic applies in our model as well: even if the prior is extremely in favor of one product, with positive probability there will be a consumer who will get a sufficiently strong signal to tilt the consumption decision towards the a priori inferior product. Thus, under strategic pricing and unbounded signals, asymptotic learning holds.

The learning results in our model depart from those of the canonical model when signals are bounded. In the herding model there is always a positive probability that the suboptimal alternative will eventually be chosen by all agents. However, intuition suggests that when prices are endogenized they serve to prevent such a herding phenomenon. Hypothetically, once society stops learning and a herd develops on the product of one firm, the other will lower its 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 referred to here as *vanishing likelihood*.

When signals are bounded the posterior belief of any agent, given his signal, is bounded away from zero and one for any interior prior. The proportion of agents whose posterior lies within ε of the boundaries of the posterior distribution obviously shrinks to zero as ε goes to zero. We say that signals exhibit *vanishing likelihood* if the *density* of consumers at the boundaries of the posterior belief distribution goes to zero.

Consumers who receive signals that induce such extreme posterior beliefs, those close to the boundaries, are those that are likely to go against a herd and purchase the less popular product. We refer to such consumers as *non-conformists*. With this interpretation in mind the property of 'vanishing likelihood' serves as a measure of the prevalence of non-conformism. More particularly we associate vanishing likelihood with a negligible level of non-conformism while signals that do not exhibit vanishing likelihood are associated with significant non-conformism.

When society herds each agent follows in the footstep of her predecessors, and thus intuitively, one expects non-conformism (when signals do not exhibit vanishing likelihood) to induce learning. Our main result shows that the opposite occurs - in the presence of strategic pricing asymptotic learning holds if and only if signals have the vanishing likelihood property.

The intuition behind our main result is as follows. With negligible non-conformism whenever a herd forms, the popular firm expects the next consumer to conform with high probability. Therefore, the popular firm's optimal action is to exploit this and set a high price. The popular firms' high price gives the less popular firm an opportunity to compete by proposing a low price, and in the rare case the consumer has an extreme signal, sell to him. The consumer's deviation from the herd entails additional social learning. In contrast, when non-conformism is significant the leading firm optimally lowers its price in order to increase its market share, namely to sell to the non-conformists. By sufficiently Lower its price, the popular firm eventually drives the other firm out of the market and learning stops.

1.1 Schumpeterian Growth

It is widely agreed that innovation and the evolution of technology constantly propel the economy forward. New technologies replace older ones and may improve product quality, reduce production costs and often completely disrupt an industry.

However not every innovation entails improvement. Arguably such innovations will naturally be driven out of the market while only the successful ones prevail. This argument forms the basis of the *Evolutionary Economics* strand that dates back to Marx, Veblen and Schumpeter.

In his seminal work, Schumpeter [27] described the process of economic growth, which he refers to as "Capitalism", as an evolutionary process which is shaped by "gales of creative destruction". Some of his contemporaries argue that large and profitable firms are the source of innovation and so regulation protecting them was essential to R&D investments. In contrast, Schumpeter argues that incumbent firms, anticipating innovation by potential entrants, invest in R&D to stay ahead of the game. Therefore such regulation is unwarranted, and may even be detrimental. However, such profitable incumbents may also use their power to drive innovation away by lowering prices. This is true in particular when it is hard to identify which innovation forms an improvement and which does not.

Does the evolutionary process guarantee that the economy will successfully separate the wheat from the chaff? This question becomes more acute with the accelerated pace of innovation witnessed in the last 2 decades [26].

Our theoretical results shed light on this issue and relate the outcome of the evolutionary process with the market structure. Our model shows that whenever the proportion of non conformist consumers (often referred to as 'early adopters' in the context of technological revolutions) is insignificant, a phenomenon captured by the technical notion of vanishing likelihood, the evolutionary process successfully sieves the better technologies. However, whenever this proportion is significant then the evolutionary process may fail and policies to support entry may be warranted in order to sustain Schumpeterian growth.

We demonstrate the validity of our observation with the following two case studies:

A tale of vanishing likelihood: In the mid 1990s, with around 10 percent market share, Barnes and Noble (BN) was a clear leader in the US book selling market (See [14],[16]). BN rose to power by perfecting the shopping experience for buyers and through aggressive discounts, $20-30\%$, over cover prices [17].

In 1995 BN faced a new type of competition - Amazon's online retail. In hindsight this is a classical example for "creative destruction", however, back then the success of online retail was ambiguous. In the late 90's the Internet was new and was definitely not conceived as a retail shopping channel [11]. Uncertainty as to the security of on-line payments was immense ([29], [12]) and many considered instantaneous gratification, absent in online shopping, as central to the shopping experience.

To overcome these challenges, Amazon offered 40 percent discounts over cover prices, while BN kept maintained its 20-30 percent discount. Early adopters of the new technology started shopping on-line, followed by more conservative shoppers. 14 years later Amazon was a clear market leader while BN closed most of its shops. With a 16.5% market share and over $\$400B$ in trade in the 2016 US market ([25]), online retail shopping is a proven superior innovation.

We note that BN, with its 28 percent operational profitability at the time ([9]), could have lowered prices and driven Amazon out of the market. Was it rational to maintain its high profit or should have BN offered additional discounts given the information available at the time? We revisit this issue after the next case study.

A tale of non-vanishing likelihood: In the game consoles market the quality of a product is determined by a variety of measurable determinants (e.g., graphic capabilities, CPU speed, RAM) and non measurable ones (e.g., design, gameplay and game titles). Back in late 1990s' game consoles offered no connected gaming so network effects and externalities were less significant in console adoption.

The *Dreamcast* game console, developed at the end of the 1990s, was

Sega's second attempt to restore its place as an industry leader.[2] The market leader at the time was Sony with its *Playstation I* console. In contrast with BN's strategy of ignoring Amazon, Sony preempted the launch of Dreamcast by offering a 30 percent discount on its own console a month before the release of Dreamcast and one year before releasing its own next-gen product (see [30], [32]). As a result, Dreamcast did not manage to penetrate the market and Sega stopped making and selling this console less than two years after its introduction.[3]

In retrospect Dreamcast was acknowledged as the superior technology and some even consider it one of the best consoles ever developed. This reasonably priced and developer friendly technology contained many futuristic features such as "network game-play" and 64-bit high resolution graphics.[4] Some of these novel features did not reach the market until more than half a decade later.[5] In Schumpeter's framework, growth had been delayed.

Vanishing vs. non-vanishing likelihood: Why did the incumbent firm price aggressively, thus interfering with the evolutionary process, in one case and not in the other? What is the primary distinction that explains this? One explanation that resonates with our model is due to the different nature of the demand side of both markets. Buyers of game consoles are highly engaged and actively seek information regarding new releases and so we witness more diverse opinions by buyers and a substantial proportion of non-conformism. In contrast, book buyers, who do show interest and are engaged when it comes to the decision of which title to buy, are less inquisitive when it comes to the actual shopping experience. Only a marginal proportion would not conform to habits. Thus the marginal proportion of innovators, captured by our notion of vanishing likelihood, is significant in the console market and is insignificant in the book market. For Sony it was rational to reduce prices in order to maintain the non-conformist market share. For BN, on the other hand, the size of the non-conformist market share coupled with the ambiguity of the success of the new online shopping

[2]Sega's previous console, the "Saturn", had failed miserably due to its "awful gameplay and inferior design" - see [31].

[3]Indeed, Tadashi Takezaki, a former executive at Sega, points out two main factors related to the failure of Dreamcast: (1) Consumers had been skeptic about Sega's abilities to produce a viable product following the aforementioned "Saturn" fiasco; and (2) Sony's aggressive pricing ([15]).

[4]To this day, over a decade an a half after its initial release, second hand consoles are being traded on eBay and new Dreamcast titles are being released See [10], which serves as a testimony to its technical superiority.

[5]On-line multiplayer interface had been introduced by the Playstation III in 2006, the same year the controls with motion detection first appeared in the Nintendo Wii.

paradigm, was insufficient to forego profits.[6]

1.2 Related Literature

Our work primarily contributes to the models of social learning when agents act sequentially initiated by Bikhchandani, Hirshleifer, and Welch [6] and Banerjee [4]. Their primary contribution is to point out the possibility of information cascades and market failure when signals are bounded. Smith and Sørensen [28] characterize the information structure that entail such potential market failure. In these papers and a many of the follow ups prices are assumed exogenous and fixed throughout. The primary departure of our model from this strand is by introducing endogenous pricing. We associate a favorite firm (product) with each state of nature and allow for the firms to set prices dynamically, based on the information available in the market.

Avery and Zemsky [3] incorporate dynamic pricing into herding models. They consider a single firm whose product value is associated with the (unknown) state of nature. Instead of offering the product at a fixed price, as in the earlier papers, they assume the price is set dynamically to be the expected value of the product conditional on all the information available publicly. Their primary interest is to study financial markets and so they have a market maker in mind which uses all the publicly available data to set prices. In contrast, we assume the firms themselves set the prices.

Our model is reminiscent of the model introduced in Bose et al. [7, 8]. Whereas our model has a duopoly competing for the consumer at each stage, Bose et al study a monopoly who competes against some outside option. Whereas our model has a general signal distribution (possibly unbounded) the Bose at al model studies a finite signal space ([8] is even more particular as it considers a binary signal space with symmetric distribution). In both models consumers, sharing a common value but with a private signal, arrive sequently and make their consumption decision based on their predecessors decisions and past prices. These papers show that herding is inevitable, even if the monopoly adjusts prices. In fact, when the public belief is sufficiently in favor of the monopoly's product then the monopolist will price low enough to attract all consumers, invariably of the signal they receive. Consequently learning stops. This result parallels our observation regarding

[6] One could suggest that the non-conformism of the game console market is captured by the notion of unbounded signals whereas the book market exhibits bounded signals. However, this explanation is unsatisfactory. The reason is that it would yield that learning and creative destruction should have been observed in the console market and not in the book market, contrary to the evidence.

the favorite firm deterring its competitor whenever the vanishing likelihood condition is not satisfied. The binary case allows the authors to investigate the threshold for the public belief as a function of the informativeness of the signals.

Roughly speaking the two main distinctions between our work and [7, 8] is the treatment of a monopoly vs. a duopoly and a finite signals space vs. arbitrary signals. The latter is more important as it allows us to discover the conditions under which pricing eliminates herding and efficiency prevails even if signals are bounded. This condition, vanishing likelihood, is never satisfied in the earlier paper.

Moscarini and Ottaviani [22] study the duopoly case and their paper focuses on the single stage interaction with 2 firms and a single knowledgeable consumer. In fact, it is exactly the model $\Gamma(\mu)$ we study in Section 3.1, however restricted to a binary and symmetric signal space. Unsurprisingly, whenever the prior belief is above (or below) some threshold, all equilibria in their model form a *deterrence equilibrium* (see definition 5), where one firm prices out the other firm. Clearly the emergence of a deterrence equilibrium will imply that learning stops in the repeated model. The authors go on and provide comparative statics over the threshold public belief for which learning stops as a function of the informativeness of the signal (here is where they leverage the restricted signal space). As signals become more informative the thresholds move to the extremes.

