

Competitive Information Disclosure in Search Markets

Simon Board

University of California, Los Angeles

Jay Lu

University of California, Los Angeles

Buyers often search across sellers to learn which product best fits their needs. We study how sellers manage these search incentives through their disclosure strategies (e.g., product trials, reviews, and recommendations) and ask how competition affects information provision. If sellers can observe the beliefs of buyers or can coordinate their strategies, then there is an equilibrium in which sellers provide the “monopoly level” of information. In contrast, if buyers’ beliefs are private, then there is an equilibrium in which sellers provide full information as search costs vanish. Anonymity and coordination thus play important roles in understanding how advice markets work.

I. Introduction

Starting with Stigler (1961), there has been a large literature examining the incentives of buyers to search for better prices. This paper studies the

We thank the editor and three anonymous referees for many excellent comments. We are grateful to John Asker, Alessandro Bonatti, Emir Kamenica, Moritz Meyer-ter-Vehn, Marek Pycia, and Asher Wolinsky for comments, as well as seminar audiences at Austin, Arizona State, Berkeley, Bonn, Chicago, Columbia Industrial Organization Conference, Duke, Econometric Society World Congress, Gerzensee, Harvard, Massachusetts Institute of Technology, Miami, Microsoft, Munich, New South Wales, North Carolina, Riverside, Queensland, San Diego, Society for Economic Dynamics, SocCal Network Economics Conference,

Electronically published September 7, 2018
[*Journal of Political Economy*, 2018, vol. 126, no. 5]
© 2018 by The University of Chicago. All rights reserved. 0022-3808/2018/12605-0003\$10.00

incentives of buyers to search for better information and asks how sellers manage these incentives through their disclosure strategies. We consider a setting in which all sellers offer the same products and choose what information to disclose. Buyers then engage in costly random search across sellers, trying to learn which product they prefer.

There are many environments that share these features. When one is shopping online for a TV, many websites provide product information and customer reviews and also provide links to Amazon, earning a commission for each sale.¹ The websites therefore wish to steer customers toward expensive TVs but must also provide enough information to prevent customers from going elsewhere. Similarly, when a customer is looking to invest his savings, a financial adviser gathers information about his financial needs and provides advice but may stop soliciting information if the customer expresses a preference for high-fee mutual funds.² And when a patient is looking into an elective medical procedure, a doctor can order tests that both inform the patient and potentially guide him toward more lucrative procedures. In all these settings, a buyer can receive advice from many potential sellers, and one would hope that competition enables buyers to become fully informed. Yet this is often not the case. For example, Clemens and Gottlieb (2014) find that small changes in Medicare fees lead to large changes in patients' decisions, implying that patients' information depends on the incentives of their doctors.

In this paper, we study the role of competition in information provision and show that its effectiveness depends on the information structure of the market. We first show that competition is ineffectual if sellers can perfectly observe buyers' beliefs or perfectly correlate their disclosures via a coordination device. In either case, there is an equilibrium in which all sellers choose the monopoly disclosure strategy, manipulating buyers to purchase the most profitable rather than the most suitable product. However, if sellers cannot fine-tune their disclosure strategies, then there is an equilibrium in which they fully reveal all the information as search costs vanish. Intuitively, a buyer can visit many sellers and accumulate more and more information, forcing any given seller to match the market and provide full information.

In the model, there are a large number of sellers ("she") who sell the same set of products. A large number of prospective buyers ("he") approach sellers in order to learn which product best suits their needs. Each

Stanford Institute for Theoretical Economics, University of Technology Sydney, and Washington St. Louis.

¹ For example, *techradar.com*, *4k.com*, *cnet.com*, and *rtings.com*.

² Mullainathan, Noeth, and Schoar (2012) document that when a customer holds a high-fee "returns-chasing portfolio, advisers were significantly more supportive than for either the company stock or the index portfolio" (12).

seller chooses how much information to disclose by releasing a signal that may increase or decrease a buyer's assessment of the product, as in Kamenica and Gentzkow (2011).³ After receiving the signal and updating his belief, the buyer chooses whether to buy a product, exit, or pay a search cost and randomly sample another seller. Different products yield different profits for a seller, so she tries to steer the buyer toward products that are more profitable, taking into account that the buyer can always move on to a competitor.

We study how the disclosure by sellers depends on their information, as captured by noise parameters (α, β) . When a seller is approached by a buyer, she observes a signal equal to the buyer's belief with noise scaled by α . She also has access to a coordination device with probability $1-\beta$ or a noisy permutation with probability β . The α -parameter captures the degree to which a seller can condition her disclosure on the buyer's belief. For example, this represents the quality of past browsing data on a customer's cookies or the information a financial adviser can deduce from an investor's portfolio. The β -parameter captures the degree to which a seller can condition her disclosure on a coordination device. For example, websites recommending TVs may use common marketing material provided by Amazon or third-party rating agencies (e.g., Energy Star) or may collect their own independent customer reviews.

In Section III, we show that if sellers can either perfectly observe beliefs ($\alpha = 0$) or perfectly coordinate their disclosures ($\beta = 0$), then the monopoly disclosure strategy is always an equilibrium. In either case, if all sellers use the monopoly strategy, then no buyer receives additional information from searching a second time. The buyer thus has no incentive to continue searching, and since the strategy maximizes profits, no seller will defect. This monopoly strategy may be very undesirable for buyers: if there is a single product, then it provides the buyer with no valuable information.

When sellers can observe buyers' beliefs ($\alpha = 0$), this monopoly equilibrium is also unique in a wide range of environments. For example, if there is a single product, the search cost implies that any one seller can provide a little less information than the market, iteratively pushing the equilibrium toward the monopoly outcome. The crucial condition is that, given any nonmonopoly strategy, there are a series of local deviations that lead to monopoly. More precisely, we show that the monopoly strategy is the unique equilibrium if every *nonmaximal belief* (i.e., where profit differs from the maximum over all possible states) is *improvable* (i.e., a monopolist wishes to release some information). In addition to the single-

³ More formally, a seller can choose any distribution of posteriors that average to the prior. Kamenica and Gentzkow show that a monopolist's optimal solution corresponds to the concavification of her profit function.

product example, this also covers cases in which there are two horizontally differentiated products or many “niche” products.

In Section IV, we show that if sellers observe a noisy signal of buyers’ beliefs ($\alpha > 0$), then there exists a sequence of equilibria that converges to full information as search costs vanish (i.e., full information is a *limit equilibrium*). In equilibrium, buyers purchase from the first seller they visit but have the option to continue searching and mimic a new buyer. If all other sellers provide some information and search costs are small, then a buyer can threaten to visit many of them and accumulate a large amount of information. This forces the current seller to provide (almost) full information. Anonymity is therefore a powerful force in enabling buyers to become well informed.

We then study the uniqueness of equilibria when $\alpha > 0$, focusing on equilibria in which buyers purchase from the first seller.⁴ If all sellers have access to a perfect coordination device ($\beta = 0$), then both monopoly and full information are limit equilibria. However, if coordination is noisy ($\beta > 0$), full information is the unique limit equilibrium if every *conditional belief* (i.e., the marginal of the prior on a subset of states) is *fully improvable* (i.e., a monopolist prefers full information to no information). Table 1 summarizes our existence results. In the case of our example 1 (which features a single product and two states), these are the only equilibria if the buyer is initially skeptical, preferring not to buy in the absence of information.

Comparing these cases, the key is whether sellers can choose strategies that discriminate between new and old buyers. When sellers can observe buyers’ beliefs, it is optimal for them to provide less information to old buyers who are already informed, undercutting the incentive to search and undermining competition. Similarly, when sellers can coordinate, they can use a disclosure strategy that is useful for a new buyer but uninformative for an old one. In contrast, when sellers cannot discriminate, the option of going to a competitor forces sellers to provide all the pertinent information as search costs vanish. This finding is important for applications. Our results provide a theoretical foundation for the notion that tracking consumers can make them “exploitable” and suggest that common marketing material or industry training courses can help financial advisers implicitly collude to the detriment of investors.⁵

Finally, in Section V, we explore several extensions. First, we argue that our results are robust to heterogeneous buyers. Second, we consider a case in which sellers can observe buyer shopping history but not beliefs

⁴ Since all sellers sell identical products, it is natural to look at such equilibria. We discuss this assumption in Sec. II.

⁵ For example, see “Google to Pay \$17 Million to Settle Privacy Case,” *New York Times* (November 18, 2013) or “Facebook Tries to Explain Its Privacy Settings but Advertising Still Rules,” *New York Times* (November 13, 2014).

TABLE 1
SUMMARY OF LIMIT EQUILIBRIA

COORDINATION NOISE	BELIEF OBSERVABILITY NOISE	
	$\alpha = 0$	$\alpha > 0$
$\beta = 0$	Monopoly	Monopoly and full information
$\beta > 0$	Monopoly	Full information

NOTE.—This table describes our existence results for limit equilibria, given the two noise parameters (α , β). In the case of example 1, these are the only equilibria if the buyer prefers not to buy in the absence of information.

and show that since sellers can still discriminate between new and old buyers, equilibria are monopolistic. Third, we argue that buyers create a positive externality when they become anonymous (e.g., by deleting cookies), thus providing a rationale for regulating tracking programs.

Literature.—The paper contributes to the literature on search by allowing sellers to choose information disclosure strategies. In the benchmark model in which sellers choose prices, Diamond (1971) shows that all sellers charge the monopoly price. Intuitively, the search cost allows any one seller to raise her price slightly above the prices set by others without losing customers.⁶ We show that when sellers choose disclosure strategies, the analysis depends on the information structure: when sellers can perfectly observe buyers' beliefs, the logic is analogous to that in Diamond's model; however, when their observations are imperfect, competition induces sellers to fully reveal the information. In contrast to the Diamond model, when a seller provides information to a buyer, this changes his beliefs and changes how he acts at subsequent sellers.

The paper also contributes to a growing literature on information disclosure with competition based on the Kamenica-Gentzkow framework (see also Aumann and Maschler 1995; Rayo and Segal 2010). Gentzkow and Kamenica (2017b) consider a general model in which multiple firms simultaneously release signals and provide conditions under which competing firms release more information than colluding firms. Gentzkow and Kamenica (2017a) study senders who release "coordinated" signals, characterize the equilibrium, and show that competition increases information. Hoffmann, Inderst, and Ottaviani (2014) suppose that heterogeneous sellers simultaneously release information to win over a customer and show that competition increases information disclosure. Li and Norman (2018a, 2018b) show that this result may fail if sellers move sequen-

⁶ This result continues to hold if sellers make multiple offers to each buyer (Board and Pycia 2014). However, it can break down if sellers sell heterogeneous products (Wolinsky 1986), buyers have multi-unit demand (Zhou 2014; Rhodes 2015), or buyers precommit to a number of searches (Burdett and Judd 1983).

tially, use independent signals, or use mixed strategies. In all these models, a buyer receives the market information and then chooses his action; in our model, the buyer must pay a search cost in order to acquire more information.⁷

We study sellers who choose disclosure strategies and obtain payoffs that depend on buyers' actions. For example, a website receives a commission if a buyer purchases a TV from Amazon, a financial adviser receives a fee if she sells a mutual fund, and a doctor is reimbursed for procedures by Medicare. In contrast, Anderson and Renault (2006) study a monopolist selling a single good who can choose both information and prices (see also Lewis and Sappington 1994; Johnson and Myatt 2006). They show that the seller should reveal whether the buyer's value exceeds the cost and charge a price equal to the buyer's expected value; this implements the efficient allocation and fully extracts from the buyer. More generally (e.g., if there are two goods), the monopoly solution will be inefficient, and the forces we study will affect the effectiveness of search in enhancing efficiency.

Finally, there is a large literature that studies costly acquisition of information. For example, Wald (1947) and Moscarini and Smith (2001) consider a decision maker who can purchase independent signals at a constant marginal cost. In contrast, the information provided to our buyer is chosen endogenously by strategic sellers.

II. Model Setup

Overview.—There is a unit mass of sellers offering the same set of products. Time is discrete, with a unit mass of buyers entering each period; each buyer searches for one product. When a buyer approaches a seller, the seller observes a signal about the buyer's belief and decides how much information to disclose. The buyer then chooses whether to buy, exit, or continue searching. Our results will depend on two parameters: α , which determines the accuracy of the sellers' signals, and β , which determines the sellers' ability to coordinate their information disclosures.

⁷ There are other related papers. Forand (2013) studies a model of information revelation with directed search. Au (2015) and Ely (2017) study dynamic information disclosure by a monopolistic seller. In addition, there are a variety of papers concerning the revelation of information when the senders are informed. One literature considers verifiable information ("persuasion games"). Here, Milgrom and Roberts (1986) find conditions under which competition leads to full revelation, and Bhattacharya and Mukherjee (2013) consider a model with multiple senders. A second literature supposes that information is unverifiable ("cheap talk"). Here, Crawford and Sobel (1982) characterize the structure of communication, and Battaglini (2002) shows how multiple senders can increase the information provided when senders have opposed preferences. Also related, Inderst and Ottaviani (2012) consider a model with a single seller and study the choice of commission rates chosen by upstream sellers.

Information.—A buyer is initially uncertain about a finite set of *payoff states* S that are relevant for his decision. There is also an independent coordination device uniformly distributed on $[0, 1]$ that sellers can use to coordinate their disclosures.⁸ Let $\Omega := S \times [0, 1]$ be the overall state space and $\Delta\Omega$ be the set of all possible beliefs for the buyer. Buyers start with an initial belief $\mathbf{p}_0 \in \Delta\Omega$, where the *payoff belief* $p_0 \in \Delta S$ is the marginal of \mathbf{p}_0 on S and puts positive probability on every $s \in S$. Denote a generic prior by \mathbf{p} and a generic payoff prior by p . When there are only new buyers in the market, we will use \mathbf{p}_0 and \mathbf{p} synonymously.

Buyers are randomly matched with sellers. When a buyer approaches a seller, the seller will have a prior η_0 about the buyer’s belief.⁹ The seller then observes a private signal about the buyer’s belief, $\mathbf{y} = (1 - \alpha)\mathbf{p} + \alpha\xi$, where the random variable $\xi \in \mathbb{R}^Q$ has full support. The parameter $\alpha \in [0, 1]$ captures the noise of the signal, so $\alpha = 0$ corresponds to the case in which sellers can perfectly observe buyers’ beliefs. Let Y be the space of private signals and η_y be her inference about the buyer’s belief after receiving the signal $\mathbf{y} \in Y$.

Sellers’ strategy.—A seller chooses how much information to disclose to a buyer based on the private signal \mathbf{y} . With probability $1 - \beta$, the seller chooses a *general signal structure* $\sigma_y^g : \Omega \rightarrow \Delta\Theta$ over a signal space Θ that allows her to coordinate with other sellers. With probability β , the seller does not have access to the coordination device and can send only an *independent signal structure* $\sigma_y^i : S \rightarrow \Delta\Theta$. We denote the collection of signal structures by $\sigma^g = (\sigma_y^g)_{y \in Y}$ and $\sigma^i = (\sigma_y^i)_{y \in Y}$ and call them *disclosure policies*. A seller’s *strategy* is then $\sigma = (\sigma^g, \sigma^i)$.

