Market Power and Asymmetric Learning in Product Markets*

Ali Kakhbod† Giacomo Lanzani‡

MIT MIT

Abstract

We study how market structures, along with asymmetries in learning technologies, affect trade in a product market. In this market, a new product of unknown quality is introduced to challenge a pre-existing product of known quality. We show that market efficiency (the first-best) is achieved both under monopoly and competition if buyers are symmetric in their learning process. Instead, if buyers are asymmetric, only a monopolistic market in which the seller of the old product also sells the new one is efficient. We identify inefficiency as a learning externality that consumptions of the unknown quality product by one buyer generates for the other buyers. The equilibrium inefficiency has two important features: (i) Efficiency at the top: the threshold for starting to serve the best learners (i.e., to enter into a Beta phase) remains the optimal one; (ii) Non-monotonicity: distortions are not monotone in the extent of the asymmetry. Importantly, if sellers can offer take-it-or-leave-it multilateral contracts, the distortion disappears. Finally, we explore the robustness of our results under different assumptions about the ability to price discriminate and other market structures.

Keywords: Market structure, Market power, Trading, Learning.

JEL Classification: C7

---

*This version August 2020. We thank Daron Acemoglu, Ricardo Caballero, Glenn Ellison, Drew Fudenberg, Kevin Li and Stephen Morris for great and very helpful comments. Giacomo Lanzani gratefully acknowledges the financial support of the Guido Cazzavillan scholarship.

†Department of Economics, Massachusetts Institute of Technology (MIT), E52-391, 50 Memorial Drive, Cambridge MA 02142, USA. Email: akakhbod@mit.edu.

‡Department of Economics, Massachusetts Institute of Technology (MIT), E52-391, 50 Memorial Drive, Cambridge MA 02142, USA. Email: lanzani@mit.edu.

1 Introduction

E-commerce (e.g., Amazon, ebay, Alibaba) provides easy access for sellers to compete and sell new products.[1] They are also platforms for buyers to publicly provide their reviews from their experiences. These reviews are important sources of information to learn about the quality or adoption potentials of new products, affecting not only competition but also the pricing of pre-existing (old) products.[2]

In those markets, the large number of available reviews of individual experiences implies that the belief about the new product is approximately the same for all market participants. However, a crucial form of heterogeneity across agents still persists: feedbacks and reviews are often in different forms (from a brief 1 to 5 scale review to detailed text comments) and agents may differ in the accuracy of the reviews provided. This important heterogeneity could be due to asymmetry between buyers in how deeply they experience new products (e.g., old and young generations).

Given this prevalent heterogeneity in reviews, how do reviews affect price competition and the learning of Bayesian market participants? When does a buyer decide to buy a new product? Does competition improve welfare? What type of distortions may arise with competition? Moreover, what policies are effective in improving market efficiency? In this paper, we offer answers to these questions, focusing on buyers' heterogeneity in providing reviews.

To study these questions, we consider a market where a new product of unknown quality is introduced to challenge a pre-existing (old) product of known quality. Two states correspond to the possible quality levels of the new product. The true state is initially unknown, and the market participants are Bayesian agents who can gradually learn the quality of the new product through reviews of the buyers experiencing the new product. Importantly, we allow the buyers to be heterogeneous in the (expected) information

[1]In 2018, consumers spent $517.36 billion online with U.S. merchants, up 15 percent from 2017. E-commerce made up 14.3 percent of total U.S. retail sales in 2018. For example, Amazon alone accounts for 40 percent of U.S. online retail, and Amazon accounted for 43.3 percent of e-commerce gains in the United States in 2019.

[2]User reviews, particularly on sites like Amazon, mean a great deal to shoppers. "*A product that has just one review is 65% more likely to be purchased than a product that has none*", according to Power Reviews CEO Matt Moog. He added that one-third of online shoppers refuse to purchase products that have not received positive feedback from customers.

about the quality of the new product that they generate by consuming it. This is a reduced form representation of the different accuracy of feedback they provide. Since their use of the product generate more learning about the quality of the good, we refer to the agent who provide more accurate signals as the good or better learners. We assume that the learning ability of each buyer is an exogenous parameter. We think that this assumption is approximately correct in many instances of online commerce, where the differences in the amount of feedback provided seem mostly related to individual attitudes instead of strategic considerations.[3]

To analyze the impact of market power in the model, we consider the efficient consumption pattern as well as the decentralized market outcomes under both a monopolistic firm serving both products and competition between two firms, each specialized in one of the products.

We first establish that both the efficient and the decentralized market outcomes feature a sequence of belief thresholds. When there is little confidence in the new product being good, only the good learners consume it. Over time if the market belief about the new product is sufficiently improved by the reviews of good learners, then worse learners also start buying the new product. In other words, all solutions feature a *Beta phase* in which only the best learners experience the new (unknown) product. We explicitly derive the Beta phase and the expected length of the Beta phase in terms of the endogenous model parameters.

To analyze the effect of market power, we consider the efficiency properties of different market outcomes, with particular attention to the comparison between monopoly and competition. The main finding is that the relative welfare performance of monopolistic and duopolistic market structures crucially relies on the learning technology across buyers. We show that both market structures lead to efficiency when buyers are symmetric in the way they process information (i.e., homogeneous learning technologies). In contrast, however, when buyers are asymmetric, competition is no longer efficient, whereas the monopolistic market remains so. The lower welfare resulting from competition may seem counterintuitive; the intuition behind these results is that in dynamic markets, bilateral

[3]As a consequence, in our model sellers can enhance information production only by targeting the best learners. We think that having sellers trying to endogenously increase the feedback produced by the agent may be important in some applications. However, we do not follow this route in this paper, because we are interested in isolating the informational externalities due to the learning asymmetries.

contracting between two parties produces learning externalities on the other market actors, proportional to their value of information. In monopolistic markets, the monopolist optimal pricing makes this value 0. Instead, under competition, a part of these positive externalities is appropriated by the potential buyers not involved in the transaction, and therefore is not internalized in prices. To make this more concrete, the situation is analogous to the scenario in which a new retailer starts to sell its product on Amazon. The new entrant wants to offer discounts to the subset of consumers that can provide detailed reviews of the new product. Intuitively, the amount of discount is increasing in the future market power of the entrant in case of success. Our results highlight that competition may induce objectively suboptimal discounting strategies.

Our model also yields novel implications in terms of the structure of the inefficiency induced by competition. Particularly, as in the first best, the equilibrium behavior features a threshold structure. For a low level of confidence, only the best learners use the new product, while worse learners move to the new product as the confidence level increases. This equilibrium features efficiency at the top: the equilibrium threshold in beliefs to start serving the very first buyers, i.e., the best learners, is the same as the first-best. This implies that all the new products that are sufficiently promising (i.e., a priori market belief is high enough for them) are given a shot. However, competition distorts the threshold for moving out from the Beta phase and start serving the entire market.

We further investigate the comparative statics of the above inefficiency. Even if asymmetries in the learning technology are necessary to have an inefficient market outcome, the size of the distortion is not monotone in the amount of heterogeneity.

Finally, we consider a possible solution to the distortions induced by competition. Indeed, we show that the introduction of multilateral contracts leads to an efficient equilibrium outcome. Precisely, we increase the commitment power of the sellers by allowing them to make take-it-or-leave-it offers to multiple market participants. These offers may require buyer 1 to pay for having the product consumed by buyer 2. Buyer 1 may accept such an offer because of the information produced by buyer 2 upon consumption. We prove that if such contracts are feasible, the decentralized outcome is efficient regardless of the heterogeneity in the learning technologies.

1.1 Related Literature

Dynamic pricing has a rich history.[4] In general, time-varying prices may have different reasons. It might be due to inability of the firms to commit to future actions (e.g. Conlisk, Gerstner and Sobel (1984), Sobel (1991)), or due to learning new experience goods (e.g. Caminal and Vives (1999), Bergemann and Välimäki (2000)), or the result of the inability of boundedly rational buyers to pay immediate attention to price changes (e.g. Radner, Radunskaya and Sundararajan (2014), Bordalo, Gennaioli and Shleifer (2016), Bordalo, Gennaioli and Shleifer (2017)), or information diffusion due to word of mouth effect Ajorlou, Jadbabaie and Kakhbod (2018).[5] Scarcity of the products with regard to the number of buyers (e.g. Gallego and van Ryzin (1994), Gershkov and Moldovanu (2009)), network externalities (e.g. Cabral, Salant and Woroch (1999)), stochastic incoming demand (e.g. Board (2008)), and time-varying values of buyers (e.g. Garrett (2013), Stokey (1979,9)) are among other causes suggested in the literature for varying prices over time. In contrast to all the works above, the reason for time-varying prices is to learn about consumers' valuations. Particularly, Bergemann and Välimäki (2000) study the duopolistic case with symmetric consumers. The main twist with respect to this paper is that our focus is on the heterogeneity between consumers. In particular, here, consumers differ in the precision of their reviews of the product. This is crucial and leads to rich predictions about how sellers discount and price discriminate between consumers based on their learning technologies. Moreover, it allows us to explore different questions like the relative efficiency performance of monopoly and competition, and how the inefficiency depends on the heterogeneity of the buyers.

Our paper is linked with works that study big data and the use-based evolution of beliefs about the quality of a product. Related questions to this type of belief dynamics have been addressed in different frameworks in several important papers (e.g., Bolton and Harris (1999), Décamps, Mariotti and Villeneuve (2006), Papanastasiou, Bimpikis and Savva (2018)). For example, Park (2001) observed possible linkage between learn-

[4] Talluri and van Ryzin (2004) and Phillips (2005) provide an extensive review of this topic.

[5] These models are typically either two-sided or one-sided. For example, Ifrach et al. (2019) and Yu, Debo and Kapuscinski (2013) consider two-sided learning models where buyers and sellers both learn the true value of a new product through consumer experiences. Papanastasiou, Bakhshi and Savva (2013) and Crapis et al. (2017) analyze one-sided learning models when the firm knows the product quality, buyers report their experiences and subsequent customers learn from these reports.

ing asymmetries and efficiency, but he does not study neither when it is likely nor the form of the inefficiency. The idea that platforms can aggregate information is linked to the literature on markets for big data (e.g., Admati and Pfleiderer (1986) and Begenau, Farboodi and Veldkamp (2018)) and mechanisms for pricing information (e.g., Anton and Yao (2002), Babaioff, Kleinberg and Paes Leme (2012), Eső and Szentes (2007), and Eliaz, Eilat and Mu (2019)).[6] In contrast to these important works, we consider how big data (availability of information through heterogeneous sources) affect welfare, market power (monopoly and competition), trading volume, Beta phase, and the nature of arising distortions in product markets. We further present policies that can effectively reduce distortions. In this regard, this paper also relates to the body of works on heterogeneous learning in financial markets. However, the nature of the asymmetry is different, because most of the attention has been dedicated to heterogeneity in beliefs. Most notably, Scheinkman and Xiong (2003) study asset prices, trading volume, and price volatility during episodes of asset price bubbles, Gennaioli and Shleifer (2018) study how investors and policymakers assign irrationally and inaccurately low probabilities to disaster outcomes leading to financial fragility, Caballero and Simsek (2019) study effect of monetary policy on financial stability, Veronesi (2019) considers general distributions of households' risk tolerance and beliefs about long-term growth.

Our paper is also related to the growing literature studying the role of online platforms sharing information, see Kremer, Mansour and Perry (2014). For example, Acemoglu et al. (2019), like us, single out an externality induced by some consumers on others. However, the rationale there is that there is partial overlapping in the private information of the different consumers, and the information provided by one consumer depresses the value of the one of the other.

Finally, this paper is also related to the growing literature on studying innovation, strategic pricing and externalities: for example, strategic information exchange (e.g., Acemoglu, Bimpikis and Ozdaglar (2014)) optimal static pricing under presence of local externalities (e.g., Sundararajan (2008), Hartline, Mirrokni and Sundararajan (2008), Candogan, Bimpikis and Ozdaglar (2012), Jadbabaie and Kakhbod (2019)), optimal advertising (e.g., Galeotti and Goyal (2009)), and experimentation with technology innovation

[6]See Acquisti, Taylor and Wagman (2016), Bergemann and Bonatti (2019) for an excellent surveys of different aspects of this literature.

(e.g., Acemoglu, Bimpikis and Ozdaglar (2011)).[7]

The rest of the paper proceeds as follows. Section 2 introduces our formal model, and Section 3 studies the first best consumption allocation. In Section 4 we move to the analysis of the decentralized outcome and present our main results. Section 5 propose several extensions of the basic model, and explore the robustness of our findings. Finally, Section 6 concludes. All the proofs are in the Appendix.

2 Model

We consider a product market where buyers (consumers) face two indivisible products. Product a is a known and established commodity that creates a known flow of payoff for buyers. However, product b is recently introduced and its true (expected flow payoff) value is unknown to both sellers (retailers) and buyers. That is, the consumption utility of product b depends on an unknown state

$$\theta \in \{\ell, h\},$$

where the states denote expected flow utility of the new product.

There are $M \in \{1, 2\}$ sellers and $n \geq 2$ possibly asymmetric buyers. When $M = 1$ a profit maximizing **monopolist** sells both of the products. In the **oligopoly** structure, i.e., $M = 2$, two different sellers compete strategically to sell the products, i.e., one seller sells the new product and the other one sells the established product. In this case, we will label each of the seller as the product he markets. Although most of the analysis focus on the comparison between duopoly and monopoly, Proposition 11 shows how our results about the relative inefficiency of competition are confirmed in a setting with multiple competitors.[8]

The key elements of the model are as follows.

[7] See also Su (2007) for dynamic pricing with strategic customers. However, the price paths are found to be monotone in these works.

[8] A different market structure in which one of the two sellers markets both products is explored in Section 5.

2.1 Buyers asymmetry and flow payoffs

We assume that each product sold in time t will survive in $[t, t+dt)$ and generates the following flow of payoff for its buyer(s). Precisely, at $[t, t+dt)$ the established product a creates

$$dC_{ai}(t) = \mu_a dt,$$

monetary value to its buyer, and in state $\theta \in \{h, \ell\}$ the new product b generates the flow of payoff

$$dC_{bi}(t) = \theta dt + \sigma_i dZ_{it},$$

for buyer i (if he owns it), where Z_{it}, $i = 1, 2, \cdots, n$, are independent standard Brownian motions (BMs) (Wiener processes). We assume that the problem is not trivial, that is,

$$\ell < \mu_a < h,$$

and that the unknown product induces a non-negative flow of utility in both states, i.e., $\ell \geq 0$. Therefore, the state θ determines the objectively preferable product, that is the same for every buyer. The state is initially unknown to all the market participants, and the buyers and the sellers share a common prior $\Pr\{\theta = h\} = \pi_0$ at time 0, i.e., the initial time of offering for the product of unknown quality.

Importantly, we let $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n$. This means that buyers are possibly **asymmetric** in processing information if they buy the risky product b. In this regard, we order buyers from the most noisy observations of buyer 1 to the least noisy observations of buyer n. We say that a buyer is a better learner the lower σ is. Clearly, when $\sigma_i = \sigma$ for all $i = 1, 2, \cdots, n$, then buyers are **symmetric**. To isolate the effect of heterogeneous learning technology (i.e., heterogenous σ_i), we assume that the buyers are otherwise identical, and in particular, they share the same valuation for the product of unknown quality in both states of the world.[9]

This heterogeneity in the learning technology for the new product is natural for different reasons. A first simple motivation is that buyers significantly differ in terms of

[9] Our main results would continue to hold as long as the valuation of the different buyers are correlated. However, if they were independent, the information externality we single out below would disappear, and the equilibrium under competition would be efficient.

the depth of their experience for the product: even if all of them care about using the best product, some consumers may use more frequently some ancillary features, generating knowledge about the product more rapidly. The fact that those different experiences translate into the public belief is particularly realistic in current markets with frequent and detailed feedback about consumer experiences through surveys or posted reviews.

2.2 Trading volume and payoffs

At each time period $[t, t+dt)$, a buyer at most uses (experiences) one product. Hence, at the beginning of period t the order of buyer i is in $\{a, b, \emptyset\}$, where $\emptyset$ means that i does not hold any product in $[t, t+dt)$. We denote by $\xi_{ik}(t)$ the (measurable) allocation process such that $\xi_{ik}(t) = 1$ if buyer i purchases product k in period t, and it is $\xi_{ik}(t) = 0$ otherwise. Therefore, the trading volume of product $k \in \{a, b\}$ at time t is

$$\mathrm{Vol}_k(t) = \sum_{i=1}^{n} \xi_{ik}(t).$$

Sellers have all the bargaining power, i.e., offers are in take-it-or-leave-it forms. At the beginning of period t, the price of product k posted by its seller is $p_k(t)$, for $k \in \{a, b\}$.

Both buyers and sellers are risk-neutral and forward-looking. They discount payoffs exponentially at a shared rate $\rho > 0$. Therefore, the payoff of buyer i is given by

$$U_i^B = \mathbf{E}\left[\int_0^\infty \rho e^{-\rho t} \underbrace{\xi_{ik}(t)}_{\text{order}} \Big(\underbrace{dC_{ki}(t)}_{\text{flow gain}} - \underbrace{p_{ki}(t)}_{\text{payment}} \, dt \Big) \right], \tag{1}$$

where $\mathbf{E}[\cdot]$ denotes the expectation operator. Without loss of generality, we normalize the production cost to 0 so that the payoff of the sellers is equal to the total revenues they obtain from the products they sell.

Importantly, how we compute these revenues depends on the market structure (monopoly v.s. oligopoly). Below we present the expected discounted payoffs in the two cases considered in the paper.

Monopoly. When there is a unique seller of both products his payoff is given by

$$U_m = \mathbf{E}\left[\int_0^\infty \rho e^{-\rho t}\underbrace{\left(\sum_{i=1}^n \xi_{ia}(t)p_{ia}(t) + \sum_{i=1}^n \xi_{ib}(t)p_{ib}(t)\right)}_{\text{overall time } t \text{ monopoly profit (sale)}} dt\right]. \tag{2}$$

Under competition, the objective function of the two competing sellers is analogous, but it takes into account that each of them benefits only from his sales.

Oligopoly. Under oligopoly, the payoff of seller $k \in \{a,b\}$ is given by

$$U_k^S = \mathbf{E}\left[\int_0^\infty \rho e^{-\rho t}\underbrace{\sum_{i=1}^n \xi_{ik}(t)p_{ik}(t)}_{\substack{\text{time } t \text{ seller } k\\ \text{profit (sale)}}} dt\right]. \tag{3}$$

2.3 Belief dynamics

At each time t all the data about the buyers' flow of payoffs are public information. Therefore, even when the amount of information produced by different buyers is different, there is a unique market belief about the type of the unknown product. Formally, it is denoted as

$$\pi_t := \Pr\left\{\theta = h | \mathcal{F}_t\right\},$$

where $\mathcal{F}_t$ is the (filtration generated by all the) information available up to time t. The following lemma characterizes the dynamics of the market belief in terms of the (endogenous) trading volume and learning technologies (i.e., σ_i) of buyers.

Lemma 1. *[Belief Evolution] We have*

$$d\pi_t = \pi_t(1-\pi_t)(h-\ell)\sqrt{\sum_{i=1}^n \frac{\xi_{ib}(t)}{\sigma_i^2}}\, dZ_t$$

where Z_t is a standard Wiener process. In particular, in the case of symmetric buyers we have

$$d\pi_t = \frac{\pi_t(1-\pi_t)(h-\ell)}{\sigma}\sqrt{Vol_b(t)}dZ_t.$$

Moreover, π is a continuous martingale with respect to $\{\mathcal{F}_t, t \geq 0\}$ and a strong Markov process that is symmetric in time.

Proof. See Appendix. □

The dynamics are quite intuitive; beliefs are more volatile when the information provided by the consumers is less precise, and when the market participants are less sure to start with (i.e., π is closer to $\frac{1}{2}$).

2.4 Learning Progression

How does market belief improve over time? What factors are important? In this environment we can characterize the expected improvement in the market belief about unknown state θ to a target belief β obtained when only an arbitrary subset of buyers $M = \left\{m_1, ..., m_j\right\}$ use product b for the fixed span of time T starting from confidence π_0. More formally,

$$MO\left(\pi_0, \beta, T, m_1, ..., m_j\right) = \mathbf{E}_{\pi_0}\left[\max\{\pi_T - \beta, 0\} | \forall t \in [0,T], \forall i \in \{1,n\}, \xi(t) = 1_{i \in M}\right].$$

The above expression introduces a natural way to measure the expected improvement in market optimism (MO) due to the experiences of buyers $\left\{m_1, ..., m_j\right\}$ of the unknown quality product b up to time T.

The following result characterizes $MO\left(\pi_0, \beta, T, m_1, ..., m_j\right)$ in terms of the learning abilities of $\left\{m_1, ..., m_j\right\}$, i.e., $\sigma_{m_1}, \cdots, \sigma_{m_j}$, the target belief β, and the horizon T.

Proposition 1. *Let buyers $\left\{m_1, ..., m_j\right\}$ use the risky product b in the time interval $[0,T]$. Then, the expected progression in the market belief π_t to the target belief β is explicitly given by:*

$$MO\left(\pi_0, \beta, T, m_1, ..., m_j\right) = (1-\beta)\pi_0\Phi(\lambda_1) - \beta(1-\pi_0)\Phi(\lambda_0)$$

where

$$\lambda_1 = \frac{1}{(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{T}{\sigma_{m_i}^2}}}\left[\ln\left(\frac{\frac{\pi_0}{1-\pi_0}}{\frac{\beta}{1-\beta}}\right)+\frac{(h-\ell)^2}{2}\left(\sum_{i=1}^{j}\frac{T}{\sigma_{m_i}^2}\right)\right]$$

$$\lambda_0 = \frac{1}{(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{T}{\sigma_{m_i}^2}}}\left[\ln\left(\frac{\frac{\pi_0}{1-\pi_0}}{\frac{\beta}{1-\beta}}\right)-\frac{(h-\ell)^2}{2}\left(\sum_{i=1}^{j}\frac{T}{\sigma_{m_i}^2}\right)\right]$$

and $\Phi(\cdot)$ *denotes the CDF of a standard normal random variable.*

Proof. See Appendix. □

To obtain intuition about the result, suppose $\beta = \pi_0$. The above proposition implies the following result characterizing the expected progression in the market belief π_t to the initial market belief π_0. Moreover, it immediately shows intuitive comparative statics on this measure. Particularly, and somewhat expectedly, MO is increasing in the horizon time T, the learning quality of buyer m_k, and the number of buyers experiencing the product b.

Corollary 1. *Let buyers* $\{m_1,...,m_j\}$ *use the risky product* b *in the time interval* $[0,T]$. *Then, the expected progression in the market belief* π_t *from the initial market belief* π_0 *is:*

$$MO(\pi_0,\pi_0,T,m_1,...,m_j) = \pi_0(1-\pi_0)\left(2\Phi\left(\frac{(h-\ell)}{2}\sqrt{\sum_{i=1}^{j}\frac{T}{\sigma_{m_i}^2}}\right)-1\right). \tag{4}$$

Moreover, $MO(\pi_0,\pi_0,T,m_1,...,m_j)$ *is increasing in the horizon time* T, *the learning quality of buyer* m_k. *i.e.,* $\frac{1}{\sigma_{m_k}}$, *and the number of buyers experiencing the product* b.