Our main result about the one shot game, Theorem 2, argues that learning stops whenever the *vanishing likelihood* condition does not hold. As this condition can never hold for a finite signal space the result in [22] follows as a corollary. The main take-home message from comparing our work with [22] is that learning is not determined by the level of informativeness of the signals but rather the vanishing likelihood condition. In particular signal distributions that satisfy vanishing likelihood need not be highly informative. The restriction to a binary model, in this case, is misleading. We believe that a similar distinction is valid to the monopoly setting of Bose et al. [7, 8].

Mueller-Frank [23, 24] introduces a pair of models with dynamic pricing of a monopoly [24] and a duopoly [23]. The model is very similar to ours with the distinction that the firms have the informational advantage and know the true state of the world. Mueller-Frank does not characterize the informational conditions that entail learning as we do. Rather, he studies the connection between welfare and learning and shows that learning is not sufficient for welfare maximization. It is worth noting that in our model, in contrast, learning is necessary and sufficient for welfare maximization.

The paper is organized as follows. Section 2 presents the model and the main theorem for the case where firms are myopic. Section 3 gives the proof of the main result. Section 4 is an extension of our model to the case where firms are farsighted. Section 5 concludes.

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 to 1 and the value of the inferior product to 0. In every time period t the two firms first set (non-negative) prices $(\tau_t^0, \tau_t^1) \in [0, 1]^2$ for their product. Then consumer t receives a private signal and must decide whether to buy product 1, product 0, or neither product. Formally, the action set of every consumer is $A = \{0, 1, e\}$, where the action $a = i$ corresponds to the decision to buy from firm i and the action $a = e$ corresponds to the decision not to buy either product. The payoff of every consumer t, given the price vector (τ_0^t, τ_1^t) as a function of the realized state ω, is

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

For simplicity we assume that both firms have no marginal cost of production at every given time period t. Hence, firm i's stage payoff, given a price vector (τ_0^t, τ_1^t), can be described as a function of the consumer's decision as follows

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

We assume that the state ω is drawn at stage $t = 0$ according to a commonly known prior distribution, such that $P(\omega = 0) = \mu = 1 - P(\omega = 1)$. The state ω is unknown both to 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[7], $h_t \in H_t = ([0, 1]^2)^{t-1} \times (\{0, 1, e\})^{t-1}$, and a private signal $s_t \in S$ (where S is some abstract measurable signal space)[8]. The firms observe only the realized history $h_t \in H_t$ at every time t

[7]Alternatively, Consumers and firms see only the histoy of consumer actions and infer the corresponding equilibrium prices.

[8]One may think of an alternative formulation in which firms and consumer only see the history of actions and calculate the sequence of corresponding equilibrium prices at each stage.

and receive no private information. Conditional on the state of the world ω, the signals are independently drawn according to a probability measure F_ω. We assume throughout that F_0 and F_1 are mutually absolutely continuous with respect to each other.[9] The prior μ and the functions F_0 and F_1 are common knowledge among the consumers and the firms.

We let $\mathcal{A} \subset \{0, 1, e\}^{[0,1]^2 \times S}$ 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 into a consumption decision. A *(pure) strategy for consumer t* in is a measurable function $\sigma^t : H_t \to \mathcal{A}$ that maps every history $h_t \in H_t$ and signal $s_t \in S$ to a decision rule. A strategy for firm i is a sequence $\overline{\tau}_i = (\tau_i^t)_{t \geq 1}$ such that for every time t, $\tau_i^t : H_t \to [0, 1]$ is a measurable function. We let $H = \cup_{t \geq 1} H_t$ be the set of all finite histories and let $\overline{\sigma} = (\sigma^t)_{t \geq 1}$ be the strategy of the consumers. We note that the strategy profile $(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)$ together with the prior μ and the functions F_0 and F_1 induce a probability distribution $\mathbf{P}_{(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)}$ over $\Omega \times H \times S^\infty$.

Let $\mu_t = \mathbf{P}_{(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)}(\omega = 0 | h_t)$ be the probability that the state is 0 conditional on the realized history h_t whenever this history is well defined. We call μ_t *the public belief at time t*. The following observation regarding the sequence of public beliefs, $\{\mu_t\}_{t=1}^\infty$ is straightforward.

Observation 1. $\{\mu_t\}_{t=1}^\infty$*, is a martingale. Thus, by the martingale convergence theorem, it must converge almost surely to a limit random variable* $\mu_\infty \in [0, 1]$.

We further note that a strategy profile $(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)$, a time t, a pair of prices $(\tau_0, \tau_1) \in [0, 1]^2$, and a decision rule $\sigma \in \mathcal{A}$ define, conditional on every history $h_t \in H_t$ that is realized with positive probability, an expected payoff $\Pi_i^t(\tau_0, \tau_1, \sigma | h_t)$ for every firm i, and an expected payoff $U_t(\tau_0, \tau_1, \sigma | h_t)$ to consumer t. We can now define the notion of Bayesian Nash equilibrium for myopic firms.

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

- For every $\tau \in [0, 1]$,

$$\Pi_i^t(\tau_0^t(h_t), \tau_1^t(h_t), \sigma_t(h_t) | h_t) \geq \Pi_i^t(\tau, \tau_{-i}^t(h_t), \sigma_t(h_t) | h_t).$$

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

$$U_t(\tau_0, \tau_1, \sigma_t(h_t) | h_t) \geq U_t(\tau_0^t, \tau_1^t, \sigma | h_t).$$

[9] 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 with this assumption the probability of a fully revealing signal, for which the posterior probability is either 0 or 1, is zero.

In words, a strategy profile $(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)$ constitutes a myopic Bayesian-Nash equilibrium if for every time t and every history h_t that is realized with positive probability, $\tau_t^i(h_t)$ maximizes the conditional expected stage payoff to every firm i and $\sigma_t(h_t)$ 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; 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 price 0 at 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 she buys product 0 whenever $\mu \geq \frac{1}{2}$ and product 1 if $\mu < \frac{1}{2}$. 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 $\sigma_t(h_t)$ be optimal with respect to *every* price vector (τ_0, τ_1), and not just with respect to $(\tau_0^t(h_t), \tau_1^t(h_t))$.

2.1 Characterization of Asymptotic Learning

We now turn to analyze asymptotic information aggregation of the sequential game. As is common in the literature, we define *asymptotic learning* as follows.

Definition 2. Let $\mu \in (0, 1)$ be the initial prior and consider a myopic Bayesian Nash equilibrium σ of the corresponding game. We say that *asymptotic learning* holds for σ if the belief martingale sequence converges to a point belief assigning probability 1 to the true state.

Thus, when asymptotic learning holds, it must be the case that the consumers and the firms eventually learn the superior product. In our case, even for a very strong public belief in favor of one firm, it is not a priory clear that the strong firm will dominate the market as the weak firm can always lower its price. We show, however, in Lemma 6 that the probability of buying from the superior firm converges to one when asymptotic learning occurs. Whenever asymptotic learning doesn't hold, only one firm prevails (from some time on all consumers buy from one firm but are not 100% certain that it is the superior one). As a result, there is a positive probability that the prevailing firm is the inferior one.

The main goal of our paper is to provide a characterization of asymptotic learning under strategic pricing in terms of the signal distribution. Such a characterization is provided by Smith and Sørensen [28] for the standard herding model. We start by presenting the formal definition of *bounded and unbounded signals* due to Smith and Sørensen [28].

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 observing the signal s, when the prior is $\mu = 0.5$.

Let $G_\omega(x) = F_\omega(\{s \in S | p(s) < x\})$, $\omega = 0, 1$, be the two cumulative distribution functions of the random variable $p(s)$ induced by the two probability distributions, F_ω, $\omega = 0, 1$, over S. The support[10] of G_ω is the interval $[\bar{\alpha}, \underline{\alpha}]$, where $\underline{\alpha} = \inf_{x\in[0,1]} G_0(x)$ and $\bar{\alpha} = \sup_{x\in[0,1]} G_0(x)$.

Definition 3. The signals are called *unbounded* if $\underline{\alpha} = 0$ and $\bar{\alpha} = 1$. Signals are *bounded* if $\underline{\alpha} > 0$ and $\bar{\alpha} < 1$.

In words, signals are 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 unbounded signals and fails under bounded signals.

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

Our analysis shows that in the sequential model with strategic pricing unbounded signals yield asymptotic learning. Bounded signals, however, *do not* necessarily lead to failure of asymptotic learning. The analysis of this case hinges on the level of conformism among consumers, as we now turn to explain.

The driving force in Smith and Sørensen's [28] characterization is the *overturning principle.* The overturning principle states that for any prior either action is played with positive probability in equilibrium. This implies that independent of the strength of a herd on one action, there is always a positive probability that the subsequent agent will be a *non-conformist* and overturn the herd.

Let us refer to a consumer as a non-conformist if his private signal is strong enough to sway his decision against the current majority. The bigger this majority is, the stronger the required countervailing signal is, and so the probability of being a non-conformist decreases and converges to zero. The next definition of *vanishing likelihood* relates to the probability density of consumer receiving an extreme signal and hence it provides an indication of the prevalence of non-conformists.

[10] Recall that F_0 and F_1 are mutually absolutely continuous and so they have the same support.

Definition 4. Signals exhibit *vanishing likelihood* if $g_1(\underline{\alpha}) = g_0(\bar{\alpha}) = 0$.

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

Theorem 1. *If signals are unbounded, then asymptotic learning holds for any prior $\mu \in (0,1)$ in any myopic Bayesian Nash equilibrium. If signals are bounded, and exhibit vanishing likelihood, then asymptotic learning holds for any prior $\mu \in (0,1)$ in any myopic Bayesian Nash equilibrium. If signals are bounded and do not exhibit vanishing likelihood, then asymptotic learning fails for any prior $\mu \in (0,1)$ in any myopic Bayesian Nash equilibrium.*

3 Proof of the Main Result

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 (0 is realized with probability μ and state 1 with probability $1-\mu$), then the two firms post a price simultaneously, and then the single consumer decides, based on his private signal and the vector of prices, whether or not to buy from any of the firms. The following observation is a direct implication of Definition 1.

Observation 2. *A strategy profile $(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)$ constitutes a myopic Bayesian Nash equilibrium if and only if for every time t and every history $h_t \in H_t$, that is realized with positive probability with respect to $\mathbf{P}_{(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)}$, the tuple $(\sigma^t(h_t), \tau_0^t(h_t), \tau_1^t(h_t))$ is a subgame perfect equilibrium (SPE) of $\Gamma(\mu_t)$.*

We note that equilibrium behavior is determined by the realized history h_t only via the public belief μ_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.