After receiving a signal $\theta \in \Theta$ from the seller, a buyer updates his belief to form posterior $\mathbf{q} \in \Delta\Omega$. Given a buyer with belief \mathbf{p} , a signal structure σ_y^g induces a posterior distribution $K_{\mathbf{p}}^{\sigma_y^g}$ for the buyer such that the average posterior coincides with the prior, $\int_{\Delta\Omega} \mathbf{q} K_{\mathbf{p}}^{\sigma_y^g}(d\mathbf{q}) = \mathbf{p}$. Given a seller strategy σ and a signal realization \mathbf{y} , a buyer with belief \mathbf{p} faces a distribution of posteriors $K_{\mathbf{p}}^{\sigma_y} := (1 - \beta)K_{\mathbf{p}}^{\sigma_y^i} + \beta K_{\mathbf{p}}^{\sigma_y^g}$. When sellers observe buyers’ beliefs ($\alpha = 0$), we can drop the subscript $\mathbf{y} = \mathbf{p}$ on the strategy and use the condensed notation $K_{\mathbf{p}}^{\sigma}$.

Buyers’ decisions.—After a buyer updates his belief to form a posterior $\mathbf{q} \in \Delta\Omega$, he chooses whether to (i) buy a product from the seller, (ii) exit the market, or (iii) pay a search cost $c > 0$ and pick a new seller at random. Each seller offers the same set of products $\{0, 1, \dots, I\}$. Product $i \geq 1$ de-

⁸ This formulation draws on Green and Stokey (1978) and Gentzkow and Kamenica (2017a).

⁹ In equilibrium, the seller’s prior η_0 will be consistent with the strategies of other sellers in the market.

livers utility $u_i(s)$ to the buyer in state $s \in S$ and profit $\pi_i > 0$ to the seller.¹⁰ We assume that different products yield different profits, that is, $\pi_i \neq \pi_j$ for all $i \neq j$. Product 0 is the exit option normalized so that u_0 is the zero vector and $\pi_0 = 0$.

If a buyer with belief $\mathbf{q} \in \Delta\Omega$ decides to stop shopping, then let $i^*(q) = \operatorname{argmax}_i q \cdot u_i$ denote his optimal product.¹¹ We can define the indirect utility vector of the buyer by $u(q) := u_{i^*(q)}$ and the profit of the seller by $\pi(q) := \pi_{i^*(q)}$. If a buyer decides to continue shopping, then he pays the search cost and arrives at a new seller. The game then proceeds as above.

Value functions.—Sellers choose their strategies to maximize profit. All sellers have identical optimization problems with the buyer’s belief \mathbf{p} being the only relevant state variable. We look for an equilibrium in symmetric Markov strategies.¹² If all other sellers choose a strategy σ , then the continuation value of a buyer with belief \mathbf{q} is

$$V_c(\sigma, \mathbf{q}) = -c + \mathbb{E}_{\mathbf{q}} \left[\int_{\Delta\Omega} \max\{r \cdot u(r), V_c(\sigma, \mathbf{r})\} K_{\mathbf{p}}^{\sigma}(\mathbf{d}\mathbf{r}) \right],$$

where $r \cdot u(r)$ is the buyer’s expected utility when stopping and the expectation $\mathbb{E}_{\mathbf{q}}$ is taken over all realizations of \mathbf{y} . A buyer purchases or exits when his belief falls in a *stopping set* given by

$$\mathbf{Q}_c(\sigma) = \{\mathbf{q} \in \Delta\Omega \mid q \cdot u(q) \geq V_c(\sigma, \mathbf{q})\}.$$

We will also let $Q_c(\sigma) \subset \Delta S$ denote the corresponding stopping payoff beliefs.

Best responses.—Suppose other sellers in the market use strategy σ . The current seller begins with belief η_0 about the buyer’s prior \mathbf{p} and, after receiving the realization $\mathbf{y} \in Y$, updates her belief to form a posterior $\eta_{\mathbf{y}}$. She then chooses a strategy $\tilde{\sigma}$ that solves

$$\max_{\tilde{\sigma}} \int_{\Delta\Omega} \int_{\mathbf{Q}_c(\sigma)} \pi(q) K_{\mathbf{p}}^{\tilde{\sigma}}(\mathbf{d}\mathbf{q}) \eta_{\mathbf{y}}(\mathbf{d}\mathbf{p}).$$

Geometrically, given \mathbf{p} , the seller’s optimal profits are given by the concavification of π over beliefs in the stopping set $\mathbf{Q}_c(\sigma)$. If $\tilde{\sigma}$ solves this problem, then we say it is *optimal* given σ .

We will usually be interested in the case in which sellers have degenerate beliefs about the distribution of buyers they face, either because they

¹⁰ As in Rayo and Segal (2010), a seller’s profit is pinned down by the product chosen by the buyer.

¹¹ In equilibrium, ties are resolved in the seller’s favor. Similarly, a buyer purchases when indifferent between stopping and continuing.

¹² That is, the strategy depends on neither calendar time nor the identity of the seller. This assumption makes the analysis simpler but is not required for our motivating example (see n. 17).

can observe buyers' beliefs or because they believe all buyers are new. Letting δ be the Dirac measure, then we have the following lemma:

LEMMA 1. If a seller knows a buyer's belief, in that $\eta_y = \delta_p$ for all y , then it is optimal for her to sell with probability one.

The proof is in Section A.1 of the appendix.

Intuitively, suppose a seller sends a buyer to a belief outside his stopping set with positive probability. Since she has access to arbitrary signals, she can increase her profits by disclosing full information to this buyer whenever he would have waited, while holding constant the information she sends in all other situations.

If $\alpha = 0$, a seller observes buyers' beliefs, meaning that $\eta_y = \delta_p$. Lemma 1 then implies that she sells with probability one. If $\alpha > 0$, a seller observes a noisy signal of the true belief. If the current seller believes that she faces only new buyers, then her prior is degenerate, $\eta_0 = \delta_{p_0}$. Since the noise ξ has full support, her prior is consistent with any realization of the signal y , and Bayes's rule implies that her posterior coincides with the prior, $\eta_y = \delta_{p_0}$, for any realization of y . Lemma 1 thus implies that if the seller is correct about facing a new buyer, then she sells with probability one.

Equilibrium.—In a symmetric Markov perfect equilibrium σ of this game, (i) a seller's strategy is optimal given that other sellers also use this strategy, and (ii) sellers' priors η_0 are consistent with buyers' behavior. Given the discussion above, we consider a particular type of Markov perfect equilibrium in which buyers purchase from the first seller; when clear, we refer to these *immediate purchase equilibria* simply as *equilibria*. When beliefs are observable ($\alpha = 0$), this assumption is without loss. When beliefs are imperfectly observable ($\alpha > 0$), then on path, it is optimal for a seller to sell immediately if all other sellers sell immediately.

When $\alpha > 0$, it is natural to ask whether equilibria with delay also exist. In online appendix S.2, we show that if there is no coordination, then selling immediately is a weak best response to any Markov equilibrium strategy; moreover, under some conditions (e.g., normal noise, ξ), it is a strict best response, meaning that equilibria with delay cannot exist. Unfortunately, lack of tractability makes it hard to say more when there is partial coordination. Formally, this restriction does not affect theorems 1–3 but does mean that the uniqueness result in theorem 4 applies only to immediate purchase equilibria.

The discussion highlights the critical difference between perfectly observable beliefs ($\alpha = 0$) and imperfectly observable beliefs ($\alpha > 0$). In the former case, sellers can see buyers' beliefs and are thus correct both on and off path, $\eta_y = \delta_p$. In the latter case, sellers think they are facing new buyers, $\eta_y = \delta_{p_0}$, which is correct on path but not off path if a buyer chooses to search. As we will see, this distinction generates the disparity between the "monopoly" results in Section III and the "full-information" results in Section IV.

Remarks.—As in Kamenica and Gentzkow (2011), we assume that the seller has no private information about the payoff state. For example, when a website provides information about TVs, it does not know the buyer's preferences. When a financial adviser tries to help an investor, she learns about his needs and has a fiduciary obligation to give the correct advice conditional on this information. And when a doctor orders a test for a patient, she does not know the outcome before it arrives.¹³ In addition, the signal structure is assumed to be observed by both parties (e.g., the number of reviews on a website, the information acquired by the financial adviser, or the tests ordered by the doctor).¹⁴

Sometimes we will consider equilibria without seller coordination. When sellers choose policies $\sigma^s = \sigma^i$ that are completely independent, we call the strategy σ *simple* and equate the strategy with the policy. If a strategy delivers monopoly profits to the seller, then we call it a *monopoly strategy*; the set of monopoly strategies coincides with the optimal disclosure solution in Kamenica and Gentzkow (2011). Since there is a single seller, we can ignore the coordination device, meaning that a monopoly strategy is automatically a simple strategy on payoff beliefs. An equilibrium in which all sellers use a monopoly strategy is called a *monopoly equilibrium*.

III. Monopoly Equilibria

In this section we show that the monopoly strategy is an equilibrium if sellers can perfectly observe buyers' beliefs ($\alpha = 0$) or can perfectly coordinate ($\beta = 0$). For the former case, we then provide conditions under which monopoly is the unique equilibrium. We first illustrate the main forces through a simple single-product example.

A. Motivating Example

Example 1 (Single product).—A seller has a single product to sell, there are two payoff states $\{L, H\}$, and the buyer wishes to buy the product if state H is more likely. For example, a consumer considers upgrading to a new TV but does not know whether he will value the new technology. A sale gen-

¹³ Gentzkow and Kamenica (2017c) provide a different interpretation: they show that the outcome of the test can be privately observed by the seller if sellers' messages are verifiable.

¹⁴ If the buyer observes signal structures before visiting, then he can direct his search toward the most informative, leading to Bertrand competition and full-information revelation. While some companies develop reputations for revealing information (e.g., Best and Quigley 2016), it is generally hard to describe an entire signal structure to customers before they have visited (e.g., when compared to describing a price).

erates profits $\pi = 1$ for the seller. Payoffs to the buyer are $u(L) = -1$ and $u(H) = 1$, so the buyer prefers to buy if $p = \Pr(H) > 1/2$.

If there were a single seller, then the monopoly strategy from Kamenica and Gentzkow (2011) provides just enough information to persuade the buyer to buy. In other words, the monopoly strategy σ^* is given by¹⁵

$$K_p^{\sigma^*} = \begin{cases} \delta_p & \text{if } p \in [\frac{1}{2}, 1] \\ (1 - 2p)\delta_{\{0\}} + 2p\delta_{\{1/2\}} & \text{if } p \in [0, \frac{1}{2}). \end{cases}$$

This is illustrated in figure 1A. If the buyer starts with a prior $p \geq 1/2$, then the seller provides no information and the buyer buys; if $p < 1/2$, then the seller sends the buyer to posteriors 0 and $1/2$ and the buyer buys with probability $2p$. For shorthand, we will sometimes denote this binary signal by the condensed notation $p \rightarrow \{0, 1/2\}$.

If sellers perfectly observe buyers' beliefs ($\alpha = 0$), then the monopoly strategy is an equilibrium. If all sellers choose this strategy, then a buyer learns nothing from a second seller after leaving the first seller and thus, given the search cost, buys immediately. Since the seller makes monopoly profits, she has no incentive to deviate.

More surprising, the monopoly strategy is the unique equilibrium of the game. To gain some intuition, fix $b > 1/2$ and suppose that all sellers provide a simple binary signal $p \rightarrow \{0, b\}$ as illustrated in figure 1B. Call this strategy σ and observe that it induces a stopping set $Q_c(\sigma) = [0, [b/(2b - 1)]c] \cup [b(1 - c), 1]$ as shown in the lower panel.¹⁶ Now, consider a seller who deviates and uses the less informative simple strategy $p \rightarrow \{0, b(1 - c)\}$. Since there is a strictly positive search cost, a buyer who receives this signal will not subsequently search and this deviation strictly raises the seller's profits. As we prove below, this logic implies that whenever sellers provide more information than the monopoly strategy, there is a profitable deviation, analogous to the classic result of Diamond (1971).¹⁷

Similarly, if sellers can perfectly coordinate ($\beta = 0$), then the coordinated monopoly strategy is an equilibrium. Under such coordination, all sellers provide the same signal, so a buyer will purchase after visiting the first seller. Since this strategy yields monopoly profits, no seller has any incentive to deviate. In contrast to the case of observed beliefs ($\alpha = 0$), this equilibrium

¹⁵ Technically, there is a set of monopoly strategies that deliver monopoly profits to the seller. Here, σ^* is the "one-shot" monopoly strategy defined in Sec. III.B.

¹⁶ Note that if $b(1 - c) \leq 1/2$, then $Q_c(\sigma) = [0, 1]$ and the seller can trivially deviate to the monopoly strategy.

¹⁷ In fact, in this example we can show an even stronger result: the monopoly strategy is the only rationalizable strategy. If we start with a stopping set $Q_0 = \{0\} \cup \{1\}$, then no matter what other sellers do, a buyer purchases if his belief falls within $Q_1 = [0, c] \cup [1 - c, 1]$. Iterating $n \geq -\log(2)/\log(1 - c)$ rounds, $[1/2, 1] \subset Q_n$ and the monopoly strategy is the only undominated strategy remaining.

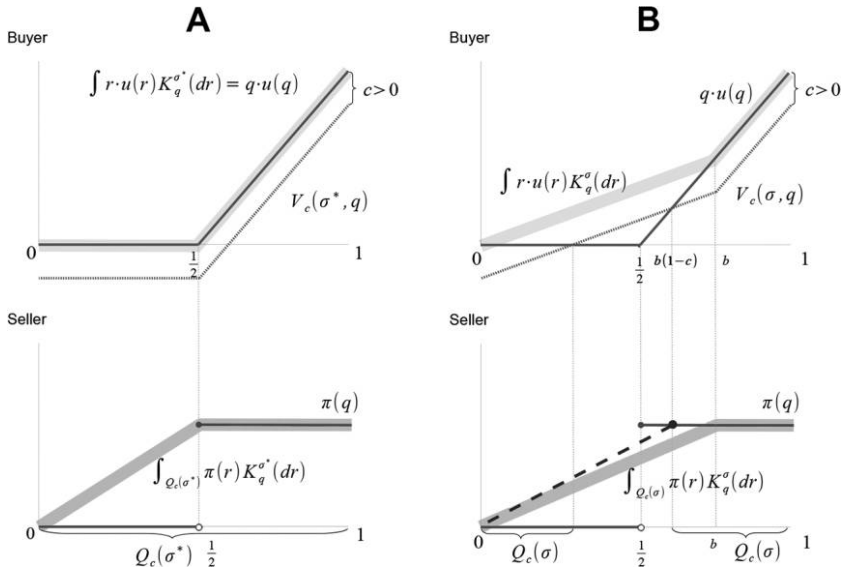


FIG. 1.—Single product (monopoly equilibrium). This figure illustrates example 1 when sellers observe buyers' beliefs ($\alpha = 0$). Panel A shows that the monopoly strategy is an equilibrium. The top panel shows the buyer's payoffs: the bold line is the payoff from stopping $q \cdot u(q)$, the shaded line is the expected payoff $\int_{\Delta S} r \cdot u(r) K_q^{\sigma^*}(dr)$ under the monopoly strategy σ^* , and the dotted line is the buyer's value function $V_c(\sigma^*, q)$. The bottom panel shows the seller's payoffs: the bold line is the profit function $\pi(q)$ and the shaded line is the expected profit $\int_{Q_c(\sigma^*)} \pi(r) K_q^{\sigma^*}(dr)$. Panel B shows that a more informative strategy cannot be an equilibrium since a seller can raise her profits by providing a little less information, as illustrated by the dashed line. Color version available as an online enhancement.

is not unique: if all other sellers fail to use the coordination device, then the current seller has no incentive to use it either.