Next, we leverage the previous results on the learning process of the agents for a *given* pattern of consumption to study optimal patterns of consumption. First, we study the optimal consumption pattern for a planner who wants to maximize the sum of the utility of the market participants. Then, we move to consider the decentralized equilibrium outcome that arises when each single market participant best replies to the strategy of the opponents, and we explore the difference between these two situations.

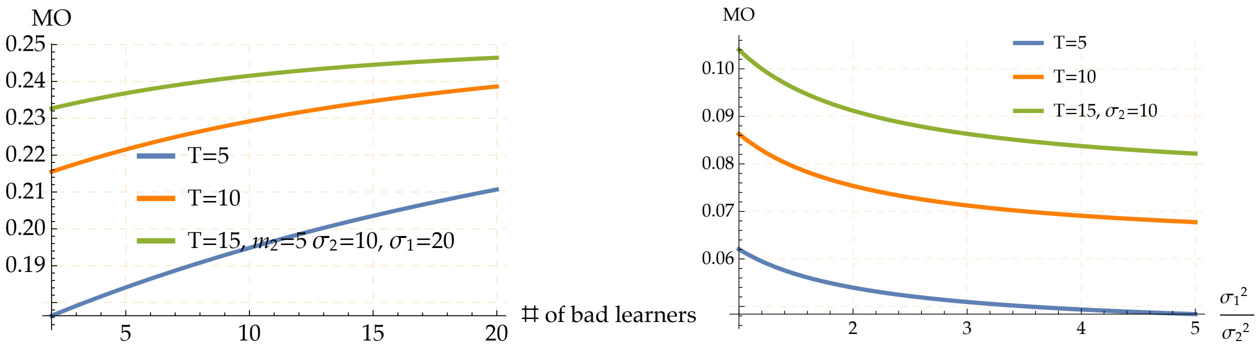

Figure 1: The left panel $MO(\pi_0, \pi_0, T, m_1, ..., m_j)$ is increasing in T and the number of buyers experiencing the product b. The right panel shows $MO(\pi_0, \pi_0, T, m_1, ..., m_j)$ is increasing in the quality of learning.

3 The first best— efficient strategies

The first best formulation. In this section, we consider the social welfare-maximizing strategies; that is, we specify strategies that maximize the sum of the utilities of all market participants. Given buyers' and sellers' payoffs (see (1)-(3)), the payments cancel out each other in the welfare-maximization problem. As a result, the objective function is the discounted sum of the consumption utility of the buyers:

$$\max_{\substack{\xi_i:\ i=1,2,\cdots,n,\\ p_k,\ k\in\{a,b\}}} \left(\sum_{i=1}^{n} U_i^B + \sum_{k\in\{a,b\}} U_k^S \right) = \max_{\xi_i:\ i=1,2,\cdots,n} \mathbf{E}\left[\sum_{i=1}^{n} \sum_{k\in\{a,b\}} \int_0^\infty \rho e^{-\rho t} \xi_{ik}(t) dC_k(t) \right].$$

Therefore, efficiency only depends on the consumption of each agent, regardless of the transfers. Given that the system is time-invariant, the optimal ξ_t only depends on the belief π_t, and the maximization can be mapped into a stopping time problem (see, e.g., Øksendal (2003)). That is, the efficient allocation is pinned down by a simple sequence of cutoffs on the market belief that we denote it by π_{fb} so that each consumer buys product a if and only if $\pi_t < \pi_{\text{fb,i}}$. In other words:

$$\xi_{ib}(t) = 1_{\{\pi_t \geq \pi_{\text{fb,i}}\}} \quad \forall\ i = 1, 2 \cdots, n, \tag{5}$$

where $1_{\{A\}}$ is the indicator function on A. With this, the Hamilton-Jacobi-Bellman (HJB) equation for this problem is given by :

$$W(\pi) = \max_{\xi_i:\ i=1,2,\cdots,N} \left\{ \sum_{i=1}^{N} (\xi_{ia}\mu_a + \xi_{ib}\mathbf{E}_\pi[\mu_b]) + W''(\pi) \sum_{i=1}^{N} \xi_{ib} \frac{g(\pi,h,\ell)}{\rho\sigma_i^2} \right\}$$

where $g(\pi,h,\ell) = \left((h-\ell)\pi(1-\pi)\right)^2$.

3.1 Symmetric buyers

In the case of symmetric buyers, by imposing the standard value matching and smooth pasting conditions (see, e.g., Dixit (1993)), we can obtain an explicit formula for the optimal cutoff π_{fb} such that every consumer buys product b if and only if $\pi \geq \pi_{\text{fb}}$. The following result summarizes.

Proposition 2. *The first-best (social welfare) maximizing cutoff is given by*

$$\pi_{\text{fb}} = \frac{(\mu_a - \ell)(\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}} - 1)}{(\ell + h) - 2\mu_a + (h-\ell)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}. \tag{6}$$

Moreover, π_{fb} is increasing in μ_a, σ^2, and ρ. It is decreasing in h and n.

Proof. See Appendix. □

It is interesting how the first-best cutoff changes with the fundamentals. First, since the known product acts as an outside option, a higher μ_a is easily seen to induce a higher π_{fb}. On the other hand, a larger h increases the value of choosing alternative b through two channels. First, it increases the instantaneous value given a particular belief; second, it increases the learning value by making $(h-\ell)$ larger. Therefore, it unambiguously induces a lower π_{fb}. The effect of a larger ℓ is instead ambiguous: it reduces the value of experimentation, but it makes the instantaneous reward of choosing b larger.

The effects of the information processing technology, the discount factor, and the number of buyers are unambiguous and intuitive. The larger is σ^2 (or ρ), the less attractive experimentation and the higher is π_{fb}. Conversely, more patient buyers stop experimenting at a lower π_{fb}. Finally, notice that the existence of the public-platform makes

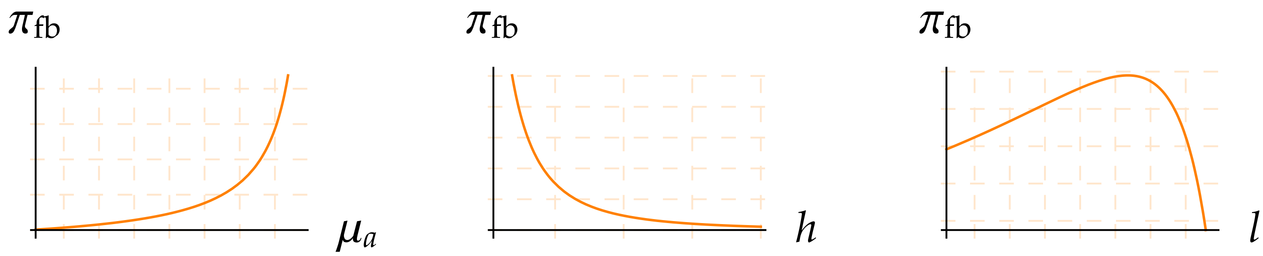

Figure 2: The first-best belief threshold π_{fb} is increasing in the drift μ_a of the known product, and decreasing in the maximum feasible drift h and non-monotone in the minimum feasible drift ℓ of the unknown product.

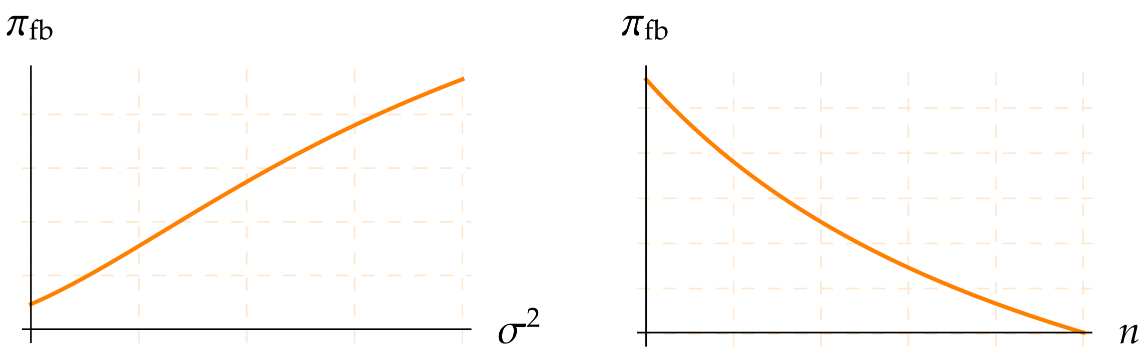

Figure 3: The first-best belief threshold π_{fb} is increasing in the volatility σ^2 of the unknown product and decreasing in the number of buyers n.

the benefit of learning linearly increasing in the number of buyers, and therefore a larger n induces a lower π_{fb}.

3.2 Welfare

We can use the derived optimal strategy and cutoff value to compute the total welfare of the agents. The following proposition provides an explicit formula.

Proposition 3. *The first-best social welfare is equal to $nW^*(\pi)$ where*

$$W^*(\pi) = \begin{cases} \mu_a & \text{if } 0 \le \pi \le \pi_{\text{fb}}; \\ \mathbf{E}_\pi[\theta] + \varphi\left[\pi^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}(1-\pi)^{\frac{1}{2}\left(\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}+1\right)}\right] & \text{if } \pi_{\text{fb}} < \pi \le 1. \end{cases}$$

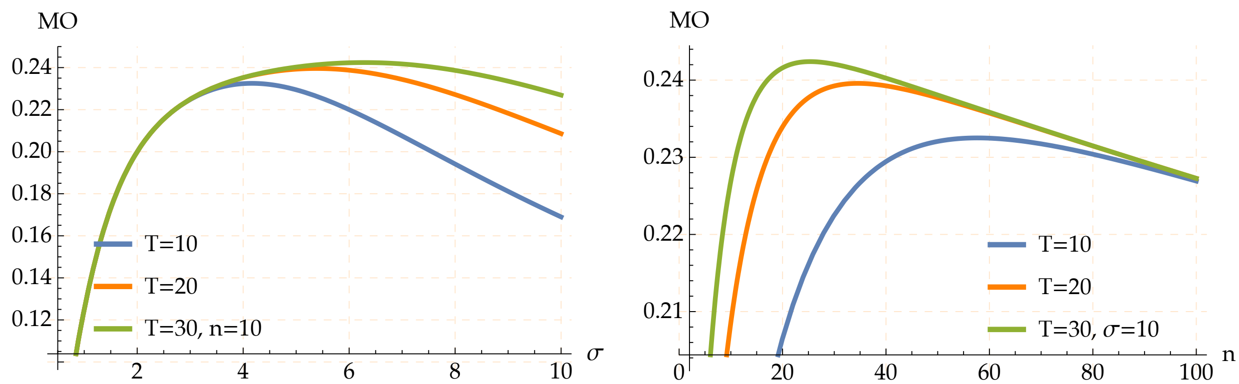

Figure 4: These panels show how the market optimism $MO(\pi_{\text{fb}}, \pi_{\text{fb}}, T)$ evolves with changing the volatility σ and the number of of buyers for different horizons $T = 10, 20$ and 30. As is shown market optimism is non-monotone in changing σ and n.

and

$$\varphi := 2\frac{h-\mu_a}{\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}-1}\left(\frac{\pi_{\text{fb}}}{1-\pi_{\text{fb}}}\right)^{\frac{1}{2}\left(\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}+1\right)}.$$

Moreover, W^* is strictly convex in $\pi_{\text{fb}} < \pi \leq 1$.

Proof. See Appendix. □

Intuitively, the welfare function has a first flat part before the optimal cutoff. The per-consumer value there it is equal to the flow of payoff guaranteed by the known product bought by all the consumers. After the cutoff the value increases in the probability assigned to the high quality of the unknown product. There, the convexity of the welfare is due to the known product that acts as an outside option whenever the market participants start to be too confident that the unknown product has a low quality.

Next, we move to the more interesting case of different learning technologies for the buyers. There, even if we cannot compute the optimal (sequence) of cutoffs, we can study some qualitative properties of the welfare-maximizing consumption pattern that we can later use to draw comparisons with the decentralized market outcome.

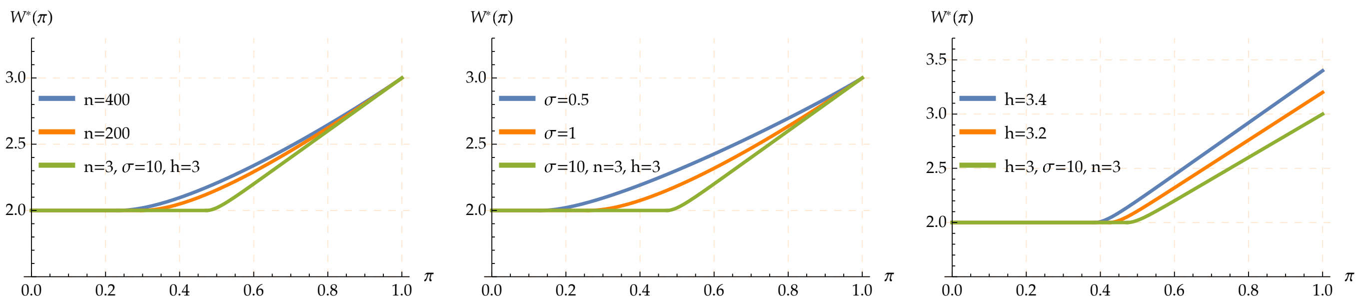

Figure 5: The first-best social welfare $nW^*(\pi)$ is weakly convex. In addition, $W^*(\pi)$ is strictly convex in $\pi_{\text{fb}} < \pi \leq 1$, where π_{fb} is at the kinks in these panels. Moreover, the panels (from left to right, respectively) show that the social welfare increases in the number of buyers n, the learning quality of buyers $\frac{1}{\sigma}$, and the extent of the high type h of product b. The other fixed parameters are $\ell = 1, \mu_a = 2, \rho = 5$.

3.3 Asymmetric buyers

We consider the case of 2 buyers. The extension to n buyers is straightforward and does not provide additional insights. Recall that we have $\sigma_1 > \sigma_2$. It is useful to consider the average (per consumer) first-best welfare, that is given by

$$W(\pi) = \max_{\xi_{1b}, \xi_{2b}} \left\{ \mu_a + \sum_{i=1}^{2} \frac{\xi_{ib}}{2} \left(\mathbf{E}_\pi [\mu_b] - \mu_a + \frac{g\big(\pi, h, \ell\big) W''(\pi)}{\rho \sigma_i^2} \right) \right\}. \tag{7}$$

Clearly, a policy ξ_{1b}^*, ξ_{2b}^* is optimal if and only if it maximizes the average utility:

$$\xi_{1b}^*, \xi_{2b}^* \in \arg\max_{\xi_{1b}, \xi_{2b}} \left\{ \mu_a + \sum_{i=1}^{2} \frac{\xi_{ib}}{2} \left(\mathbf{E}_\pi [\mu_b] - \mu_a + \frac{g\big(\pi, h, \ell\big) W''(\pi)}{\rho \sigma_i^2} \right) \right\}.$$

Our line of attack consists in conjecturing a solution of the form:

$$\xi_{ib} = \begin{cases} \xi_{1b} = \xi_{2b} = 1 & \text{if } \frac{\mu_a - \mathbf{E}_\pi[\mu_b]}{2} \leq \frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_1^2} W''(\pi), \\ \xi_{2b} = 1 = 1 - \xi_{1b} & \text{if } \frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_1^2} W''(\pi) \leq \frac{\mu_a - \mathbf{E}_\pi[\mu_b]}{2} < \frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_2^2} W''(\pi), \\ \xi_{1b} = \xi_{2b} = 0 & \text{if } \frac{\mu_a - \mathbf{E}_\pi[\mu_b]}{2} > \frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_2^2} W''(\pi). \end{cases}$$

The structure of this candidate optimal policy is simple: There are two thresholds $\pi_{\text{fb},1}$ and $\pi_{\text{fb},2}$ with $0 < \pi_{\text{fb},2} < \pi_{\text{fb},1} < 1$, with both buyers purchasing the risky product b when the market belief is sufficiently high, no buyer purchasing b when the market belief is sufficiently low, and only the buyer with the best learning technology purchasing the risky product for intermediate beliefs. The intuitive reason why the threshold for the better learner is lower is simple. First, note that information is valuable for overall welfare because it allows better consumption choice for the consumers (this mathematically translates into the convexity of the value function). Second, the higher signal precision of the better learner implies that she can trade-off exploitation in favor of information generation at more favorable terms, and therefore it is optimal to start to do so for more pessimistic beliefs.

Given the conjectured policy function, we derive the value function, and we check that it is indeed in cutoff strategies. The following result summarizes.

Lemma 2. *The first best policy is in cutoff strategies with* $\pi_{\text{fb},2} < \pi_{\text{fb},1}$.

In the next sections of the paper, we compare these first-best cutoffs with the ones obtained under the two different competition structures presented above: a monopolist selling both products and the competition between two different sellers. Before doing so, we explore the determinants and comparative statics of the particular transition period, called Beta phase, that it is peculiar of our model with asymmetries in the learning technologies.

3.4 Beta phase

In the Beta phase, only the subset of best learners buys the new product. In this section, we explicitly characterize the expected length of the Beta phase and the probability of leaving the Beta phase. What is the expected waiting time of leaving the Beta phase? Leaving the Beta phase means either the market belief becomes sufficiently optimist (i.e., $\pi > \pi_{\text{fb},1}$) so that all buyers start buying the new product, or the market belief sufficiently reduces (i.e., $\pi < \pi_{\text{fb},2}$) for which all buyers discard the new product as a failure. What factors affect the expected length of the Beta phase? Moreover, how do these factors affect the probability of leaving the Beta phase? The following proposition explicitly addresses these questions.

Proposition 4. *Let $\pi_{\text{fb},2} < \pi_0 < \pi_{\text{fb},1}$. Define $\sigma(y) := \frac{y(1-y)(h-\ell)}{\sigma_2}$. Then,*

$$\begin{aligned}\mathbf{E}_{\pi_0}[\text{Beta phase}] &= \mathbf{E}_{\pi_0}[\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}] \\ &= \frac{\pi_{\text{fb},1} - \pi_0}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \int_{\pi_{\text{fb},2}}^{\pi_0} (y - \pi_{\text{fb},2}) \frac{2dy}{\sigma^2(y)} + \frac{\pi_0 - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \int_{\pi_0}^{\pi_{\text{fb},1}} (\pi_{\text{fb},1} - y) \frac{2dy}{\sigma^2(y)}.\end{aligned}$$

Particularly, $\lim_{\pi_0 \searrow \pi_{\text{fb},2}} \partial_{\pi_0} \mathbf{E}_{\pi_0}[\text{Beta phase}] > 0$ *and* $\lim_{\pi_0 \nearrow \pi_{\text{fb},1}} \partial_{\pi_0} \mathbf{E}_{\pi_0}[\text{Beta phase}] < 0$. *Moreover*

- *The probability of discarding the new product as a failure is*

$$\begin{aligned}Pr_{\pi_0}\{\text{discarding}\} &= Pr_{\pi_0}\{\pi\Big(\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}\Big) = \pi_{\text{fb},2}\} \\ &= \frac{\pi_{\text{fb},1} - \pi_0}{\pi_{\text{fb},1} - \pi_{\text{fb},2}}.\end{aligned}$$

- *The probability that the new product serves the whole market*

$$\begin{aligned}Pr_{\pi_0}\{\text{serving the whole market}\} &= Pr_{\pi_0}\{\pi\Big(\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}\Big) = \pi_{\text{fb},1}\} \\ &= \frac{\pi_0 - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}}.\end{aligned}$$

Particularly, $\partial_{\pi_0} Pr_{\pi_0}\{\textit{discarding}\} < 0$, *and* $\partial_{\pi_0} Pr_{\pi_0}\{\textit{serving the whole market}\} > 0$.

Proof. See Appendix. □

The above result explicitly characterizes the expected Beta phase and the probabilities that the new product either serves the whole market or it is discarded as a failure in terms of the endogenous thresholds $\pi_{\text{fb},1}$ and $\pi_{\text{fb},2}$. Moreover, it produces intuitive comparative statics based on the initial market belief about the new product. Particularly, the expected length of the Beta phase increases in the initial market belief when market belief is initially sufficiently small (i.e., $\lim_{\pi_0 \searrow \pi_{\text{fb},2}} \partial_{\pi_0} \mathbf{E}_{\pi_0}[\text{Beta phase}] > 0$) and the expected length of the Beta phase decreases in the initial market belief when the market belief is initially sufficiently large (i.e., $\lim_{\pi_0 \searrow \pi_{\text{fb},2}} \partial_{\pi_0} \mathbf{E}_{\pi_0}[\text{Beta phase}] > 0$). Moreover, the probability of discarding the new product as a failure decreases in the initial market belief and the probability that the new product starts to serve the whole market increases in the initial market belief.

We now have a complete first-best benchmark that we can compare to the decentralized outcome we characterize in the next section.

4 Analysis: Decentralized Outcome

We aim to specify equilibrium strategies when buyers are symmetric and asymmetric in their learning technology σ_i. Notably, we aim to discover how the market structure and competition between sellers affect learning, trading volume, and efficiency.

To analyze this model we restrict our attention to Markov perfect equilibria. Given the timing of offers, the pricing strategy of seller k is a measurable function from the belief space to the real numbers $p_{k,i} : [0,1] \to \mathbb{R}^N$, and the purchasing strategy of buyer i is a pair of measurable functions $\xi_i = (\xi_{ia}, \xi_{ib}) : [0,1] \times \mathbb{R}^2 \to \{0,1\}^2$.

We state the relevant equilibrium notion for the case of competition, an analogous definition that takes into account that the choice variable of the monopolist has two dimensions is used for the study of the monopoly case.

Definition 1. *A collection of strategies (ξ^*, p^*) is a Markov Perfect Equilibrium if*

$$\begin{array}{rcll} \forall k & \in & \{1,2\} \forall i \in \{1,2\}, \forall p_{k,i} \in \mathbb{R}, \forall \pi & U_k^S(p^*, \xi^*, \pi) \geq U_k^S\left(p_{k,i}, p^*_{-k,i}, \xi^*, \pi\right) \\ \forall i & \in & \{1,2\}, \forall \xi_i \forall \pi & U^B(p^*, \xi^*, \pi) \geq U^B\left(p^*, \xi^*_{-i}, \xi_i, \pi\right). \end{array}$$

Under both market structures (monopoly, competition), we allow the seller to discriminate between buyers depending on their learning technology. This is a natural assumption in situations in which the sellers can provide discounts to the buyers that report their experiences with a particular product in a more detailed way through surveys or reviews. Therefore, the equilibrium choice of the monopolist and each of the two sellers under competition are respectively equal to $(p_{m,a,1}, p_{m,a,2}, p_{m,b,1}, p_{m,b,2})$ under monopoly, $(p_{a,1}, p_{a,2})$, and $(p_{b,1}, p_{b,2})$ under competition.