3.1 Analysis of $\Gamma(\mu)$

We begin by studying the consumer's best-reply strategy in $\Gamma(\mu)$. We denote the consumer's posterior belief after observing the signal $s_t = s$ by $p_\mu(s)$. It follows readily from Bayes rule that

$$p_\mu(s) = \frac{\mu p(s)}{\mu p(s) + (1-\mu)(1-p(s))}. \tag{3}$$

The bounds $\underline{\alpha}$ and $\bar{\alpha}$, together with Equation (3), imply that $p_\mu(s) \in [\underline{\alpha}_\mu, \bar{\alpha}_\mu]$ with probability one, where:

$$\underline{\alpha}_\mu = \frac{\mu\underline{\alpha}}{\mu\underline{\alpha} + (1-\mu)(1-\underline{\alpha})} \text{ and } \bar{\alpha}_\mu = \frac{\mu\bar{\alpha}}{\mu\bar{\alpha} + (1-\mu)(1-\bar{\alpha})} \tag{4}$$

Fix a price vector $\tau = (\tau_0, \tau_1)$ and note that the consumer optimizes her expected utility if and only if she follows the following strategy:

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

We note that in every perfect Bayesian equilibrium of the game $\Gamma(\mu)$ the strategy σ^* constitutes an (almost surely) unique strategy of the consumer.

Note further that every price vector (τ_0, τ_1) induces two possible market scenarios: *the full market scenario*, where under $\sigma^*(\mu, s, \tau)$ the consumer always buys from one firm or another for almost all signal realizations, and *a non-full market scenario*, where $\sigma^*(\mu, s, \tau) = e$ for some positive measure of signals $s \in S$.

We can infer from (5) that the consumer buys from firm 0 whenever

$$p_\mu(s) - \tau_0 \geq (1 - p_\mu(s)) - \tau_1.$$

and the market is full or whenever

$$p_\mu(s) - \tau_0 \geq 0.$$

and the market is not full.

Given a prior μ and a pair of prices τ_0, τ_1, we let $v_\mu(\tau_0, \tau_1)$ be the threshold in terms of the private belief above which firm zero is chosen. That is, choosing firm zero 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:

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

One can easily see that $v_\mu(\tau_0, \tau_1)$ is a continuous function of (μ, τ_0, τ_1).

We can therefore write the expected utility of firm 0 in the game $\Gamma(\mu)$ for the price vector is τ as follows:

$$\begin{aligned} &\Pi_0(\tau, \mu) = \\ &(\mu\,(1 - G_0(v_\mu(\tau_0, \tau_1)) + (1-\mu)\,(1 - G_1(v_\mu(\tau_0, \tau_1)))\,\tau_0 = \\ &[1 - (\mu G_0\,(v_\mu(\tau_0, \tau_1)) + (1-\mu) G_1\,(v_\mu(\tau_0, \tau_1)))]\,\tau_0 \end{aligned} \tag{7}$$

A similar equation can be derived for $\Pi_1(\tau, \mu)$, the profit of firm 1.

In what follows we make a distinction between two forms of perfect Bayesian equilibria of the game $\Gamma(\mu)$: *a deterrence equilibrium*, where only a single firm sells with positive probability, and *a non-deterrence equilibrium*, where both firms sell with positive probability. That is,

Definition 5. *A deterrence equilibrium (DE)* in $\Gamma(\mu)$ is a Bayesian SPE, $(\tau_0^*, \tau_1^*, \sigma^*)$, in which there exists a unique firm i such that:

$$Pr_F(\{s|\sigma^*(\mu, s, \tau^*) = i\}) \neq 0.$$

A non-deterrence equilibrium (NDE) is an equilibrium that is not DE.

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

Theorem 2. *Let τ be a Bayesian sub-game perfect equilibrium of the game $\Gamma(\mu)$:*

1. *If signals are bounded and $g_0(\underline{\alpha}) > 0$, then for some high enough prior, $\mu_0 \in (0,1)$, whenever $\mu > \mu_0$ firm 1 is deterred.*
2. *Symmetrically, if signals are bounded and $g_1(\bar{\alpha}) > 0$, then for some low enough prior $\mu_1 \in (0,1)$, whenever $\mu < \mu_1$) firm 0 is deterred.*
3. *If the signals are unbounded or they exhibit the vanishing likelihood property, then no firm is deterred.*

The proof of Theorem 2 as well as the complete analysis of this stage game is relegated to Appendix A. We next explain the logic behind the proof of Theorem 2 for bounded signal distribution. Consider the case where signals exhibit vanishing likelihood and assume that the prior $\mu < 1$ is very close to 1 and thus strongly in favor of firm 0. Since the proportion of non-conformist consumers is vanishing, in equilibrium firm 0 is better off neglecting those non-conformist consumers who are in favor of firm 1. This fact leaves a margin for a small portion of non-conformists to buy product 1 and so implies that firm 1 is not deterred. If, however, the proportion of those non-conformist consumers is non-vanishing, then firm 0 is better off pricing aggressively and pushing firm 1 out of the market. As a result, a deterrence equilibrium holds.

3.2 Proof of Theorem 1

We next introduce the formal proof of Theorem 1, based on Theorem 2.

Proof of Theorem 1. First we show that if there is no vanishing likelihood, then the martingale of the public belief must converge to an interior point. Assume signal likelihoods are non-vanishing, that is, $g_0(\underline{\alpha}) > 0$. Therefore, by Theorem 2, there exists μ_0 such that $\forall \mu \in (\mu_0, 1)$ there is a unique Bayesian subgame perfect equilibrium of $\Gamma(\mu)$ in which the consumer almost surely chooses firm zero (firm one is deterred by firm zero). This implies that if $\mu_t \in (\mu_0, 1)$, then $\mu_{t+1} = \mu_t$ with probability 1. Now assume to the contrary that the state of the world is $\omega = 0$ and that asymptotic learning holds and so $\lim_{t\to\infty} \mu_t = 1$.[11] Therefore, there exists a time $\hat{t}$ for which $\mu_{\hat{t}} \in (\mu_0, 1)$. This entails that $\mu_t = \mu_{\hat{t}}$ for every $t > \hat{t}$, which yields a contradiction as $\mu_\infty = 1$.[12]

Next we show that if vanishing likelihood holds, then the martingale of the public belief converges to a limit belief in which the true state is assigned probability 1. We start by proving that for every interval $[a, b] \subseteq (0, 1)$ there exists $c > \underline{\alpha}$ and $d < \bar{\alpha}$ such that if $\mu \in [a, b]$, then, for every subgame perfect equilibrium price vector $\tau^*(\mu)$ of $\Gamma(\mu)$, it holds that $v_\mu(\tau^*(\mu)) \in [c, d]$. Assume by way of contradiction that there exists a sequence of $\{\mu_k\} \subseteq [a, b]$ and a corresponding sequence of equilibrium prices $\{\tau^*(\mu_k)\}$ such that $lim_{k\to\infty} v_\mu(\tau^*(\mu_k)) = \underline{\alpha}$.

It follows from Lemma 4 (see appendix A) that $\tau^*(\mu_\infty)$ is a subgame perfect equilibrium price vector of $\Gamma(\mu_\infty)$, and that $v_\mu(\tau^*(\mu_\infty)) = \underline{\alpha}$. Hence, $\tau^*(\mu_\infty)$ is a deterrent SPE of $\Gamma(\mu_\infty)$. Since, by assumption $\mu_\infty \in [a, b] \subset (0, 1)$, and since vanishing likelihood holds, we have reached a contradiction to Theorem 2.

We conclude by proving that if the public belief μ_t lies in some interior interval $[a, b] \subset (0, 1)$ then μ_{t+1} must lie at a distance of at least $\varepsilon > 0$ from μ_t with a probability of at least $\delta > 0$, where ε and δ depend only on the interval $[a, b]$. We show this by examining the likelihood ratio. Conditional on the consumer choosing firm zero at time t we have that:

$$\frac{\mu_{t+1}}{1 - \mu_{t+1}} = \frac{\mu_t}{1 - \mu_t} \frac{1 - G_0(v_\mu(\tau^*(\mu_t))}{1 - G_1(v_\mu(\tau^*(\mu_t))} \tag{8}$$

Since $v_\mu(\tau^*(\mu_t)) \in [c, d] \subseteq (0, 1)$ and the distribution $G_0(\cdot)$ first-order stochastic dominates over $G_1(\cdot)$ (see Lemma A1 in Acemoglu et al. [1]), we know that $\frac{1-G_0(v_\mu(\tau^*(\mu_t))}{1-G_1(v_\mu(\tau^*(\mu_n))} \geq y > 1$. Since $\mu_t \in [a, b]$ it readily follows that there exists $\varepsilon > 0$ such that $\mu_{t+1} > \mu_t + \varepsilon$, conditional on the consumer choosing firm zero at time t. Since $v_\mu(\tau^*(\mu_t)) \in [c, d] \subset (\underline{\alpha}, \bar{\alpha})$, this event occurs with a probability of at least $\delta = 1 - G_1(v_\mu(\tau^*(\mu_t)) > 0$.

By observation 1, the limit $\mu_\infty = \lim_{t\to\infty} \mu_t$ exists and by the above

[11] If asymptotic learning occurs, then for all finite t, $\mu_t < 1$.

[12] Note that $\mu_t < 1$ for all finite t.

argument $\mu_\infty \in \{0,1\}$ with probability 1. This shows that asymptotic learning holds. $\square$

4 Social Learning and Farsighted Firms

In this section we show that our main result carries forward 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 at every time period, except that now, each firm tries to maximize its discounted sum of the stream of payoffs. We follow Maskin and Tirole ([19] – [21]) and analyze the *Markov perfect Nash equilibria* (MPE) of the corresponding Bayesian repeated game.

For initial prior μ and a pair of prices $\tau = (\tau_1, \tau_2)$, let $q_i(\tau, \mu)$ be the probability that the optimal action of the consumer is action $i \in \{0, 1, e\}$ and let $\mu_i(\mu, \tau)$ be the posterior public belief given that the prior is μ and the consumer chooses action $i \in \{0, 1, e\}$. A strategy of firm i in this non-myopic game is a non-negative sequence of prices $\tau_i^* = (\tau_i^t)_{t \in \mathbb{N}}$ such that for every time t, the mapping $\tau_i^t : H_t \to [0,1]$ is measurable. $\tau_i^t(\ )$ determines the price of firm i at time t as a function of the history. Every strategy profile $\tau^* = (\tau_1^*, \tau_2^*)$, initial prior μ, and a discount factor $\delta > 0$ define an expected payoff to every firm i:

$$W_{i,\mu(\tau)} = (1-\delta) \sum_{t=0}^{\infty} \delta^t \Pi_i(\tau_1^t(h_t), \tau_2^t(h_t), \mu_t)$$

where $\Pi_i(\cdot)$ is the expected revenue per period as defined by equation (7).