B. Monopoly Is an Equilibrium

We now return to the general model, with multiple states and multiple products. We wish to show that if sellers can perfectly observe buyers' beliefs ($\alpha = 0$) or perfectly coordinate ($\beta = 0$), then monopoly is always an equilibrium.

First, we need a preliminary result. Consider a monopolist who uses a strategy σ , and let $A^\sigma \subset \Delta S$ be the set of absorbing beliefs in which no information is provided, that is, $K_p^\sigma = \delta_p$ for all $p \in A^\sigma$. We say that such a strategy is one-shot if it takes the buyer only to absorbing beliefs, that is, $K_p^\sigma(A^\sigma) = 1$ for all $p \in \Delta S$. This means the buyer gets no more information if he were to return to the monopolist under σ .

LEMMA 2. There exists a one-shot monopoly strategy.

Proof. Kamenica and Gentzkow (2011) show that a monopolist’s profit, which we denote by Π^* , equals the concavification of the profit function π . Let R be the set of beliefs in which the two functions coincide, that is, $\Pi^*(q) = \pi(q)$. Consider the following strategy σ^* : the seller uses any optimal strategy if $p \notin R$ and releases no information if $p \in R$. First observe that σ^* takes the buyer to posteriors in R . The reason is that the optimal strategy at any prior p must send the buyer to beliefs where no information is optimal; otherwise, the seller should incorporate this information in her initial signal. Second, using the definition of σ^* , beliefs in R are absorbing so σ^* is one-shot. Since σ^* yields monopoly profits, it is thus a one-shot monopoly strategy. QED

We now present our first main result: monopoly is always an equilibrium. Note that this also implies the existence of an equilibrium.

THEOREM 1. Suppose sellers can perfectly observe beliefs ($\alpha = 0$) or perfectly coordinate ($\beta = 0$). There exists a monopoly equilibrium.

Proof. First, suppose that $\alpha = 0$ and all sellers use a simple one-shot monopoly strategy σ^* , which exists by lemma 2. By definition, this strategy sends buyers to absorbing beliefs, that is, $K_p^{\sigma^*}(A^{\sigma^*}) = 1$. Since buyers receive no information at absorbing beliefs, they stop, that is, $A^{\sigma^*} \subset Q_c(\sigma^*)$. Hence the seller achieves her monopoly profit and has no incentive to deviate, meaning that σ^* is an equilibrium.

Now suppose $\beta = 0$ and let σ^{**} be the perfectly coordinated strategy corresponding to σ^* . If a buyer were to return to the market, then he would receive no information, that is, $K_q^{\sigma^{**}} = \delta_q$ for all q in the support of $K_p^{\sigma^{**}}$ for any $p \in \Delta\Omega$. Hence the buyer prefers to stop rather than continue, that is, $q \in Q_c^{\sigma^{**}}$ for all q in the support of $K_p^{\sigma^{**}}$ for any $p \in \Delta\Omega$. Since the seller cannot do better than achieving monopoly profits, she will not deviate, meaning that σ^{**} is an equilibrium. QED

Intuitively, if all sellers use a one-shot monopoly strategy, a buyer receives only one round of information and therefore purchases from the first seller. Since all sellers make monopoly profits, they have no incentive to deviate. This logic relies on sellers’ ability to give information to new buyers but not to old buyers, either because sellers can observe buyers’ beliefs or because they can perfectly coordinate their disclosure strategies.

C. Uniqueness of the Monopoly Equilibrium

When sellers can perfectly coordinate ($\beta = 0$), there is a monopoly equilibrium even if buyers’ beliefs are imperfectly observed ($\alpha > 0$). However, if all other sellers use independent signals, then no one seller can coordinate by herself. As we show in Section IV, this means that full-information equilibria also exist as search costs vanish.

In contrast, when buyers’ beliefs are perfectly observed ($\alpha = 0$), the monopoly equilibrium is “often” unique. Before establishing this result, we

first provide an example of a nonmonopoly equilibrium. This motivates the sufficient condition for the uniqueness result. For the remainder of this section, we thus assume that beliefs are perfectly observed ($\alpha = 0$).¹⁸

Example 2 (Vertical differentiation).—Suppose there are two vertically differentiated products and two payoff states L and H representing the buyer’s taste for quality (e.g., Kamenica 2008, sec. IIA), with $p = \Pr(H)$. The first product is a cheap, low-quality TV yielding utility $u_1 = (-\frac{1}{3}, \frac{2}{3})$; the second is an expensive, high-quality TV yielding utility $u_2 = (-1, 1)$. Thus, the buyer prefers no TV if $p \in [0, \frac{1}{3})$, the cheap TV if $p \in [\frac{1}{3}, \frac{2}{3})$, and the expensive TV if $p \in [\frac{2}{3}, 1]$. Assume that the expensive TV is more profitable for the seller, that is, $\pi_1 = 1$ and $\pi_2 = \frac{3}{2}$. The monopoly strategy σ^* is given by

$$K_p^{\sigma^*} = \begin{cases} (1 - 3p)\delta_{\{0\}} + 3p\delta_{\{\frac{1}{3}\}} & \text{if } p \in [0, \frac{1}{3}) \\ (2 - 3p)\delta_{\{\frac{1}{3}\}} + (3p - 1)\delta_{\{\frac{2}{3}\}} & \text{if } p \in [\frac{1}{3}, \frac{2}{3}) \\ \delta_p & \text{if } p \in [\frac{2}{3}, 1]. \end{cases}$$

By theorem 1, this monopoly strategy is an equilibrium, as illustrated in figure 2A. One can interpret this strategy as the seller recommending the high-quality TV for high beliefs, recommending either the high- or low-quality TV for intermediate beliefs, and recommending either the low-quality TV or no TV for low beliefs.

When $c > 0$ is small enough, there is also a second simple equilibrium σ in which sellers provide more information,

$$K_p^\sigma = \begin{cases} \frac{2 - 3p}{2}\delta_{\{0\}} + \frac{3p}{2}\delta_{\{\frac{2}{3}\}} & \text{if } p \in [0, \frac{2}{3}) \\ \delta_p & \text{if } p \in [\frac{2}{3}, 1] \end{cases}$$

as illustrated in figure 2B. This gives rise to the stopping set $Q_c(\sigma) = [0, \underline{\varepsilon}_c] \cup [\frac{2}{3} - \bar{\varepsilon}_c, 1]$ for some $\underline{\varepsilon}_c > 0$ and $\bar{\varepsilon}_c > 0$. This strategy differs from the monopoly strategy by never recommending the low-quality TV. As shown in figure 2B, the presence of the search cost allows the seller to provide a little less information than her competitors. However, there is no local deviation that is profitable: if a seller sends the buyer to $p \rightarrow \{0, \frac{2}{3} - \bar{\varepsilon}_c\}$, then the seller’s profit would drop. There is also no profitable global deviation: if the seller sends the buyer to $p \rightarrow \{0, \frac{1}{3}\}$ as under the monopoly strategy, then the buyer with belief $\frac{1}{3}$ would refuse to buy, correctly anticipating that other sellers will provide significantly more information.

¹⁸ When the search cost c is sufficiently high, monopoly is trivially the unique equilibrium. All of the results in this section apply to any positive search cost, including those arbitrarily small. As one might expect, the set of equilibria is decreasing in the search cost; indeed, we show in online app. S.1 that given any simple equilibrium for $c > 0$, there is an equilibrium for all $c' < c$ with the same profits.

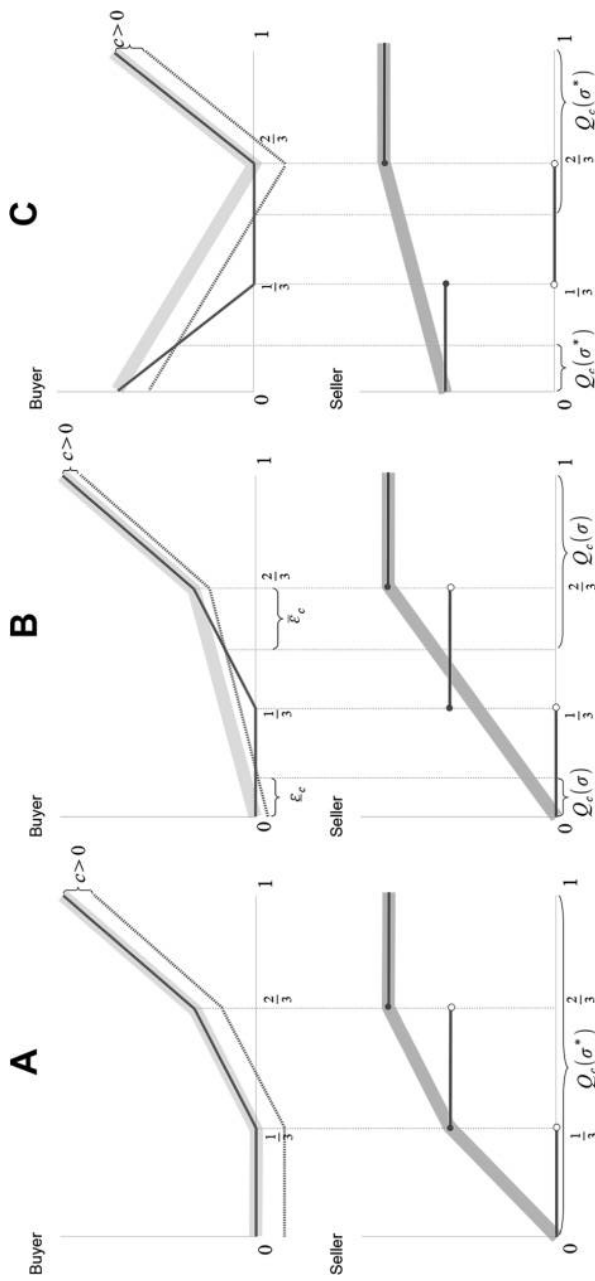


FIG. 2.—Vertical and horizontal differentiation. Panels A and B show the two equilibria that arise with vertical differentiated products, monopoly and nonmonopoly, respectively. Panel C shows the unique equilibrium that arises with horizontally differentiated products; this coincides with the monopoly strategy. For a description of the different lines, see the caption for figure 1. Color version available as an online enhancement.

We now provide a sufficient condition for monopoly to be the unique equilibrium. We say that a belief $\mathbf{p} \in \Delta\Omega$ is *maximal* if $\pi(\mathbf{p}) = \max_{s \in S_p} \pi(1_s)$, where $S_p \subset S$ is the support of the payoff belief p . We say that \mathbf{p} is *nonmaximal* if the equality fails. A belief \mathbf{p} is *improvable* if a monopolist would like to provide some nontrivial information, that is, $\Pi^*(p) > \pi(\mathbf{p})$, where $\Pi^*(p)$ is the monopoly profit. We first present a preliminary lemma.

LEMMA 3. Suppose $\alpha = 0$. If any nonmaximal belief is improvable, then in any equilibrium, a buyer stops at all maximal beliefs.

The proof is in the appendix, Section A.2.

To illustrate what it means for nonmaximal beliefs to be improvable, consider figure 3A, where there are three products $\{1, 2, 3\}$ with $\pi_1 > \pi_2 > \pi_3$. The set $H_i \subset \Delta S$ denotes the set of payoff beliefs in which the buyer favors product i . Note that $1_s \in H_i$ for all $i \in \{1, 2, 3\}$, so the most profitable product is favored in state s_1 , the second most profitable in state s_2 , and the third most profitable in state s_3 . The shaded dark regions represent the set of maximal beliefs. This example satisfies the condition in lemma 3: given any belief p that is not maximal, a monopolist would choose to release some information (i.e., p is improvable). In particular, the monopoly strategy can be decomposed in two steps. First, it releases a binary news signal about s_1 , taking the prior $p \rightarrow \{q_1, q_{23}\}$, where q_1 just persuades the buyer to buy the first product and q_{23} is on the edge between s_2 and s_3 . Second, the strategy releases a binary signal about s_2 such that $q_{23} \rightarrow \{q_2, 1_{s_3}\}$, where q_2 just persuades the buyer to purchase product 2.¹⁹

To understand lemma 3, suppose by contradiction that there is a maximal payoff belief in H_1 that is not in the stopping set $Q_c(\sigma)$ of some equilibrium strategy σ . A buyer with a belief on the boundary of $Q_c(\sigma) \cap H_1$ gets no information since he purchases the most profitable product. By a continuity argument, we can find a belief close enough to the boundary but not in the stopping set such that the buyer gets very little information. Since the search cost is strictly positive, this small amount of information is not enough to incentivize the buyer to shop, so the buyer stops at that belief, yielding a contradiction. This shows that the buyer stops at any belief in H_1 , and the proof follows by applying the same argument to all maximal beliefs.

Given this lemma, we have the following uniqueness result:

THEOREM 2. Suppose $\alpha = 0$. Any equilibrium is monopoly if every nonmaximal belief is improvable.

Proof. If other sellers use some equilibrium strategy σ , the current seller can send the buyer to any belief in the stopping set $Q_c(\sigma)$. Suppose she uses the one-shot monopoly strategy σ^* , which exists by lemma 2. This sends the buyer to beliefs in its absorbing set A^* , which are unimprovable and thus,

¹⁹ If p is maximal, then there is some freedom about the monopoly strategy; but since the buyer's decision is the same (he buys the product), such differences are payoff-irrelevant.




FIG. 3.—Uniqueness of monopoly equilibrium. Panel A shows the monopoly strategy when all nonmaximal beliefs are improvable. The shaded dark regions (H_1 , 1_s , and the segment from q_2 to 1_s) represent the set of maximal beliefs. In this case, maximal beliefs are in the stopping set of any equilibrium. Panel B shows that other equilibria can exist when the sufficient condition fails. The belief p is nonmaximal and is unimprovable when the second product is sufficiently profitable. For small ϵ , there is a nonmonopoly equilibrium in which all sellers “skip over” product 2 and send the buyer $p \rightarrow \{q_1, 1_s\}$. Color version available as an online enhancement.

by the premise, maximal. By lemma 3, these beliefs must be in the stopping set $Q_\epsilon(\sigma)$. Hence the seller can attain her monopoly profits, as required. QED

One can understand the uniqueness result in terms of “local deviations” as in example 1 or figure 3A. Whenever the monopolist sends the buyer only to beliefs that are maximal, then no matter what the other sellers do, one seller can always provide a little less information and creep closer to the monopoly outcome. This logic does not apply when a monopolist wishes to use products in the middle of the probability space as in example 2 or figure 3B.

The following corollary illustrates a wide range of cases in which the sufficient condition in theorem 2 holds.