Section 5 explores the robustness of our results to the price discrimination hypothesis. For the moment, we notice that the negative result we are going to obtain in case of competition is only reinforced by the assumption that the seller can apply price discrimination. Indeed, it is well known that even in static markets, the combination of market power for the seller, asymmetric consumers, and impossibility to discriminate between consumers create inefficiencies, the most prominent example being a monopolistic market with a continuous and strictly increasing demand function. However, such inefficiencies are usually avoided when the seller can discriminate. Instead, we are going to show that in dynamic markets with learning, discrimination is not enough to eliminate inefficiencies.

We start with the analysis of the monopolistic equilibrium.

4.1 Monopoly

We start by proving that the revenue-maximizing policy of a monopolist is *efficient* independently of the learning technologies. To prove this result, we first derive the buyers and the monopolist's HJB equations.

Recall that we are assuming that the sellers, and in this particular case the monopolist, have all the bargaining power with offers in take-it-or-leave-it forms. Therefore, the HJB equation of a buyer i captures the comparison between the two products at the posted prices:[10]

$$v_i(\pi) = \max\Bigg\{ \underbrace{\mu_a - p_{a,i}(\pi)}_{\text{flow gain of } a} + \underbrace{g\Big(\pi,h,\ell\Big)\sum_{j\neq i}\frac{\xi_{jb}\Big(\pi,p_{a,j},p_{b,j}\Big)}{2\rho\sigma_j^2}v_i''(\pi)}_{\text{learning gain from others buying } b},$$

$$\underbrace{\mathbf{E}_\pi[\mu_b] - p_{b,i}}_{\text{flow gain of } b} + \underbrace{g\Big(\pi,h,\ell\Big)\Bigg(\sum_{j\neq i}\frac{\xi_{jb}\Big(\pi,p_{a,j},p_{b,j}\Big)}{2\rho\sigma_j^2} + \frac{1}{2\rho\sigma_i^2}\Bigg)v_i''(\pi)}_{\text{learning gain from } i \text{ and others buying } b}\Bigg\}. \tag{8}$$

Each term of the above HJB equation has two parts: If buyer i buys the product a of known quality then $\mu_a - p_a$ is the instant (expected) flow payoff, and $g\Big(\pi,h,\ell\Big)\sum_{j\neq i}\frac{\xi_{jb}(t)}{2\rho\sigma_j^2}v_i''(\pi)$ is

[10] We note that the HJB solution to continuous functions can also be extended to the weaker viscosity solutions as well, see Appendix A.

the expected continuation payoff (which is due to learning). The buyer can learn even if he does not choose the unknown product b, since the platform shares the reports from other buyers. A similar decomposition holds when buyer i buys the risky product b of unknown quality, but then the trading volume of the risky product b is increased by 1 and his (expected) flow payoff becomes $\mathbf{E}_\pi[\mu_b] - p_b$.

The monopolist's HJB equation can be obtained similarly in terms of the behavior of the buyer:

$$w_m(\pi) = \sup_{p_a, p_b} \left\{ \sum_{i=1}^{n} \left(\frac{\xi_{ib}(\pi, p_{a,i}, p_{b,i})}{2\rho\sigma_i^2} p_{b,i} + \frac{\xi_{ia}(\pi, p_{a,i}, p_{b,i})}{2\rho\sigma_i^2} p_{a,i} + g\big(\pi, h, \ell\big) \frac{\xi_{ib}(\pi, p_{a,i}, p_{b,i})}{2\rho\sigma_i^2} w_m''(\pi) \right) \right\}. \tag{9}$$

Next, we explore the case of symmetric and asymmetric learning technologies. Under the monopolistic structure we are currently analyzing the two cases lead to similar welfare conclusions, but we keep them separate because of the critical difference they feature under competition.

4.1.1 Symmetric buyers

We start characterizing the equilibrium prices posted by the monopolist in a cutoff equilibrium when buyers are symmetric.

Lemma 3. *In every equilibrium with symmetric π_m^*, the prices are as follows. If $\pi < \pi_m^*$*

$$p_a = \mu_a \qquad \textit{and} \qquad p_b \geq \mathbf{E}_\pi[\theta] - \mu_a. \tag{10}$$

If $\pi \geq \pi_m^*$

$$p_a \geq \mu_a - \mathbf{E}_\pi[\theta] \qquad \textit{and} \qquad p_b = \mathbf{E}_\pi[\theta]. \tag{11}$$

To prove this result, as a preliminary observation, notice that the problem of the monopolist reduces to the choice between either selling the product of known quality or selling the one whose quality is unknown. Indeed, given this choice, there is no reason to charge less than the maximal willingness to pay of buyers, since the informational content generated by the use of the product is not affected by the price. Therefore, if $\pi < \pi_m^*$ we have $p_a = \mu_a$. Since by the definition of cutoff equilibrium the monopolist is selling the product of known quality at those beliefs, it immediately follows from the value function of the buyer that $p_b \geq \mathbf{E}_\pi[\theta] - \mu_a + v''(\pi) = \mathbf{E}_\pi[\theta] - \mu_a$. Similarly, if $\pi \geq \pi_m^*$ the monopolist sets a price of product b equal to the willingness to pay $p_b = \mathbf{E}_\pi[\theta] + v''(\pi)$, and a price of product a such that product b is sold. That is, $p_a \geq \mu_a - p_b = \mathbf{E}_\pi[\theta] - v''(\pi)$. However, if we plug these prices in the value function of the buyers, we obtain that v is identically equal to 0, and so is its second derivative v''. This, together with the previously computed prices, imply the result.

Next, how do we find the threshold π_m^*? Given the pricing strategy of the monopolist characterized in Lemma 3, we know the value function of the monopolist for beliefs below the threshold. Therefore, to find the threshold, we combine a smooth pasting and value matching conditions with the second-order ODE given by the diffusion process derived in Lemma 1. The main take away is that the monopolist who sells both products chooses which product to deliver to the market using the same belief threshold as in the welfare-maximizing benchmark, i.e., $\pi_m^* = \pi_{\text{fb}}$. As a result, a monopoly achieves efficiency, summarized by the following proposition.

Proposition 5. *The followings hold: (i) Any symmetric equilibrium is specified by a cutoff π_m^*. (ii) All the symmetric equilibria are efficient (i.e., welfare-maximizing), and we have*

$$\pi_m^* = \pi_{\text{fb}}.$$

Proof. See Appendix. □

The above result is not surprising because a monopolist with the power to make take-it-or-leave-it offers is able to extract all the surplus from symmetric buyers. An important question is: how robust is this result? Interestingly, we next show that this result does not depend neither on the symmetry, nor on the ability to price discriminate accordingly to the learning technology of the buyers.

4.1.2 Asymmetric Buyers

The following result shows under monopoly efficiency is still achieved even when buyers' learning technologies are heterogenous.

Proposition 6. *When the buyers are asymmetric, the equilibrium under a monopolistic market structure remains efficient.*

Proof. See Appendix. □

The monopolist can make take-it-or-leave-it offers and can perfectly discriminate the buyers. Therefore, it is optimal to sell the products at a price equal to their static expected utility flow. Hence, under this pricing strategy, the total profit is equal to total welfare. Therefore the allocation strategy of the products that maximizes profits becomes welfare-maximizing as well.

Of course, other fairness concerns may arise since, in a monopolistic market, the inefficiency is eliminated, but the entire surplus accrues to the monopolist. More interestingly, we are going to see that efficiency is preserved even when the monopolist cannot use first-degree price discrimination.

4.1.3 Beyond Price Discrimination

A reasonable concern might be that the above *no distortion result* about the monopolistic market structure with asymmetric buyers may be driven by the fact that the monopolist is allowed to use first-degree price discrimination. To show it is *not* the case, we prove an analogous result in the case in which the monopolist cannot discriminate.

Definition 2. *We say that the monopolist cannot discriminate if the strategy space is restricted to satisfy* $p_{k,1} = p_{k,2}$ *for all* $k \in \{a, b\}$.

We notice that our result on the efficiency of the monopolist with asymmetric buyers holds under this restriction (i.e., she can choose different prices for the two products, but those prices cannot be learning technology dependent).

Given the above definition, the following result holds.

Proposition 7. *When the buyers are asymmetric and the monopolist cannot discriminate, there is an efficient equilibrium. If the monopolist is forced to use the same price for both products equilibrium is not efficient.*

Proof. See Appendix. □

This result highlights a key difference between asymmetries in the learning technologies and asymmetries in the valuation of the new product (i.e., heterogeneous parameters ℓ and h across buyers). In the latter case, it is well known that the incentive compatibility of the buyers induces inefficient outcomes under competition. In our model of asymmetric learning technologies this does not happen.

The intuition behind the result is as follows. The willingness to pay of the consumers for product a is the same for both buyers and equal to μ_a. Instead, their willingness to pay for good b at belief π is potentially different: it is equal to the instant expected flow of utility $\mathbf{E}_\pi[\theta]$ plus the value of learning (i.e. v_i'') multiplied by the amount of information produced by the buyer. Even if $\mathbf{E}_\pi[\theta]$ is common across all the agents, differences in the learning parts may create incentive compatibility issue. However, in our proof we show that the monopolist can always obtain the total surplus by setting the price of the products equal to their expected flow of utility. Indeed, when the monopolist uses such a pricing strategy, the agent has zero value of information (i.e, v'' is the function constant at zero), and therefore they have the same willingness to pay, eliminating any incentive compatibility issue. Finally, its worth noting that from the inspection of the proofs of Propositions 6 and 7, we observe that the only substantial difference is that when the monopolist cannot discriminate, the efficient equilibrium is not such that the buyers have *strict* incentives to choose one product over the other.[11]

Next, we show that the above efficiency results is crucially different under a duopolistic market structure.

[11] However, it is obvious that if every type of price differentiation is banned and the monopolist is forced to use the same price for the two products distortion may arise.

4.2 Competition

In the case of duopolistic competition between the sellers, the value function of seller k is the solution to the following HJB equation:

$$w_k(\pi) = \sup_{p_k} \left\{ \sum_{i=1}^{n} \xi_{ik} p_{k,i} + g\Big(\pi, h, \ell\Big) \sum_{i=1}^{n} \frac{\xi_{ib}(t)}{2\rho\sigma_i^2} w_k''(\pi) \right\}, \qquad k \in \{a, b\}. \tag{12}$$

We next show that in sharp contrast to the monopolistic market structures, symmetric and asymmetric markets will have very different welfare implications. Once again, we consider separately the case of symmetric and asymmetric buyers.

4.2.1 Symmetric buyers

When the buyers and the pricing strategy are symmetric the HJB equation of seller k simplifies to:

$$w_k(\pi) = \sup_{p_k} \left\{ \underbrace{p_k \mathrm{Vol}_k}_{\text{flow gain}} + \underbrace{\mathrm{Vol}_b \frac{g\Big(\pi, h, \ell\Big)}{2\rho\sigma^2} w_k''(\pi)}_{\text{learning gain from product } b} \right\}, \qquad k \in \{a, b\}. \tag{13}$$

The right-hand-side of the HJB equation has two terms: the first term is the expected flow payoff $p_k \mathrm{Vol}_k$ (given that the volume of seller k's sale is Vol_k) and the second term is his continuation payoff that depends on Vol_b (i.e., volume of seller b's sale) via $\mathrm{Vol}_b \frac{g\left(\pi,h,\ell\right)}{2\rho\sigma^2} w_k''(\pi)$.

We start with a preliminary observation on the implications of equilibrium.

When $\pi = 0$, there is certainty in the low quality of product b, and the environment reduces to a static Bertrand competition between two sellers with the same 0 marginal cost but different quality of the products. The standard argument for price competition guarantees the following equilibrium behavior.

Lemma 4. *At $\pi = 0$ all buyers buy from seller a:*

$$Vol_a(0) = 2, \qquad Vol_b(0) = 0,$$

and the equilibrium prices are given by

$$p_a(0) - p_b(0) = \mu_a - \ell,$$

with

$$\ell - \mu_a \leq p_b(0) \leq 0.$$

Proof. See Appendix. □

In what follows, we are going to focus on the case in which $p_b(0) = 0$, since $p_b < 0$ is the pathological equilibrium in which seller b charges a price lower than his marginal cost (without any learning advantages, since beliefs are settled once and for all) only because he knows he will not sell to anyone.

Next, we characterize the pricing strategy of the sellers that can be sustained in a symmetric cutoff equilibrium.

Lemma 5. *In every equilibrium with symmetric π^*, the prices are as follows.*
If $\pi < \pi^*$:

$$p_a = \mu_a - \mathbf{E}_\pi[\theta] \qquad \text{and} \qquad p_b = 0. \tag{14}$$

If $\pi \geq \pi^*$:

$$p_a = p_b + \mu_a - \mathbf{E}_\pi[\theta] - \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} v''(\pi), \tag{15}$$

with

$$p_b \in \left[-\frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_b''(\pi),\ \mathbf{E}_\pi[\theta] - \mu_a + \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} \Big(v''(\pi) + w_a''(\pi) \Big) \right]. \tag{16}$$

Proof. See Appendix. □

When the belief is below the cutoff π^* seller a is serving the entire market, and the higher the perceived quality of the alternative product b, the lower the price he is able to ask. When the belief is higher than the threshold, the difference between the prices of the two products is pinned down by the competition and the indifference condition of the buyers. To make the buyers indifferent between the two alternative products, product a has to be discounted by both the difference in current expected payoffs and the forgone learning opportunity.

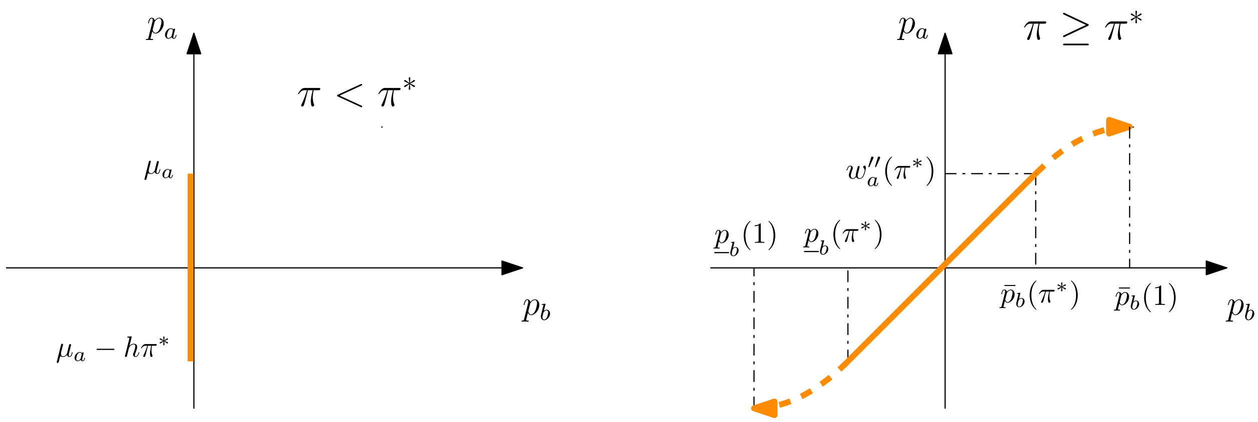

Figure 6: This figure plots equilibrium prices when $\pi < \pi^*$ and $\pi \geq \pi^*$. It is assumed that $\ell \sim 0$.

More interestingly, a multiplicity of pricing behavior can arise. The minimal price that seller b can ask is negative, and it is equal to the learning component of his value function w_b. The idea is that the seller is willing to sell below his cost hoping to prove to the market that his product is good, so to enjoy positive rent in the future. Instead, the maximal price is composed by three terms. First, it depends on the perceived difference in the quality of the products: the maximal price that he can charge is (linearly) increasing in $\mathbf{E}_\pi[\theta] - \mu_a$. Second, seller b can charge slightly higher prices because of the learning generated by the consumption of product b. This learning value is proportional to the variance of the signal, and it can be decomposed in to further parts: the higher price that a buyer is willing to pay for learning about quality of the product (i.e., $v''(\pi)$) and the lower competition exerted by seller a because he benefits from learning too (i.e., $w_a''(\pi)$).

Remark 1 (Equilibrium selection)**.** *Given the above multiplicity in prices, a multiplicity of consumption patterns arises as well. Since our focus is on the different efficiency properties of* *monopoly and competition, we want to consider the* minimal departure *from monopoly that is induced by having multiple competing sellers. Therefore, in what follows, we consider the equilibrium that is* more favorable to the sellers. *The reason is that a monopoly is a situation in which the surplus accruing to the (unique) seller is maximal, and therefore the minimal departure is obtained by focusing on the equilibrium that maximizes the sellers' profits when a competitor is added. Again, since we are going to highlight the difference between monopoly and competition, our findings will be more surprising the less we depart from the monopoly with our selection.*

Moving forward we derive the cutoff π^* under competition. Given the structure of the symmetric equilibrium, when $\pi \in [0, \pi^*]$ each buyer's value function is given by

$$v(\pi) = \mathbf{E}_\pi[\theta],$$

that follows from (8) and Lemma 5. In addition, due to (12) and the fact that $\mathrm{Vol}_b = 0$, the value function of each seller is

$$\begin{aligned} w_a &= np_a = n\Big(\mu_a - \mathbf{E}_\pi[\theta]\Big), \\ w_b &= 0. \end{aligned}$$

To characterize the equilibrium cutoff π^*, we combine the second order ODE (12) with a series of smooth pasting and value matching conditions at the cutoff value for the value function of seller a. Most importantly, given the equilibrium selection above (see Remark 1), we show that $\pi^* = \pi_{\text{fb}}$. As a consequence, efficiency is obtained under competition when buyers are symmetric. The following proposition summarizes.

Proposition 8. *The followings hold: (i) Any symmetric equilibrium is specified by a cutoff π^*. (ii) The equilibrium with the highest sellers' profits is efficient, (i.e., welfare-maximizing). (iii) The consumer surplus is strictly higher than under monopoly. (iv) The value function of the two sellers is convex, and the value function of the buyers is concave.*

Proof. See Appendix. □

It is important to understand the economic forces driving the result. At a first sight, the result does not seem surprising. Standard reasoning from static markets tell us that,

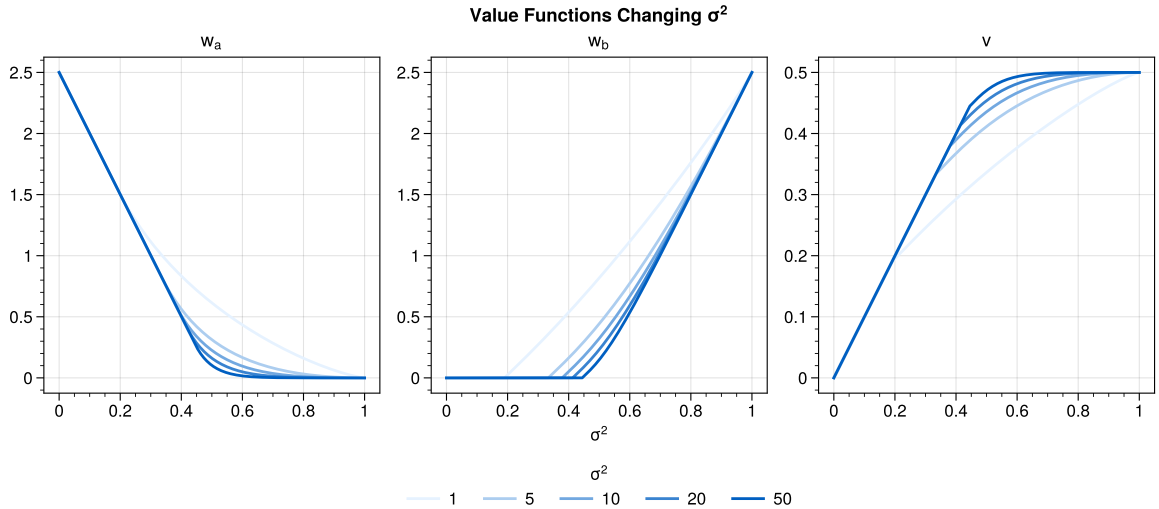

Figure 7: This figure plots the value functions w_a, w_b and v when σ^2 changes, fixing other parameters to $n = 5, h = 1, \ell = 0, \mu = .5$ (explicit characterizations of the value functions are provided in the proof of Proposition 8). As shown in the figure, with increasing σ^2 the cutoff π^*, expectedly, moves to the right (i.e., it increases).

since the buyers are symmetric, there is no reason for the sellers to use price discrimination. In static markets, it is well-known that the absence of incentives to price discrimination (or the possibility for the seller to perfectly discriminate) is sufficient to guarantee that the market power of the seller does not induce inefficiencies. One may think that the same is happening here too. Our next result below shows that this is not the case: If buyers are asymmetric, the possibility to use price discrimination does not amend inefficiencies. Indeed, in markets with learning externalities, efficiency is obtained only if additionally the seller is able to internalize the learning externality of the other market participants. Equation (16) and the proof of Proposition 8 show that when buyers are symmetric, this is the case. Seller b can benefit from the higher willingness to pay of the consumers (i.e., $v''(\pi){>}0$) and the softer competition of seller a (i.e., $w_a''(\pi){>}0$).

Importantly, in the next section we show that asymmetries in the learning technologies preclude efficient internalization of these externalities.

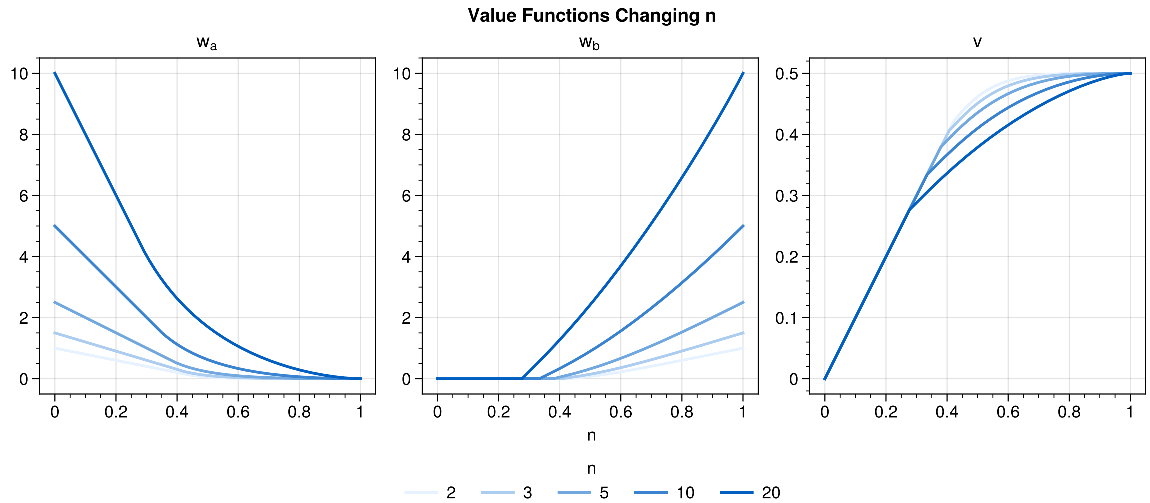

Figure 8: This figure plots the value functions w_a, w_b and v when n changes, fixing other parameters to $\sigma^2 = 10, h = 1, \ell = 0, \mu = .5$ (explicit characterizations of the value functions are provided in the proof of Proposition 8). As shown in the figure, with increasing n the cutoff π^*, expectedly, moves to the left (i.e., it decreases).