In the non-myopic case we view the set of all prior beliefs $[0,1]$ as our state space. The transition probabilities given a prior belief μ are determined by the above $q_j's$ to the three potential states μ_j for $j \in \{0, 1, e\}$.

Definition 6. A strategy τ_i^* of firm i in the non-myopic sequential game is called *Markovian* (see Bergemann and Välimäki [5]) if there exists a mapping $\sigma_i : [0,1] \to [0,1]$ such that

$$\tau_i^t(h_t) = \sigma_i(\mu_t) \text{ , for every firm } i \in \{0,1\}. \tag{9}$$

A pair of Markovian strategies $\tau^* = (\tau_0^*, \tau_1^*)$ comprises a *Markov perfect equilibrium* (MPE) if for every initial prior μ the profile τ^* is a Nash equilibrium of the repeated sequential consumption game.[13]

[13]We henceforth identify σ with τ^*.

In (9) we explicitly demand that the firms strategies at time t depend only on the *current state*, which is the public belief μ_t. The Markovian property implies that for every time t, the continuation payoff to firm i depends only on μ_t. Therefore, for any Markovian strategy profile τ^* and initial prior $\mu \in [0,1]$, we can simplify the notation and write $W_{i,\mu(\tau)} = W_i(\mu)$, where $W_i(\mu)$ is the expected payoff to firm i in the sequential game with initial prior μ. Using this notation, we get that the continuation payoff to firm i, conditional on the history h_t, is simply $W_i(\mu_t)$. It follows that for every Markovian strategy profile τ^* and $\mu \in [0,1]$,

$$W_i(\mu) = (1-\delta)q_i(\mu, \tau^*(\mu))\tau_i^*(\mu) + \delta\big(\sum_{j\in\{0,1,e\}} q_j(\mu, \tau^*(\mu))W_i(\mu_j(\tau_i^*(\mu))\big) \tag{10}$$

Note that the payoff to firm i comprises two parts: $q_i(\mu, \tau^*(\mu))\tau_i^*(\mu)$ is the myopic expected payoff to firm i, and

$$\sum_{j\in\{0,1,e\}} q_j(\mu, \tau^*(\mu))W_i^{\tau^*}(\mu_j(\tau_i^*(\mu))$$

is the expected future payoff to firm i. The sum comprises three summands, each of which determines the continuation payoff to firm i subject to the consumer decision that determines the new state. That is, with probability $q_j(\mu, \tau^*(\mu))$ the consumer chooses alternative $j \in \{0,1,e\}$, which yields a posterior belief μ_j and a continuation payoff $W_i(\mu_j)$. It follows from the definition that τ^* is a MPE iff for every firm i and prior μ the price $\tau^*(\mu)_i$ maximizes the expected payoff to firm i on the right-hand side of equation (10).

In the following two theorems we show that our main result in the myopic case carries forward to the non-myopic case. The proofs are relegated to Appendix C.

Theorem 3. *Asymptotic learning holds in every MPE, for every discount factor $\delta < 1$, if the signal distributions exhibit the vanishing likelihood property.*

By Theorem 3, Vanishing likelihood is a sufficient condition for asymptotic learning[14]. If we further restrict attention to MPE τ^* with continuation payoffs W_i that are Lipschitz-continuous in the prior μ, which we denote by "LMPE", then our results for the myopic case carry over completely. I.e. for LMPE, "Vanishing likelihood" is a sufficient and necessary condition for asymptotic learning to hold.

[14]Note that the property where asymptotic learning occurs in a setting where signals are bounded carries over from the myopic case with very little additional structure and assumptions.

Theorem 4. *Asymptotic learning holds in every LMPE, for every discount factor $\delta < 1$, if and only if the signal distributions exhibit the vanishing likelihood property.*

5 Discussion

In the classical models of social learning, the consumers' utility from each alternative is fixed. In that setting, when signals are bounded, there is always a positive probability that the inferior product will prevail (see Banerjee [4] and Bikhchandani et al. [6]). However, when signals are unbounded there are always non-conformists consumers who go against the herd and purchase the product most others won't (see Smith and Sørensen [28]). These non-conformists are instrumental for the aforementioned information aggregation when the consumers' choice is between products with fixed prices. However, our setting involves strategic pricing that alters these results, as each of the firms can now lower its price and attract consumers, even when the prior belief is biased against it, or price out its competitor if the prior belief is in its favor. We ask two questions: When are these pricing strategies optimal? And what implications do these strategies have for the manner in which markets aggregate information? We find that the proportion of non-conformist consumers plays a significant role in answering these questions. An intuitive extension from the model with fixed prices suggests that more non-conformists implies more social learning. However, our main finding is the exact opposite: social learning occurs only when the number of non-conformist consumers is small. This is the condition we refer to as "vanishing likelihood."

We study the conditions under which markets in which firms are engaged in a pricing competition enable or hinder social learning. We do so by introducing a simple setting of duopolistic pricing competition. We first study a simplified model where firms are myopic and prove that in this setting, when signals are bounded, social learning occurs if and only if the signal distributions exhibit the vanishing likelihood property. We then extend these results to a version of the model with forward-looking firms that maximize their expected discounted future revenue stream.

The rationale behind this counterintuitive result is uncovered when analyzing the firms' incentives in the stage game. As society learns one of the firms, say firm zero, emerges as the better one. At that stage the new consumer, prior to receiving a signal, assigns a high probability to firm zero having the superior product. In other words, the stage game begins with a biased prior towards firm zero, which now wants to exploit this nearmonopolistic status and set a high price. The only reason not to do so is when the next consumer is very likely to receive a strong signal that firm one is superior and consequently not conform with its predecessors. This argument can be ignored by firm zero when the probability of this event is low enough, which is exactly captured by our notion of 'vanishing likelihood'. Therefore when signals exhibit vanishing likelihood the popular firm ignores non-conformist consumers and foregoes this market share by setting prices high. Firm one sets prices low and wins over the consumer in the rare event that he is is a non-conformist, thus breaking the herd phenomenon. Notice that no matter how small the probability of this is in the stage game, when we go back to the repeated game it eventually happens with probability one.

While this work contributes to the literature of social learning, the vanishing likelihood property and its effect on firms' strategic behavior has interesting implications for market behavior and in particular on market entry and the adoption of new technologies. We study these aspects in a companion paper (Arieli et al. [2]).

References

[1] Acemoglu, D., Daleh, M. A., Lobel, I., and Ozdaglar, A. Bayesian Learning in Social Networks. *Review of Economic Studies 78* (2011), 1201–1236.

[2] Arieli, I., Koren, M., and Smorodinsky, R. Predatory Pricing and Information Aggregation in Markets with a Common Value. *Available at SSRN 2825009* (2016).

[3] Avery, C., and Zemsky, P. Multidimensional Uncertainty and Herd Behavior in Financial Markets. *The American Economic Review 88*, 4 (1998), 724–748.

[4] Banerjee, A. A Simple Model of Herd Behavior. *The Quarterly Journal of Economics 107*, 3 (1992), 797–817.

[5] Bergemann, D., and Välimäki, J. Learning and Strategic Pricing. *Econometrica 64*, 5 (1996), 1125–1149.

[6] Bikhchandani, S., Hirshleifer, D., and Welch, I. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy 100*, 5 (1992), 992–1026.

[7] Bose, S., Orosel, G., Ottaviani, M., and Vesterlund, L. Dynamic monopoly pricing and herding. *The RAND Journal of Economics 37*, 4 (2006), 910–928.

[8] Bose, S., Orosel, G., Ottaviani, M., and Vesterlund, L. Monopoly Pricing in the Binary Herding Model. *Economic Theory 37*, 2 (2008), 203–241.

[9] Brooker, K. Beautiful Dreamer Jeff Bezos has slashed costs, imposed discipline, and worked like hell to turn Amazon into a lean, efficient retailer. Yet he still wants investors to see his company as something more, as the paragon company of the Internet revolution., 2000.

[10] Brooks, T. Did You Know The Sega Dreamcast Still Gets New Releases?, 2016.

[11] Cleland, K. WEB LACKS TRANSACTIONS, BUT INFLUENCES DECISIONS | News - AdAge. *AdvertisingAge* (apr 1995).

[12] Crisp, B. C., Jarvenpaa, S. L., and Todd, P. A. Individual Differences and Internet Shopping Attitudes and Intentions. *Information research : an international electronic journal.* (1995).

[13] Duffie, D., Malamud, S., and Manso, G. Information Percolation in Segmented Markets. *Journal of Economic Theory 153*, 1 (2014), 1–32.

[14] Ghemawat, P. Leadership Online (A): Barnes & Noble vs. Amazon.com. *Harvard Business School Cases 44*, 0 (2004), 19.

[15] Gifford, K. Why did the Dreamcast fail? Sega's marketing veteran looks back, 2013.

[16] Greco, A. N. *The book publishing industry.* Routledge, 2013.

[17] Grimes, W. Book War: Shops vs. Superstores; As Chains Grow, Struggling Small Stores Stress Expertise, 1995.

[18] Herrera, H., and Hörner, J. Biased Social Learning. *Games and Economic Behavior 80* (2013), 131–146.

[19] Maskin, E., and Tirole, J. A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs. *Econometrica 56*, 3 (1988), 549–569.

[20] Maskin, E., and Tirole, J. A Theory of Dynamic Oligopoly, II: Price Competition , Kinked Demand Curves , and Edgeworth Cycles. *Econometrica 56*, 3 (1988), 571–599.

[21] Maskin, E., and Tirole, J. Markov Perfect Equilibrium: I. Observable Actions. *Journal of Economic Theory 100*, 2 (2001), 191–219.

[22] Moscarini, G., and Ottaviani, M. ECONOMIC MODELS OF SOCIAL LEARNING. In *Decisions, Games and Markets.* Springer US., 1997, pp. 265–298.

[23] Mueller-frank, M. Market Power, Fully Revealing Prices and Welfare. 2012.

[24] Mueller-frank, M. Price Effieciency and Welfare. 2016.

[25] National Retail Federation. National Retail Federation Forecasts Holiday Sales to Increase 3.6% | National Retail Federation, 2016.

[26] OECD. *OECD Science, Technology and Industry Scoreboard 2015.* OECD Science, Technology and Industry Scoreboard. OECD Publishing, oct 2015.

[27] Schumpeter, J. A. *Capitalism , Socialism & Democracy.* Harper & Brothers, 1942.

[28] Smith, B. Y. L., and Sørensen, P. Pathological Outcomes of Observational Learning. *Econometrica 68*, 2 (2000), 371–398.

[29] Smith, R. G. PAYING THE PRICE ON THE INTERNET: FUNDS TRANSFER CRIME IN CYBERSPACE. In *Internet Crime* (Melbourne, 1998), the Australian Institute of Criminology.

[30] Snodgrass, J. BBC News | Sci/Tech | Sega Dreamcast to spark price war, 1999.

[31] The Video Game Critic'. The Video Game Critic's Sega Saturn Console Review page.

[32] Wikia. Video Game Sales Wiki - Price cuts, 2017.