COROLLARY 1. Suppose $\alpha = 0$. All equilibria are monopoly if

- i. there is one product,
- ii. there are two “horizontally differentiated” products in that $u_1(s) \geq 0$ implies $u_2(s) \leq 0$, and
- iii. products are “niche” in that $p \cdot u_i \geq 0$ implies $p \cdot u_j \leq 0$ for any $j \neq i$.

The proof is in the appendix, Section A.3.

Corollary 1 provides sufficient conditions for uniqueness. Part i considers the case of a single product, extending example 1 to any finite number of states. Part ii considers two “horizontally differentiated” products

as illustrated in figure 2C, where in each state, buyers want only one product or the other, although for intermediate beliefs they may want both. Finally, part iii considers a large number of “niche” products as illustrated in figure 3A, where buyers want only one product at any given belief.

It is notable that the improvability condition in theorem 2 involves only the monopoly strategy and is therefore stricter than required. For example, consider a reversed version of example 2 in which the low-quality product is more profitable than the high-quality product, and $p \in (0, \frac{1}{3})$. While this does not satisfy the sufficient condition, monopoly is still the unique equilibrium since any seller can raise her profits by providing slightly less information than the market. More generally, one could show uniqueness by checking whether profits drop along each path of local deviations; this depends on how buyer-optimal products are ordered in belief space.

IV. Full-Information Equilibria

In this section, we suppose that sellers imperfectly observe buyers’ beliefs ($\alpha > 0$). Our first result is that, as search costs vanish, there is a sequence of equilibria that converges to full information. We then derive conditions under which all equilibria reveal full information as search costs vanish, in stark contrast to the results in the last section.

Given that there are only new buyers on path, sellers have a degenerate prior $\eta_0 = \delta_{p_s}$, ignore the noisy signal \mathbf{y} , and thus have degenerate beliefs $\eta_y = \delta_{p_s}$. We can therefore drop the subscript \mathbf{y} and let σ^g and σ^i denote the general and independent signal structures, respectively, which we will refer to as policies without loss of generality.

We now describe how buyers update their payoff beliefs when they receive independent signals from a seller. Let $r(\theta)$ be the posterior of a buyer with prior p after observing a signal $\theta \in \Theta$, and let $r_q(\theta)$ denote the posterior of a buyer with prior q who observes the same signal. As shown by Alonso and Camara (2016), the posteriors of the two buyers are related as follows:

$$r_q(\theta) = \phi_q(r(\theta)),$$

where $\phi_q : \Delta S \rightarrow \Delta S$ is a mapping from the posteriors of buyer p to the posteriors of buyer q satisfying Bayes’s rule,

$$[\phi_q(r)](s) = \frac{q(s) \frac{r(s)}{p(s)}}{\sum_{s'} q(s') \frac{r(s')}{p(s')}}. \quad (1)$$

Intuitively, for any two buyers, the proportionality of their likelihood ratios for any two states remains constant when we update via Bayes’s rule.

That is, if buyer q starts off twice as optimistic as buyer p , then buyer q will remain twice as optimistic as p after any signal realization.

This section will be interested in convergence results as search costs get small. A *limit equilibrium* is a sequence of equilibria σ_n and associated search costs $c_n \rightarrow 0$ such that the buyer's equilibrium payoffs converge. We say a strategy is *fully informative* if it has full support on all degenerate posteriors, and a limit equilibrium is *full-information* if the buyer's equilibrium payoffs converge to that of a fully informative strategy.

A. *Motivating Example*

Example 1 (continued).—As before there are two payoff states $\{L, H\}$, and the buyer's initial prior is $p = \Pr(H)$. We first abstract from coordination and assume that sellers use simple strategies. We claim that when $p < 1/2$, the only equilibrium strategy provides full information as $c \rightarrow 0$. Intuitively, since the monopoly strategy provides some information, as the cost vanishes, a buyer can obtain a large number of signals at low cost and become (almost) fully informed. Hence the current seller must provide (almost) full information in order to beat this outside option and make a sale.

Since sellers use independent policies that do not depend on y , the stopping sets are of the form $Q_c(\sigma) = [0, a] \cup [b, 1]$, where $a \leq 1/2 \leq b$. The reason is that information has no value at the boundaries and a buyer's value function is convex and increasing in the posterior q . As a result, the seller's optimal strategy is a binary signal $p \rightarrow \{0, b\}$ as shown in figure 4A. Moreover, a buyer with prior b will have posteriors $b \rightarrow \{\phi_b(0), \phi_b(b)\}$. Using equation (1), we get

$$\begin{aligned} \phi_b(b) &= \frac{\frac{b}{p}}{\frac{b}{p} + (1-b)\frac{1-b}{1-p}} && \text{with probability } \frac{b}{\phi_b(b)}, \\ \phi_b(0) &= 0 && \text{with probability } 1 - \frac{b}{\phi_b(b)}. \end{aligned}$$

By the definition of the stopping set $Q_c(\sigma)$, the buyer with belief b must be indifferent between stopping and searching again, that is,

$$2b - 1 = \frac{b}{\phi_b(b)} [2\phi_b(b) - 1] - c.$$

Rearranging, this becomes

$$b^2 - [1 + p - c(1 - p)]b + p = 0. \tag{2}$$

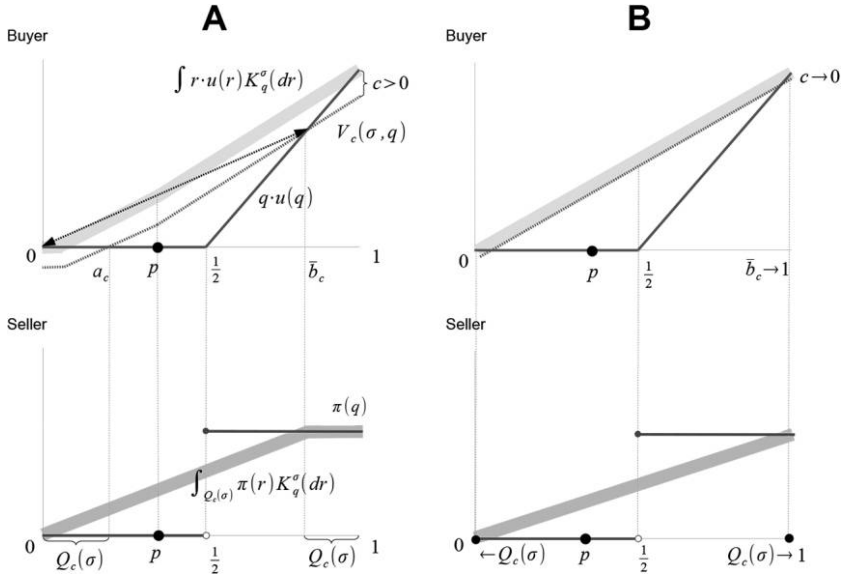


FIG. 4.—Single product (full-information equilibrium). This figure illustrates example 1 when sellers receive noisy signals about buyers’ beliefs ($\alpha > 0$) and use simple strategies. Panel A shows the unique equilibrium, which provides more information than a monopolist. Panel B shows that as the search cost vanishes, that is, $c \rightarrow 0$, this equilibrium converges to full information. Color version available as an online enhancement.

Since $p < \frac{1}{2}$, for small search cost c , only the larger root $b = \bar{b}_c$ of this quadratic is greater than $\frac{1}{2}$ and leads to a sale. Hence, the unique equilibrium is that sellers provide a signal $p \rightarrow \{0, \bar{b}_c\}$, which induces a stopping set $Q_c(\sigma) = [0, \bar{a}_c] \cup [\bar{b}_c, 1]$. Moreover, as $c \rightarrow 0$, the value of searching increases and $\bar{b}_c \rightarrow 1$, as shown in figure 4B. This means that the seller provides full information.

When $p \geq \frac{1}{2}$, a monopolist would provide no information and there are two limit equilibria as search costs vanish: full information and no information. First observe that “no information” is an equilibrium for any search cost: if sellers provide no information, the buyer will not search; and since “no information” is the monopoly strategy, no seller will defect. There are two further equilibria that correspond to the two roots of the quadratic in equation (2), which are denoted by $\{\underline{b}_c, \bar{b}_c\} \subset [\frac{1}{2}, 1]$. As $c \rightarrow 0$, $\underline{b}_c \rightarrow p$, which corresponds to no information, whereas $\bar{b}_c \rightarrow 1$, which corresponds to full information. In summary, full information is always a limit equilibrium and is the unique limit equilibrium if the monopolist provides any information.

So far, we have abstracted from coordination. No matter the level of β , there is always an equilibrium in which sellers ignore the coordination device, inducing the full-information limit equilibrium above. Moreover, if

coordination is imperfect ($\beta > 0$) and the monopolist provides information ($p < 1/2$), then full information is the only limit equilibrium. Intuitively, no matter what information the coordinating policy provides, there is at least proportion β of potential sellers who provide nontrivial independent information. Therefore, as the search cost vanishes, a buyer can obtain a large number of signals at low cost and become almost fully informed.

This example shows that when sellers imperfectly observe buyers' beliefs, equilibria can be much more informative than the monopoly strategy. Crucially, when sellers can perfectly observe buyers' beliefs, a buyer who receives information from one seller obtains no useful information if he were to go back to the market. However, when the observation of beliefs is imperfect, the buyer can pretend to be uninformed and receive more information from another seller in the market. The possibility of taking this outside option enables competition to function.

B. Full Information Is a Limit Equilibrium

We now return to the general model with multiple states and multiple products and show that there always exists a full-information limit equilibrium.

THEOREM 3. Suppose $\alpha > 0$. There exists a full-information limit equilibrium.

The proof is in the appendix, Section B.2.

The basic intuition is as follows. Suppose sellers ignore the coordination device and use simple strategies. Although a buyer visits only one seller on the equilibrium path, he always has the option of visiting other sellers and obtaining new information. As the search cost vanishes, the buyer can thus threaten to visit a growing number of sellers at a shrinking total cost. This means that if each seller provides some information about each state, the buyer will become fully informed. Hence the current seller has to match the market and also provide full information.²⁰

We now provide an overview of the proof. Fix a search cost $c > 0$ and suppose all other sellers use a simple strategy σ . This generates a value function $V_c(\sigma, \cdot)$ for the buyer and a corresponding stopping set $Q_c(\sigma)$. Faced with this stopping set, the current seller has a set of optimal simple strategies that we denote by $\varphi_c(\sigma)$. In other words, the mapping $\sigma \mapsto \varphi_c(\sigma)$ is the best-response correspondence. The set of all simple strategies is a convex compact space. If the best-response correspondence φ_c is non-empty, is convex-valued, and has a closed graph, then a direct application

²⁰ Note that theorem 3 does not say that the seller must provide full information but that the limit equilibrium strategy must give the buyer his full-information payoff. Intuitively, a buyer who makes the same decision in two different states will not pay a search cost to distinguish them.

of the Kakutani-Fan-Glicksberg fixed-point theorem implies that an equilibrium exists. Unfortunately, φ_c may not have a closed graph: as σ_n changes, the set of stopping beliefs $Q_c(\sigma_n)$ can increase, meaning that there are products available to the seller in the limit that are not available along the sequence, resulting in profits that may jump up discontinuously.²¹

To construct an equilibrium, we consider strategies that provide a lot of information. For example, consider figure 5A, where the D_ε regions represent payoff beliefs such that the buyer's stopping payoff is ε -away from that of full information. The first step is to observe that if all sellers use strategies with support in D_ε , then when search costs are small, the best response also has support in D_ε . Intuitively, if other sellers provide a lot of information, then buyers have high continuation values and the stopping sets are inside D_ε . Second, we show that φ_c has a closed graph. When ε is small, D_ε intersects only with beliefs in which products are chosen under full information. This means that as other sellers change their strategies, this does not affect the set of products that are available to the current seller. Together, these properties imply that an equilibrium exists with a payoff that is ε -close to that of full information. We can then take $\varepsilon \rightarrow 0$ as $c \rightarrow 0$ as required.

C. Uniqueness of the Full-Information Limit Equilibrium

Theorem 3 establishes that there always exists a full-information limit equilibrium. In this section, we characterize the set of limit equilibria and provide a sufficient condition under which the full-information limit equilibrium is unique.²²

To illustrate our results, consider the example illustrated in figures 5B and 5C, where there is a single product and p is the buyer's initial prior over the three payoff states $S = \{s_1, s_2, s_3\}$. For any event $E \subset S$, let $p_E \in \Delta S$ denote its *conditional belief* given E , that is, $p_E(s) = p(s)/p(E)$ for $s \in E$ (and $p_E(s) = 0$ for $s \notin E$). Figure 5B shows a limit equilibrium that is partitional and provides less than full information. Here, sellers provide independent information about s_1 only, so a buyer's posterior lies on the line connecting the conditional beliefs $p_{\{s_1\}}$ and $p_{\{s_2, s_3\}}$. From the current seller's perspective, a buyer near $p_{\{s_2, s_3\}}$ purchases her product, so the seller has no incentive to provide information about s_2 versus s_3 . As search costs go to zero, buyers thus learn for sure whether s_1 is true or not but learn nothing about s_2 versus s_3 ; this limit equilibrium is thus characterized by the partition $\{\{s_1\}, \{s_2, s_3\}\}$ and yields the buyer lower payoffs than full information.

²¹ For example, suppose that in example 2 there is a sequence of simple strategies σ_n such that $Q_c(\sigma_n) = [0, \frac{1}{3} - \varepsilon_n] \cup [\frac{2}{3}, 1]$. For all $\varepsilon_n > 0$, the optimal strategy is $p \rightarrow \{0, \frac{2}{3}\}$, whereas the optimal strategy in the limit is no information.

²² Theorem 1 shows that if there is perfect coordination ($\beta = 0$), then monopoly is an equilibrium; our discussion thus assumes that $\beta > 0$.

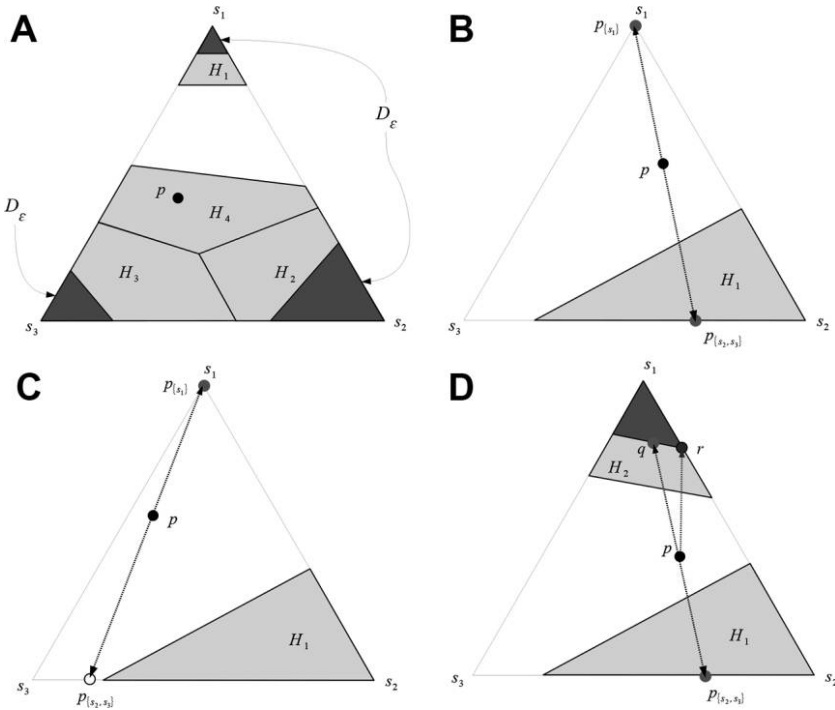


FIG. 5.—Limit equilibria. Panel A shows the construction of a full-informatory limit equilibrium in the proof of theorem 3. It also shows that this limit is unique as every nontrivial conditional belief is fully improvable. Panel B shows a limit equilibrium in which the buyer learns only about s_1 . Note that $p_{\{s_2, s_3\}}$ is not fully improvable since the buyer purchases the most profitable product possible at that conditional belief. Panel C shows the case in which learning only about s_1 is not a limit equilibrium as $p_{\{s_2, s_3\}}$ is fully improvable. By theorem 4, the unique limit equilibrium is full-informatory. Panel D shows an example in which $p_{\{s_2, s_3\}}$ is not fully improvable, but learning only about s_1 is still not a limit equilibrium. Color version available as an online enhancement.