4.2.2 Asymmetric buyers

First, we show that the decentralized outcome induced by competition is no longer efficient with asymmetric buyers, and we then explore the nature of the inefficiency.

The HJB equation for the buyers is almost the same as in the symmetric case, with the only difference that the learning component involved in the trade-off between the two products is now idiosyncratic:

$$v_i(\pi) = \max\left\{\mu_a - p_{a,i} + g\big(\pi, h, \ell\big) \sum_{j \neq i} \frac{\xi_{jb}}{2\rho\sigma_j^2} v_i''(\pi), \right.$$
$$\left. \mathbf{E}_\pi[\mu_b] - p_{b,i} + g\big(\pi, h, \ell\big) \sum_{j=1}^{2} \frac{\xi_{jb}}{2\rho\sigma_j^2} v_i''(\pi) \right\}.$$

Price competition between sellers imposes indifference between the arguments of the above maximization. Indeed, in the right-hand-side of the above HJB, if the second argument was larger, a profitable deviation for seller b would be to slightly increase $p_{b,i}$, collecting higher per unit revenues and selling to the same number of buyers. An analogous profitable deviation obtains for seller a if the first argument was strictly larger than the second. Therefore:

$$p_{b,1} - p_{a,1} = \mathbf{E}_{\pi}\left[\mu_b\right] - \mu_a + \frac{1}{2\rho\sigma_1^2} g\Big(\pi, h, \ell\Big) v_1''(\pi), \tag{17}$$

for buyer 1 and

$$p_{b,2} - p_{a,2} = \mathbf{E}_{\pi}\left[\mu_b\right] - \mu_a + \frac{1}{2\rho\sigma_2^2} g\Big(\pi, h, \ell\Big) v_2''(\pi), \tag{18}$$

for buyer 2.

Thus, equations (17) and (18) together give us the markup on the expected values that seller b is able to apply due to the learning component.

We next consider the optimal pricing problem faced by the sellers.

The HJB equations of the sellers become more complicated objects since they now control a two dimensional variable (prices). For seller b:

$$w_b(\pi) = \max_{p_{b,1}, p_{b,2}} \left\{ \begin{array}{c} 0, \\ p_{b,2} + \frac{1}{2\rho\sigma_2^2} g\Big(\pi, h, \ell\Big) w_b''(\pi), \\ p_{b,1} + \frac{1}{2\rho\sigma_1^2} g\Big(\pi, h, \ell\Big) w_b''(\pi), \\ p_{b,2} + p_{b,1} + \frac{\sigma_1^2+\sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} g\Big(\pi, h, \ell\Big) w_b''(\pi) \end{array} \right\}$$

where the four terms allow for the different combinations of buyers of product b. In particular, the value function of seller b is equal to his total revenues from good and bad learners plus the learning induced by the consumption of his product. Similarly for seller a we have:

$$w_a(\pi) = \max_{p_{a,1}, p_{a,2}} \left\{ \begin{array}{c} \frac{\sigma_1^2+\sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} g\Big(\pi, h, \ell\Big) w_a''(\pi) \\ p_{a,2} + \frac{1}{2\rho\sigma_1^2} g\Big(\pi, h, \ell\Big) w_a''(\pi), \\ p_{a,1} + \frac{1}{2\rho\sigma_2^2} g\Big(\pi, h, \ell\Big) w_a''(\pi), \\ p_{a,2} + p_{a,1} \end{array} \right\}$$

with the only difference being that here buyers contribute to learning when they do *not* buy from a.

To prove the next result we will plug the indifference conditions of the buyers into

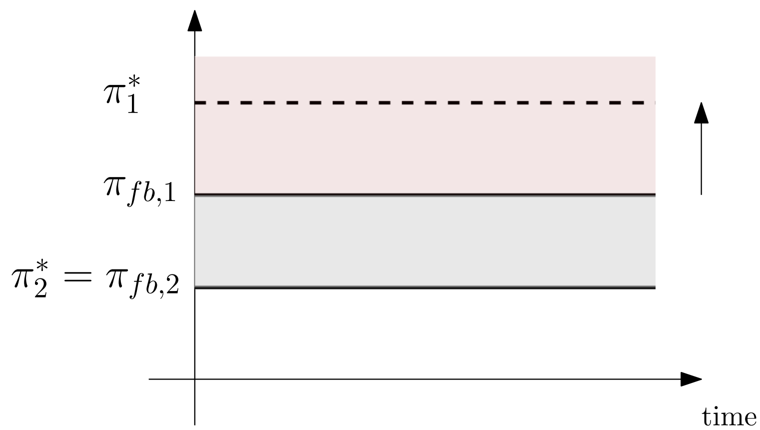

Figure 9: This panel shows while completion does not lead to the first best outcome, there is efficiency at the top.

the HJB equation of seller b to prove the discrepancy in the threshold.

Proposition 9. *When the buyers are asymmetric, the equilibrium with the highest sellers' profit is not efficient. However, there is efficiency at the top, therefore,*

$$\pi_2 = \pi_{\text{fb},2}.$$

Finally, in this equilibrium the good learner has no incentive to mimic the bad learner.

Proof. See Appendix. □

The inefficiency at π_1 (with respect to $\pi_{\text{fb},1}$) follows from not taking into account the learning externality that consumption by bad learners induces over good learners. Indeed, when the seller b serves a buyer with the bad learning technology, an informational gain is produced for every market participant in the economy. The optimal price-setting by seller b internalizes the learning gain for his and the buyer, and competition incorporates the learning externality of seller a into the price. However, without using multilateral contracts (a theme that we explore in Section 5), there is no way in which the benefit over good learners is internalized. This issue does not arise in the symmetric case because at the unique threshold (the same for the smallest threshold π_2^* here) no buyer has a strictly positive value of information.

The efficiency at the top part of the result highlights the way in which distortions arise. Competition does not affect the number and quality of the innovation that are given a shot, that is, the threshold for being tested by the fraction of the market that produce better (more precise) information about the product (i.e., during the Beta phase). However, what is affected is the confidence in the quality of the product that is required to start to serve the entire market.

The final part of the result is also important, because it implies that the equilibrium does not require that the seller is able to discriminate the buyers accordingly to their learning technology. Indeed, good learners will self-select in buying product b during the Beta phase given the prices posted by the sellers.

We notice that the extent of inefficiency is not monotone in the difference between the learning technologies of buyers. Proposition 8 already guarantees that efficiency holds in the case of symmetric buyers, and the next result shows that the distortion disappears also in the limit where one of the two buyers does not produce any valuable information (through his experience) about the product's quality.

Proposition 10. *Fix σ_2. The equilibrium with a buyer that does not generate any information about the quality of the product is efficient:*

$$\lim_{\sigma_1 \to \infty} (\pi_1 - \pi_{\text{fb},1}) = 0$$

Proof. See Appendix. □

The intuition behind the previous result is that if serving the general public does not produce additional information about the product, there is no learning externality that has to be considered when deciding if the product is ready for the entire market.

Finally, we note that the non-monotonicity in the difference in learning technologies implied by the previous proposition is a robust feature of the model, and it continues to arise even with multiple levels of learning technologies. However, what is lost in the more general case is the stark conclusion that if the worst learners become completely uninformative, then efficiency is fully restored. Indeed, as long as there are two types of learners with their variance of the signal strictly between zero and plus infinity, competition induces some inefficiency.

5 Discussion

In this section we first explore the robustness of our results to an alternative market structure in which the seller of the known product can start to compete by directly selling the risky product. Then, we show that inefficiency is not a concern if the seller can offer multilateral contracts.

5.1 Different market structure

Multiple sellers of product b Most of our analysis focused on duopolistic competition. A natural question is what happens if the new product b can be produced by multiple, and not just one, competitors. In particular, suppose that there are $m > 2$ sellers making take-it-or-leave-it offers to the consumers, with seller 1 selling the established product a, and all the other sellers marketing product b.

Proposition 11. *When there is one seller that can selling product a, and $m-1$ sellers selling product b, and buyers are asymmetric, the equilibrium is not efficient, and the amount of experimentation is too low:*

$$\forall i \quad \pi_i = \pi_{\text{myopic}} < \pi_{\text{fb},i}.$$

Direct competition in the new product market What if the incumbent (old seller) diversifies and along with the new seller starts selling product b as well? Does this restore efficiency?

It may be the case that, with introducing the new unknown product b by the new seller, the incumbent (i.e, the seller of the known quality product a) starts competing directly by selling b as well. In this case, one can hope that direct competition in the market for this new product would restore efficiency. The next result shows that this is *not* the case, and our inefficiency of competition result is robust to this different nature of competition.

Proposition 12. *When there is a seller that can serve both markets, one seller that can serve only market b, and buyers are asymmetric, the equilibrium is not efficient, and the amount of experimentation is too low:*

$$\forall i \quad \pi_i = \pi_{\text{myopic}} < \pi_{\text{fb},i}.$$

Proof. See Appendix. □

The result shows that our baseline model was the most favorable for the welfare properties of competition. In situations in which the incumbent can start to compete in both market, the problem almost reduces to static competition, and even at the top distortion arises; some promising new product are not even brought to the Beta phase.

5.2 Multilateral Contracts

There is a way, however, to eliminate distortion and maintain competition. This is achieved by increasing the commitment power of the sellers. More precisely, suppose that now seller $k \in \{a, b\}$ can commit to offering a multilateral contract of the following form: one-unit of product k will be delivered to buyer $i \in \{1, 2\}$ if buyer i makes a transfer $t^i_{i,k}$ to seller k and buyer $j \neq i$ makes a transfer $t^j_{i,k}$ to seller k.

That is, we allow the seller to ask a buyer to pay for the fact that the product is delivered to another consumer. We notice that the buyer who does not receive the product may have an incentive to pay because of the learning externality that is generated by the use of product b. Although there are service markets in which a similar structure may be implemented in the form to a subscription to a platform that shares buyers' experiences, we think that in most of the cases assuming such a commitment power is unwarranted. Still, if we allow for this possibility, the competition outcome becomes efficient.

Proposition 13. *When seller b is allowed to use multilateral contracts, the equilibrium is efficient.*

Proof. See Appendix. □

6 Conclusion

We study the interaction between the market structure (monopoly vs oligopoly) and asymmetry in learning technologies in a dynamic product market. In this market, a new product of unknown quality competes against an established one, with Bayesian learning induced by the consumption of the unknown product. We establish that the optimal policy (both the efficient and the market solution) is in terms of a sequence of belief thresholds, with a Beta phase in which only a subset of best learners is served.

We analyze the learning progression in the market belief, the expected time before the product moves from the Beta phase to serving the entire market, and consider the efficiency implications of different market structures.

Under a monopolistic market structure in which the same seller sells the new product and the old product, the resulting equilibrium always achieves efficiency. In sharp contrast, however, if two different sellers compete in marketing the products, efficiency is achieved only if the buyers are symmetric in their learning technologies.

We identify the inefficiency as a learning externality that consumption of the product by one buyer generates for the other buyers. The equilibrium inefficiency has two features: (i) Efficiency at the top: the threshold for starting to serve the best learners (i.e., to enter into a Beta phase) remains the optimal one. (ii) Non-monotonicity: distortions are not monotone in the extent of the asymmetry.

We show that in markets in which the sellers can use multilateral contracts, the distortion disappears. Finally, we consider the robustness of our results under different assumptions about the ability to price discriminate and other market structures.

Our results shed light on how prevalent heterogeneity of reviews in online marketplaces (e.g. Amazon, ebay, Alibaba) affect pricing, learning, competition, welfare, and adoption of new products.

Appendix: Omitted proofs

Lemma 1. *[Belief Evolution] We have*

$$d\pi_t = \pi_t(1-\pi_t)(h-\ell)\sqrt{\sum_{i=1}^{n}\frac{\xi_{ib}(t)}{\sigma_i^2}}dZ_t$$

where Z_t is a standard Wiener process. In particular, in the case of symmetric buyers we have

$$d\pi_t = \frac{\pi_t(1-\pi_t)(h-\ell)}{\sigma}\sqrt{Vol_b(t)}dZ_t.$$

Moreover, π is a continuous martingale with respect to $\{\mathcal{F}_t, t \geq 0\}$ and a strong Markov process that is symmetric in time.

Proof of Lemma 1. Define $d\mathfrak{C}_{bi} := \frac{\theta}{\sigma_i}dt + dZ_{it}$, for $i = 1, \cdots, n$, where $\theta \in \{h, \ell\}$. So, $\left(d\mathfrak{C}_{bi}\right)^2 = \frac{\theta^2}{\sigma_i^2}(dt)^2 + dZ_{it}^2 + 2dZ_{it}dt = dt$. Recall that $dZ_{it} \perp dZ_{jt}$, for all $i \neq j$. We denote as $m_i(t)$ the consumer with the i-th worst learning technology that is consuming product b at time t. Then, by Bayes' rule, we have

$$\pi_{t+dt} = \frac{\pi_t \Pr\left(\frac{h}{\sigma_{m_1(t)}}, \frac{h}{\sigma_{m_2(t)}}, \cdots, \frac{h}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right)}{\pi_t \Pr\left(\frac{h}{\sigma_{m_1(t)}}, \frac{h}{\sigma_{m_2(t)}}, \cdots, \frac{h}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right) + (1-\pi_t)\Pr\left(\frac{\ell}{\sigma_{m_1(t)}}, \frac{\ell}{\sigma_{m_2(t)}}, \cdots, \frac{\ell}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right)} \tag{19}$$

where for $\theta \in \{h, \ell\}$:

$$\begin{aligned}
\Pr\left(\frac{\theta}{\sigma_{m_1(t)}}, \frac{\theta}{\sigma_{m_2(t)}}, \cdots, \frac{\theta}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right) &= \prod_{i=1}^{\mathrm{Vol}_b(t)} \frac{1}{\sqrt{2\pi dt}} \exp\left(\frac{-1}{2dt}\left(d\mathfrak{C}_{bm_i(t)} - \frac{\theta}{\sigma_{m_i(t)}}dt\right)^2\right) \\
&= \left(\frac{1}{\sqrt{2\pi dt}}\right)^{\mathrm{Vol}_b(t)} \exp\left(\frac{-1}{2dt}\sum_{i=1}^{\mathrm{Vol}_b(t)}\left(d\mathfrak{C}_{bm_i(t)} - \frac{\theta}{\sigma_{m_i(t)}}dt\right)^2\right).
\end{aligned} \tag{20}$$

Using (19) we also have

$$\begin{aligned}
d\pi_t &= \pi_{t+dt} - \pi_t \\
&= \pi_t(1-\pi_t)\frac{\Pr\left(\frac{h}{\sigma_{m_1(t)}}, \frac{h}{\sigma_{m_2(t)}}, \cdots, \frac{h}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right) - \Pr\left(\frac{\ell}{\sigma_{m_1(t)}}, \frac{\ell}{\sigma_{m_2(t)}}, \cdots, \frac{\ell}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right)}{\pi_t \Pr\left(\frac{h}{\sigma_{m_1(t)}}, \frac{h}{\sigma_{m_2(t)}}, \cdots, \frac{h}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right) + (1-\pi_t)\Pr\left(\frac{\ell}{\sigma_{m_1(t)}}, \frac{\ell}{\sigma_{m_2(t)}}, \cdots, \frac{\ell}{\sigma_{m_{\mathrm{Vol}_b}(t)}}\right)}.
\end{aligned} \tag{21}$$

To simplify (21) we note that for $i = 1, \cdots, \mathrm{Vol}_b(t)$:

$$\frac{-1}{2dt}\left(d\mathfrak{C}_{bm_i(t)} - \frac{\theta}{\sigma_{m_i(t)}}dt\right)^2 = \frac{\theta}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{\theta}{\sigma_{m_i(t)}}\right)^2 dt - \frac{1}{2},$$

where the equality follows because $\left(d\mathfrak{C}_{bm_i(t)}\right)^2 = dt$. Using the above equality, plugging

(20) into (21) implies that

$$d\pi_t = \pi_t(1-\pi_t)\frac{\exp\left(\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{h}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{h}{\sigma_{m_i(t)}}\right)^2 dt\right)\right) - \exp\left(\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{\ell}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{\ell}{\sigma_{m_i(t)}}\right)^2 dt\right)\right)}{\pi_t \exp\left(\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{h}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{h}{\sigma_{m_i(t)}}\right)^2 dt\right)\right) + (1-\pi_t)\exp\left(\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{\ell}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{\ell}{\sigma_{m_i(t)}}\right)^2 dt\right)\right)} \tag{22}$$

Using Taylor expansion (removing the higher order terms), we further have for $\theta \in \{h, \ell\}$:

$$\begin{aligned}\exp\left(\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{\theta}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{\theta}{\sigma_{m_i(t)}}\right)^2 dt\right)\right) &= 1 + \sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{\theta}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{\theta}{\sigma_{m_i(t)}}\right)^2 dt\right) \\ &\quad + \frac{1}{2}\left(\sum_{i=1}^{\text{Vol}_b(t)}\frac{\theta}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{1}{2}\left(\frac{\theta}{\sigma_{m_i(t)}}\right)^2 dt\right)^2 \\ &= 1 + \sum_{i=1}^{\text{Vol}_b(t)}\frac{\theta}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}\end{aligned} \tag{23}$$

where the last equality follows because $dZ_{it}dZ_{jt} = 0$ for $i \neq j$ and $(dt)^k = 0$, for $k > 1$. Next, plugging (23) into (22) implies

$$\begin{aligned}d\pi_t &= \pi_t(1-\pi_t)\frac{\sum_{i=1}^{\text{Vol}_b(t)}\frac{(h-\ell)}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}}{1+\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{\pi_t h+(1-\pi_t)\ell}{\sigma_{m_i(t)}}\right)d\mathfrak{C}_{bm_i(t)}} \\ &= \pi_t(1-\pi_t)\left(\sum_{i=1}^{\text{Vol}_b(t)}\frac{(h-\ell)}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}\right)\left(1-\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{\pi_t h+(1-\pi_t)\ell}{\sigma_{m_i(t)}}\right)d\mathfrak{C}_{bm_i(t)}\right) \\ &= \pi_t(1-\pi_t)(h-\ell)\left(\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{1}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)} - \frac{\pi_t h+(1-\pi_t)\ell}{\sigma^2_{m_i(t)}}dt\right)\right)\end{aligned} \tag{24}$$

where the second equality follows by the Taylor expansion of the denominator, and the last equality follows because $d\mathfrak{C}_{bm_i(t)}d\mathfrak{C}_{bm_j(t)} = 0$ when $i \neq j$ and $(d\mathfrak{C}_{bm_i(t)})^2 = dt$. Finally, note that

$$\mathbf{E}_{\pi_t}\left[\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{1}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}-\frac{\pi_t h+(1-\pi_t)\ell}{\sigma^2_{m_i(t)}}dt\right)\right]=\sum_{i=1}^{\text{Vol}_b(t)}\mathbf{E}_{\pi_t}\left[\frac{1}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}-\frac{\pi_t h+(1-\pi_t)\ell}{\sigma^2_{m_i(t)}}dt\right]=0$$

$$\mathbf{Var}_{\pi_t}\left[\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{1}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}-\frac{\pi_t h+(1-\pi_t)\ell}{\sigma^2_{m_i(t)}}dt\right)\right]=\sum_{i=1}^{\text{Vol}_b(t)}\frac{1}{\sigma^2_{m_i(t)}}dt$$

Therefore,

$$\sum_{i=1}^{\text{Vol}_b(t)}\left(\frac{1}{\sigma_{m_i(t)}}d\mathfrak{C}_{bm_i(t)}-\frac{\pi_t h+(1-\pi_t)\ell}{\sigma^2_{m_i(t)}}dt\right)\sim\sqrt{\sum_{i=1}^{\text{Vol}_b(t)}\frac{1}{\sigma^2_{m_i(t)}}}dZ_t$$

where Z_t is the standard BM. Therefore,

$$d\pi_t=\pi_t(1-\pi_t)(h-\ell)\sqrt{\sum_{i=1}^{\text{Vol}_b(t)}\frac{1}{\sigma^2_{m_i(t)}}}dZ_t=\pi_t(1-\pi_t)(h-\ell)\sqrt{\sum_{i=1}^{n}\frac{\xi_{ib}(t)}{\sigma^2_i}}dZ_t,$$

finishing the proof.

□

Proposition 1. *Let buyers $\{m_1,...,m_j\}$ use the risky product b in the time interval $[0,T]$. Then, the expected progression in the market belief π_t to the target belief β is explicitly given by:*

$$MO\left(\pi_0,\beta,T,m_1,...,m_j\right)=(1-\beta)\pi_0\Phi(\lambda_1)-\beta(1-\pi_0)\Phi(\lambda_0)$$

where

$$\lambda_1=\frac{1}{(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{T}{\sigma^2_{m_i}}}}\left[\ln\left(\frac{\frac{\pi_0}{1-\pi_0}}{\frac{\beta}{1-\beta}}\right)+\frac{(h-\ell)^2}{2}\left(\sum_{i=1}^{j}\frac{T}{\sigma^2_{m_i}}\right)\right]$$

$$\lambda_0=\frac{1}{(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{T}{\sigma^2_{m_i}}}}\left[\ln\left(\frac{\frac{\pi_0}{1-\pi_0}}{\frac{\beta}{1-\beta}}\right)-\frac{(h-\ell)^2}{2}\left(\sum_{i=1}^{j}\frac{T}{\sigma^2_{m_i}}\right)\right]$$

and $\Phi(\cdot)$ denotes the CDF of a standard normal random variable.

Proof of Proposition 1. To prove the proposition we first need to prove the following lemma.

Lemma 6. *Let buyers $\{m_1, ..., m_j\}$ use the risky product b in the time interval $[0, T]$. Define $U := 1-\pi$ and $L := \pi$. Given the result of Lemma 1, let π solves $d\pi = \pi(1-\pi)(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}dZ$ where Z is a BM under the probability measure $\mathbb{P}$, for a given π_0.*[12] *This is equivalent to the followings:*

Part 1. There is a standard BM, Z^U, so that the process $\gamma_y := \frac{L}{U}$ solves $\frac{d\gamma_y}{\gamma_y} = (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}dZ^U$, and $\gamma_{y_0} > 0$.

Part 2. There is a standard BM, Z^L, so that the process $\gamma_z := \frac{U}{L}$ solves $\frac{d\gamma_z}{\gamma_z} = (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}dZ^L$ and $\gamma_{z0} > 0$.

Proof of Lemma 6. Note that whenever is clear the dependence of the process to time is removed, for ease of notation.