A Proofs of the Stage Game.

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

In this stage game, the decision of the consumer can be described by the simple threshold strategy σ^* described in Section 4. The analysis of equilibrium behavior will therefore focus on analyzing the firms' behavior. The expected profit of firm zero is described in equation (7). The functions $(G_\omega)_{\omega=0,1}$ are differentiable on $(\underline{\alpha}, \bar{\alpha})$ and so, whenever $v_\mu(\tau_0, \tau_1) \in (\underline{\alpha}, \bar{\alpha})$, the payoff functions $\Pi_i(\tau, \mu)$ are differentiable. The partial derivative $\frac{\partial \Pi_0((\tau_0,\tau_1),\mu)}{\partial \tau_0}$ is given by:

$$\begin{aligned} &\frac{\partial \Pi_0((\tau_0, \tau_1), \mu)}{\partial \tau_0} = \\ &[1 - (\mu G_0\left(v_\mu(\tau_0, \tau_1)\right) + (1-\mu) G_1\left(v_\mu(\tau_0, \tau_1)\right))] \\ &\quad - \tau_0 \frac{\partial v_\mu(\tau_0, \tau_1)}{\partial \tau_0} \left[\mu g_0(v_\mu(\tau_0, \tau_1)) + (1-\mu) g_1(v_\mu(\tau_0, \tau_1))\right] \end{aligned} \tag{11}$$

Equation (11) will be crucial for the following equilibrium analysis.

Lemma 1. *Let τ_0^*, τ_1^* be an equilibrium of the game $\Gamma(\mu)$. The following conditions hold:*

1. *If no firm is deterred, then $\frac{\partial \Pi_i(\tau_i,\tau_i^*)}{\partial \tau_i}|_{\tau_i^*} = 0$ for any firm $i = 0, 1$.*
2. *If firm one is deterred, then $\frac{\partial \Pi_0(\tau_0,\tau_1^*)}{\partial \tau_0}|_{\tau_0^*} \leq 0$.*
3. *If firm zero is deterred, then $\frac{\partial \Pi_1(\tau_1,\tau_0^*)}{\partial \tau_1}|_{\tau_1^*} \leq 0$.*

Proof. Proof of 1: We note that since the equilibrium is not deterrent we have that $v_\mu(\tau_0^*, \tau_1^*) \in (\underline{\alpha}, \bar{\alpha})$. The lemma follows from standard first-order conditions.

Proof of 2: Assume to the contrary that τ^* is DE and that $\frac{\partial \Pi_0(\tau_i,\tau_i^*)}{\partial \tau_0}|_{\tau_0^*} > 0$. Then firm zero can increase its profit by increasing its price by ε, a contradiction.

Proof of 3: Symmetric to the proof of 2. □

We next show that any DE of the game $\Gamma(\mu)$ must be a full market equilibrium.

Proposition 1. *If (τ_0^*, τ_1^*) is a DE, then $Pr_F(\{s | \sigma^*(\mu, s, \tau^*) = e\}) = 0$.*

In words, whenever one firm has zero probability of selling, the other firm sells with probability one.

Proof. Without loss of generality assume that no buyer buys from firm 1 and so its profit must be zero. Assume to the contrary that $Pr_F(\{s|\sigma^*(\mu, s, \tau^*) = e\}) > 0$. Recall that $\sigma^*(\mu, s, \tau^*) = e \iff 1 - \tau_1^* < p_\mu(s) < \tau_0^*$ (this follows from the structure of σ^*; see equation (5)). Hence

$$Pr_F(\{s|1 - \tau_1^* < p_\mu(s) < \tau_0^*\}) > 0.$$

Therefore there must exist $\bar{\tau}_1$ such that $Pr_F(\{s|p_\mu(s) < 1 - \bar{\tau}_1\}) > 0$ and $1 - \bar{\tau}_1 < \tau_0^*$, which implies that $\bar{\tau}_1$ is positive. If firm 1 deviates from τ_1^* to $\bar{\tau}_1$, then there is positive probability that agents will buy from firm one. Thus firm one's expected profit is positive. In contradiction to the equilibrium constraint. □

Another key property of a DE is given in the following lemma.

Lemma 2. *Let $\tau^* = (\tau_0^*, \tau_1^*)$ be a deterrence equilibrium (DE) in the game $\Gamma(\mu)$. If firm zero controls the market, then $\underline{\alpha}_\mu \geq \frac{1}{2}$ (and if firm one controls the market, 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 her) it must be the case that the consumer's posterior belief assigns a probability of at most 0.5 that i is the superior firm.

Proof. Assume to the contrary that $\underline{\alpha}_\mu < \frac{1}{2}$ and that τ^* is a DE where firm 1 gets no buyer. Clearly $\Pi_1(\tau^*) = 0$. From Proposition 1 we know that the market is full and hence that $\underline{\alpha}_\mu - \tau_0^* \geq 0$. Consider a deviation of firm 1 to the price $\tau_1 = \tau_0^*$. Note that the set $\{s \in S|p_\mu(s) \in (\underline{\alpha}_\mu, \frac{1}{2})\}$ has positive probability and that for any signal in this set buyers prefer firm 1 to firm 0 (and hence also to the action e). Therefore, at the price $\tau_1 = \tau_0^*$ firm 1 has a positive profit. Thus it is profitable for firm 1 to deviate in contradiction to the equilibrium assumption. The proof of the case where firm zero receives no buyers is symmetric, and thus omitted. □

The following corollary suggests that whenever signals are unbounded there is no DE. In fact all equilibria are NDE.

Corollary 1. *If signals are unbounded then there is no DE in $\Gamma(\mu)$.*

Proof. As signals are unbounded $\bar{\alpha} = 0$ and $\underline{\alpha} = 1$, we get that $\underline{\alpha}_\mu = 0$ and $\bar{\alpha}_\mu = 1$. The proof now follows from Lemma 2. □

An outstanding question is under what conditions does the DE equilibrium with bounded signals exist. We now turn to show that this depends on whether or not signals exhibit vanishing likelihood. We first show that, in a DE, it is always the case that there is a unique price vector where firm one sets its price to zero and firm zero sets its price such that its least favorable consumer becomes indifferent between buying from firm zero or receiving the other firm's product for free.

Proposition 2. *Let $\tau^* = (\tau_0^*, \tau_1^*)$ be a deterrence equilibrium (DE) in game $\Gamma(\mu)$. Then*

$$\tau^* = \begin{cases} (2\underline{\alpha}_\mu - 1, 0) & \text{if } \underline{\alpha}_\mu \geq \frac{1}{2} \\ (0, 1 - 2\bar{\alpha}_\mu) & \text{if } \bar{\alpha}_\mu \leq \frac{1}{2} \end{cases} \tag{12}$$

Proof. Without loss of generality, assume that $\underline{\alpha}_\mu \geq \frac{1}{2}$. The other case follows from symmetry considerations. By Lemma 2 we conclude that in this case firm one is deterred. This in turn implies that $\underline{\alpha}_\mu - \tau_0^* \geq 0$ and $\underline{\alpha}_\mu - \tau_0^* \geq 1 - \underline{\alpha}_\mu - \tau_1^*$.

Next we prove that $\tau_0^* = 2\underline{\alpha}_\mu - 1 + \tau_1^*$. Assume to the contrary that $\underline{\alpha}_\mu - \tau_0^* > 1 - \underline{\alpha}_\mu - \tau_1^*$ and consider the following three cases: (1) $\underline{\alpha}_\mu - \tau_0^* > 0 \geq 1 - \underline{\alpha}_\mu - \tau_1^*$ – firm zero has a profitable deviation by raising its price to $\tau_0' = \underline{\alpha}_\mu$; (2) $\underline{\alpha}_\mu - \tau_0^* = 0 > 1 - \underline{\alpha}_\mu - \tau_1^*$ – firm one has a profitable deviation by lowering its price to $\tau_1 = 1 - \underline{\alpha}_\mu - \varepsilon$ for some small enough ε; (3) $\underline{\alpha}_\mu - \tau_0^* > 1 - \underline{\alpha}_\mu - \tau_1^* \geq 0$ – firm zero has a profitable deviation to $\tau_0' = 2\underline{\alpha}_\mu - 1 + \tau_1^*$.

Since $\tau_0^* = 2\underline{\alpha}_\mu - 1 + \tau_1^*$, the consumer with the lowest signal is indifferent between both products and weakly prefers either product to action e. If $\tau_1^* > 0$, firm one can slightly lower its price and attract buyers with positive probability. This would entail a positive expected profit and hence be a profitable deviation. Yet this is impossible in equilibrium and therefore $\tau_1^* = 0$, which in turn implies $\tau_0^* = 2\underline{\alpha}_\mu - 1$. □

Lemma 2 argued that the condition $\underline{\alpha}_\mu > \frac{1}{2}$ is necessary for a DE (in which firm 1 is deterred) to exist. We now turn to study the implications of this condition.

Lemma 3. *If $\tau^* = (\tau_0^*, \tau_1^*)$ is a Bayesian SPE of $\Gamma(\mu)$, then $\Pi_0(\tau^*) \geq 2\underline{\alpha}_\mu - 1$. Furthermore, if τ^* is a non-deterrence equilibrium, then $\tau_0^* > 2\underline{\alpha}_\mu - 1$.*

Proof. It is easy to see that for a price of $\tau_0 = 2\underline{\alpha}_\mu - 1$ firm zero attracts all consumers regardless of $\tau_1^* \geq 0$. Hence, the price $\tau_0 = 2\underline{\alpha}_\mu - 1$ guarantees a profit of $\underline{\alpha}_\mu - 1$ for firm 0. Therefore, in SPE it must be the case that $\Pi_0(\tau^*) \geq 2\underline{\alpha}_\mu - 1$. To prove the second part of the lemma, we note as

τ^* is a non-deterrence equilibrium, there is a positive probability that the consumer buys from firm 1. By the first part of the lemma firm 0's profit must be at least that of a deterrence equilibrium, $\Pi_0(\tau^*) \geq 2\underline{\alpha}_\mu - 1$, in which it wins with probability 1 and thus $\tau_0^* > 2\underline{\alpha}_\mu - 1$. $\square$

Until now we have shown that a deterrence equilibrium never occurs when signals are unbounded (Corollary 1) or when the prior is not biased (Lemma 2). We show next that if the signals exhibit VL, then for every prior $\mu \in (0,1)$, all equilibria of $\Gamma(\mu)$ are non-deterrent (even when signals are bounded).