In contrast, in figure 5C, we will see that full information is the only limit equilibrium. In particular, the partition $\{\{s_1\}, \{s_2, s_3\}\}$ is no longer an equilibrium. Intuitively, a buyer with belief close to $p_{\{s_2, s_3\}}$ does not buy, so the seller can strictly increase her profits by providing information about s_2 versus s_3 ; as long as sellers provide some information in this dimension, then as search costs vanish, a buyer can visit many sellers and fully learn the state. Hence, a single seller must match the market and the unique limit equilibrium is full-informatory.

We now formally characterize limit equilibria in terms of partitions. We call a strategy *partitional* if it corresponds to revealing to the buyer a partition \mathcal{E} of the payoff state space S . Given any independent policy σ , we say a partition \mathcal{E}_σ is *induced by* σ if for every event $E \in \mathcal{E}_\sigma$, $s, s' \in E$ iff the pol-

icy does not distinguish between them, that is, $\sigma(s) = \sigma(s')$. Finally, observe that since the space of independent policies is compact,²³ we can consider convergent sequences, $\sigma_n \rightarrow \sigma$, where we call σ its limit.

LEMMA 4. Suppose $\alpha > 0$ and $\beta > 0$. In any limit equilibrium, buyers' payoffs exceed those from the partition induced by the limit of the independent policies. Moreover, payoffs are equal if $\beta = 1$.

The proof is in Section B.3 of the appendix.

When sellers cannot coordinate ($\beta = 1$), lemma 4 says that all limit equilibria are partitional. Intuitively, if sellers use independent policies that are informative about state s versus s' , then as the search cost vanishes, a buyer can visit a large number of sellers, receive a large number of independent signals, and perfectly distinguish between the two states in the limit. In order to make a sale, any one seller must therefore match the market by providing this information immediately to the buyer, meaning that any limit equilibrium must be partitional.

However, there are a couple of subtleties with this intuition. First, as the search cost vanishes, it may be the case that the information provided by σ_n shrinks so that the partitions induced by each σ_n are richer than the partition induced by the limit policy σ . Even though a buyer can obtain more and more signals for a given search expenditure, lemma 4 says that his payoff is determined by the coarser partition corresponding to σ . Intuitively, sellers provide enough information so that the buyer searches only once on path, and his limit payoff is determined by the limit policy σ .²⁴

Second, when sellers can coordinate ($\beta < 1$), limit equilibria may not be partitional. This occurs when sellers' independent policies do not distinguish between states s and s' , while the coordinating signal does, so the buyer learns about the two states from only one draw of the coordinating signal. This can be an equilibrium because if a seller defected and provided independent information about s versus s' , then the buyer would continue searching for the coordinated signal.²⁵ In such a case, lemma 4 says that buyers still earn at least as much as from learning the partition induced by the limit of sellers' independent policies.

We now provide a sufficient condition for full information to be the unique limit equilibrium. We say that a payoff belief p is *nontrivial* if there are at least two products that a buyer with belief p could potentially buy;

²³ See lemma 9 in Sec. B.2 in the appendix.

²⁴ As an illustration, consider a reversed version of example 2 in which the low-quality good is more profitable than the high-quality good. Suppose $p > \frac{2}{3}$ and sellers provide vanishing amounts of information as $c_n \rightarrow 0$, so that $Q_{c_n}(\sigma_n) \cap [\frac{1}{3}, \frac{2}{3}] = \emptyset$. For any positive search cost, $\mathcal{E}_{\sigma_n} = \{\{L\}, \{H\}\}$, whereas in the limit, $\mathcal{E}_{\sigma} = \{\{L, H\}\}$.

²⁵ Again, consider the reversed version of example 2 from n. 24. When $p > \frac{2}{3}$ and the cost is small, there is a nonpartitional limit equilibrium in which the coordinating policy is a monopoly strategy (i.e., $p \rightarrow \{\frac{2}{3}, 1\}$), while the independent policy reveals no information. A buyer purchases after receiving either of these signals; in both cases, a subsequent signal provides no additional value.

that is, $\pi(1_s)$ is not constant for all $s \in S_p$. Recall from Section III that a belief is improvable if a monopolist prefers to release some information at that belief. We say that a payoff belief p is *fully improvable* if a monopolist prefers to release full information rather than no information, that is, $\bar{\Pi}(p) > \pi(p)$, where $\bar{\Pi}(p) := \sum_s p(s)\pi(1_s)$ is the full-information profit.

THEOREM 4. Suppose $\alpha > 0$ and $\beta > 0$. Any limit equilibrium is full-information if every nontrivial conditional belief is fully improvable.

This is proved in Section B.4 of the appendix.

Intuitively, if a single seller would like to release a little information along some dimension, a buyer can visit many sellers and accumulate this information, thus forcing a seller to release full information to make a sale. To prove this result, note that lemma 4 implies that buyers' payoffs are bounded below by the limit of the independent policies. We would thus like to take an equilibrium in which sellers provide no independent information about state s versus s' and show that any current seller would like to deviate and provide some information. If the stopping set for the current seller includes all beliefs on the segment connecting s and s' , then the result would require only that a monopolist would want to release some information at the conditional belief $p_{\{s,s'\}}$ (i.e., $p_{\{s,s'\}}$ is improvable). However, this is complicated by the two "subtleties" discussed after lemma 4. First, the partitions induced by the independent policies can be discontinuous in the limit; so even though the limit stopping set $Q_0(\sigma)$ may contain the entire segment of beliefs, each $Q_{c_n}(\sigma_n)$ may not (see n. 24). Second, with coordination, limit equilibria may not even be partitional, so $Q_0(\sigma)$ itself may not include the segment (see n. 25). Nevertheless, a seller can always provide full information at $p_{\{s,s'\}}$. Thus, theorem 4 states that if every conditional belief is fully improvable, then any limit equilibrium is full-information.

In terms of the examples, example 1 (with $p < 1/2$) and figures 5A and 5C satisfy full improvability, so full information is the only limit equilibrium. On the other hand, figure 5B does not satisfy full improvability and full information is not the unique limit equilibrium.²⁶

It is notable that the full improvability condition in theorem 4 considers a monopolist's incentives only at conditional beliefs and is therefore stricter than required. Indeed, a partitional strategy that provides less than full information can be ruled out in other ways. First, as suggested by the above intuition, the partition may contain a conditional belief that is im-

²⁶ One can also ask whether monopoly is still an equilibrium with $\alpha > 0$. When there is no coordination ($\beta = 1$), monopoly is generically a limit equilibrium iff the monopolist provides no information. This follows from lemma 4 and the fact that a monopolist generically does not provide full information along any dimension. When there is coordination ($\beta < 1$), n. 25 illustrates that there can still be equilibria whereby the coordinated policy is monopoly. However, in order for this to be an equilibrium, lemma 4 requires that the independent policy releases no information.

provable (rather than fully improvable).²⁷ Second, even if no conditional belief is improvable, there may be no sequence of equilibria that converges to the strategy.²⁸

Taken together, theorems 1–4 show how market outcomes depend on whether sellers can observe buyers' beliefs or coordinate their disclosures. Under perfect observation of beliefs or perfect coordination, monopoly is an equilibrium. Intuitively, sellers can discriminate between new and old buyers, allowing them to implicitly collude. However, under imperfect observation and imperfect coordination, full information is an equilibrium and, if conditional beliefs are fully improvable, the only equilibrium. Even though buyers purchase from the first seller on path, the option to anonymously mimic an uninformed buyer and receive more information forces sellers to compete against each other and provide much more information to buyers.

V. Extensions

In this section we discuss a number of extensions and applications. For simplicity, we discuss all the issues in the context of the two-state, single-product example (example 1) and abstract from the possibility of coordination; that is, we set $\beta = 1$.

The model assumes that buyers have a common prior p_0 . We first argue that the spirit of our results is not affected if priors were heterogeneous. To see this, suppose that the prior is $p_0 \in \{p_L, p_H\}$ with probabilities $\{f_L, f_H\}$. If sellers can perfectly observe buyers' beliefs, they can condition their disclosure strategy on the buyer's type. Hence, monopoly is the unique equilibrium for any search cost exactly as in Section III.A. In contrast, if sellers imperfectly observe buyers' beliefs, then full information is a limit equilibrium as in Section IV.A. This is easiest to see if buyers' beliefs are private ($\alpha = 1$), so disclosure policies are independent of \mathbf{y} and buyers' acceptance sets are of the form $Q_c(\sigma) = [0, a] \cup [b, 1]$ for $a \leq \frac{1}{2} \leq b$. If other sellers provide a lot of information, then $b \approx 1$ and it is optimal for the current seller to use a strategy that sells to both types, that is, $p_L \rightarrow \{0, b\}$.²⁹ Replacing p with p_L in Section IV.A, this means that a buyer

²⁷ For example, consider example 2 with $p \in (\frac{1}{3}, \frac{2}{3})$ and no coordination ($\beta = 1$). Even though p is not fully improvable, the no-information limit equilibrium can be ruled out because a seller can always earn higher profits by sending the buyer to $p \rightarrow \{p - \varepsilon, 1 - \varepsilon\}$ for small $\varepsilon > 0$.

²⁸ For example, consider fig. 5D with three payoff states and two goods. Given the partition $\mathcal{E} = \{\{s_1\}, \{s_2, s_3\}\}$, the seller has no incentive to provide information about s_2 vs. s_3 since $p_{(s_2, s_3)}$ is unimprovable. However, away from the limit, the seller wishes to provide information about s_2 vs. s_3 . In particular, the seller would like to send the buyer to posterior r rather than q , meaning that there is no limit equilibrium that corresponds to \mathcal{E} .

²⁹ If the seller sells to both types, then its profits are bounded below by $f_L p_L + f_H p_H$, which can be attained by releasing full information. If the seller sells to p_H only, then its profits are $f_H p_H / b$, which is smaller when $b \approx 1$.

can become fully informed as $c \rightarrow 0$ and the equilibrium converges to full information. Moreover, when $p_H < 1/2$, full information is the only limit equilibrium. Intuitively, if there are lots of p_L types, then the seller always sells to both types, that is, $p_L \rightarrow \{0, b\}$ as above. If there are lots of p_H types, then the low types may stay in the market for multiple periods; nevertheless, the seller always sells to p_H types immediately, and so the signal is at least as informative as $p_H \rightarrow \{0, b\}$.³⁰ In either case, buyers become fully informed as search costs vanish and sellers provide full information in the limit.

The model also assumes that sellers observe signals of buyers' beliefs, leading to the dichotomy between the cases with perfect observability ($\alpha = 0$) and imperfect observability ($\alpha > 0$). We now claim that the monopoly results of Section III.A carry over to a model in which sellers observe which sellers a buyer has visited in the past and their signal structures but not the resulting belief. First, one can see that there is an equilibrium in which the first seller provides the monopoly level of information and subsequent sellers provide no information. Given such strategies, a buyer purchases from the first seller who makes monopoly profits. If a buyer deviates and searches, then any subsequent seller knows that the buyer's belief is already in the acceptance set and thus provides no more information. Second, under the assumptions of example 1, this monopoly strategy is the unique equilibrium. Given an acceptance set of the form $Q_c(\sigma) = [0, a_1] \cup [b_1, 1]$, the first seller uses a binary signal, $p \rightarrow \{0, b_1\}$. If the buyer approaches seller 2, then she knows that the buyer's belief q lies in $\{0, b_1\}$. If $q = 0$, seller 2 cannot affect the buyer's posterior, so she should assume that $q = b_1$. Given an acceptance set $Q_c(\sigma) = [0, a_2] \cup [b_2, 1]$, seller 2 thus uses a binary signal $b_1 \rightarrow \{0, b_2\}$. Crucially, seller 2 would use exactly the same strategy if she could observe the belief of the buyer. The same logic applies to subsequent sellers, so the equilibria of the game coincide with those of the game in Section III.A. Given that the monopoly strategy is the only rationalizable outcome (see n. 17), it is also the unique equilibrium here. Intuitively, the logic behind the monopoly outcome relies on the seller being able to discriminate between new and old buyers; seeing the history of the buyer's shopping behavior is sufficient for this.

As a final extension, one might think that buyers can affect whether their beliefs are observed by sellers. This may mean browsing the internet anonymously or being tight-lipped when a financial adviser asks about their portfolio. By deleting their cookies, anonymous buyers exert a positive externality on others. Intuitively, when faced with an anonymous

³⁰ To see this, the first step is to show that the seller uses a binary signal of the form $p_H \rightarrow \{0, r\}$; this follows from the fact that "bad news" that lowers beliefs never leads to a sale, so the seller should send the buyer to posterior $q = 0$. The second step is to show the seller to p_H immediately; if this did not happen, then a seller can provide "two rounds" of information and raise her profits.

buyer, a seller provides more information, anticipating that he can return to the market and pose as a new buyer. This extra information benefits both new and anonymous buyers. However, if there is a small cost to becoming anonymous, then there is no equilibrium in which all buyers choose to do so. If this were the case, then sellers would reveal enough information to prevent an anonymous buyer from searching again; hence any anonymous buyer obtains the same payoff as one who can be tracked, while the former has to pay the fixed cost. This externality may provide a rationale for regulating online tracking, encouraging sellers to provide more information to their customers.

VI. Conclusion

This paper studies a market in which buyers search for better information about products and sellers choose how much information to disclose. This is particularly applicable to markets in which intermediaries provide advice and are compensated via commission, such as when websites recommend products, financial advisers advise on mutual funds, and doctors suggest different medical treatments. This also applies to areas such as organizational economics and political economy. For example, consider a CEO (“buyer”) who is interested in investing in a new innovation and sequentially approaches potential project managers (“sellers”) to investigate the market potential and report back. While the CEO wants to maximize company profits, the managers want the firm to make the investment and be chosen to lead the project.