Given Lemma 1, applying Ito's lemma gives

$$d\gamma_y = \gamma_y(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\left[dZ + (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\pi_t dt\right]$$

Define Z^U by $dZ^U = dZ + (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\pi_t dt$ with $Z_0^U = 0$. We next show that Z^U is indeed a standard BM. Note that Z is a BM under the probability measure $\mathbb{P}$, and $\mathbb{P}_U$ and $\mathbb{P}_L$ are probability measures derived from P. Hence,

$$\frac{d\mathbb{P}_U}{d\mathbb{P}} = \frac{U_T}{U_0} = \frac{1-\pi_T}{1-\pi_0}$$

[12] Note that the solution π to the SDE is unique both in strong and weak sense, see, e.g., section 5.2 in Øksendal (2003).

Moreover, $\frac{d[1-\pi_t]}{1-\pi_t} = -(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\pi_t dZ$. Therefore,

$$1-\pi_t = (1-\pi_0)e^{-(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\int_0^t \pi_s dZ_s - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\right)^2\int_0^t \pi_s^2 ds}$$

Moreover,

$$\frac{d\mathbb{P}_U}{d\mathbb{P}} = e^{-(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\int_0^T \pi_s dZ_s - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\right)^2\int_0^T \pi_s^2 ds}$$

So, given the above relation, **Girsanov theorem** shows that Z^U is indeed a $\mathbb{P}_U$-BM.

Next, we argue that Part 1 proves Part 2. Let Z^U be a $\mathbb{P}_U$ BM. Since $\gamma_z = \gamma_y^{-1}$, by Ito's lemma, we get

$$d\gamma_z = (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\gamma_z\left[-dZ^U + (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}dt\right].$$

Moreover,

$$\frac{d\mathbb{P}_L}{d\mathbb{P}_U} = \frac{\frac{L_T}{L_0}}{\frac{U_T}{U_0}} = \frac{\gamma_{y_T}}{\gamma_{y_0}} = e^{(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}Z_T^U - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\right)^2 T},$$

therefore, by **Girsanov theorem**, the process Z^L defined by $dZ^L = -dZ^U + (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}dt$ and $Z_0^L = 0$ is a $\mathbb{P}_L$-BM.

Finally, we show the converse holds as well. That is, the implication of Part 2 proves that $d\pi = \pi(1-\pi)(h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}dZ$ where Z is a BM under $\mathbb{P}$. Suppose that Z^L is defined as in the Lemma. Then, Ito's lemma gives

$$d\pi_t = (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}\pi_t(1-\pi_t)\left[-dZ_t^L + (h-\ell)\sqrt{\sum_{i=1}^{j}\frac{1}{\sigma_{m_i}^2}}(1-\pi_t)dt\right].$$

Let Z be the process defined by $dZ = -dZ^L + (h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}(1-\pi_t)dt$ with $Z_0 = 0$. Since,

$$\frac{d\pi_t}{\pi_t} = -(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}(1-\pi_t)dZ^L + \left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 (1-\pi_t)^2 dt,$$

we have

$$\pi_t = \pi_0 e^{-(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\int_0^t (1-\pi_s)dZ_s^L + \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 \int_0^t (1-\pi_s)^2 ds}$$

and

$$\frac{d\mathbb{P}}{d\mathbb{P}_L} = \frac{\pi_0}{\pi_T} = e^{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\int_0^t (1-\pi_s)dZ_s^L - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 \int_0^t (1-\pi_s)^2 ds}.$$

Therefore, by **Girsanov theorem**, Z is a BM under $\mathbb{P}$. □

To finish the proof we also need another lemma, described below.

Lemma 7. *Define the event $\mathcal{A} := \{\omega \in \Omega : \pi_T(\omega) > \beta\}$. Then,*

$$\mathbf{E}_{\pi_t}[\max\{\pi_T - \beta, 0\}] = (1-\beta)L_t \cdot P_L\{\mathcal{A}|\mathcal{F}_t\} - \beta U_t \cdot P_U\{\mathcal{A}|\mathcal{F}_t\}$$

Proof of Lemma 7. Note that

$$\pi_T - \beta = (1-\beta)L_T - \beta U_T$$

Therefore, we have

$$\begin{aligned}\mathbf{E}_{\pi_t}[\max\{\pi_T - \beta, 0\}] &= (1-\beta)\mathbf{E}\left[L_T \mathbf{1}_{\mathcal{A}}|\mathcal{F}_t\right] - \beta\mathbf{E}\left[U_T \mathbf{1}_{\mathcal{A}}|\mathcal{F}_t\right] \\ &= (1-\beta)L_t \cdot P_L\{\mathcal{A}|\mathcal{F}_t\} - \beta U_t \cdot P_U\{\mathcal{A}|\mathcal{F}_t\}\end{aligned}$$

finishing the proof. □

Next, equipped with the above lemmas, we finish the proof of the proposition. Let Z^L and Z^U be BMs as in Lemma 6, thus (using the exponential martingale formula) we have

$$\gamma_{y_t} = \gamma_{y_0} e^{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}} Z_t^U - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 t}$$

and

$$\gamma_{z_t} = \gamma_{z_0} e^{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}} Z_t^L - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 t}.$$

Therefore, sine Z^U is a BM, we have

$$\begin{aligned}
\mathbb{P}_U\{\mathcal{A}|\mathcal{F}_t\} &= P_U\left\{\gamma_{y_T} > \frac{\beta}{1-\beta} | \gamma_{y_t}\right\} \\
&= P_U\left\{\ln \frac{\gamma_{y_T}}{\gamma_{y_t}} > \ln \frac{\frac{\beta}{1-\beta}}{\gamma_{y_t}}\right\} \\
&= P_U\left\{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}(Z_T^U - Z_t^U) > \ln \frac{\frac{\beta}{1-\beta}}{\gamma_{y_t}} + \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 (T-t)\right\} \\
&= \Phi\left(\frac{1}{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\sqrt{T-t}}\left[-\ln \frac{\frac{\beta}{1-\beta}}{\gamma_{y_t}} - \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 (T-t)\right]\right) \qquad (25)
\end{aligned}$$

where the last equality follows by Lemma 6. Similarly,

$$\begin{aligned}
\mathbb{P}_L\{\mathcal{A}|\mathcal{F}_t\} &= P_L\left\{\gamma_{z_T} > \frac{1-\beta}{\beta} | \gamma_{y_t}\right\} \\
&= P_L\left\{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}(Z_T^L - Z_t^L) > \ln \frac{\frac{1-\beta}{\beta}}{\gamma_{z_t}} + \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 (T-t)\right\} \\
&= \Phi\left(\frac{1}{(h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\sqrt{T-t}}\left[-\ln \frac{\frac{\beta}{1-\beta}}{\gamma_{y_t}} + \frac{1}{2}\left((h-\ell)\sqrt{\sum_{i=1}^{j} \frac{1}{\sigma_{m_i}^2}}\right)^2 (T-t)\right]\right) \qquad (26)
\end{aligned}$$

where the last equality follows by Lemma 6. Equations (25) and (26) along with Lemma

7 finish the proof of the first part of the Proposition (note that the proposition is stated for when $t = 0$). □

Proposition 2. *The first-best (social welfare) maximizing cutoff is given by*

$$\pi_{\text{fb}} = \frac{(\mu_a - \ell)(\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}} - 1)}{(\ell + h) - 2\mu_a + (h - \ell)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}. \tag{6}$$

Moreover, π_{fb} is increasing in μ_a, σ^2, and ρ. It is decreasing in h and n.

Proposition 3. *The first-best social welfare is equal to $nW^*(\pi)$ where*

$$W^*(\pi) = \begin{cases} \mu_a & \text{if } 0 \le \pi \le \pi_{\text{fb}}; \\ \mathbf{E}_\pi[\theta] + \varphi\left[\pi^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}(1-\pi)^{\frac{1}{2}\left(\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}+1\right)}\right] & \text{if } \pi_{\text{fb}} < \pi \le 1. \end{cases}$$

and

$$\varphi := 2\frac{h - \mu_a}{\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}} - 1}\left(\frac{\pi_{\text{fb}}}{1 - \pi_{\text{fb}}}\right)^{\frac{1}{2}\left(\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}+1\right)}.$$

Moreover, W^* is strictly convex in $\pi_{\text{fb}} < \pi \le 1$.

Proof of Propositions 2 and 3. The (average) first-best welfare is given by

$$W(\pi) = \max_{\xi_{1b},\ldots,\xi_{nb}}\left\{\frac{(n - \sum_{i=1}^n \xi_{ib})\mu_a + \sum_{i=1}^n \xi_{ib}\mathbf{E}_\pi[\mu_b]}{n} + \frac{g(\pi,h,\ell)W''(\pi)}{n\rho\sigma^2}\left(\sum_{i=1}^n \xi_{ib}\right)\right\}.$$

The policy $\xi^*_{1b}, \ldots, \xi^*_{nb}$ is optimal if and only if

$$\xi^*_{1b}, \ldots, \xi^*_{nb} \in \arg\max_{\xi_{1b},\ldots,\xi_{nb}}\left\{\frac{(n - \sum_{i=1}^n \xi_{ib})\mu_a + \sum_{i=1}^n \xi_{ib}\mathbf{E}_\pi[\mu_b]}{n} + \frac{g(\pi,h,\ell)W''(\pi)}{n\rho\sigma^2}\left(\sum_{i=1}^n \xi_{ib}\right)\right\}.$$

Rearranging the terms

$$W(\pi) = \max_{\xi_{1b},\ldots,\xi_{nb}} \left\{ \mu_a + \left[\frac{g(\pi,h,\ell)W''(\pi)}{n\rho\sigma^2} - \frac{\mu_a - \mathbf{E}_\pi[\theta]}{n} \right] \left(\sum_{i=1}^{n} \xi_{ib} \right) \right\}. \tag{27}$$

Given the structure of (27) it is clear that the optimal strategy has the following form:

$$\begin{cases} \xi_{1b} = \ldots = \xi_{nb} = 1 & \text{if } \frac{\mu_a - \mathbf{E}_\pi[\mu_b]}{n} \leq \frac{g(\pi,h,\ell)}{n\rho\sigma^2} W''(\pi), \\ \xi_{1b} = \ldots = \xi_{nb} = 0 & \text{otherwise.} \end{cases}$$

We next derive the value function. This is done by obtaining it in the two regions/cases defined by the unique threshold π_{fb}. We consider two cases:

- Case 1: If $\pi \geq \pi_{\text{fb}}$, then following Wronskian approach of second order ODEs (Zaitsev and Polyanin (2002)) we have

$$\begin{aligned} W(\pi) = n\mathbf{E}_\pi[\theta] &+ \zeta_1 \left[\pi^{\frac{1}{2}\left(1 - \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} (1-\pi)^{\frac{1}{2}\left(1 + \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} \right] \\ &+ \zeta_2 \left[\pi^{\frac{1}{2}\left(1 + \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} (1-\pi)^{\frac{1}{2}\left(1 - \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} \right]. \end{aligned}$$

- Case 2: If $\pi < \pi_{\text{fb}}$, we trivially have

$$W(\pi) = \mu_a.$$

We note that W is bounded by nh on $\pi_{\text{fb}} < \pi \leq 1$, thus $\zeta_2 = 0$ (note that otherwise $\lim_{\pi \to 1} (1-\pi)^{\frac{1}{2}\left(1 - \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}$ would explode in this region). Hence,

$$W(\pi) = \zeta_1 \left[\pi^{\frac{1}{2}\left(1 - \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} (1-\pi)^{\frac{1}{2}\left(1 + \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} \right] + n\mathbf{E}_\pi[\theta]. \tag{28}$$

Now, the value matching and smooth pasting for value function give

$$W(\pi_{\text{fb}}) = n\mu_a, \tag{29}$$

$$W'(\pi_{\text{fb}}) = 0. \tag{30}$$

First, using (28) and (29) we have

$$\zeta_1 = \frac{n(\mu_a - \mathbf{E}_{\pi_{\text{fb}}}[\theta])}{(\pi_{\text{fb}})^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}(1-\pi_{\text{fb}})^{\frac{1}{2}\left(1+\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}}.$$

But if we combine this with

$$0 = \frac{w_m'(\pi_{\text{fb}})}{w_m(\pi_{\text{fb}})}, \tag{31}$$

we obtain that

$$\pi_{\text{fb}} = \frac{\mu_a - \ell + (\ell - \mu_a)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}{2\mu_a - (\ell + h) + (\ell - h)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}. \tag{32}$$

Hence (due to (6)) $\pi_{\text{fb}} = \pi_{\text{fb}}$.

There are two unknowns that we have to identify to characterize the welfare-maximizing behavior: φ and π_{fb}. To pin down φ, π_{fb} we use the value matching $W(\pi_{\text{fb}}) = \mu_a$ and the smooth pasting conditions $W'(\pi_{\text{fb}}) = 0$ at the boundary condition $\pi = \pi_{\text{fb}}$.

Given these constraints it is easy to verify that

$$\pi_{\text{fb}} = \frac{(\mu_a - \ell)(\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}} - 1)}{(\ell + h) - 2\mu_a + (h - \ell)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}$$

and

$$\varphi = 2\frac{h - \mu_a}{\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}} - 1}\left(\frac{\pi_{\text{fb}}}{1 - \pi_{\text{fb}}}\right)^{\frac{1}{2}\left(\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}+1\right)}$$

finishing the proof. □

Moreover, to have enough equations to identify all the unknowns needed to pin down the value function, we start by proving the necessity of the super-contact condition:

Lemma 8. *The following holds for the average welfare*

$$W_1''(\pi_{\text{fb},1}) = W_0''(\pi_{\text{fb},1}).$$

Proof. We prove the statement considering two separate cases and obtaining a contradiction in both of them.

Case 1: Let $W_1''(\pi_{\text{fb},1}) < W_0''(\pi_{\text{fb},1})$. In this case the contradiction is obtained by showing that a failure of the super-contact condition would imply that the value function is nondifferentiable. First recall that the optimality condition gives

$$\frac{g(\pi_{\text{fb},1},h,\ell)}{2\rho\sigma_1^2}W_1''(\pi_{\text{fb},1}) \le \frac{\mu_a - \mathbf{E}_{\pi_{\text{fb},1}}[\theta]}{2} \le \frac{g(\pi_{\text{fb},1},h,\ell)}{2\rho\sigma_1^2}W_0''(\pi_{\text{fb},1}).$$

Suppose now that $\frac{\mu_a - \mathbf{E}_{\pi_{\text{fb},1}}[\theta]}{2} = \frac{g(\pi_{\text{fb},1},h,\ell)}{2\rho\sigma_1^2}W_0''(\pi_{\text{fb},1})$. Then, when π tends to $\pi_{\text{fb},1}$ from above we have

$$\begin{aligned} W(\pi) &= \mathbf{E}_{\pi_{\text{fb},1}}[\theta] + \left(\frac{\sigma_1^2+\sigma_2^2}{\sigma_1^2\sigma_2^2}\right) g(\pi_{\text{fb},1},h,\ell)\frac{W_0''(\pi_{\text{fb},1})}{2\rho} \\ &= \frac{\mu_a + \mathbf{E}_{\pi_{\text{fb},1}}[\theta]}{2} + g(\pi_{\text{fb},1},h,\ell)\frac{W_0''(\pi_{\text{fb},1})}{2\rho\sigma_2^2} \end{aligned}$$

and when π tends to $\pi_{\text{fb},1}$ from below we have

$$W(\pi) = \frac{\mu_a + \mathbf{E}_{\pi_{\text{fb},1}}[\theta]}{2} + g(\pi_{\text{fb},1},h,\ell)\frac{W_1''(\pi_{\text{fb},1})}{2\rho\sigma_2^2}.$$

However, by assumption $W_1''(\pi_{\text{fb},1}) < W_0''(\pi_{\text{fb},1})$ thus $W(\pi)$ will not be continuous at $\pi_{\text{fb},1}$, which is a contradiction. Hence, we must have

$$\frac{g(\pi_{\text{fb},1},h,\ell)}{2\rho\sigma_1^2}W_1''(\pi_{\text{fb},1}) \le \frac{\mu_a - \mathbf{E}_{\pi_{\text{fb},1}}[\theta]}{2} < \frac{g(\pi_{\text{fb},1},h,\ell)}{2\rho\sigma_1^2}W_0''(\pi_{\text{fb},1}).$$

Next, since $W_0''(\cdot)$ and $W_1''(\cdot)$ are continuous, there exists $\epsilon > 0$ so that when $\pi \in \left(\pi_{\text{fb},1} - \epsilon, \pi_{\text{fb},1}\right)$ then

$$\frac{\mu_a - \mathbf{E}_\pi[\theta]}{2} < \frac{g(\pi,h,\ell)}{2\rho\sigma_1^2} W_0''(\pi).$$

$$W_1''(\pi) < W_0''(\pi).$$

However, recall that $W_0(\pi_{\text{fb},1}) = W_1(\pi_{\text{fb},1})$ (the value matching condition), $W_0'(\pi_{\text{fb},1}) = W_1'(\pi_{\text{fb},1})$ (the smooth pasting condition), and $W_1''(\pi_{\text{fb},1}) < W_0''(\pi_{\text{fb},1})$ (by assumption) and $W_0(\cdot)$ and $W_1(\cdot)$ are convex.

Thus, by integrating we have that for all $\pi \in \left(\pi_{\text{fb},1} - \epsilon, \pi_{\text{fb},1}\right)$:

$$W_0(\pi) > W_1(\pi)$$

which is a contradiction for buyer 1's optimality condition when $\pi \in (0, \pi_{\text{fb},1}]$.
Case 2: Let $W_1''(\pi_{\text{fb},1}) > W_0''(\pi_{\text{fb},1})$. We know that for $\pi \in (\pi_{\text{fb},1}, 1]$ we have

$$\frac{\mu_a - \mathbf{E}_\pi[\theta]}{2} \leq \frac{g(\pi,h,\ell)}{2\rho\sigma_1^2} W_0''(\pi).$$

Moreover, since $W_0''(\cdot)$ and $W_1''(\cdot)$ are both continuous, there exists $\epsilon > 0$ such that for every $\pi \in (\pi_{\text{fb},1} - \epsilon, \pi_{\text{fb},1} + \epsilon)$ we have

$$\frac{\mu_a - \mathbf{E}_\pi[\theta]}{2} < \frac{g(\pi,h,\ell)}{2\rho\sigma_1^2} W_1''(\pi),$$

which is a contradiction when $\pi \in [0, \pi_{\text{fb},1})$. □

Lemma 2. *The first best policy is in cutoff strategies with* $\pi_{\text{fb},2} < \pi_{\text{fb},1}$.

Proof of Lemma 2. We consider the restriction of the second order differential equation to the three regions created by the posited cutoffs. Following Wronskian approach of second order ODEs (Zaitsev and Polyanin (2002)) we have

- Case 1: If $\pi \in [\pi_{\text{fb},2}, \pi_{\text{fb},1})$ then

$$W_1(\pi) = \frac{1}{2}\{\mu_a + \mathbf{E}[\mu_b]\} + \xi_1 \pi^{\frac{1}{2}(\lambda+1)}(1-\pi)^{-\frac{1}{2}(\lambda-1)} + \xi_2 \pi^{-\frac{1}{2}(\lambda-1)}(1-\pi)^{\frac{1}{2}(\lambda+1)}$$

with $\lambda = \sqrt{1 + \frac{8\rho\sigma_2^2}{(h-\ell)^2}}$.

- Case 2: If $\pi \geq \pi_{\text{fb},1}$, then

$$W_0(\pi) = \mathbf{E}_\pi[\mu_b] + \xi_0 \pi^{-\frac{1}{2}(\bar{\lambda}-1)} (1-\pi)^{\frac{1}{2}(\bar{\lambda}+1)}$$

where $\bar{\lambda} = \sqrt{1 + \frac{8\rho}{(h-\ell)^2 \left(\frac{\sigma_1^2+\sigma_2^2}{\sigma_1^2\sigma_2^2}\right)}}$.

- Case 3: If $\pi < \pi_{\text{fb},2}$, we trivially have

$$W(\pi) = \mu_a.$$

There are five unknowns that we have to identify to characterize the welfare-maximizing behavior: $\xi_0, \xi_1, \xi_2, \pi_{\text{fb},1}, \pi_{\text{fb},2}$. Therefore, we are going to use five conditions. First, we have the standard value matching

$$W_1(\pi_{\text{fb},1}) = W_0(\pi_{\text{fb},1}) \text{ and } \mu_a = W_1(\pi_{\text{fb},2})$$

and smooth pasting

$$W_1'(\pi_{\text{fb},1}) = W_0'(\pi_{\text{fb},1}) \text{ and } 0 = W_1'(\pi_{\text{fb},2}).$$

Combining this with Lemma 8 we can check that the value function has the desired shape.

□

Proposition 4. *Let $\pi_{\text{fb},2} < \pi_0 < \pi_{\text{fb},1}$. Define $\sigma(y) := \frac{y(1-y)(h-\ell)}{\sigma_2}$. Then,*

$$\begin{aligned}\mathbf{E}_{\pi_0}[\text{Beta phase}] &= \mathbf{E}_{\pi_0}[\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}] \\ &= \frac{\pi_{\text{fb},1} - \pi_0}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \int_{\pi_{\text{fb},2}}^{\pi_0} (y - \pi_{\text{fb},2}) \frac{2dy}{\sigma^2(y)} + \frac{\pi_0 - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \int_{\pi_0}^{\pi_{\text{fb},1}} (\pi_{\text{fb},1} - y) \frac{2dy}{\sigma^2(y)}.\end{aligned}$$

Particularly, $\lim_{\pi_0 \searrow \pi_{\text{fb},2}} \partial_{\pi_0} \mathbf{E}_{\pi_0}[\text{Beta phase}] > 0$ *and* $\lim_{\pi_0 \nearrow \pi_{\text{fb},1}} \partial_{\pi_0} \mathbf{E}_{\pi_0}[\text{Beta phase}] < 0$. *Moreover*

- *The probability of discarding the new product as a failure is*

$$Pr_{\pi_0}\{\text{discarding}\} = Pr_{\pi_0}\{\pi\Big(\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}\Big) = \pi_{\text{fb},2}\}$$
$$= \frac{\pi_{\text{fb},1} - \pi_0}{\pi_{\text{fb},1} - \pi_{\text{fb},2}}.$$

- *The probability that the new product serves the whole market*

$$Pr_{\pi_0}\{\text{serving the whole market}\} = Pr_{\pi_0}\{\pi\Big(\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}\Big) = \pi_{\text{fb},1}\}$$
$$= \frac{\pi_0 - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}}.$$

Particularly, $\partial_{\pi_0} Pr_{\pi_0}\{discarding\} < 0$, and $\partial_{\pi_0} Pr_{\pi_0}\{serving\ the\ whole\ market\} > 0$.