Proposition 3. *If the signal distribution exhibits vanishing likelihood, then for every $\mu \in (0,1)$ there is no deterrence equilibrium in $\Gamma(\mu)$.*

Proof. Assume that τ^* is a deterrence equilibrium in which firm 1 is deterred. By Lemma 2 and Proposition 2: $\underline{\alpha}_\mu > \frac{1}{2}$ and $(\tau_0^*, \tau_1^*) = (2\underline{\alpha}_\mu - 1, 0)$. Since τ^* is DE, from Lemma 1 we get $\frac{\partial \Pi_0(\tau_0,0)}{\partial \tau_0}|_{\tau_0 = 2\underline{\alpha}_\mu - 1} \leq 0$. To prove the proposition we compute $\frac{\partial \Pi_0(\tau_0,0)}{\partial \tau_0}|_{\tau_0 = 2\underline{\alpha}_\mu - 1}$ explicitly and show that when vanishing likelihood holds, this expression is bounded away from zero.

From Proposition 1 we know that for the given price vector the market is full. We use *the log-likelihood ratio transformation* (see, e.g., Smith and Sørensen [28], Herrera and Hørner [18], and Duffie et al. [13]). This transformation additively separates between the effect of pricing and the effect of the prior. The log-likelihood ratio of state 0 conditional on observing the signal s is defined as $\log(\frac{p(s)}{1-p(s)})$; we denote the corresponding pair of CDFs by

$$\bar{G}_\omega(x) = F_\omega(\{s \in S | \log(\frac{p(s)}{1-p(s)}) < x\}),\ \omega = 0, 1$$

and the corresponding density functions by $\bar{g}_\omega$. By equation (4), if the bounds of the signal distribution are $\underline{\alpha}$ and $\bar{\alpha}$, then the bounds of the log likelihood ratio distribution are $\log(\frac{\underline{\alpha}}{1-\underline{\alpha}})$ and $\log(\frac{\bar{\alpha}}{1-\bar{\alpha}})$, respectively. In particular, the functions $(\bar{G}_\omega)_{\omega=0,1}$ are differentiable on $(\log(\frac{\underline{\alpha}}{1-\underline{\alpha}}), \log(\frac{\bar{\alpha}}{1-\bar{\alpha}}))$. The log likelihood of the posterior belief is therefore

$$\log(\frac{p_\mu(s)}{1-p_\mu(s)}) = \log(\frac{\mu}{1-\mu}) + \log(\frac{p(s)}{1-p(s)}). \tag{13}$$

In a similar manner we can now transform the indifference signal and denote $\bar{v}_\mu(\tau) = \log(\frac{v_\mu(\tau)}{1-v_\mu(\tau)})$. From our analysis of the consumer's optimal strategy σ^*, we have that

$$\bar{v}_\mu(\tau_0, \tau_1) \equiv \begin{cases} \log(\frac{\frac{1+\tau_0-\tau 1}{2}}{1-\frac{1+\tau_0-\tau 1}{2}}) - \log(\frac{\mu}{1-\mu})] & \text{if the market is full} \\ \log(\frac{\tau_0}{1-\tau_0}) - \log(\frac{\mu}{1-\mu}) & \text{if the market is not full.} \end{cases} \tag{14}$$

By equation (14), we can describe firm zero's expected profit as follows:

$$\Pi_0(\tau_0,\tau_1) = \big(\mu(1-\bar{G}_0(\bar{v}_\mu(\tau_0,\tau_1)) + (1-\mu)(1-\bar{G}_1(\bar{v}_\mu(\tau_0,\tau_1))\big)\tau_0 \tag{15}$$

From equation (14) note that

$$\bar{v}_\mu(2\underline{\alpha}_\mu - 1, 0) = log(\frac{\underline{\alpha}}{1-\underline{\alpha}}) \tag{16}$$

Plugging equation (16) into equation (15) and calculating its derivative we get

$$\begin{aligned} &\frac{\partial \Pi_0(\tau_0,0)}{\partial \tau_0}|_{\tau_0=2\underline{\alpha}_\mu-1} = \\ &1-(2\underline{\alpha}_\mu-1)\left(\frac{\partial \bar{v}_\mu(\tau_0,0)}{\partial \tau_0}|_{2\underline{\alpha}_\mu-1}\right)(\mu\bar{g}_0(\underline{\alpha})+(1-\mu)\bar{g}_1(\underline{\alpha})) \end{aligned} \tag{17}$$

Since vanishing likelihood holds we have that both $g_1(\underline{\alpha}) = 0$, and $\bar{g}_1(\underline{\alpha}) = 0$. Therefore, to complete the proof, we need to show that $\bar{g}_0(\underline{\alpha}) = 0$. Assume to the contrary that $\bar{g}_1(\underline{\alpha}) < \bar{g}_0(\underline{\alpha})$, then, since the distributions are absolutely mutually continuous, there exists ε such that $\int_{\underline{\alpha}}^{\underline{\alpha}+\varepsilon} \bar{g}_0(s)ds > \int_{\underline{\alpha}}^{\underline{\alpha}+\varepsilon} \bar{g}_1(s)ds \Rightarrow \bar{G}_0(\underline{\alpha}+\varepsilon) > \bar{G}_1(\underline{\alpha}+\varepsilon)$. This is a contradiction as $\bar{G}_0$ FOSD over $\bar{G}_1$.

Therefore, if $\bar{g}_1(\underline{\alpha}) = 0$, it must be the case that $\bar{g}_0(\underline{\alpha}) = 0$. Thus, from equation (17), $\frac{\partial \Pi_0(\tau_0,0)}{\partial \tau_0}|_{\tau_0=2\underline{\alpha}_\mu-1} = 1$, and we have reached the desired contradiction. □

A.2 Equilibrium Analysis for Biased Priors

We now consider the case where signals are bounded; in contrast to the assumptions of Proposition 3, signals do not exhibit vanishing likelihood. It turns out that if, say $g_0(\underline{\alpha}) > 0$, then when μ is sufficiently close to 1 all SPE of $\Gamma(\mu)$ are DE. We study first the case where $\mu = 1$.

Observation 3. *If τ^* is an equilibrium of $\Gamma(1)$, then:*

1. *$\tau_0^* = 1$.*
2. *τ^* is a deterrence equilibrium for firm zero.*
3. *$v_\mu(\tau^*) = \underline{\alpha}$.*

Proof. Since $\mu = 1$, it follows that $\underline{\alpha}_\mu = 1$. As a result, the consumer will almost surely choose product zero for all non-negative price vectors such that $\tau_0 \leq 1$. Firm zero can therefore dominate the market with a price 1. □

Lemma 4. *The equilibrium correspondence that maps every prior μ into the set of SPE prices (τ_0^*, τ_1^*) of $\Gamma(\mu)$, is upper semi-continuous.*

Proof. The proof is standard and therefore omitted. $\square$

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 2. *Let $\{\mu_k\}_{k=0}^{\infty} \subset (0,1)$ be a sequence of priors that converges to 1, and let $\{\tau_0^*(\mu_k)\}_{k=1}^{\infty}$ be any appropriate sequence such that for every k, $\tau_0^*(\mu_k)$ is an SPE price of firm 0 in $\Gamma(\mu_k)$. It holds that*

$$\lim_{k\to\infty} \tau_0^*(\mu_k) = 1.$$

Corollary 2 readily follows from Observation 3 and Lemma 4.

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

Corollary 3. *Let $\{\mu_k\}_{k=1}^{\infty} \subseteq (0,1)$ be a sequence of priors such that $\lim_{k\to\infty} \mu_k = 1$. Let $(\tau_0^*(\mu_k), \tau_1^*(\mu_k))$ be some SPE price pair for the game $\Gamma(\mu_k)$. Then*

$$lim_{k\to\infty} v_{\mu_k}(\tau_0^*(\mu_k), \tau_1^*(\mu_k)) = \underline{\alpha}.$$

Proof. Assume the corollary is false and that there exists some convergent subsequence which limit is not $\underline{\alpha}$. For this subsequence there is a sub-subsequence for which the sequence of price pairs, $\{(\tau_0^*(\mu_k), \tau_1^*(\mu_k))\}_{k=1}^{\infty}$, converges to some limit (τ_0^*, τ_1^*). Without loss of generality we may assume that the sub-subsequence is actually the original sequence and

$$\exists \lim_{k\to\infty} v_{\mu_k}(\tau_0^*(\mu_k), \tau_1^*(\mu_k)) = \beta > \underline{\alpha}$$

(note that β cannot be smaller than $\underline{\alpha}$ as the threshold is always in the interval $[\underline{\alpha}, \bar{\alpha}]$).

Therefore, for some $\tilde{k}$ and any $k > \tilde{k}$:

$$v_{\mu_k}(\tau_0^*(\mu_k), \tau_1^*(\mu_k)) > \underline{\alpha} + \frac{\beta - \underline{\alpha}}{2}$$

and so the probability of consumer k purchasing from firm zero is bounded above by

$$Pr_{\mu_k}\left(p(s) > v_{\mu_k}(\tau_0^*(\mu_k), \tau_1^*(\mu_k))\right) < Pr_{\mu_k}\left(p(s) > \underline{\alpha} + \frac{\beta - \underline{\alpha}}{2}\right) < 1.$$

In addition, firm zero's price in $\Gamma(\mu_k)$ is $\tau_0^*(\mu_k) \leq 1$ and so and all $k > \tilde{k}$ the profit of firm zero is uniformly bounded away from 1.

On the other hand, a deviation by firm zero to the price $\tau_0 = 2\underline{\alpha}_{\mu_k} - 1$ will guarantee a profit of $2\underline{\alpha}_{\mu_k} - 1$. By equation (4) $\lim_{k\to\infty} 2\underline{\alpha}_{\mu_k} - 1 = 1$ which implies that the price $2\underline{\alpha}_{\mu_k} - 1$ is a profitable deviation for firm zero in the game $\Gamma(\mu_k)$ for any $k > \tilde{k}$ and hence a contradiction. $\square$

Corollary 3 is stronger than Observation 3 as it shows that when the prior approaches 1, the probability that the consumer will purchase product 0 approaches 1.

The following lemma is a counterpart of Corollary 2. Unlike Corollary 2 it does not follow immediately from Observation 3 and Lemma 4 since $(1, \tau)$ is an SPE price vector of $\Gamma(1)$ for any $\tau \in [0, 1]$.

Lemma 5. *Let $\{\mu_k\}_{k=0}^{\infty} \subset (0, 1)$ be a sequence of priors that converges to 1 and let $\tau_1^*(\mu_k)$ be an SPE price of firm 1 in $\Gamma(\mu_k)$, then $\lim_{k\to\infty} \tau_1^*(\mu_k) = 0$.*

Proof. Assume by contradiction that there exists a convergent subsequence for which the limit is some $\delta > 0$. Without loss of generality we may assume this subsequence is actually the original sequence, i.e. $\lim_{k\to\infty} \tau_1^*(\mu_k) = \delta$. Therefore there exists $\bar{k}$ such that $\tau_1^*(\mu_k) > \frac{\delta}{2}$ for all $k > \bar{k}$. In addition, as $\lim_{k\to\infty}(\mu_k) = 1$ then $\lim_{k\to\infty} \underline{\alpha}_{\mu_k} = 1$ (recall equation (4)). Therefore, there exists a $\tilde{k}$ such that $1 - \underline{\alpha}_{\mu_k} < \frac{\delta}{2}$ for all $k > \tilde{k}$. Cobining the last two observation implies that for any $k > max\{\bar{k}, \tilde{k}\}$, $\tau_1^*(\mu_k) > 1 - \underline{\alpha}_{\mu_k}$ and no consumer buys product one in $\Gamma(\mu_k)$ (recall equation (5)). This, in turn, implies that the expected profit of firm one is zero.