It has been widely recognized that if there is a single seller, then the advice she provides will be distorted in her interest (e.g., high-commission products). However, one might think that competition is a substitute for regulation: intuitively, a buyer can visit many sellers and buy from the one who gives the best advice. Since competition is quite extensive among websites, financial advisers, and doctors, this would suggest that regulation would be unnecessary. Our analysis suggests that the impact of competition in advice markets depends on the information structure of the market. We show that even if there are many competing sellers, the monopoly disclosure strategy is an equilibrium if sellers can observe buyers’ beliefs or can successfully coordinate their disclosures. In either case, sellers can provide less information to old buyers than to new ones, discouraging buyers from searching and undermining competition. On the other hand, when sellers cannot fine-tune their disclosures, there is an equilibrium in which sellers provide full information as search costs vanish. When buyers search for information, they can visit many sellers and become better informed. This forces any given seller to match the market and provide close to full information in order to make a sale.

On the normative side, our results suggest that the effectiveness of competition depends on the details of the market. Thus, a regulator might be able to enhance the role of competition without directly regulating commissions (e.g., the United Kingdom’s ban on commissions for financial advisers). For example, regulators can make it harder to condition advice on buyers’ beliefs by requiring financial advisers to describe a strategy for clients before they see the customer’s current portfolio. Alternatively, regulators could make coordination harder by restricting the common marketing materials provided by mutual funds, increasing the dispersion of opinions. These forces are particularly powerful online, where sellers are increasingly able to calibrate information disclosure to the needs and experiences of the customer. Our results highlight the role that anonymity can play in fostering competition between recommender systems and illustrate how tracking technology can undermine the competitive mechanism and support collusive equilibria.

Appendix

A. *Proofs for Section III*

1. Proof of Lemma 1

Lemma 1 says that if a seller knows a buyer’s belief, that is, $\eta_y = \delta_p$, then it is optimal for her to sell with probability one. In other words, the stopping set $\mathbf{Q}_c(\sigma)$ is almost sure under the measure $K_p^{\sigma_y}$ for any possible realization of \mathbf{y} given \mathbf{p} . We will prove a stronger version of this below, which will be useful for proving subsequent results.

LEMMA 5. Suppose $\tilde{\sigma}$ is optimal given σ . If $\alpha = 0$, then for all $\mathbf{p} \in \Delta\Omega$, $K_p^{\tilde{\sigma}}(\mathbf{Q}_c(\sigma)) = 1$. Moreover, if $\sigma = \tilde{\sigma}$ is an equilibrium, then

$$V_c(\sigma, \mathbf{p}) = -c + \int_{\Delta\Omega} q \cdot u(q) K_p^{\sigma}(dq). \tag{A1}$$

If $\alpha > 0$ and $\eta_0 = \delta_{p_0}$, then the same holds for $\mathbf{p} = \mathbf{p}_0$.

Proof. Let $\tilde{\sigma}$ be optimal given σ and $\mathbf{Q} := \mathbf{Q}_c(\sigma)$ be the corresponding stopping set. First, consider the case in which $\alpha = 0$ so the seller knows the buyer’s belief, that is, $\eta_y = \delta_p$ for any buyer with belief $\mathbf{p} = \mathbf{y}$. Let $\mu = K_p^{\tilde{\sigma}}$ be the distribution on posteriors induced by $\tilde{\sigma}$, so profit is given by $\int_{\mathbf{Q}} \pi(r) \mu(d\mathbf{x})$. We will show that the buyer will always choose to purchase from the seller, that is, $\mu(\mathbf{Q}) = 1$.

Suppose otherwise, so μ puts strictly positive weight outside \mathbf{Q} . Consider the following deviation: provide an independent draw of full information to all buyers who would have ended up outside \mathbf{Q} under μ . Call this new signal $\hat{\sigma}$ and let $\nu = K_p^{\hat{\sigma}}$ be the corresponding posterior distribution. The seller’s profit from using ν is

$$\int_{\mathbf{Q}} \pi(r) \nu(d\mathbf{x}) = \int_{\mathbf{Q}} \pi(r) \mu(d\mathbf{x}) + \int_{\Delta\Omega \setminus \mathbf{Q}} \bar{\Pi}(r) \mu(d\mathbf{x}),$$

where $\bar{\Pi}(r)$ is the seller's profit from providing full information for a buyer with payoff belief r . We now show that the second term on the right-hand side is strictly positive. For any \mathbf{r} that is not in \mathbf{Q} , the buyer pays the search cost $c > 0$ in order to acquire some information. Hence, after being given full information, he must purchase with strictly positive probability (else he should have just stopped at \mathbf{r}). Since all products have positive profits, this means that the seller strictly benefits by providing full information at \mathbf{r} , that is $\bar{\Pi}(r) > 0$. By assumption, μ puts strictly positive weight on beliefs outside \mathbf{Q} , so the expected profit from using $\hat{\sigma}$ is strictly larger than from $\tilde{\sigma}$, contradicting the fact that $\tilde{\sigma}$ is optimal. Hence, μ must put all its weight in \mathbf{Q} and the buyer stops with probability one. Moreover, if $\sigma = \tilde{\sigma}$ is an equilibrium, then we have

$$\begin{aligned} V_c(\sigma, \mathbf{p}) &= -c + \int_{\Delta\Omega} \max\{q \cdot u(q), V_c(\sigma, \mathbf{q})\} K_{\mathbf{p}}^{\sigma}(d\mathbf{q}) \\ &= -c + \int_{\Delta\Omega} q \cdot u(q) K_{\mathbf{p}}^{\sigma}(d\mathbf{q}) \end{aligned}$$

as desired.

Finally, consider the case of $\alpha > 0$ and $\eta_0 = \delta_{\mathbf{p}_0}$. Recall from the text that from the full-support assumption, this implies that for the initial prior \mathbf{p}_0 , $\eta_y = \delta_{\mathbf{p}_0}$ for all realizations of y . Everything now follows as above but for $\mathbf{p} = \mathbf{p}_0$. QED

2. Proof of Lemma 3

Let $\alpha = 0$ and suppose that any nonmaximal belief is improvable. Consider some equilibrium σ and let

$$\Pi(\mathbf{p}) := \int_{\mathbf{Q}(\sigma)} \pi(q) K_{\mathbf{p}}^{\sigma}(d\mathbf{q})$$

be the corresponding profit from the equilibrium strategy. Fix some \mathbf{p} that is maximal, that is,

$$\pi(p) = \max_{s \in S_p} \pi(1_s) = \pi(1_s) = \pi_k,$$

where π_k and $s \in S_p$ are the corresponding profit and state. We will now show that \mathbf{p} is also a stopping belief under σ , that is, $\mathbf{p} \in \mathbf{Q}_c(\sigma)$.

Let H_k be the set of payoff beliefs in which product k is chosen under no information. Consider beliefs between \mathbf{p} and some degenerate belief at s , that is,

$$\mathbf{p}_{\lambda} := (1 - \lambda)\mathbf{p} + \lambda 1_s$$

for $\lambda \in [0, 1]$, and let p_{λ} denote the corresponding payoff belief. Note that since Π is concave and upper semicontinuous (see Aliprantis and Border 2006, lemma 17.30), $\Pi(\mathbf{p}_{\lambda})$ as a function of $\lambda \in [0, 1]$ is also concave and upper semicontinuous. This implies that $\Pi(\mathbf{p}_{\lambda})$ as a function of λ is continuous on $[0, 1]$.

We will now use a continuity argument as $\lambda \rightarrow 0$ to show that the buyer stops at \mathbf{p} . By the contrapositive of our assumption, every belief that is not improvable must

be maximal. This means that π_k is the highest attainable profit in ΔS_p .³¹ Note that the convexity of H_k implies that $p_\lambda \in H_k$ as both $p \in H_k$ and $1_s \in H_k$. Since π_k is the highest attainable profit in ΔS_p , π_k is also the monopoly profit at $p_\lambda \in H_k$, that is, $\Pi^*(p_\lambda) = \pi_k$. This implies that

$$\Pi(\mathbf{p}_\lambda) \leq \Pi^*(p_\lambda) = \pi_k.$$

Finally, define

$$\bar{\lambda} := \inf\{\lambda \in [0, 1] \mid \Pi(\mathbf{p}_\lambda) = \pi_k\}$$

and $\bar{\mathbf{p}} := \mathbf{p}_{\bar{\lambda}}$. Note that by the continuity of $\Pi(\mathbf{p}_\lambda)$, $\Pi(\bar{\mathbf{p}}) = \pi_k$.

We will now show that $\bar{\lambda} = 0$ so $\Pi(\bar{\mathbf{p}}) = \pi_k$. Suppose otherwise so we can find an increasing sequence $\lambda_n \nearrow \bar{\lambda}$. Let $\mathbf{p}_n := \mathbf{p}_{\lambda_n}$, so by continuity

$$\lim_n \Pi(\mathbf{p}_n) = \Pi(\bar{\mathbf{p}}) = \pi_k.$$

Since π_k is the highest attainable profit, it must be that $\lim_n K_{\mathbf{p}_n}^\sigma(H_k) = 1$. By the definition of $\bar{\lambda}$, $\Pi(\mathbf{p}_n) < \pi_k$. Thus, \mathbf{p}_n cannot be in the stopping set $\mathbf{Q}_c(\sigma)$ since otherwise the seller can achieve greater profits by providing no information. By the definition of $\mathbf{Q}_c(\sigma)$, it must be that

$$p_n \cdot u_k < V_c(\sigma, \mathbf{p}_n) = -c + \int_{\Delta\Omega} q \cdot u(q) K_{\mathbf{p}_n}^\sigma(d\mathbf{q}),$$

where the last expression follows from equation (A1). Rearranging terms yields

$$p_n \cdot u_k - \left[\int_{\{q \in \Delta\Omega \mid q \in H_k\}} q K_{\mathbf{p}_n}^\sigma(d\mathbf{q}) \right] \cdot u_k < -c + \int_{\{q \in \Delta\Omega \mid q \notin H_k\}} q \cdot u(q) K_{\mathbf{p}_n}^\sigma(d\mathbf{q}).$$

Note that both the left expression and the second term in the right expression vanish as $K_{\mathbf{p}_n}^\sigma(H_k) \rightarrow 1$. Since $c > 0$, we can find some $\lambda_n < \bar{\lambda}$ such that the above inequality is not satisfied, implying $\mathbf{p}_n \in \mathbf{Q}_c(\sigma)$, which yields a contradiction. Thus, $\bar{\lambda} = 0$, so $\Pi(\bar{\mathbf{p}}) = \Pi(\mathbf{p}) = \pi_k$.

Finally, since π_k is the most profitable product in ΔS_p and $\Pi(\bar{\mathbf{p}}) = \pi_k$, it must be that $K_{\bar{\mathbf{p}}}^\sigma(H_k) = 1$. Hence, information has no effect on the buyer's purchasing behavior, so the buyer might as well stop, that is, $\bar{\mathbf{p}} \in \mathbf{Q}_c(\sigma)$. QED

3. Proof of Corollary 1

We will show that in all three cases, nonmaximal beliefs are improvable and then apply theorem 2. Let Π^* denote the monopoly profit function. Note that part i follows trivially from either ii or iii. For part ii, consider two "horizontally differentiated" products with $\pi_1 > \pi_2$. We will prove this by contradiction. Suppose there is a belief \mathbf{p} that is neither maximal nor improvable, so

³¹ Suppose otherwise and let $\pi_j > \pi_k$ be the highest attainable profit in ΔS_p . Let $q \in \Delta S_p$ be such that $\pi(q) = \pi_j$. Clearly q is not improvable, so it must be maximal, which contradicts the definition of π_k .

$$\Pi^*(p) = \pi_2 < \pi_1 = \pi(1_s)$$

for some $s \in S_p$. We will show that it is a strictly profitable deviation for the seller to provide a perfect signal about s . Since $1_s \in H_1$, $u_1(s) \geq 0$ and “horizontal differentiation” implies that $1_s \cdot u_2 = u_2(s) \leq 0$. Let \mathbf{r} be the conditional belief knowing that s is not true, that is,

$$\mathbf{p} = p(s)\mathbf{1}_s + [1 - p(s)]\mathbf{r},$$

where $\mathbf{1}_s$ is the conditional belief knowing that s is true. Note that $1_s \cdot u_2 \leq 0$ and $p \cdot u_2 \geq 0$ imply that $r \cdot u_2 \geq 0$. Hence, the buyer chooses some profitable product at \mathbf{r} so $\pi(r) \geq \pi_2$. This implies that the seller can earn a strictly higher payoff by providing a perfect signal about state s , contradicting the fact that $\Pi^*(p) = \pi_2$. Hence, all nonmaximal beliefs are improvable.

The argument for part iii is analogous. We prove it again by contradiction. Suppose there is a belief \mathbf{p} that is neither maximal nor improvable, so

$$\Pi^*(p) = \pi_i < \pi_1 = \pi(1_s),$$

where π_1 is the statewise optimal product in S_p and $s \in S_p$. Again, consider the deviation of providing a perfect signal about s and let \mathbf{r} be the conditional belief knowing that s is not true, so $\mathbf{p} = p(s)\mathbf{1}_s + [1 - p(s)]\mathbf{r}$. By “nicheness” at 1_s , $1_s \cdot u_1 \geq 0$ implies $1_s \cdot u_i \leq 0$. Since $p \cdot u_i \geq 0$, this implies that $r \cdot u_i \geq 0$ from the definition of \mathbf{r} . By “nicheness” at r , this means that $r \cdot u_j \leq 0$ for all $j \neq i$. Hence, the profit at r is at least that of π_p , so by the same argument as before, the seller can obtain a strictly higher payoff by providing a perfect signal about s , yielding a contradiction. QED

B. Proofs for Section IV

This section includes the proofs for theorems 3 and 4. In subsection B.1, we first introduce some preliminary notation and results that will be useful in subsections B.2 and B.3 when we prove the theorems.

1. Preliminaries

We first establish some useful preliminary notation and results. First, note that when $\alpha > 0$, then all sellers ignore their private signal \mathbf{y} . Thus, σ^s and σ^i are independent of \mathbf{y} ; so without loss of generality, we let σ^s and σ^i denote the general and independent signal structures. The value function of the buyer is now given by

$$V_c(\sigma, \mathbf{q}) = -c + \int_{\Delta\Omega} \max\{r \cdot u(r), V_c(\sigma, \mathbf{r})\} [(1 - \beta)K_q^{\sigma^s} + \beta K_q^{\sigma^i}] (d\mathbf{r}).$$

Moreover, if a seller uses a simple strategy, then she provides only a single independent signal structure, which we will simply refer to as the strategy.

When sellers use a simple strategy σ , we can work only with payoff beliefs and a buyer with the initial prior p_0 has posterior beliefs distributed according to $K_{p_0}^\sigma$. Hence, we can associate a simple strategy σ with its posterior distribution $\mu = K_{p_0}^\sigma$.