Proof of Proposition 4. We first note that the corresponding comparative statistics are immediate from the explicit characterizations of $\mathbf{E}_{\pi_0}[\text{Beta phase}]$, $Pr_{\pi_0}\{\text{serving the whole market}\}$ and $Pr_{\pi_0}\{\text{serving the whole market}\}$ and the fact that the endogenous $\pi_{\text{fb},2}$ and $\pi_{\text{fb},2}$ do not depend on π_0 (see Lemma 2).

Next, we explicitly derive $\mathbf{E}_{\pi_0}[\text{Beta phase}]$, $Pr_{\pi_0}\{\text{serving the whole market}\}$ and $Pr_{\pi_0}\{\text{serving the whole market}\}$.

To prove this statement we make a use of the following known result. For ease of notation let us define $\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1}) = \inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}$.

Theorem [Extended Feynman-Kac Formula]. *Let $\Phi(x), f(x), F(x)$, $x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}]$, be continuous functions (f is non-negative). Let $u(x)$, $x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}]$ be a solution to*

$$\frac{\sigma^2(x)}{2} u''(x) - (\lambda + f(x))u(x) = -\lambda\Phi(x) - F(x), \quad x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}]$$

and $u(\pi_{\text{fb},2}) = \Phi(\pi_{\text{fb},2})$ and $u(\pi_{\text{fb},1}) = \Phi(\pi_{\text{fb},1})$ then

$$u(x) = \mathbf{E}_x\left[\Phi\Big(\pi_{\tau\wedge\mathcal{H}(\pi_{\text{fb},2},\pi_{\text{fb},1})}\Big)e^{-\int_0^{\tau\wedge\mathcal{H}(\pi_{\text{fb},2},\pi_{\text{fb},1})}} f(\pi_s)ds + \int_0^{\tau\wedge\mathcal{H}(\pi_{\text{fb},2},\pi_{\text{fb},1})} F(\pi_s)e^{-\int_0^s f(\pi_r)dr}ds\right]$$

where τ is random variable with the density $\lambda e^{-\lambda t}\mathbf{1}_{t\in[0,\infty)}$.

The proof of the theorem follows by a simple extension of the celebrated Feynman-Kac formula, omitted.

Using this theorem (that we call it *Extended Feynman-Kac Formula*) we prove the proposition. To prove the proposition we use a lemma and 2 corollaries of *Extended Feynman-Kac Formula*.

First, in the following lemma we show that $\mathbf{E}_{\pi_0}[\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})] < \infty$.

Lemma 9. $\mathbf{E}_{\pi_0}[\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})] < \infty$.

Proof. The proof follows from the *Extended Feynman-Kac Formula*. To show it, consider a family of functions $\{u_\lambda(x) : x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}]\}_{\lambda \geq 0}$ that are solution to the following $\lambda-$parametric problem:

$$\frac{\sigma^2(x)}{2} u''(x) - \lambda u(x) = -1, \quad x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}] \tag{33}$$

and $u(\pi_{\text{fb},2}) = u(\pi_{\text{fb},1}) = 0$. From the *Extended Feynman-Kac Formula* it follows that $u_\lambda(x) = \mathbf{E}_x[\tau \wedge \mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})]$ for $\lambda > 0$. Next, we argue that $\sup_{\lambda > 0} u_\lambda(x) \leq u_0(x)$, where $u_0(x)$ solves (33) when $\lambda = 0$.

Next, since $\lim_{\lambda \to 0} \tau = \infty$ thus $\lim_{\lambda \to 0} \tau \wedge \mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1}) = \mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})$. Therefore $\mathbf{E}_{\pi_0}[H(\pi_{\text{fb},2}, \pi_{\text{fb},1})] < \infty$, finishing the proof. □

Next, we present two useful corollaries.

Corollary 1. *Let $f(x)$ and $F(x)$, $x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}]$, be continuous functions and $f(x)$ be non-negative. Let the function Φ be defined only at two points $\pi_{\text{fb},2}$ and $\pi_{\text{fb},1}$. Then the function*

$$q(x) = \mathbf{E}_x\left[\Phi(\pi_{\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})}) e^{-\int_0^{\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})} f(\pi_s) ds} + \int_0^{\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})} F(\pi_s) e^{-\int_0^s f(\pi_r) dr} ds\right] \tag{34}$$

is the solution of the following problem

$$\frac{\sigma^2(x)}{2} q''(x) - f(x) q(x) + F(x) = 0, \quad x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}], \tag{35}$$

and $q(\pi_{\text{fb},2}) = \Phi(\pi_{\text{fb},2})$ and $q(\pi_{\text{fb},1}) = \Phi(\pi_{\text{fb},1})$.

The proof of this corollary follows directly from *Extended Feynman-Kac Formula* by assuming $\lambda = 0$, replacing $u(x)$ with $q(x)$.

Corollary 2. *The solution of the problem*

$$\frac{\sigma^2(x)}{2} q''(x) + F(x) = 0, \quad x \in [\pi_{\text{fb},2}, \pi_{\text{fb},1}],$$

$q(\pi_{\text{fb},2}) = \Phi(\pi_{\text{fb},2})$ *and* $q(\pi_{\text{fb},1}) = \Phi(\pi_{\text{fb},1})$ *has the following form*

$$\begin{aligned} q(x) = {} & \frac{\pi_{\text{fb},1} - x}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \left(\Phi(\pi_{\text{fb},2}) + \int_{\pi_{\text{fb},2}}^{x} (y - \pi_{\text{fb},2}) \frac{2F(y)}{\sigma^2(y)} dy \right) \\ & + \frac{x - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \left(\Phi(\pi_{\text{fb},1}) + \int_{x}^{\pi_{\text{fb},1}} (\pi_{\text{fb},1} - y) \frac{2F(y)}{\sigma^2(y)} dy \right). \end{aligned}$$

The proof of this corollary is directly followed from *Extended Feynman-Kac Formula*.

Using the above two corollaries, we have $\Pr_{\pi_0}\{\pi\Big(\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}\Big) = \pi_{\text{fb},2}\} = \frac{\pi_{\text{fb},1} - \pi_0}{\pi_{\text{fb},1} - \pi_{\text{fb},2}}$

and $\Pr_{\pi_0}\{\pi\Big(\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}\Big) = \pi_{\text{fb},1}\} = \frac{\pi_0 - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}}$. These results follow from the above corollaries by assuming $F = f = 0, \Phi(\pi_{\text{fb},2}) = 1$ and $\Phi(\pi_{\text{fb},1}) = 0$.

In addition

$$\begin{aligned} & \mathbf{E}_{\pi_0}[\inf\{t : \pi_t \notin (\pi_{\text{fb},2}, \pi_{\text{fb},1})\}] = \mathbf{E}_{\pi_0}[\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})] = \\ & = \frac{\pi_{\text{fb},1} - \pi_0}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \int_{\pi_{\text{fb},2}}^{\pi_0} (y - \pi_{\text{fb},2}) \frac{2dy}{\sigma^2(y)} + \frac{\pi_0 - \pi_{\text{fb},2}}{\pi_{\text{fb},1} - \pi_{\text{fb},2}} \int_{\pi_0}^{\pi_{\text{fb},1}} (\pi_{\text{fb},1} - y) \frac{2dy}{\sigma^2(y)} \end{aligned}$$

which is followed by the above corollaries by assuming $F = 1, f = 0, \Phi(\pi_{\text{fb},2}) = \Phi(\pi_{\text{fb},1}) = 0$ (implying $q(\pi_0) = \mathbf{E}_{\pi_0}[\mathcal{H}(\pi_{\text{fb},2}, \pi_{\text{fb},1})]$ is the solution to (35)).

By these results, the proof of the proposition is now complete.

□

Lemma 3. *In every equilibrium with symmetric* π_m^**, the prices are as follows. If* $\pi < \pi_m^*$

$$p_a = \mu_a \qquad \textit{and} \qquad p_b \geq \mathbf{E}_\pi[\theta] - \mu_a. \tag{10}$$

If $\pi \geq \pi_m^*$

$$p_a \geq \mu_a - \mathbf{E}_\pi[\theta] \qquad \textit{and} \qquad p_b = \mathbf{E}_\pi[\theta]. \tag{11}$$

Proof of Lemma 3. The proof is in the text. □

Lemma 5. *In every equilibrium with symmetric* π^**, the prices are as follows.*
If $\pi < \pi^*$*:*

$$p_a = \mu_a - \mathbf{E}_\pi[\theta] \qquad \textit{and} \qquad p_b = 0. \tag{14}$$

If $\pi \geq \pi^*$*:*

$$p_a = p_b + \mu_a - \mathbf{E}_\pi[\theta] - \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} v''(\pi), \tag{15}$$

with

$$p_b \in \left[-\frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_b''(\pi),\ \mathbf{E}_\pi[\theta] - \mu_a + \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} \Big(v''(\pi) + w_a''(\pi) \Big) \right]. \tag{16}$$

Proof of Lemma 5. Consider a symmetric equilibrium with the cutoff π^*, thus for all $\pi > \pi^*$ all buyers submit their orders to seller b. Let $v(\cdot)$ denote a buyer's value function in this (symmetric) equilibrium. Let us consider $\pi > \pi^*$. Due to (8), for any $\pi > \pi^*$ we have

$$\mu_a - p_a + (n-1) \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} v''(\pi) \leq \mathbf{E}_\pi[\theta] - p_b + n \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} v''(\pi). \tag{36}$$

At the equilibrium, due to price competition between sellers, (36) holds with equality. Indeed, if the right hand side was larger, it would be profitable for seller b to slightly increase p_b, collecting higher per unit revenues and selling to the same number of buyers.

As a result, we must have

$$\underbrace{\Big(\mu_a - p_a\Big) - \Big(\mathbf{E}_\pi[\theta] - p_b\Big)}_{\substack{\text{opportunity cost of} \\ \text{choosing the risky product "}b\text{"}}} = \underbrace{\frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} v''(\pi)}_{\text{benefit of information}}$$

Hence, rearranging implies that

$$p_a - p_b = \mu_a - \mathbf{E}_\pi[\theta] - \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} v''(\pi). \tag{37}$$

Moreover, (12) implies that for any $\pi > \pi^*$

$$w_a(\pi) = n\frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_a''(\pi) \geq np_a \tag{38}$$

where the equality follows because when $\pi^* < \pi \leq 1$ then $\text{Vol}_a = 0$ and $\text{Vol}_b = n$, and the inequality follows because of the optimality condition of the value function in (12). Thus

$$p_a \leq \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_a''(\pi). \tag{39}$$

Similarly, we have

$$w_b(\pi) = np_b + n\frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_b''(\pi) \geq 0.$$

Thus,

$$p_b + \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_b''(\pi) \geq 0. \tag{40}$$

Using (37) and putting together (39) and (40) specify regions for the optimal prices

$$p_a \in \left[\mu_a - \mathbf{E}_\pi[\theta] - \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2}\Big(v''(\pi) + w_b''(\pi)\Big), \ \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_a''(\pi)\right], \tag{41}$$

and

$$p_b \in \left[-\frac{g\big(\pi,h,\ell\big)}{2\rho\sigma^2} w_b''(\pi),\ \mathbf{E}_\pi[\theta] - \mu_a + \frac{g\big(\pi,h,\ell\big)}{2\rho\sigma^2}\Big(v''(\pi) + w_a''(\pi)\Big) \right]. \tag{42}$$

Let us consider $\pi < \pi^*$. Due to (8), we have

$$\mu_a - p_a \geq \mathbf{E}_\pi[\theta] - p_b + \frac{g\big(\pi,h,\ell\big)}{2\rho\sigma^2} v''(\pi) \tag{43}$$

As before, in equilibrium (43) holds as an equality because of the competition between sellers. But since v is flat in that region, the result follows. □

Proposition 5. *The followings hold: (i) Any symmetric equilibrium is specified by a cutoff π_m^*. (ii) All the symmetric equilibria are efficient (i.e., welfare-maximizing), and we have*

$$\pi_m^* = \pi_{\text{fb}}.$$

Proof of Proposition 5. First, notice that the continuation value of each market participant is always nonnegative since the all have a strategy that guarantees a deterministic zero payoff. At the same time, observe that the pricing strategies of Lemma 3 for some cutoff π_m, it is optimal for the buyers to use the strategies

$$\xi_{i,a}(\pi, p_{i,a}, p_{i,b}) = 1 \quad \text{if and oly if } \mu_a - p_{i,a} = \max\{\mu_a - p_{i,a}, \mathbf{E}_\pi[\mu_b] - p_{i,b}, 0\} \tag{44}$$

$$\xi_{i,b}(\pi, p_{i,a}, p_{i,b}) = 1 \quad \text{if and oly if } \mathbf{E}_\pi[\mu_b] - p_{i,b} = \max\{\mu_a - p_{i,a}, \mathbf{E}_\pi[\mu_b] - p_{i,b}, 0\}. \tag{45}$$

Moreover, the induced expected discounted utility for the buyers is equal to 0. Therefore, by setting the cutoff equal to the welfare-maximizing one, the monopolist obtains the first-best welfare. Since we have noted that the continuation utilities of all market participants have to be nonnegative, using that cutoff is optimal for the monopolist. □

Proposition 6. *When the buyers are asymmetric, the equilibrium under a monopolistic market structure remains efficient.*

Proof of Proposition 6. The proof follows the same lines of Lemma 3 and Proposition 5. The unique nontrivial decision for the monopolist is who to serve with product b. Indeed, as in the proof of Lemma 3, it is always optimal to ask for a price equal to the willingness to pay of the consumers, and therefore the consumer value function is identically 0. But if v_a and v_b are equal to 0, we have that by setting the cutoff at the first best, the monopolist is able to get the entire surplus. □

Proposition 7. *When the buyers are asymmetric and the monopolist cannot discriminate, there is an efficient equilibrium. If the monopolist is forced to use the same price for both products equilibrium is not efficient.*

Proof of Proposition 7. For the equilibrium where the monopolist cannot discriminate, notice that in the proof of Proposition 6 (see also Lemma 3) that a pricing strategy that would have worked in constructing the welfare-maximizing equilibrium is

$$p_a = \mu_a \qquad \textit{and} \qquad p_b = \mathbf{E}_\pi[\theta]. \tag{46}$$

It is immediate to see that under this pricing strategy the buyers are always indifferent between the two products. Therefore, by letting

$$\xi_{ib}(t) = 1_{\{\pi_t \geq \pi_{i,fb}\}} \tag{47}$$

we obtain an equilibrium that is welfare-maximizing. The result where the monopolist is forced to use the same price for both products is trivial, because by Lemma 2 the first-best features two different thresholds. □

Proposition 8. *The followings hold: (i) Any symmetric equilibrium is specified by a cutoff π^*. (ii) The equilibrium with the highest sellers' profits is efficient, (i.e., welfare-maximizing). (iii) The consumer surplus is strictly higher than under monopoly. (iv) The value function of the two sellers is convex, and the value function of the buyers is concave.*

Proof of Proposition 8. Part (i) follows from the text. Since for $\pi > \pi^*$ it holds that

$$w_a(\pi) = n\frac{g\big(\pi, h, \ell\big)}{2\rho\sigma^2} w_a''(\pi) \tag{48}$$

following Wronskian approach of second order ODEs (Zaitsev and Polyanin (2002)) we have

$$w_a(\pi) = \zeta_1 \left[\pi^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} (1-\pi)^{\frac{1}{2}\left(1+\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} \right] + \zeta_2 \left[\pi^{\frac{1}{2}\left(1+\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} (1-\pi)^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} \right].$$

We note that $w_a(\pi)$ is bounded on $\pi^* < \pi \leq 1$, thus $\zeta_2 = 0$ (note that otherwise $\lim_{\pi\to 1}(1-\pi)^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}$ would explode in this region). Hence,

$$w_a(\pi) = \zeta_1 \left[\pi^{\frac{1}{2}\left(1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} (1-\pi)^{\frac{1}{2}\left(1+\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)} \right]. \tag{49}$$

Now, notice that by Lemma 5, we have that the utility of the consumers is decreasing in w_a'' in the case of the highest possible price. By the previous equation, we have that w_a'' is larger when w_a' is smaller. Therefore, the best possible scenario for the seller is the one in which w_a' is as small as possible. However, as it is well known, the viscosity solution cannot have a concave kink, and therefore w_a' is minimal when w_a satisfies the smooth pasting condition at π^*. Now, the value matching and smooth pasting for seller a at π^* give

$$w_a(\pi^*) = n\Big(\mu_a - \mathbf{E}_{\pi^*}[\theta]\Big), \tag{50}$$

$$w_a'(\pi^*) = n\,\frac{\partial}{\partial \pi}\Big(\mu_a - \mathbf{E}_{\pi}[\theta]\Big)\Big|_{\pi=\pi^*} = n(\ell - h). \tag{51}$$

Given (49) we have

$$\frac{w_a'(\pi^*)}{w_a(\pi^*)} = \left(\frac{1-\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}}{\pi^*} - \frac{1+\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}}{1-\pi^*} \right)\frac{1}{2} = \frac{\ell - h}{\mu_a - \mathbf{E}_{\pi^*}[\theta]}, \tag{52}$$

where the last equality follows from (50) and (51). Substituting $\mathbf{E}_{\pi^*}[\theta] = \pi^* h + (1-\pi^*)\ell$

and then solving (52) with respect to π^* implies

$$\pi^* = \frac{\mu_a - \ell + (\ell - \mu_a)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}{2\mu_a - (\ell + h) + (\ell - h)\sqrt{1 + 8\frac{\sigma^2\rho}{n(h-\ell)^2}}}. \tag{53}$$

Hence (due to (6)) $\pi^* = \pi_{\text{fb}}$, finishing the proof of (ii)

Moreover, from (50) we obtain that $\zeta_1 = n\frac{2(h-\mu_a)}{\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}-1}\left(\frac{\pi^*}{1-\pi^*}\right)^{\frac{1}{2}\left(1+\sqrt{1+8\frac{\sigma^2\rho}{n(h-\ell)^2}}\right)}$. As a result, $w_a''(\pi) = W''(\pi)$ for $\pi > \pi^*$(see Proposition 3), proving the convexity of w_a.

To simplify the notation, in the rest of the proof, set

$$g = \frac{2\sigma^2\rho}{(h-l)^2} \quad \text{and} \quad s = \sqrt{1 + 8\frac{\sigma^2\rho}{n(h-l)^2}} = \sqrt{1 + \frac{4g}{n}}.$$

It follows that

$$\begin{aligned} w_a'(\pi) &= \left(\frac{1-s}{2\pi} - \frac{1+s}{2(1-\pi)}\right) w_a(\pi) \\ w_a''(\pi) &= \frac{s^2-1}{4\pi^2(1-\pi)^2} w_a(\pi) = \frac{2\rho\sigma^2}{n(h-l)^2\pi^2(1-\pi)^2} w_a(\pi) = \frac{g}{n}\frac{w_a(\pi)}{\pi^2(1-\pi)^2} \end{aligned}$$

Recall that by Lemma 5 v is defined by the differential equation

$$\begin{aligned} v(\pi) = \mu + \frac{(h-l)^2\pi^2(1-\pi)^2}{2\rho\sigma^2}((n-1)v''(\pi) - w_a''(\pi)) &= \mu + \frac{n-1}{g}\pi^2(1-\pi)^2 v''(\pi) - \frac{w_a(\pi)}{n} \\ &= \mu + \frac{n-1}{g}\pi^2(1-\pi)^2 v''(\pi) - W\pi^{\frac{1-s}{2}}(1-\pi)^{\frac{1+s}{2}} \end{aligned}$$

with initial condition $v(\pi^*) = h\pi^* + l(1-\pi^*)$ and the stipulation that v is bounded on $(\pi^*, 1)$. We shall see presently that there is a unique v satisfying the bounded condition.

The function $f(\pi) = \pi^{\frac{1-s}{2}}(1-\pi)^{\frac{1+s}{2}}$ obeys the differential equation

$$f''(\pi) = \frac{s^2-1}{4}f(\pi) = \frac{g}{n}f(\pi) \qquad \text{on } (\pi^*, 1).$$

Therefore, the function $\tilde{v}(\pi) = v(\pi) - \mu - A_1 f(\pi)$ where

$$A_1 = \frac{n-1}{g}\frac{g}{n}A_1 - W, \quad \text{i.e.,} \quad A_1 = -nW = -\frac{2n(h-\mu)}{s-1}\left(\frac{\pi^*}{1-\pi^*}\right)^{\frac{1+s}{2}},$$

satisfies the differential equation

$$\tilde{v}(\pi) = C\pi^2(1-\pi)^2(\tilde{v})''(\pi), \quad C = \frac{n-1}{g}.$$

Setting

$$t = \sqrt{1+\frac{4}{C}} = \sqrt{1+\frac{4g}{n-1}},$$

then the solution space for $\tilde{v}$ is parameterized by two constants A_2 and C_2

$$\tilde{v}(\pi) = A_2\pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}} + C_2\pi^{\frac{1+t}{2}}(1-\pi)^{\frac{1-t}{2}}.$$

Since $t > 1$, it is clear $\tilde{v}(\pi)$ being bounded forces $C_2 = 0$. Putting everything together, and using $-nWf(\pi) = w_a(\pi)$, the solution to $v(\pi)$ is

$$v(\pi) = \mu - w_a(\pi) + A_2\pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}}$$

where A_2 is chosen to satisfy the initial condition $v(\pi^*) = \pi^*h + (1-\pi^*)l$. In particular, the definition of π^* implies

$$w_a(\pi^*) = n(\mu - (\pi^*h + (1-\pi^*)l) = n(\mu - v(\pi^*))$$

Therefore,

$$v(\pi^*) = \mu - n\mu + nv(\pi^*) + A_2(\pi^*)^{\frac{1-t}{2}}(1-\pi^*)^{\frac{1+t}{2}} \implies A_2 = \frac{(n-1)(\mu - v(\pi^*))}{(\pi^*)^{\frac{1-t}{2}}(1-\pi^*)^{\frac{1+t}{2}}}.$$

It will prove convenient to set

$$\Delta = \mu - v(\pi^*) = \mu - (\pi^*h + (1-\pi^*)l),$$

so that

$$A_2 = \frac{(n-1)\Delta}{(\pi^*)^{\frac{1-t}{2}}(1-\pi^*)^{\frac{1+t}{2}}}.$$

From the closed form solution of $v(\pi)$, we immediately get (iii), and we see that $v(\pi)$ is concave and that $\lim_{x\to 1} v(\pi) = \mu$. Furthermore, we have

$$v(\pi) + w_a(\pi) = \mu + A_2 \pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}},$$

so that

$$\pi^2(1-\pi)^2(v''(\pi) + w_a''(\pi)) = A_2 \frac{t^2-1}{4} \pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}} = A_2 \frac{g}{n-1} \pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}}.$$

Recall that by Lemma 5 w_b is defined via the differential equation

$$w_b(\pi) = n(\pi h + (1-\pi)l - \mu) + \frac{n}{g}\pi^2(1-\pi)^2(v''(\pi) + w_a''(\pi)) + \frac{n}{g}\pi^2(1-\pi)^2 w_b''(\pi).$$