In addition it must be the case that $\tau_0^*(\mu_k) \geq \underline{\alpha}_{\mu_k}$. To see this note that equation (5) implies that whenever $\tau_1^*(\mu_k) > 1 - \underline{\alpha}_{\mu_k}$, then for all $\tau_0 \leq \underline{\alpha}_{\mu_k}$, the consumer almost surely buys from firm zero. Thus, by equation (7), every $\tau_0 < \underline{\alpha}_{\mu_k}$ is dominated by $\tau_0 = \underline{\alpha}_{\mu_k}$.

It is now possible for firm one to deviate to some price, τ_1', satisfying $1 - \underline{\alpha}_{\mu_k} - \tau_1' > 0 \geq \underline{\alpha}_{\mu_k} - \tau_0^*(\mu_k)$. There exists some $\epsilon > 0$ such that any agent whose signal is in the interval $[\underline{\alpha}_{\mu_k}, \underline{\alpha}_{\mu_k} + \epsilon]$ will prefer to buy from firm one and so firm one's profit will be positive. Therefore firm one has a profitable deviation and we reach the desired contradiction. □

Finally, we present the complementary condition of Proposition 3 for the case of non-vanishing likelihood. We show that when the signal distribution likelihood is non-vanishing and the prior is sufficiently biased in favor of one firm, there is a unique equilibrium in which the a priori disadvantageous firm is deterred.

Proposition 4. *If $g_0(\underline{\alpha}) > 0$, then $\exists \mu_0 \in (0, 1)$, such that $(2\underline{\alpha}_\mu - 1, 0)$ is the unique equilibrium in $\Gamma(\mu)$ for all $\mu > \mu_0$. Symmetrically, if $g_1(\bar{\alpha}) > 0$, then $\exists \mu_1 \in (0, 1)$, such that $(0, 2\bar{\alpha}_\mu - 1)$ is the unique equilibrium in $\Gamma(\mu)$, for all $\mu < \mu_1$.*

Proof. We prove only the first part of the proposition as the second part follows from symmetry considerations. Assume that $g_0(\underline{\alpha}) > 0$. Assume to the contrary that there exists a sequence $\{\mu_k\}_{k=1}^{\infty}$ such that $\lim_{k\to\infty} \mu_k = 1$

and, for all k, there exists a non-deterrence equilibrium in the corresponding $\Gamma(\mu_k)$. Now recall the derivative of the profit function for firm zero:

$$\frac{\partial \Pi_0(\tau)}{\partial \tau_0} = A(\mu,\tau) - \tau_0 \frac{\partial v_\mu(\tau)}{\partial \tau_0} B(\mu,\tau). \tag{18}$$

where

$$A(\mu,\tau) = 1 - (\mu G_0(v_\mu(\tau_0,\tau_1)) + (1-\mu)G_1(v_\mu(\tau_0,\tau_1))) \tag{19a}$$
$$B(\mu,\tau) = \mu g_0(v_\mu(\tau_0,\tau_1)) + (1-\mu)g_1(v_\mu(\tau_0,\tau_1)). \tag{19b}$$

Note that for all $\mu \in (0,1)$ and for all $\tau \in [0,1]^2$: $0 \leq A(\mu,\tau) \leq 1$ and $0 \leq B(\mu,\tau)$. Also note that $B(\mu,\tau)$ is a composite of continuous functions of μ, and hence is continuous in μ. From Corollary 3 we know that $\lim_{k\to\infty} v_{\mu_k}(\tau_0^*(\mu_k)), \tau_1^*(\mu_k)) = \underline{\alpha}$. Therefore, when the vanishing likelihood property does not hold,

$$\lim_{k\to\infty} B(\mu_k, \tau^*(\mu_k)) = g_0(\underline{\alpha}) > 0$$

As in the proof of Proposition 3, we can use the log likelihood transformation to additively separate the effect of pricing on the consumer indifference signal. Using equation (14) we can calculate the appropriate derivative of the threshold function:

$$\frac{\partial \bar{v}_\mu(\tau)}{\partial \tau_0}|_{\tau^*(\mu)} = \begin{cases} \frac{2}{1-(\tau_0^*(\mu)-\tau_1^*(\mu))^2} & \text{if the market is full} \\ \frac{1}{\tau_0^*(\mu)(1-\tau_0^*(\mu))} & \text{if the market is not full} \end{cases} \tag{20}$$

Now by Corollary 2 and Lemma 5 above $\lim_{k\to\infty} \tau_0^*(\mu_k) = 1$ and $\lim_{k\to\infty} \tau_1^*(\mu_k) = 0$. Hence

$$\lim_{k\to\infty} \frac{\partial \bar{v}_{\mu_k}(\tau)}{\partial \tau_0}|_{\tau^*(\mu_k)} = \infty \tag{21}$$

Since $\bar{v}_\mu(\tau)$ is a log-likelihood ratio transformation of $v_\mu(\tau)$, we can perform the inverse transformation $v_\mu(\tau) = \frac{e^{\bar{v}_\mu(\tau)}}{1+e^{\bar{v}_\mu(\tau)}}$, and get

$$\lim_{k\to\infty} \frac{\partial v_{\mu_k}(\tau)}{\partial \tau_0}|_{\tau^*(\mu_k)} = \lim_{k\to\infty} \frac{e^{\bar{v}_\mu(\tau)}}{(1+e^{\bar{v}_\mu(\tau)})^2} \frac{\partial \bar{v}_\mu(\tau)}{\partial \tau_0}|_{\tau^*(\mu_k)} = \infty \tag{22}$$

The latter equality holds as signals are bounded, and thus, by Corollary 3, $\lim_{k\to\infty} \frac{e^{\bar{v}_\mu(\tau)}}{(1+e^{\bar{v}_\mu(\tau)})^2} = \lim_{k\to\infty} v_\mu(\tau)(1-v_\mu(\tau)) = \underline{\alpha}(1-\underline{\alpha})$.

When combining equation (22) with equation (18) we find that

$$\lim_{k\to\infty} \frac{\partial \Pi_0(\tau)}{\partial \tau_0}|_{\tau^*(\mu_k)} = -\infty$$

On the other hand, we assume that $\tau^*(\mu_k)$ is a NDE, and hence by Lemma 1, for all $k < \infty$, the following equality holds:

$$\frac{\partial \Pi_0(\tau)}{\partial \tau_0}|_{\tau^*(\mu_k)} = 0, \tag{23}$$

a contradiction. □

The importance of the distinction between a DE and NDE in our stage game is expressed in the following lemma. It argues that the probability of buying from firm 0 is state-dependent if and only if the equilibrium being played is a NDE.

Proof of Theorem 2

Proof. The first part of Theorem 2 follows from Proposition 4. The second part of the theorem follows from Proposition 3. To show part (3) of the theorem, let $\tau = (\tau_0, \tau_1)$ be an SPE equilibrium of $\Gamma(\mu)$ for $\mu \in (0,1)$. If signals are unbounded it follows from Corollary 1 that τ is not DE. If the signals are bounded but exhibit vanishing likelihood, then Proposition 3 shows that τ is a NDE. □

B Additional Proofs for the Main Model

Lemma 6. *Let $(\overline{\sigma}, \overline{\tau}_0, \overline{\tau}_1)$ be a myopic Bayesian equilibrium. If asymptotic learning holds, then conditional on state $\omega \in \Omega$,*

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

The lemma states that if asymptotic learning holds, then the probability that consumer t buys the superior product approaches one as t goes to infinity.

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 surely. By Corollary 3 we have that $\lim_{t\to\infty} v_{\mu_t}(\tau_0^t(\mu_t), \tau_1^t(\mu_t)) = \underline{\alpha}$. Therefore,

$$\lim_{t\to\infty} \mathbf{P}_{(\overline{\sigma},\overline{\tau}_0,\overline{\tau}_1)}(\{\sigma^t(\mu_t, s, \overline{\tau}(\mu_t)) = 0\}|\omega = 0) = \lim_{t\to\infty} G_0(v_{\mu_t}(\tau_0^t(\mu_t), \tau_1^t(\mu_t))) = G_0(\underline{\alpha}) = 1.$$

□

Lemma 7. $2\underline{\alpha}_\mu - 1$ *is a strictly convex function of μ.*

Proof. Let

$$h(\mu) = 2\underline{\alpha}_\mu - 1 = 2\frac{\mu\underline{\alpha}}{\mu\underline{\alpha} + (1-\mu)(1-\underline{\alpha})} - 1. \tag{24}$$

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

$$\frac{\partial^{(2)} h(\mu)}{\partial^{(2)}\mu} = \frac{4(1-\underline{\alpha})\underline{\alpha}(2\underline{\alpha}-1)}{(\mu(1-2\underline{\alpha}) - (1-\underline{\alpha}))^3}. \tag{25}$$

The numerator of equation (25) is negative since $\underline{\alpha} < 0.5$. The denominator of (25) is also negative since $\mu \leq 1$ and therefore

$$1 - \underline{\alpha} > 1 - 2\underline{\alpha} \geq \mu(1 - 2\underline{\alpha}).$$

Hence, for every $\mu \in [0, 1]$, $\frac{\partial^{(2)} h(\mu)}{\partial^{(2)} \mu} > 0$ and $h(\mu)$ is strictly convex. □

C Proofs for the Farsighted Firms

C.1 Proof of Theorem 3

Following the same logic as in the proof of Theorem 1, we get that asymptotic learning occurs if and only if $\forall \mu \in (0, 1)$, there is no firm that choose a deterrence strategy.

Assume to the contrary that for some discount factor $\delta < 1$ and MPE τ^* there exists $\mu^d \in (0, 1)$ for which, without loss of generality, only firm zero is visited with positive probability. Since we are restricting our analysis to non-negative prices, it must be the case that $\underline{\alpha}_{\mu^d} > \frac{1}{2}$ (by Lemma 2). In addition, a similar consideration to Proposition 2 shows that $\tau^*(\mu^d) = (2\underline{\alpha}_{\mu^d} - 1, 0)$.

The signal distribution likelihood is vanishing; therefore, from Theorem 2, we know that $\forall \mu < 1$, there exists $\tau_0' > 2\underline{\alpha}_\mu - 1$, which yields a higher expected profit to firm zero in the current period. By Lemma 7, $2\underline{\alpha}_\mu - 1$ is convex in μ, and we show that firm zero's can profit by deviating to a strategy where it plays the myopic best reply, τ_0', in the current period and plays the deterrence price $2\underline{\alpha}_\mu - 1$ the subsequent periods.