Henceforth, we will use the simple strategy σ and its associated posterior distribution μ on payoff beliefs in ΔS interchangeably. We will refer to μ as the seller's strategy as well. Since signals are all independent, without loss we can define the canonical signal space as the set of posterior beliefs of buyer p . In other words, each signal realization corresponds to a posterior belief $r \in \Delta S$ of the buyer with prior p . The joint distribution of states and signals from the perspective of buyer p is given by Bayes's rule as

$$p(s)\sigma(s, dr) = \mu(dr)r(s). \tag{A2}$$

Now, a buyer with prior q has signal distribution μ_q , and for each signal realization r , his posterior belief is $\phi_q(r)$, that is

$$q(s)\sigma(s, dr) = \mu_q(dr)[\phi_q(r)](s).$$

Combining this with equation (A2) yields the following:

$$\mu_q(dr) = \mu(dr) \sum_{s \in S} q(s) \frac{r(s)}{p(s)}. \tag{A3}$$

Also recall the following expression for $\phi_q(r)$ from equation (1):

$$[\phi_q(r)](s) \left[\sum_{s' \in S} q(s') \frac{r(s')}{p(s')} \right] = q(s) \frac{r(s)}{p(s)}.$$

Consider a buyer with prior q_0 who searches t times when all sellers are using strategy μ . Let $q_t := \phi_{q_{t-1}}(r_t)$ be his posterior belief in period t given the t -period signal realization $r_t \in \Delta S$. Assuming zero search costs, his payoff at the end of the t periods is given by

$$W^t(\mu, q_0) := \int_{(\Delta S)^t} \max_i (u_i \cdot q_t) \mu_{q_0}(dr_1) \cdots \mu_{q_{t-2}}(dr_{t-1}) \mu_{q_{t-1}}(dr_t).$$

When $t = 1$, we just set $W := W^1$. We first show that W^t has a simpler expression.

LEMMA 6. For any strategy μ , belief q , and time period t ,

$$W^t(\mu, q) = \int_{(\Delta S)^t} \max_i \left(\sum_s u_i(s) q(s) \frac{r_1(s)}{p(s)} \cdots \frac{r_t(s)}{p(s)} \right) \mu(dr_1) \cdots \mu(dr_t).$$

Proof. Using the expression for μ_{q_t} from equation (A3), we obtain

$$\begin{aligned} W^t(\mu, q_0) &= \int_{(\Delta S)^t} \max_i (u_i \cdot q_t) \left[\sum_s q_{t-1}(s) \frac{r_t(s)}{p(s)} \right] \mu_{q_0}(dr_1) \cdots \mu_{q_{t-2}}(dr_{t-1}) \mu(dr_t) \\ &= \int_{(\Delta S)^t} \max_i \left(\sum_s u_i(s) q_t(s) \left[\sum_s q_{t-1}(s) \frac{r_t(s)}{p(s)} \right] \right) \mu_{q_0}(dr_1) \cdots \mu_{q_{t-2}}(dr_{t-1}) \mu(dr_t). \end{aligned}$$

Since $q_t := \phi_{q_{t-1}}(r_t)$, from equation (1), we have

$$W^t(\mu, q_0) = \int_{(\Delta S)^t} \max_i \left(\sum_s u_i(s) q_{t-1}(s) \frac{r_t(s)}{p(s)} \right) \mu_{q_0}(dr_1) \cdots \mu_{q_{t-2}}(dr_{t-1}) \mu(dr_t).$$

Repeating the above argument for each μ_{φ} we obtain the desired expression for W^t . QED

We now show that if the signal is informative about some event, then the buyer will eventually learn it. In order to prove this, recall some of the notation from Section IV.C. Recall that p_E denotes the conditional belief of event $E \subset S$ under the prior \mathbf{p} ; that is, $p_E(s) = p(s)/p(E)$ for all $s \in E$ and $p_E(s) = 0$ otherwise. Given any strategy μ with corresponding signal structure σ , let \mathcal{E}_μ denote the smallest such partition of S such that $\sigma(s) = \sigma(s')$ for every $s, s' \in E \in \mathcal{E}_\mu$. Note that from equation (A2), this implies that for $s, s' \in E \in \mathcal{E}_\mu$,

$$\mu(dr) \frac{r(s)}{r(s')} = p(s)\sigma(s, dr) \frac{1}{r(s')} = \sigma(s', dr) \frac{p(s)}{r(s')} = \mu(dr) \frac{p(s)}{p(s')},$$

so $r(s)/r(s') = p(s)/p(s')$ μ -a.s. This means that conditional posteriors are equal to the conditional priors; that is, $r_E = p_E$ μ -a.s. for all $E \in \mathcal{E}_\mu$. Let $\mu_{\mathcal{E}}$ denote the limit partitional strategy given μ , that is,

$$\mu_{\mathcal{E}} := \sum_{E \in \mathcal{E}_\mu} p(E) \delta_{p_E}.$$

We now show that the payoff of a buyer who searches indefinitely eventually converges to that of its limit partitional strategy.

LEMMA 7. For any strategy μ , $W^t(\mu, q) \rightarrow W(\mu_{\mathcal{E}}, q)$.

Proof. Fix a strategy μ and let \mathcal{E}_μ be its limit partition. Since $r(s)/p(s) = r(E)/p(E)$ for all $s \in E \in \mathcal{E}_\mu$, we can rewrite W^t from lemma 6 as

$$\begin{aligned} W^t(\mu, q) &= \int_{(\Delta S)^t} \max_i \left(\sum_E \frac{r_1(E)}{p(E)} \cdots \frac{r_t(E)}{p(E)} \sum_{s \in E} u_i(s) q(s) \right) \mu(dr_1) \cdots \mu(dr_t) \\ &= \int_{(\Delta S)^t} \max_i \left(\sum_E \frac{r_1(E)}{p(E)} \cdots \frac{r_t(E)}{p(E)} q(E) (u_i \cdot q_E) \right) \mu(dr_1) \cdots \mu(dr_t). \end{aligned}$$

Defining

$$z_t(E) := \frac{r_1(E)}{p(E)} \cdots \frac{r_t(E)}{p(E)} q(E),$$

we can rewrite this as

$$\begin{aligned} W^t(\mu, q) &= \int_{(\Delta S)^t} \max_i \left(\frac{\sum_{E'} z_t(E') (u_i \cdot q_{E'})}{\sum_{E'} z_t(E')} \right) \left[\sum_E z_t(E) \right] \mu(dr_1) \cdots \mu(dr_t) \\ &= \sum_E q(E) \int_{(\Delta S)^t} \max_i \left(\frac{\sum_{E'} z_t(E') (u_i \cdot q_{E'})}{\sum_{E'} z_t(E')} \right) \mu_E(dr_1) \cdots \mu_E(dr_t), \end{aligned}$$

where

$$\mu_E(dr) := \frac{r(E)}{p(E)} \mu(dr).$$

Consider two events $E, E' \in \mathcal{E}_\mu$ such that $E \neq E'$. We want to show that the ratio of $z_t(E')$ to $z_t(E)$ goes to zero μ_E -a.s. as $t \rightarrow \infty$. Now

$$\begin{aligned} \frac{1}{t} \log \frac{z_t(E')}{z_t(E)} &= \frac{1}{t} \log \left(\frac{r_1(E')}{p(E')} \dots \frac{r_t(E')}{p(E')} \frac{p(E)}{r_1(E)} \dots \frac{p(E)}{r_t(E)} \frac{q(E')}{q(E)} \right) \\ &= \frac{1}{t} \sum_{t' \leq t} \log \left(\frac{r_{t'}(E')}{p(E')} \frac{p(E)}{r_{t'}(E)} \right) + \frac{1}{t} \log \left(\frac{q(E')}{q(E)} \right). \end{aligned}$$

By the law of large numbers, we have

$$\frac{1}{t} \log \frac{z_t(E')}{z_t(E)} \rightarrow \int_{\Delta S} \log \left(\frac{r(E')}{p(E')} \frac{p(E)}{r(E)} \right) \mu_E(dr).$$

Since $E \neq E'$, by the definition of \mathcal{E}_μ , it must be that $r(E')/r(E)$ is different on some strictly positive μ -measure. The reason is that if $r(E')/r(E)$ is constant μ -a.s., then for all $s \in E$ and $s' \in E'$,

$$\frac{r(s)}{p(s)} = \frac{r(E)}{p(E)} = \frac{r(E')}{p(E')} = \frac{r(s')}{p(s')};$$

so $E \cup E' \in \mathcal{E}_\mu$, which contradicts the fact that \mathcal{E}_μ is the smallest such partition of S by definition. By Jensen's inequality,

$$\begin{aligned} \int_{\Delta S} \log \left(\frac{r(E')}{p(E')} \frac{p(E)}{r(E)} \right) \mu_E(dr) &< \log \left(\int_{\Delta S} \frac{r(E')}{p(E')} \frac{p(E)}{r(E)} \mu_E(dr) \right) \\ &= \log \left(\int_{\Delta S} \frac{r(E')}{p(E')} \mu(dr) \right) \\ &= \log \left(\frac{\int_{\Delta S} r(E') \mu(dr)}{p(E')} \right) = 0, \end{aligned}$$

where the strict inequality is due to the fact that $r(E')/r(E)$ is nonconstant μ_E -a.s. Thus, $(1/t) \log(z_t(E')/z_t(E))$ converges to something strictly less than zero. Hence, it must be that $\log(z_t(E')/z_t(E)) \rightarrow -\infty$ or $z_t(E')/z_t(E) \rightarrow 0$ for all $E' \neq E$. By dominated convergence, this implies that

$$\begin{aligned} \lim_t W^t(\mu, q) &= \sum_E q(E) \int_{(\Delta S)^t} \max_i \left(\lim_t \frac{\sum_{E'} z_t(E')(u_i \cdot q_{E'})}{\sum_{E'} z_t(E')} \right) \mu_E(dr_1) \dots \mu_E(dr_t) \\ &= \sum_E q(E) \int_{(\Delta S)^t} \max_i (u_i \cdot q_E) \mu_E(dr_1) \dots \mu_E(dr_t) \\ &= \sum_E q(E) \max_i (u_i \cdot q_E). \end{aligned}$$

From lemma 6 and the definition of μ_ε , we have

$$\begin{aligned} W(\mu_\varepsilon, q) &= \int_{\Delta S} \max_i \left(\sum_s u_i(s) q(s) \frac{r(s)}{p(s)} \right) \mu_\varepsilon(dr) = \sum_E p(E) \max_i \left(\sum_{s \in E} u_i(s) \frac{q(s)}{p(E)} \right) \\ &= \sum_E p(E) \max_i \left(\frac{q(E)}{p(E)} u_i \cdot q_E \right) = \sum_E q(E) \max_i (u_i \cdot q_E), \end{aligned}$$

so $\lim_t W^t(\mu, q) = W(\mu_\varepsilon, q)$. QED

2. Proof of Theorem 3

We now prove that when $\alpha > 0$, there always exists a full-information limit equilibrium. We will be considering only simple strategies, so without loss we will consider only payoff beliefs in ΔS . We will also be making use of the preliminary results in the previous section, which uses some of the notation from Section IV.C. Given any $\varepsilon > 0$, let $D_\varepsilon \subset \Delta S$ be the set of beliefs such that the buyer's stopping payoff is at most ε -away from his full-information payoff. In other words,

$$D_\varepsilon := \{q \in \Delta S | q \cdot u(q) \geq W(\bar{\mu}, q) - \varepsilon\},$$

where $\bar{\mu}$ is the fully informative strategy. Let M_ε be the set of simple strategies that put weight only in D_ε , that is, $\mu(D_\varepsilon) = 1$ for all $\mu \in M_\varepsilon$. Denote $\Pi_\varepsilon(\mu, \nu)$ as the profit of a seller for using simple strategy ν if all other sellers are using simple strategy μ , that is,

$$\Pi_\varepsilon(\mu, \nu) := \int_{Q_\varepsilon(\mu)} \pi(q) \nu(dq),$$

and define the *best-response correspondence* $\varphi_\varepsilon : M_\varepsilon \rightrightarrows M_\varepsilon$ such that for $\nu \in \varphi_\varepsilon(\mu)$,

$$\Pi_\varepsilon(\mu, \nu) \geq \Pi_\varepsilon(\mu, \nu')$$

for any other simple strategy ν' .

We first show that we can restrict ourselves to simple strategies and these best-response correspondences without loss of generality.

LEMMA 8. Suppose $\tilde{\sigma} = (\tilde{\sigma}^g, \tilde{\sigma}^i)$ is optimal given a simple strategy σ . Then the simple strategy $\tilde{\sigma}^i$ is also optimal given σ .

Proof. Since $\tilde{\sigma} = (\tilde{\sigma}^g, \tilde{\sigma}^i)$ is optimal given σ , it solves

$$\max_{\tilde{\sigma}} \int_{Q_\varepsilon(\sigma)} \pi(q) [(1 - \beta) K_p^{\tilde{\sigma}^g} + \beta K_p^{\tilde{\sigma}^i}] (dr).$$

Note that since σ is simple, we can work with payoff beliefs when buyers go back to the market. Define $\mu^g := K_p^{\tilde{\sigma}^g}$ and $\mu^i := K_p^{\tilde{\sigma}^i}$ and suppose

$$\int_{Q_\varepsilon(\sigma)} \pi(r) \mu^g(dr) > \int_{Q_\varepsilon(\sigma)} \pi(r) \mu^i(dr).$$

In this case, if the seller were to use the strategy $(\tilde{\sigma}^g, \hat{\sigma})$, where $\hat{\sigma}$ is the independent version of $\tilde{\sigma}^g$, then she would earn a strictly higher profit violating the optimality of $\tilde{\sigma}$. Thus,

$$\int_{Q_c(\hat{\sigma})} \pi(r)\mu^g(dr) \leq \int_{Q_c(\hat{\sigma})} \pi(r)\mu^i(dr).$$

Moreover, since the seller can always set $\tilde{\sigma}^g = \tilde{\sigma}^i$, it must be that

$$\int_{Q_c(\hat{\sigma})} \pi(r)\mu^g(dr) = \int_{Q_c(\hat{\sigma})} \pi(r)\mu^i(dr).$$

Thus, by using the simple strategy $\tilde{\sigma}^i$, the seller can achieve the same optimal profit. QED

The main idea is that we will construct a sequence of equilibrium payoffs that converges to the buyer’s full-information payoff as search costs vanish. For each ε , we will find some small enough search cost c such that $\varphi_c(\mu)$ is nonempty for all $\mu \in M_c$. We will then use the Kakutani-Fan-Glicksberg (KFG) theorem to show that there exists an equilibrium strategy $\mu \in M_c$. Hence, as $\varepsilon \rightarrow 0$, we can find a sequence of decreasing search costs and corresponding equilibrium strategies $\mu \in M_c$ such that the buyer’s payoff converges to $W(\bar{\mu}, p)$.

First, we show that the domain M_c satisfies the usual conditions so that we can apply KFG.

LEMMA 9. M_c is a nonempty, compact, and convex metric space.

Proof. To see why M_c is nonempty, note that the fully informative strategy $\bar{\mu}$ puts weight only on degenerate beliefs, so $\bar{\mu}(D_c) = 1$ and $\bar{\mu} \in M_c$. The convexity of M_c follows trivially. To see why M_c is compact, note that since ΔS is a compact metric space, the set of all simple strategies is also a compact metric space by theorem 15.11 of Aliprantis and Border (2006). By their theorem 3.28, it is also sequentially compact. We now show that M_c is also sequentially compact. Consider the sequence $\mu_n \in M_c$, so it has a convergent subsequence $\mu_n \rightarrow \mu$. Since μ_n are all strategies, it must be that

$$\int_{\Delta S} q\mu(dq) = \lim_n \int_{\Delta S} q\mu_n(dq) = p,$$

so μ is also a simple strategy. Moreover, since $\mu_n(D_c) = 1$ and D_c is closed, by theorem 15.3 of Aliprantis and Border (2006),

$$\mu(D_c) \geq \lim_n \sup \mu_n(D_c) = 1,$$

so $\mu \in M_c$. Thus, M_c is sequentially compact and therefore compact. QED

Next, we show that φ_c is convex-valued. The proof follows directly from the fact that expected profits are linear in strategies.