It follows from the previous calculations that the function

$$\tilde{w}_b(\pi) = w_b(\pi) - n(\pi h + (1-\pi)l - \mu) + nA_2 \pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}}$$

satisfies the equation

$$\tilde{w}_b(\pi) = \frac{n}{g}\pi^2(1-\pi)^2(\tilde{w}_b)''(\pi).$$

Since $\tilde{w}_b(\pi)$ is also bounded, the same calculation above shows that $\tilde{w}_b = nB_2\pi^{\frac{1-s}{2}}(1-\pi)^{\frac{1+s}{2}}$ for some constant B_2. Therefore,

$$w_b(\pi) = n(\pi h + (1-\pi)l - \mu) - nA_2 \pi^{\frac{1-t}{2}}(1-\pi)^{\frac{1+t}{2}} + nB_2 \pi^{\frac{1-s}{2}}(1-\pi)^{\frac{1+s}{2}}.$$

To determine B_2, we can use the initial condition $w_b(\pi^*) = 0$. Thus,

$$0 = w_b(\pi^*) = -n\Delta - n(n-1)\Delta + nB_2(\pi^*)^{\frac{1-s}{2}}(1-\pi^*)^{\frac{1+s}{2}} \implies B_2 = \frac{n\Delta}{(\pi^*)^{\frac{1-s}{2}}(1-\pi^*)^{\frac{1+s}{2}}}.$$

From this closed form solution we see that w_b is convex, concluding the proof. □

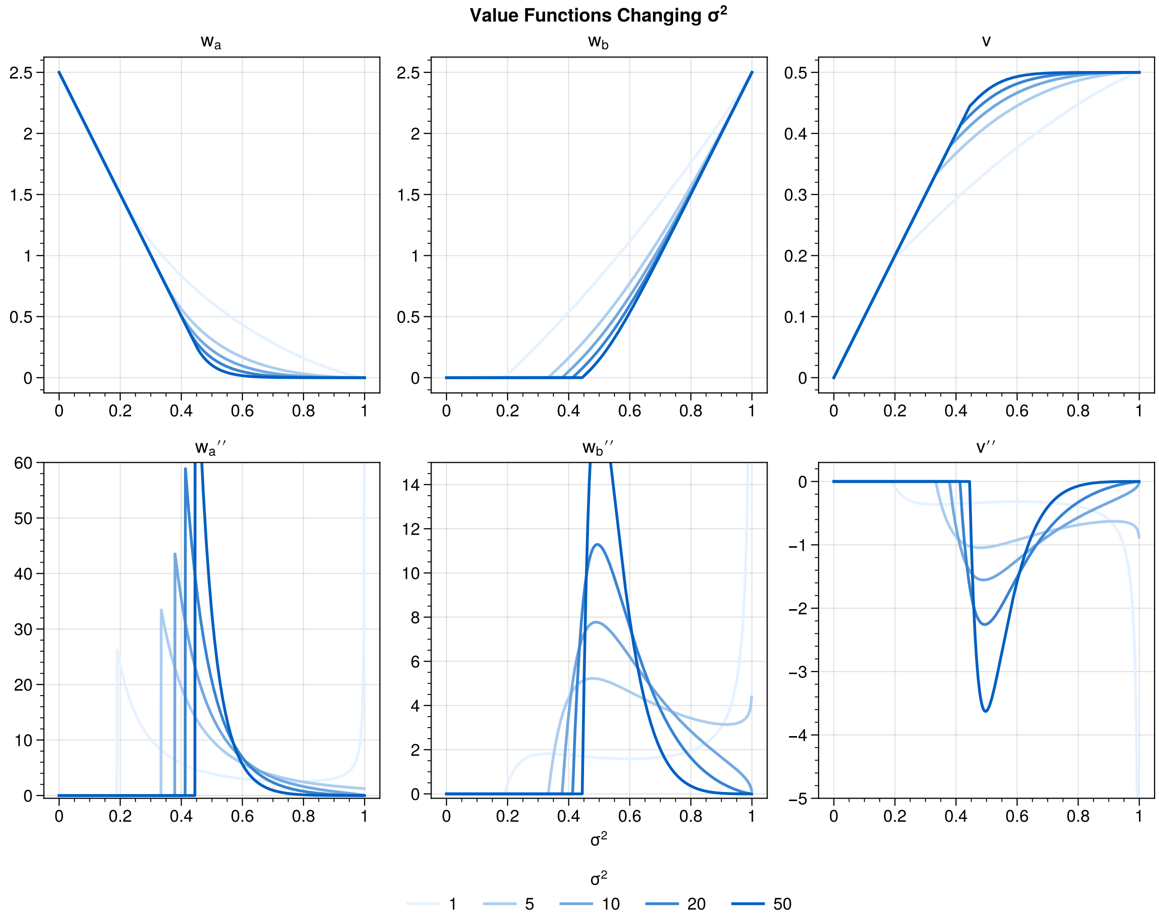

Figure 10: This figure plots the value functions w_a, w_b and v and their second derivatives w_a'', w_b'' and v'' when σ^2 changes, fixing other parameters to $n = 5, h = 1, \ell = 0, \mu = .5$ (using the explicit characterizations of the value functions in the proof of Proposition 8). As shown in the figure, with increasing σ^2 the cutoff π^*, expectedly, moves to the right (i.e., it increases). Moreover, it shows that v'' is concave (i.e., $v'' \leq 0$), w_a'' and w_b'' are convex (i.e., $w_a'' \geq 0$ and $w_b'' \geq 0$).

Proposition 9. *When the buyers are asymmetric, the equilibrium with the highest sellers' profit is not efficient. However, there is efficiency at the top, therefore,*

$$\pi_2 = \pi_{\text{fb},2}.$$

Finally, in this equilibrium the good learner has no incentive to mimic the bad learner.

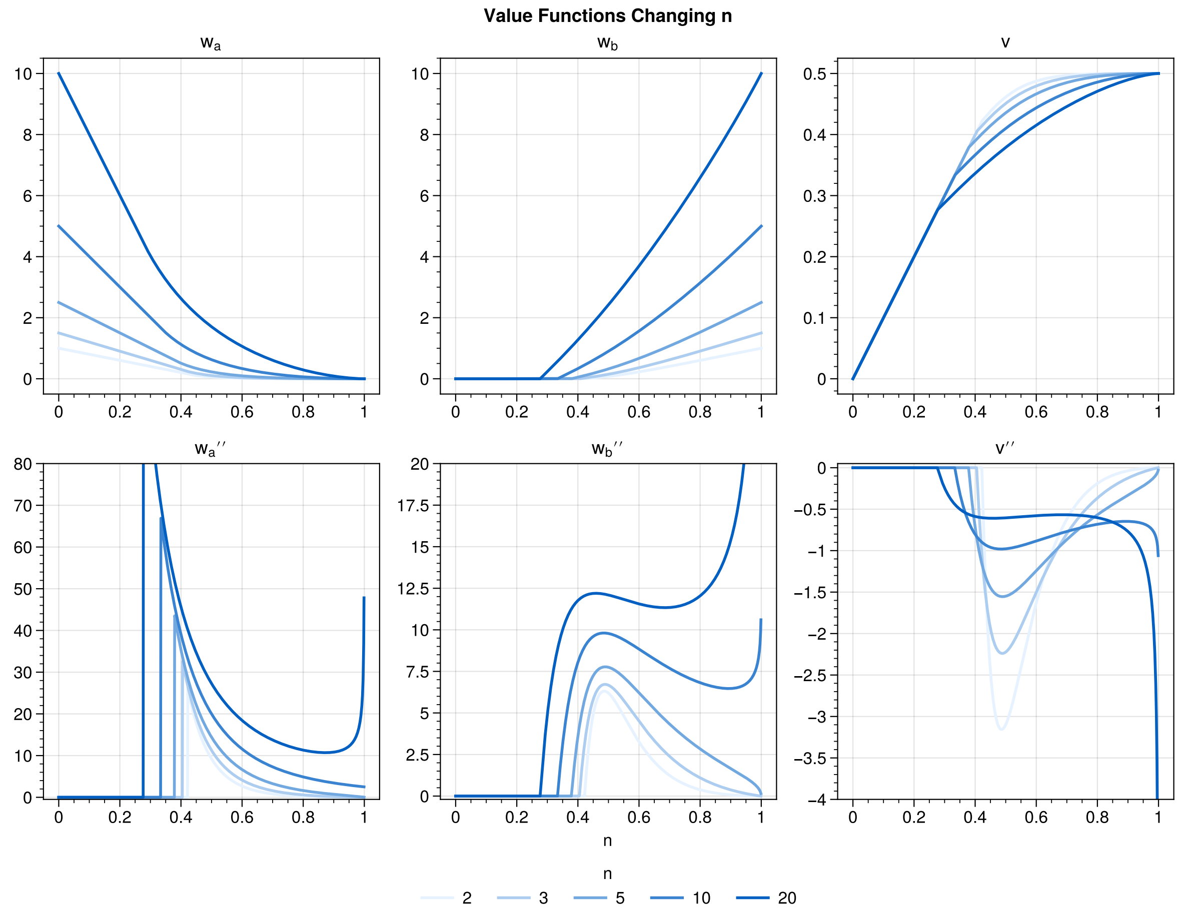

Figure 11: This figure plots the value functions w_a, w_b and v and their second derivatives w_a'', w_b'' and v'' when n changes, fixing other parameters to $\sigma^2 = 10, h = 1, \ell = 0, \mu = .5$ (using the explicit characterizations of the value functions in the proof of Proposition 8). As shown in the figure, with increasing n the cutoff π^*, expectedly, moves to the left (i.e., it decreases). Moreover, it shows that v'' is concave (i.e., $v'' \leq 0$), w_a'' and w_b'' are convex (i.e., $w_a'' \geq 0$ and $w_b'' \geq 0$).

Proof of Proposition 9. To simplify notation, for all $\pi \in (0,1)$ let

$$f(\pi) = g\Big(\pi, h, \ell\Big).$$

We first prove by contradiction that we must have

$$v_2''(\pi_{\text{fb},1}) < 0.$$

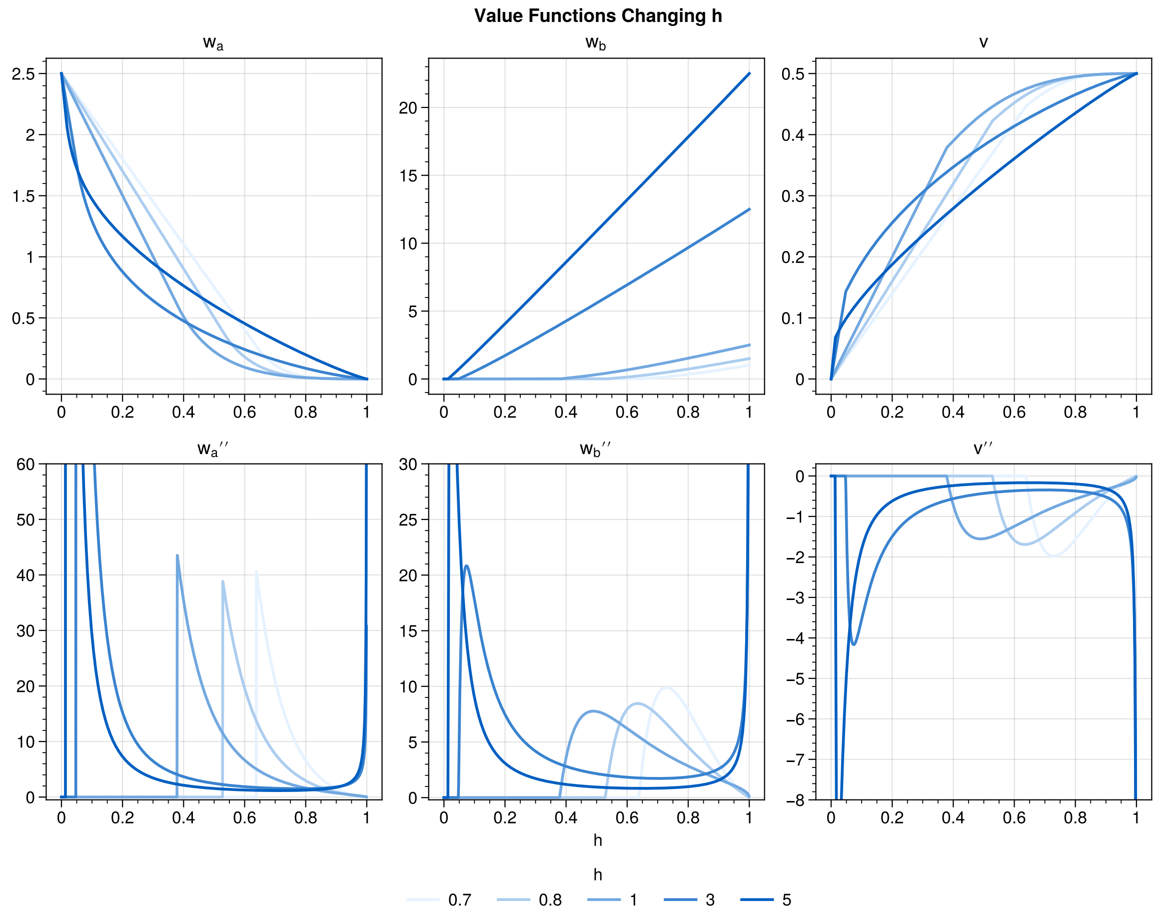

Figure 12: This figure plots the value functions w_a, w_b and v and their second derivatives w_a'', w_b'' and v'' when h changes, fixing other parameters to $\sigma^2 = 10, n = 5, \ell = 0, \mu = .5$ (using the explicit characterizations of the value functions in the proof of Proposition 8). As shown in the figure, with increasing h the cutoff π^*, expectedly, moves to the left (i.e., it decreases). Moreover, it shows that v'' is concave (i.e., $v'' \leq 0$), w_a'' and w_b'' are convex (i.e., $w_a'' \geq 0$ and $w_b'' \geq 0$).

Indeed, suppose $v''(\pi_{\text{fb},1}) \geq 0$. Since the equilibrium pricing conditions of Lemma 5 are easily seen to hold with heterogeneity as well, in the region $(\pi_{\text{fb},2}, \pi_{\text{fb},1})$ we have

$$v_2(\pi) = \mu_a - \frac{g\Big(\pi, h, \ell\Big)}{2\rho\sigma_2^2} w_a''(\pi) \implies v_2''(\pi) = -\frac{\partial^2 \left(\frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_2^2} w_a''(\pi) \right)}{\partial \pi^2}$$

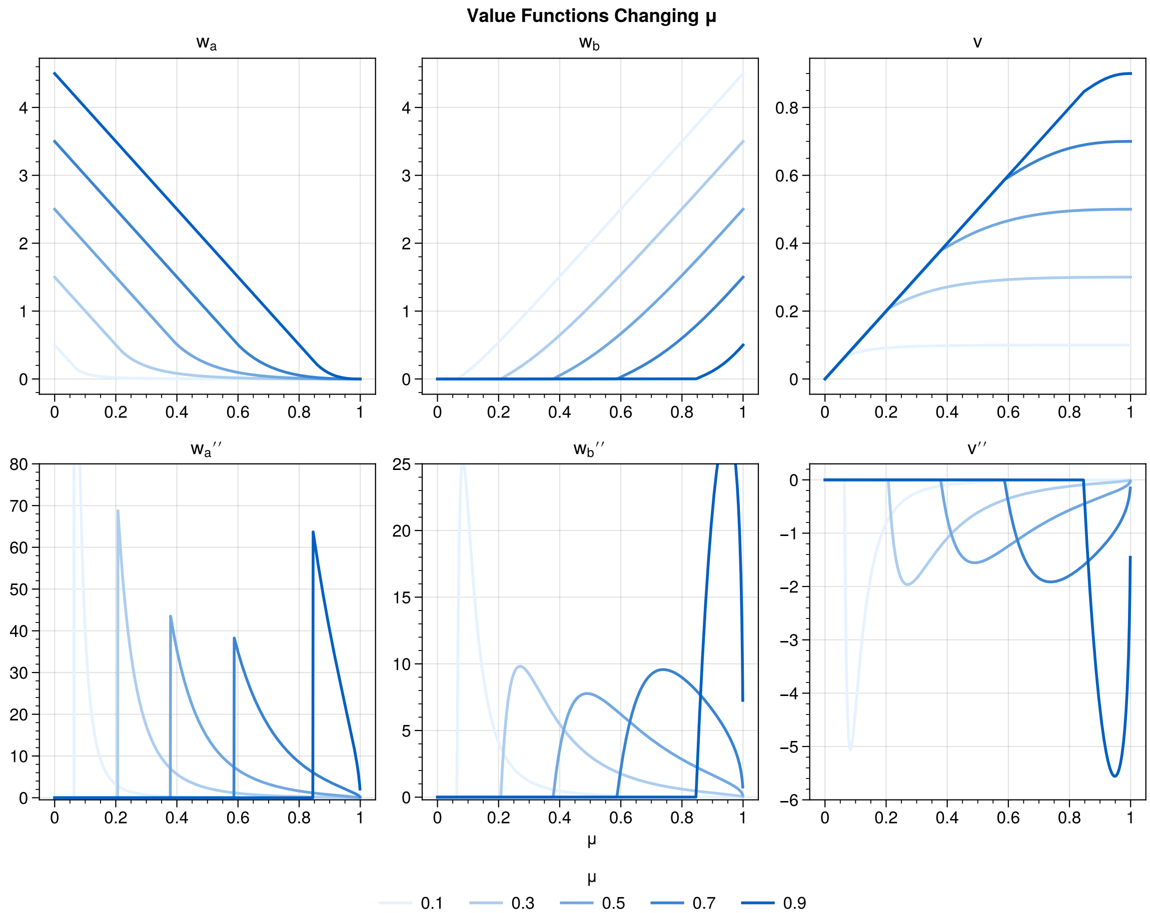

Figure 13: This figure plots the value functions w_a, w_b and v and their second derivatives w_a'', w_b'' and v'' when μ changes, fixing other parameters to $\sigma^2 = 10, n = 5, \ell = 0, h = 1$ (using the explicit characterizations of the value functions in the proof of Proposition 8). As shown in the figure, with increasing μ the cutoff π^*, expectedly, moves to the right (i.e., it increases). Moreover, it shows that v'' is concave (i.e., $v'' \leq 0$), w_a'' and w_b'' are convex (i.e., $w_a'' \geq 0$ and $w_b'' \geq 0$).

But notice that

$$w_a(\pi) = \frac{g\big(\pi, h, \ell\big)}{2\rho\sigma_2^2} w_a''(\pi) \implies w_a''(\pi) = \frac{\partial^2 \left(\frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_2^2} w_a''(\pi) \right)}{\partial \pi^2} = -v_2''(\pi) \tag{54}$$

so that $v_2''(\pi_{fb,1}) \geq 0$ implies $w_a''(\pi_{fb,1}) \leq 0$. But using $w_a(\pi) = \frac{g\left(\pi,h,\ell\right)}{2\rho\sigma_2^2} w_a''(\pi)$ again, this would imply that $w_a(\pi_{\text{fb},1}) = 0$. But this contradicts the fact that the instantaneous payoff of seller a are always weakly positive and strictly positive for π sufficiently close to 0. Then, plug equations (17) and (18) into seller's b HJB equation

$$w_b(\pi) = \max_{p_{b,1},p_{b,2}} \left\{ \begin{array}{c} 0, \\ p_{2,a} + \mathbf{E}_\pi[\mu_b] - \mu_a + \frac{1}{2\rho\sigma_2^2} f(\pi) v_1''(\pi) + \frac{1}{2\rho\sigma_2^2} f(\pi) w_b''(\pi), \\ p_{1,a} + \mathbf{E}_\pi[\mu_b] - \mu_a + \frac{1}{2\rho\sigma_1^2} f(\pi) v_1''(\pi) + \frac{1}{2\rho\sigma_1^2} f(\pi) w_b''(\pi), \\ p_{2,a} + \mathbf{E}_\pi[\mu_b] - \mu_a + \frac{1}{2\rho\sigma_2^2} f(\pi) v_1''(\pi) + p_{1,a} + \mathbf{E}_\pi[\mu_b] - \mu_a \\ + \frac{1}{2\rho\sigma_1^2} f(\pi) v_1''(\pi) + \frac{\sigma_1^2+\sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} f(\pi) w_b''(\pi). \end{array} \right\}$$

Next, we show that the equilibrium is a cutoff one, but with cutoffs different from the first best. Denote the threshold obtained under competition as $\bar{\pi}_1 > \bar{\pi}_2$, with the usual interpretation that buyer i buys from seller b if and only if $\pi \geq \bar{\pi}_i$. To find $\bar{\pi}_1$ we need to have seller a indifferent between selling only to 1 and not selling at all

$$p_{1,a} + \frac{1}{2\rho\sigma_2^2} f(\bar{\pi}_1) w_a''(\bar{\pi}_1) = \frac{\sigma_1^2 + \sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} f(\bar{\pi}_1) w_a''(\bar{\pi}_1)$$

and seller b indifferent between selling to both buyers or only to buyer 2:

$$\begin{aligned} & p_{2,a} + \mathbf{E}_{\bar{\pi}}[\mu_b] - \mu_a + \frac{1}{2\rho\sigma_2^2} f(\bar{\pi}_1) v_1''(\bar{\pi}_1) + \frac{1}{2\rho\sigma_2^2} f(\bar{\pi}_1) w_b''(\bar{\pi}_1) \\ = \quad & p_{2,a} + \mathbf{E}_{\bar{\pi}}[\mu_b] - \mu_a + \frac{1}{2\rho\sigma_2^2} f(\bar{\pi}_1) v_1''(\bar{\pi}_1) + p_{1,a} + \mathbf{E}_{\bar{\pi}}[\mu_b] - \mu_a \\ & + \frac{1}{2\rho\sigma_1^2} f(\bar{\pi}_1) v_1''(\bar{\pi}_1) + \frac{\sigma_1^2 + \sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} f(\bar{\pi}_1) w_b''(\bar{\pi}_1) \end{aligned}$$

summing the two equations we get:

$$\frac{\left(\mu_a - \mathbf{E}_{\bar{\pi}_1}[\mu_b]\right) 2\rho\sigma_1^2}{f(\bar{\pi}_1)} = v_1''(\bar{\pi}_1) + w_b''(\bar{\pi}_1) + w_a''(\bar{\pi}_1). \tag{55}$$

Similarly, we obtain

$$\frac{\left(\mu_a - \mathbf{E}_{\bar{\pi}_2}[\mu_b]\right) 2\rho\sigma_2^2}{f(\bar{\pi}_2)} = v_2''(\bar{\pi}_2) + w_b''(\bar{\pi}_2) + w_a''(\bar{\pi}_2), \tag{56}$$

whereas

$$\frac{\left(\mu_a - \mathbf{E}_{\pi_{\text{fb},1}}[\mu_b]\right) 2\rho\sigma_1^2}{f(\pi_{\text{fb},1})} = v_2''(\pi_{\text{fb},1}) + v_1''(\pi_{\text{fb},1}) + w_b''(\pi_{\text{fb},1}) + w_a''(\pi_{\text{fb},1}). \tag{57}$$

Similarly, we obtain

$$\frac{\left(\mu_a - \mathbf{E}_{\pi_{\text{fb},2}}[\mu_b]\right) 2\rho\sigma_2^2}{f(\pi_{\text{fb},2})} = v_1''(\pi_{\text{fb},2}) + v_2''(\pi_{\text{fb},2}) + w_b''(\pi_{\text{fb},2}) + w_a''(\pi_{\text{fb},2}). \tag{58}$$

However, we have shown above that $v_2''(\pi_{\text{fb},1}) \neq 0$. The proof that the cutoff for buyer 2 is equal to the first best one follows by the same kind of argument paired but observing at the end that $v_1''(\pi_{\text{fb},2}) = 0$. Finally, notice that during the Beta phase, by equation (54), the value for the good learner is μ_a, and that by mimicking the bad learner he would get $\mathbf{E}[\mu_b]$. Since in the Beta phase $\mathbf{E}[\mu_b] < \mu_a$ the final part of the result follows. □

Proposition 10. *Fix σ_2. The equilibrium with a buyer that does not generate any information about the quality of the product is efficient:*

$$\lim_{\sigma_1 \to \infty} (\pi_1 - \pi_{\text{fb},1}) = 0$$

Proof of Proposition 10. The result follows immediately by rewriting Equation as

$$\left(\mu_a - \mathbf{E}_{\bar{\pi}_1}[\mu_b]\right) 2\rho\sigma_1^2 = f(\bar{\pi}_1)(v_1''(\bar{\pi}_1) + w_b''(\bar{\pi}_1) + w_a''(\bar{\pi}_1)). \tag{59}$$

and noticing that as σ_1 goes to infinity $v_{1,\sigma}$ is converging to the value function of the buyer 1 in the equilibrium of the economy in which the belief diffusion is defined as

$$d\pi_t = \pi_t(1-\pi_t)(h-\ell)\sqrt{\frac{\xi_{2b}(t)}{\sigma_2^2}}dZ_t.$$

and that this latter value function is bounded. □

Lemma 4. *At $\pi = 0$ all buyers buy from seller a:*

$$Vol_a(0) = 2, \qquad Vol_b(0) = 0,$$

and the equilibrium prices are given by

$$p_a(0) - p_b(0) = \mu_a - \ell,$$

with

$$\ell - \mu_a \leq p_b(0) \leq 0.$$

Proof of Lemma 4. Suppose first that $p_a - p_b < \mu_a - \ell$. In that case, every buyer chooses product a. If $p_a < 0$, it would be a strictly profitable deviation for seller a to set a price equal to 0. If $p_a \geq 0$, it would be strictly preferable to sell at a slightly higher price and continue to serve the entire market. If instead $p_a - p_b < \mu_a - \ell$, it is seller b that is serving the entire market and that can always finds a profitable deviation; to $p_b = 0$ if $p_b < 0$, and to a slightly higher price if $p_b \geq 0$.