To prove this formally, recall that in the current period, the pair of prices after firm zero deviates is $(\tau_0', 0)$. Therefore, the market is full and $q_e(\mu^d, (\tau_0', 0)) = 0$. Hence, the continuation belief obtains only two values with positive probability and by the martingale property it must be the case that

$$\mu^d = \varphi\big(\mu^d, (\tau_0', 0)\big)\mu_0(\mu^d, (\tau_0', 0)) + \bigg(1 - \varphi\big(\mu^d, (\tau_0', 0)\big)\bigg)\mu_1(\mu^d, (\tau_0', 0)) \quad (26)$$

Hence a deviation of firm zero to play τ_0' now and deterrence in the future yields the following expected payoff:

$$\begin{aligned}&(1 - \delta)\varphi(\mu^d, (\tau_0', 0))\tau_0' \\ &+ \delta[\varphi(\mu^d, (\tau_0', 0))(2\underline{\alpha}_{\mu_0(\mu^d, \tau_0')} - 1) + (1 - \varphi(\mu^d, (\tau_0', 0)))(2\underline{\alpha}_{\mu_1(\mu^d, \tau_0')} - 1)]\end{aligned} \quad (27)$$

The expected payoff for firm zero from deviating to the myopic best reply τ_0' is comprised of two parts. The first part $\varphi(\mu^d, (\tau_0', 0))\tau_0'$ is the current period expected payoff and, by the choice of τ_0', is strictly larger than $2\underline{\alpha}_{\mu^d}-1$. The second part is the expected payoff from playing deterrence in subsequent periods, and is also larger than $2\underline{\alpha}_{\mu^d}-1$ due to the convexity of $2\underline{\alpha}_{\mu}-1$ proved in Lemma 7. This stands in contradiction to the assumption that τ^* is a MPE.

C.2 Proof of Theorem 4.

By Theorem 3, if the signal distribution likelihood is vanishing, then asymptotic learning holds. To complete the proof, we show that if the signal distribution does not exhibit the vanishing likelihood property, then asymptotic learning does not hold.

Assume to the contrary that the signal distribution likelihood is non-vanishing and that asymptotic learning holds for some LMPE τ^*. As asymptotic learning holds, it follows from Theorem 2 that there is no t for which there is a deterrence equilibrium at the corresponding stage game $\Gamma(\mu_t)$. As before, the probability of a consumer choosing firm zero is

$$\varphi(\mu,\tau)=\mu\Big(1-G_0\big(v_\mu(\tau)\big)\Big)+(1-\mu)\Big(1-G_1\big(v_\mu(\tau)\big)\Big).$$

For tractability we assume that the market is full; extending the analysis to the general case is straightforward and hence omitted. The contradictory assumption implies that a single deviation to the monopolist price $2\underline{\alpha}_{\mu}-1$ is not profitable. The full market assumption implies that for every μ,

$$\begin{multline}(1-\delta)(2\underline{\alpha}_{\mu}-1)+\delta W_0(\mu)\le\\ (1-\delta)\varphi(\mu,\tau^*)\tau_0^*+\delta\Big(\varphi(\mu,\tau)W_0\big(\mu_0(\mu,\tau^*)\big)+\big(1-\varphi(\mu,\tau^*)\big)W_0\big(\mu_1(\mu,\tau^*)\big)\Big)\end{multline} \tag{28}$$

Rearranging (28) yields:

$$\begin{multline}2\underline{\alpha}_{\mu}-1-\varphi(\mu,\tau^*)\tau_0^*\le\\ \frac{\delta}{1-\delta}\Bigg(\varphi(\mu,\tau^*)\Big(W_0\big(\mu_0(\mu,\tau^*)\big)-W_0(\mu)\Big)+\big(1-\varphi(\mu,\tau^*)\big)\Big(W_0\big(\mu_1(\mu,\tau^*)\big)-W_0(\mu)\Big)\Bigg)\end{multline} \tag{29}$$

$W_0(\mu)$ is Lipschitz continuous; therefore, by definition, there exists a constant $C\in\mathbb{R}_+$ such that

$$W_0\big(\mu_0(\mu,\tau)\big)-W_0(\mu)\le\frac{C}{2}(\mu_0(\mu,\tau)-\mu) \tag{30}$$

and similarly, since $\mu>\mu_1(\mu,\tau)$,

$$W_0(\mu)-W_0\big(\mu_1(\mu,\tau)\big)\le\frac{C}{2}(\mu-\mu_1(\mu,\tau)) \tag{31}$$

Next by Bayes rule

$$\mu_0(\mu,\tau) = \frac{\mu(1-G_0(v_\mu(\tau)))}{\varphi(\mu,\tau)},$$
$$\mu_1(\mu,\tau) = \frac{\mu G_0(v_\mu(\tau))}{1-\varphi(\mu,\tau)}.$$

A simple calculation shows that

$$\mu_0(\mu,\tau) - \mu = \frac{\mu(1-\mu)(G_1(v_\mu(\tau)) - G_0(v_\mu(\tau)))}{\varphi(\mu,\tau)} \tag{32}$$

$$\mu - \mu_1(\mu,\tau) = \frac{\mu(1-\mu)(G_1(v_\mu(\tau)) - G_0(v_\mu(\tau)))}{1-\varphi(\mu,\tau)} \tag{33}$$

Substituting equations (30)–(33) into (29), we get

$$2\underline{\alpha}_\mu - 1 - \varphi(\mu,\tau^*)\tau_0^* \leq \frac{\delta}{1-\delta}(G_1(v_\mu(\tau)) - G_0(v_\mu(\tau)))\mu(1-\mu)C \tag{34}$$

We plug $\varphi(\mu,\tau)$ into (34), and divide by $G_1(v_\mu(\tau)) - G_0(v_\mu(\tau))$, which is positive, and get:

$$\frac{2\underline{\alpha}_\mu - 1 - (\mu(1-G_0(v_\mu(\tau))) + (1-\mu)(1-G_1(v_\mu(\tau))))\tau_0^*}{G_1(v_\mu(\tau)) - G_0(v_\mu(\tau))} \leq \frac{\delta}{1-\delta}\mu(1-\mu)C \tag{35}$$

It is easy to see that as μ approaches 1, the right-hand side of equation (35) approaches 0. To see the desired contradiction, we now prove that the left-hand side is positive and bounded away from zero. We rearrange (35) and get

$$-\mu\tau_0^* + \frac{2\underline{\alpha}_\mu - 1 - \tau_0^*}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))} + \frac{G_1(v_\mu(\tau^*))\tau_0^*}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))} \leq \frac{\delta}{1-\delta}\mu(1-\mu)C \tag{36}$$

It is easy to see that $-\mu_t\tau_0$ is bounded below by -1. We next show that $\frac{G_1(v_\mu(\tau^*))\tau_0^*}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))} > 1$. By a standard first-order approximation it holds that for every $\epsilon > 0$ there exists $x_\epsilon > \underline{\alpha}$ such that for every $x \leq x_\epsilon$,

$$(x - \underline{\alpha})[g_i(\underline{\alpha}) - \epsilon] \leq G_i(x) \leq (x - \underline{\alpha})[g_i(\underline{\alpha}) + \epsilon].$$

By Corollary 2 it holds that $\lim_{\mu\to 1} v_\mu(\tau^*) = \underline{\alpha}$. Therefore, for every $\epsilon > 0$ there exists $\mu_\varepsilon < 1$ such that $v_\mu(\tau^*) \leq x_\epsilon$ for all $\mu \geq \mu_\varepsilon$ and hence

$$\begin{aligned}\frac{G_1(v_\mu(\tau^*))}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))} &> \frac{(g_1(\underline{\alpha}) - \varepsilon)(v_\mu(\tau^*) - \underline{\alpha}))}{(g_1(\underline{\alpha}) + \varepsilon)(v_\mu(\tau^*) - \underline{\alpha}) - (g_0(\underline{\alpha}) - \varepsilon)(v_\mu(\tau^*) - \underline{\alpha})} \\ &= \frac{g_1(\underline{\alpha}) - \varepsilon}{g_1(\underline{\alpha}) - g_0(\underline{\alpha}) + 2\varepsilon}\end{aligned} \tag{37}$$

Since G_0 FOSD G_1 (see Lemma A1 in Acemoglu et al. [1]) and the signals' likelihood is non-vanishing, we have that $g_1(\underline{\alpha}) > g_0(\underline{\alpha}) \geq 0$. Since

$\lim_{\varepsilon\to 0} \frac{g_1(\underline{\alpha})-\varepsilon}{g_1(\underline{\alpha})-g_0(\underline{\alpha})+2\varepsilon} = \frac{g_1(\underline{\alpha})}{g_1(\underline{\alpha})-g_0(\underline{\alpha})} > 1$, there exists $\mu' < 1$ such that for all $\mu \geq \mu'$ it holds that[15]

$$\frac{G_1(v_\mu(\tau^*))\tau_0^*}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))} > 1.$$

All that is left is to show is that $\frac{2\underline{\alpha}_\mu - 1 - \tau_0^*}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))}$ approaches zero as μ approaches one. The numerator of the left-hand side of (35) is negative by Lemma 3 and by stochastic dominance, while the denominator is positive. We use the approximation again:

$$\frac{2\underline{\alpha}_\mu - 1 - \tau_0^*}{G_1(v(\tau^*)) - G_0(v(\tau^*))} > \frac{2\underline{\alpha}_\mu - 1 - \tau_0^*}{(g_1(\underline{\alpha}) - g_0(\underline{\alpha}) + 2\varepsilon)(v_\mu(\tau^*) - \underline{\alpha})} \tag{38}$$

Now recall that $v_\mu(\tau) - \underline{\alpha} = v_\mu(\tau) - v_\mu(2\underline{\alpha}_\mu - 1) = v'_\mu(\tilde{\tau})(\tau_0 - 2\underline{\alpha}_\mu + 1)$ for some $\tilde{\tau} \in (2\underline{\alpha}_\mu - 1, \tau_0)$. Therefore,

$$\frac{2\underline{\alpha}_\mu - 1 - \tau_0^*}{G_1(v_\mu(\tau^*)) - G_0(v_\mu(\tau^*))} > \frac{-1}{(g_1(\underline{\alpha}) - g_0(\underline{\alpha}) + 2\varepsilon)v'_\mu(\tilde{\tau^*})} \tag{39}$$

From Lemma 5, we know that as μ approaches 1, the SPE price vector $(\tau_0^*(\mu), \tau_1^*(\mu))$ of $\Gamma(\mu)$ satisfy $\tau_0^*(\mu) \to 1$ and $\tau_1^*(\mu) \to 0$. Therefore, $v'_\mu(\tilde{\tau}) \to \infty$. As a result, the second element of (36) approaches zero and the left-hand side of equation (36) is bounded away from zero while the right-hand side approaches zero, a contradiction.

[15]By Lemma 3 we know that $\tau_0^* \to 1$ as $\mu \to 1$.