Finally, consider the case in which $p_a - p_b = \mu_a - \ell$. Since at least one of the seller is not serving the entire market, if $p_b > 0$ then one of them has the incentive to slightly undercut the other obtaining the entire market. If $p_b = 0$ and seller a is not serving the entire market, he has an incentive to offer a price slightly below $\mu_a - \ell$ and serve the entire market. If $p_b < 0$ and seller a is not serving the entire market, seller b is making negative profits and he has an incentive to switch to $p_b = 0$. Finally, if $p_b < \mu_a - \ell$ any seller who is serving a some buyer is making negative profits, and therefore he is better off by offering a price equal to 0. □

Proposition 11. *When there is one seller that can selling product a, and $m - 1$ sellers selling product b, and buyers are asymmetric, the equilibrium is not efficient, and the amount of experimentation is too low:*

$$\forall i \quad \pi_i = \pi_{\text{myopic}} < \pi_{\text{fb},i}.$$

Proof of Proposition 11. Notice that with multiple sellers of product b we now need to introduce some additional notation for the prices offered. In particular, for all $i \in \{1, 2\}$,

$p_{i,a}$ continues to denote the price asked to buyer i by the unique seller of product b, whereas for all $j \in \{1,...,m-1\}$ and $i \in \{1,2\}$, p_{i,b_j} denotes the price asked to buyer i by the j-th seller of product b. We conjecture that the equilibrium is as below, and we then check that it is correct. The pricing strategy are:

$$
\begin{aligned}
p_{i,b_j}(\pi) &= 0 \qquad\qquad\qquad\qquad \forall \pi \in [0,1], i \in \{1,2\}, j \in \{1,...,m-1\}, \\
p_{i,a}(\pi) &= \max\{\mu_a - \mathbf{E}_\pi[\mu_b], 0\} \qquad \forall \pi \in [0,1], i \in \{1,2\},
\end{aligned}
$$

The equilibrium strategy postulated for an arbitrary buyer i is to always maximize the one period payoff, to buy product a when indifferent between two products, and to accept the offer of the seller with the lowest index when indifferent between the prices offered by multiple sellers of product b. Given the pricing strategies, this means that she buys product b if and only if $\mathbf{E}_{\pi_t}[\mu_b] > \mu_a$. Indeed, denote as π_{myopic} the belief such that $\mathbf{E}_{\pi_{myopic}}[\mu_b] = \mu_a$. If $\pi \leq \pi_{myopic}$, the buyer can choose between buying the product b at price 0, and buying product a at price $\mu_a - \mathbf{E}_\pi[\mu_b]$. Since both choices induce an immediate payoff of $\mathbf{E}_\pi[\mu_b]$, the buyer is indifferent and buying a is maximizes the immediate payoff. If $\pi > \pi_{myopic}$, buyer i can decide to choose between buying the product b at price 0 (it does not matter from which seller), and buying product a at price 0. Since $\mathbf{E}_\pi[\mu_b] > \mu_a$, buying product b maximizes the immediate payoff.

Since the threshold is the same for buyers with different σ, it is immediate to see that this outcome is not efficient. We now check that it is indeed an equilibrium. It is immediate to see that the resulting value function for the j-th seller of product b is

$$w_{b_j}(\pi) = 0 \quad \forall \pi \in [0,1].$$

By the Martingale property of beliefs, we also have

$$v_i(\pi) = \rho \mathbf{E}_\pi[\mu_b].$$

Finally,

$$w_a(\pi) = \begin{cases} \mu_a - \mathbf{E}_\pi[\mu_b] & \text{if } \mathbf{E}_\pi[\mu_b] \leq \mu_a \\ 0 + \frac{\sigma_1^2+\sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} g(\pi,h,\ell)\, w_a''(\pi) & \text{if } \mathbf{E}_\pi[\mu_b] > \mu_a \end{cases}. \tag{60}$$

From these value function, it is immediate to check that the one proposed is indeed an equilibrium.

Consider first the j-th seller of product b, b_j. Since w_{b_j} is constant, at every belief the best choice for b_j is to maximize the immediate payoff. Given $(p_{i,b_{j'}})_{j' \neq j}, p_{i,a}$, and $\xi_{i,b}$, the pricing strategy $p_{i,b,E} = 0$ indeed maximize the immediate payoff. Indeed, a lower price would induce weakly negative immediate payoff, whereas with an higher price he would not sell the product, achieving an immediate payoff of 0.

For an arbitrary buyer i, $v_i''(\pi) = 0$ implies that the maximization of immediate payoffs prescribed by the equilibrium is optimal.

Finally, consider the seller of product a. If $\pi < \pi_{myopic}$, $w_a''(\pi) = 0$ and therefore she maximizes the immediate payoffs. Given that for all $j \in \{1, ..., m-1\}$, $(p_{i,b_j})(\pi) = 0$, for every buyer she can decide whether to sell the product a, at a maximal price $\mu_a - \mathbf{E}_\pi[\mu_b]$, or not to sell. Since $\mu_a - \mathbf{E}_\pi[\mu_b] > 0$, the prescribed behavior induce higher immediate payoffs. Finally, consider $\pi \geq \pi_{myopic}$. By (61), and since $w_a(\pi) \geq 0$, and $\frac{\sigma_1^2+\sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} g(\pi, h, \ell) \geq 0$, it follows that $w_a''(\pi) \geq 0$ as well. Notice that when considering what to sell to buyer i, the immediate payoffs of selling product a at the maximal price $\mu_a - \mathbf{E}_\pi[\mu_b]$ is negative, and in that case no learning is generated. Instead, by not selling product a the incumbent obtain an immediate payoff of 0, plus a weakly positive learning term $\frac{g(\pi,h,\ell)}{2\rho\sigma_i^2} w_I''(\pi)$. Therefore, the prescribed strategy is optimal. □

Proposition 12. *When there is a seller that can serve both markets, one seller that can serve only market b, and buyers are asymmetric, the equilibrium is not efficient, and the amount of experimentation is too low:*

$$\forall i \quad \pi_i = \pi_{\text{myopic}} < \pi_{\text{fb},i}.$$

Proof of Proposition 12. Notice that we need a small change in notation, because now both sellers compete in market b, and therefore we have to denote differently their price. We identify, the seller serving both market as the incumbent (I), and we denote the price she charges in market b to buyer i as $p_{i,b,I}$. Similarly, the seller who only operates in market b is identified as the entrant (E) and she charges price $p_{i,b,E}$. We conjecture that the equilibrium is as below, and we then check that it is correct. The pricing strategy are:

$$\begin{aligned} p_{i,b,E}(\pi) &= p_{i,b,I}(\pi) = 0 \quad \forall \pi \in [0,1], i \in \{1,2\}, \\ p_{i,a}(\pi) &= \max\{\mu_a - \mathbf{E}_\pi[\mu_b], 0\} \quad \forall \pi \in [0,1], i \in \{1,2\}, \end{aligned}$$

The equilibrium strategy postulated for an arbitrary buyer i is to always maximize the one period payoff, and when indifferent, buy from the incumbent (I), possibly product a. Given the pricing strategies, this means that she buys product b if and only if $\mathbf{E}_{\pi_t}[\mu_b] > \mu_a$. Indeed, denote as π_{myopic} the belief such that $\mathbf{E}_{\pi_{myopic}}[\mu_b] = \mu_a$. If $\pi \leq \pi_{myopic}$, the buyer can choose between buying the product b at price 0, and buying product a at price $\mu_a - \mathbf{E}_\pi[\mu_b]$. Since both choices induce an immediate payoff of $\mathbf{E}_\pi[\mu_b]$, the buyer is indifferent and buying a is maximizes the immediate payoff. If $\pi > \pi_{myopic}$, buyer i can decide to choose between buying the product b at price 0 (it does not matter from which seller), and buying product a at price 0. Since $\mathbf{E}_\pi[\mu_b] > \mu_a$, buying product b maximizes the immediate payoff.

Since the threshold is the same for buyers with different σ, it is immediate to see that this outcome is not efficient. We now check that it is indeed an equilibrium. It is immediate to see that the resulting value function for the entrant E is

$$w_E(\pi) = 0 \quad \forall \pi \in [0,1].$$

By the Martingale property of beliefs, we also have

$$v_i(\pi) = \rho \mathbf{E}_\pi[\mu_b].$$

Finally,

$$w_I(\pi) = \begin{cases} \mu_a - \mathbf{E}_\pi[\mu_b] & \text{if } \mathbf{E}_\pi[\mu_b] \leq \mu_a \\ 0 + \frac{\sigma_1^2 + \sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} g(\pi, h, \ell)\, w_I''(\pi) & \text{if } \mathbf{E}_\pi[\mu_b] > \mu_a \end{cases}. \tag{61}$$

From these value function, it is immediate to check that the one proposed is indeed an equilibrium. Consider first the entrant (E). Since w_E is constant, at every belief the best choice for E is to maximize the immediate payoff, something that, given $p_{i,b,I}(\pi), p_{i,a}(\pi)$, and $\xi_{i,b}(t)$ is achieved by the pricing strategy $p_{i,b,E}(\pi) = 0$. Indeed, a lower price would induce weakly negative immediate payoff, whereas with an higher price he would not sell the product, achieving an immediate payoff of 0.

For an arbitrary buyer i, $v_i''(\pi) = 0$ implies that the maximization of immediate payoffs prescribed by the equilibrium is optimal.

Finally, consider the incumbent (I). If $\pi < \pi_{myopic}$, $w_I''(\pi) = 0$ and therefore she maximizes the immediate payoffs. Given $p_{i,b,E}(\pi) = 0$, for every buyer she can decide whether to sell product a, at a maximal price $\mu_a - \mathbf{E}_{\pi_{myopic}}[\mu_b]$, or product b at a maximal price of 0. Since $\mu_a - \mathbf{E}_{\pi_{myopic}}[\mu_b] > 0$, the prescribed behavior induce higher immediate payoffs. Finally, consider $\pi \geq \pi_{myopic}$. By (61), and since $w_I(\pi) \geq 0$, and $\frac{\sigma_1^2+\sigma_2^2}{2\rho\sigma_1^2\sigma_2^2} g(\pi, h, \ell) \geq 0$, it follows that $w_I''(\pi) \geq 0$ as well. Notice that when considering what to sell to buyer i, the immediate payoffs of selling product a at the maximal price $\mu_a - \mathbf{E}_{\pi_{myopic}}[\mu_b]$ is negative, and in that case no learning is generated. Instead, by selling product b at the maximal price 0, the incumbent obtain an immediate payoff of 0, plus a weakly positive learning term $\frac{g(\pi,h,\ell)}{2\rho\sigma_i^2} w_I''(\pi)$. Therefore, the prescribed strategy is optimal. □

A Viscosity solution

Here, we argue for the buyer i's value function. Similarly, one can prove it for sellers (both for monopoly and oligopoly).

Define buyer i's value function for a given pricing strategy of the sellers as

$$v_i(\pi) = \sup_{\xi_{ik}, k \in \{a,b\}} \mathbf{E}\left[\int_0^\infty \rho e^{-\rho t}\, \xi_{ik}(t)\Big(dC_{ki}(t) - p_{ki}(t)dt\Big)\right]. \tag{62}$$

Theorem 1. *If p is an equilibrium pricing strategy for the seller then the above value function is a viscosity solution to the HJB equation in Section 4.*

Proof. To prove this result the following lemma is useful.

Lemma 10. *Suppose*

(i) that $v_i(\cdot)$ is a polynomial growth function, that is there exist constants q and m so that $|v_i(\pi)| \leq q(1+\pi^m)$

(ii) for any stopping time τ

$$v_i(\pi) = \sup_{\xi_{ik}, k\in\{a,b\}} \left[\mathbf{E} \int_0^{\tau} \rho e^{-\rho t}\ \xi_{ik}(t) \Big(dC_{ki}(t) - p_{k,i}(t)dt \Big) + \rho e^{-\rho\tau} v_i(\pi_\tau) \right].$$

Then v_i is a viscosity solution to the HJB equation in Section 4.

Proof of Lemma 10. The proof follows directly from Theorem 5.1, Chapter 8, on Fleming and Soner (2006) □

To use the above lemma we need to prove conditions (i) and (ii). The following lemma proves condition (ii) holds.

Lemma 11. *The function $v_i(\cdot)$ is continuous and satisfies the dynamic programming principle.*

Proof of Lemma 11. The proof follows directly from Lemma 2.1, Chapter 5, on Fleming and Soner (2006) □

Next, we need to show that condition (i) is also satisfied. First notice that for every $\pi \in [0,1]$, $hn \geq W(\pi)$. Moreover, under any equilibrium strategy the continuation value of each market participant is always nonnegative since the all have a strategy that guarantees a deterministic zero payoff. In turns, this implies that the the continuation value of each market participant is always weakly smaller than $W(\pi) \leq hn$. Therefore, condition $v_i(\pi) \leq hn$ and (i) is also satisfied, finishing the proof of the claim.

□

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2019. "Too much data: Prices and inefficiencies in data markets." National Bureau of Economic Research.

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar. 2011. "Experimentation, patents, and innovation." *American Economic Journal: Microeconomics*, 3(1): 37–77.

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar. 2014. "Dynamics of Information Exchange in Endogenous Social Networks." *Theoretical Economics*, 9(1): 41–97.

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The economics of privacy." *Journal of economic Literature*, 54(2): 442–92.

Admati, Anat R, and Paul Pfleiderer. 1986. "A monopolistic market for information." *Journal of Economic Theory*, 39(2): 400–438.

Ajorlou, Amir, Ali Jadbabaie, and Ali Kakhbod. 2018. "Dynamic Pricing in Social Networks: The Word-of-Mouth Effect." *Management Science*, 64(2): 971–979.

Anton, James J, and Dennis A Yao. 2002. "The sale of ideas: Strategic disclosure, property rights, and contracting." *The Review of Economic Studies*, 69(3): 513–531.

Babaioff, Moshe, Robert Kleinberg, and Renato Paes Leme. 2012. "Optimal mechanisms for selling information." 92–109.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp. 2018. "Big data in finance and the growth of large firms." *Journal of Monetary Economics*, 97: 71–87.

Bergemann, Dirk, and Alessandro Bonatti. 2019. "Markets for information: An introduction." *Annual Review of Economics*, 11: 85–107.

Bergemann, Dirk, and Juuso Välimäki. 2000. "Experimentation in markets." *Review of Economic Studies*, 67(2): 213–234.

Board, Simon. 2008. "Durable-Goods Monopoly with Varying Demand." *Review of Economic Studies*, 75(2): 391–413.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic experimentation." *Econometrica*, 67(2): 349–374.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2016. "Competition for attention." *The Review of Economic Studies*, 83(2): 481–513.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2017. "Memory, attention, and choice." National Bureau of Economic Research.

Caballero, Ricardo, and Alp Simsek. 2019. "Prudential Monetary Policy." *Working paper*.

Cabral, Luis M.B., David J. Salant, and Glenn A. Woroch. 1999. "Monopoly pricing with network externalities." *International Journal of Industrial Organization*, 17(2): 199–214.

Caminal, Roman, and Xavier Vives. 1999. "Price dynamics and consumer learning." *Journal of Economics & Management Strategy*, 8(1): 95–131.

Candogan, Ozan, Kostas Bimpikis, and Asuman Ozdaglar. 2012. "Optimal Pricing in Networks with Externalities." *Operations Research*, 60(4): 883–905.

Conlisk, John, Eitan Gerstner, and Joel Sobel. 1984. "Cyclic Pricing by a Durable Goods Monopolist." *Quarterly Journal of Economics*, 99(3): 489–505.

Crapis, Davide, Bar Ifrach, Costis Maglaras, and Marco Scarsini. 2017. "Monopoly Pricing in the Presence of Social Learning." *Management Science*, 63(11): 3586–3608.

Décamps, Jean-Paul, Thomas Mariotti, and Stéphane Villeneuve. 2006. "Irreversible investment in alternative projects." *Economic Theory*, 28(2): 425–448.

Dixit, Avinash K. 1993. *The art of smooth pasting.* Vol. 2, Taylor & Francis.

Eliaz, Kfir, Ran Eilat, and Xiaosheng Mu. 2019. "Optimal Privacy-Constrained Mechanisms."

Eső, Péter, and Balazs Szentes. 2007. "Optimal information disclosure in auctions and the handicap auction." *The Review of Economic Studies*, 74(3): 705–731.

Fleming, Wendell H, and Halil Mete Soner. 2006. *Controlled Markov processes and viscosity solutions.* Vol. 25, Springer Science & Business Media.

Galeotti, Andrea, and Sanjeev Goyal. 2009. "Influencing the influencers: a theory of strategic diffusion." *RAND Journal of Economics*, 40(3): 509–532.

Gallego, Guillermo, and Garrett van Ryzin. 1994. "Optimal dynamic pricing of inventories with stochastic demand over finite horizons." *Management Science*, 40(8): 999–1020.

Garrett, Daniel F. 2013. "Incoming demand with private uncertainty." *Working paper*.

Gennaioli, Nicola, and Andrei Shleifer. 2018. *A Crisis of Beliefs: Investor Psychology and Financial Fragility.* Princeton University Press.

Gershkov, Alex, and Benny Moldovanu. 2009. "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach." *American Economic Journal: Microeconomics*, 1(2): 168–198.

Hartline, Jason, Vahab Mirrokni, and Mukund Sundararajan. 2008. "Optimal marketing strategies over social networks." *Proc. 17th Internat. Conf. World Wide Web (WWW '08) (ACM, New York)*, 189–198.

Ifrach, Bar, Costis Maglaras, Marco Scarsini, and Anna Zseleva. 2019. "Bayesian Social Learning from Consumer Reviews." *Operations Research*, 67(5): 1209–1221.

Jadbabaie, Ali, and Ali Kakhbod. 2019. "Optimal Contracting in Networks." *Journal of Economic Theory*, 183: 1094–1153.

Kremer, Ilan, Yishay Mansour, and Motty Perry. 2014. "Implementing the "wisdom of the crowd"." *Journal of Political Economy*, 122(5): 988–1012.

Øksendal, Bernt. 2003. *Stochastic differential equations.* Springer.

Papanastasiou, Yiangos, Kostas Bimpikis, and Nicos Savva. 2018. "Crowdsourcing Exploration." *Management Science*, 64(4): 1477–1973.

Papanastasiou, Yiangos, Nitin Bakhshi, and Nicos Savva. 2013. "Social Learning from Early Buyer Reviews: Implications for New Product Launch." *Working paper*.

Park, Kichool. 2001. "Essays in Strategic Experimentation."

Phillips, Robert. 2005. *Pricing and Revenue Optimization.* Stanford Business Books.

Radner, Roy, Ami Radunskaya, and Arun Sundararajan. 2014. "Dynamic pricing of network goods with boundedly rational consumers." *Proceedings of the National Academy of Sciences*, 111(1): 99–104.

Scheinkman, Jose A, and Wei Xiong. 2003. "Overconfidence and speculative bubbles." *Journal of political Economy*, 111(6): 1183–1220.

Sobel, Joel. 1991. "Durable Goods Monopoly with Entry of New Consumers." *Econometrica*, 59(5): 1455–1485.

Stokey, Nancy L. 1979. "Intertemporal Price Discrimination." *Quarterly Journal of Economics*, 93(3): 355–371.

Stokey, Nancy L. 1981. "Rational expectations and durable goods pricing." *The Bell Journal of Economics*, 112–128.

Sundararajan, Arun. 2008. "Local network effects and complex network structure." *The B.E. Journal of Theoretical Economics*, 7(1): 1935–1704.

Su, Xuanming. 2007. "Intertemporal Pricing with Strategic Customer Behavior." *Management Science*, 53(5): 726–741.

Talluri, Kalyan T., and Garrett van Ryzin. 2004. *The Theory and Practice of Revenue Management (International Series in Operations Research & Management Science).* Springer.

Veronesi, Pietro. 2019. "Heterogeneous households under uncertainty." National Bureau of Economic Research.

Yu, Man, Laurens Debo, and Roman Kapuscinski. 2013. "Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods." *Working paper*.

Zaitsev, Valentin F, and Andrei D Polyanin. 2002. *Handbook of exact solutions for ordinary differential equations.* Chapman and Hall/CRC.