LEMMA 10. $\varphi_c(\mu)$ is convex for all $\mu \in M_c$.

Proof. Let $v_1, v_2 \in \varphi_c(\mu)$ and consider $v := av_1 + (1 - a)v_2$ for some $a \in [0, 1]$. Note that v is also a strategy and

$$\Pi_c(\mu, v) = a\Pi_c(\mu, v_1) + (1 - a)\Pi_c(\mu, v_2) \geq \Pi_c(\mu, v')$$

for any other strategy v' . Hence, $v \in \varphi_c(\mu)$, so $\varphi_c(\mu)$ is convex. QED

As an intermediary step, we show that the buyer's value function is jointly continuous in both strategies and posterior beliefs.

LEMMA 11. The buyer's value function V_c is continuous in (μ, q) and convex in q .

Proof. Fix some $c > 0$ and suppose that all sellers are using the strategy μ . If the buyer searches once, then his payoff is

$$V_c^1(\mu, q) := -c + W(\mu, q).$$

Using the expression for W from lemma 6, we see that W is continuous by corollary 15.7 of Aliprantis and Border (2006). Moreover, since support functions are convex, W is also convex in q . Hence, V_c^1 is continuous and convex in q , and define a sequence of finite-period value functions iteratively where

$$V_c^{t+1}(\mu, q) := -c + \int_{\Delta S} \max \{ \max_i \phi_q(r) \cdot u_i, V_c^t(\mu, \phi_q(r)) \} \mu_q(dr).$$

Recall from equation (A3) that

$$\mu_q(dr) = \mu(dr) \sum_s q(s) \frac{r(s)}{p(s)},$$

and also note that $\phi_q(r)$ is continuous in q . Hence, applying the same argument for V_c^1 inductively, we conclude that V_c^t is continuous and convex in q . Since V_c^t converges uniformly to the value function V_c , the latter is also continuous and convex in q . QED

The next step is to show that the best-response correspondence φ_c takes on nonempty values so we can apply KFG. While this may not be true in general, we will choose ε small enough such that this holds. Set $\bar{\varepsilon} > 0$ such that it is less than the minimal nonzero difference between the buyer's full and partitional information payoffs. In other words, choose $\bar{\varepsilon}$ such that whenever $W(\bar{\mu}, p) > W(\mu_\varepsilon, p)$ for any limit partitional strategy μ_ε , then

$$\bar{\varepsilon} < W(\bar{\mu}, p) - W(\mu_\varepsilon, p).$$

Note that such an $\bar{\varepsilon}$ always exists as there are only a finite number of partitions of S .

We now show that for all $\varepsilon \leq \bar{\varepsilon}$, we can always find a search cost c small enough such that φ_c takes on nonempty values on M_ε . The main idea is as follows. Recall that D_ε is the set of beliefs such that the buyer's payoff is ε -close to his full-information payoff. Now whenever μ sends the buyer only to beliefs in D_ε for some $\varepsilon \leq \bar{\varepsilon}$, then the definition of $\bar{\varepsilon}$ ensures that the buyer's payoff under μ approaches that of full information as search costs vanish. Hence, the buyer will stop only if he gets close to his full-information payoff, so we can find a search cost c small enough such that the stopping set $Q_c(\mu)$ is completely contained in D_ε for all $\mu \in M_\varepsilon$. Since any optimal best-response strategy $\nu \in \varphi_c(\mu)$ will not send the buyer outside his stopping set, it must be that ν also puts full weight in D_ε , so $\nu \in M_\varepsilon$ as well.

LEMMA 12. If $\varepsilon \leq \bar{\varepsilon}$, then there exists a $c_\varepsilon > 0$ such that for all $c \leq c_\varepsilon$, $\varphi_c(\mu)$ is nonempty for all $\mu \in M_\varepsilon$.

Proof. Let $\varepsilon \leq \bar{\varepsilon}$. We first show that for every $\mu \in M_\varepsilon$, the value function $V_c(\mu, q)$ converges to $W(\bar{\mu}, q)$ uniformly as $c \rightarrow 0$. For each search cost c , suppose the buyer searches t_c times, where t_c is the largest integer less than $c^{-1/2}$ so $t_c c \leq c^{1/2}$. Since the buyer can always do this, his value function is at least his payoff from visiting t_c sellers. In other words,

$$V_c(\mu, q) \geq -t_c c + W^{t_c}(\mu, q) \geq -c^{1/2} + W^L(\mu, q).$$

As $c \rightarrow 0$, $t_c \rightarrow \infty$, so $\liminf_n V_{c_n}(\mu, q) \geq W(\mu_\varepsilon, q)$ from lemma 7.

We now show that $W(\mu_\varepsilon, q) = W(\bar{\mu}, q)$. First note that since $\mu(D_\varepsilon) = 1$, the definition of D_ε means that μ -a.s.

$$\max_i r \cdot u_i \geq W(\bar{\mu}, r) - \varepsilon.$$

Note that the left-hand expression is convex in r ; so if we let $r = \sum_{E \in \mathcal{E}_\mu} r(E)r_E$, then μ -a.s.

$$\sum_{E \in \mathcal{E}_\mu} r(E) \max_i (r_E \cdot u_i) \geq \max_i \left(\sum_{E \in \mathcal{E}_\mu} r(E)r_E \right) \cdot u_i \geq W(\bar{\mu}, r) - \varepsilon.$$

Since $r_E = p_E$ μ -a.s. for all $E \in \mathcal{E}_\mu$, rearranging we have that μ -a.s.

$$\varepsilon \geq W(\bar{\mu}, r) - \sum_{E \in \mathcal{E}_\mu} r(E) \max_i (p_E \cdot u_i).$$

Note that the buyer's full-information payoff $W(\bar{\mu}, r)$ is linear in r . Hence, taking expectations with respect to μ yields

$$\varepsilon \geq W(\bar{\mu}, p) - \sum_{E \in \mathcal{E}_\mu} p(E) \max_i (p_E \cdot u_i) = W(\bar{\mu}, p) - W(\mu_\varepsilon, p).$$

By the definition of $\bar{\varepsilon}$, this implies that $W(\bar{\mu}, p) = W(\mu_\varepsilon, p)$. From lemma 6, it is straightforward to show that this implies $W(\mu_\varepsilon, q) = W(\bar{\mu}, q)$.

We thus have $\liminf_n V_{c_n}(\mu, q) \geq W(\bar{\mu}, q)$. Since $\bar{\mu}$ is fully informative, this implies that the buyer's value function converges to that of full information as search costs vanish, that is, $V_{c_n}(\mu, q) \rightarrow W(\bar{\mu}, q)$. We will use Dini's theorem (theorem 2.66 of Aliprantis and Border [2006]) to show that this convergence is uniform on the domain $M_\varepsilon \times \Delta S$. Note that $M_\varepsilon \times \Delta S$ is compact by lemma 9 and V_c is continuous by lemma 11. Since V_{c_n} is decreasing monotonically, $V_{c_n}(\mu, q) \rightarrow W(\bar{\mu}, q)$ uniformly by Dini's theorem. Hence, we can find some $c_\varepsilon > 0$ such that

$$|V_{c_n}(\mu, q) - W(\bar{\mu}, q)| < \varepsilon \quad \text{for all } (\mu, q) \in M_\varepsilon \times \Delta S.$$

This implies that for all $c \leq c_\varepsilon$,

$$\begin{aligned} Q_c(\mu) \subset Q_\varepsilon(\mu) &= \{q \in \Delta S \mid q \cdot u(q) \geq V_c(\mu, q)\} \\ &\subset \{q \in \Delta S \mid q \cdot u(q) \geq W(\bar{\mu}, q) - \varepsilon\} = D_\varepsilon. \end{aligned}$$

If ν is a best-response strategy to any $\mu \in M_\varepsilon$, then by the same argument as in lemma 5, ν must have full support in the stopping set, that is, $\nu(Q_c(\mu)) = 1$. Hence, $\nu(D_\varepsilon) \geq \nu(Q_c(\mu)) = 1$, so $\nu \in M_\varepsilon$ as desired. QED

Going forward, we assume $\varepsilon < \bar{\varepsilon}$. The last step before we can apply KFG is to show that φ_c has a closed graph. In general, this may not be true, but again we will find ε small enough such that this holds. Let $\mu_n \rightarrow \mu$, $\nu_n \in \varphi_c(\mu_n)$, and $\nu_n \rightarrow \nu$. If we can show that $\nu \in \varphi_c(\mu)$, then φ_c has a closed graph. Let $\nu^* \in \varphi_c(\mu)$ be a best-response strategy to μ . Recall that H_i is the set of beliefs in which u_i is chosen under no information, assuming ties are resolved in the seller's favor. The first step is to find ε small enough such that whenever ν^* recommends some product i , that is, $\nu^*(H_i) > 0$, then we can find some state $s \in S$ where $1_s \in H_i$ as well. For instance, in the vertical differentiation example of Section III.C, this means that the optimal strategy never recommends the low-quality TV. In figure 5A of Section IV, this means that D_ε is close enough to the vertices such that the seller will never recommend product 4.

Let Δu denote the smallest nonzero difference between buyer payoffs across all products at all degenerate beliefs. In other words, $\Delta u \leq |u_j(s) - u_i(s)|$ for all i, j and $s \in S$, where $u_j(s) \neq u_i(s)$. Define $\tilde{\varepsilon} > 0$ such that for every $\pi_j > \pi_i$,

$$\tilde{\varepsilon} < \left(\frac{\pi_j - \pi_i}{\pi_j} \right) \Delta u.$$

This is the smallest profitable profit deviation for the smallest difference in buyer payoffs. Note that this is well defined as the number of products and states are both finite. Set $\varepsilon^* := \min\{\bar{\varepsilon}, \tilde{\varepsilon}\}$.

LEMMA 13. Let $\varepsilon \leq \varepsilon^*$. Then $\nu^*(H_i) > 0$ implies $1_s \in H_i$ for some $s \in S$.

Proof. We will prove this by contradiction. Suppose $\nu^*(H_i) > 0$ for some product i and $1_s \notin H_i$ for all $s \in S$. Since $\varepsilon \leq \bar{\varepsilon}$, $\nu^* \in M_\varepsilon$, so we know that $\nu^*(H_i \cap D_\varepsilon) > 0$. Now, consider a belief $q \in H_i \cap D_\varepsilon$. We will show that providing full information at q will yield a strictly higher profit for the seller, which contradicts the fact that ν^* is a best-response strategy. Let $S_i \subset S$ denote the set of states such that product i is buyer-optimal, that is, $u_i(s) = u(1_s)$ for all $s \in S_i$. Since $1_s \notin H_i$ for all $s \in S$ by assumption, this means that for every $s \in S$, either $s \notin S_i$ or $s \in S_i$; but there is another buyer-optimal product that is more profitable for the seller, that is $\pi(1_s) \geq \pi_j > \pi_i$ for some product j . Hence, by providing full information at q , the seller will get profit

$$\sum_{s \in S} q(s) \pi(1_s) \geq \sum_{s \notin S_i} q(s) \pi(1_s) + \sum_{s \in S_i} q(s) \pi_j \geq \pi_j \sum_{s \in S_i} q(s) \tag{A4}$$

as all profits are positive. Since $q \in H_i \cap D_\varepsilon$, we know that

$$q \cdot u_i = q \cdot u(q) \geq W(\bar{\mu}, q) - \varepsilon.$$

Rearranging,

$$\varepsilon \geq \sum_s q(s) [u(1_s) - u_i(s)] = \sum_{s \notin S_i} q(s) [u(1_s) - u_i(s)] \geq (\Delta u) \sum_{s \notin S_i} q(s).$$

This implies that

$$\sum_{s \in S_i} q(s) = 1 - \sum_{s \notin S_i} q(s) \geq 1 - \frac{\varepsilon}{\Delta u}.$$

Combining this with inequality (A4), we get

$$\sum_{s \in S} q(s) \pi(1_s) \geq \pi_j \left(1 - \frac{\varepsilon}{\Delta u} \right) = \pi_j - \varepsilon \frac{\pi_j}{\Delta u} > \pi_i,$$

where the last strict inequality follows from the definition of $\tilde{\varepsilon}$. Since providing full information at q is more profitable for the seller, ν^* cannot be optimal, yielding a contradiction. QED

The next step is to show that we can find a sequence of strategies ν_n^* that approximates the seller’s optimal payoff with arbitrary precision.

LEMMA 14. Let $\varepsilon \leq \tilde{\varepsilon}^*$. Then there exists a sequence of strategies ν_n^* such that

$$\Pi_c(\mu_n, \nu_n^*) \rightarrow \Pi_c(\mu, \nu^*).$$

Proof. We are going to partition the stopping set $Q_c(\mu)$ into different regions corresponding to the different products. In other words, define $Q^i := Q_c(\mu) \cap H_i$ so all the Q^i together form a partition of $Q_c(\mu)$. For each product such that $\nu^*(Q^i) > 0$, let q^i be the ν^* -average belief in Q^i , that is,

$$q^i := \frac{1}{\nu^*(Q^i)} \int_{Q^i} q \nu^*(dq).$$

Note that $q^i \in Q^i$ since Q^i is convex. Since ν^* is a strategy, the prior p must be the ν^* -average of the beliefs q^i . Moreover, since each $\nu^*(Q)$ is strictly positive, p must be in the interior of the convex hull of all q^i .

Note that if every q^i is in every stopping set $Q_c(\mu_n)$, then a strategy with supports on q^i will give the seller exactly the optimal profit $\Pi_c(\mu, \nu^*)$, and we are done. Unfortunately, q^i may not be in $Q_c(\mu_n)$, which means that the limit profit from such a strategy may be less than $\Pi_c(\mu, \nu^*)$. To overcome this, we will construct a sequence of beliefs $q_n^i \in Q_c(\mu_n)$ such that $q_n^i \rightarrow q^i$ and the profits from a sequence of strategies with supports on q_n^i converge to $\Pi_c(\mu, \nu^*)$. We do this as follows. For each product i and signal μ_n , define $\lambda_n^i = 1$ if $q^i \in Q_n^i$ and

$$\lambda_n^i := \frac{c}{c + V_c(\mu_n, q^i) - q^i \cdot u_i}$$

otherwise. Note that in the case of the latter, $V_c(\mu_n, q^i) > q^i \cdot u_i$, so $\lambda_n^i < 1$. Finally, define $q_n^i := \lambda_n^i q^i + (1 - \lambda_n^i) 1_s$, where $1_s \in H_i$ as guaranteed by lemma 13. Since V_c is continuous by lemma 11,

$$V_c(\mu_n, q^i) \rightarrow V_c(\mu, q^i) \leq q^i \cdot u_i,$$

where the inequality follows from the fact that $q^i \in Q_c(\mu)$. By the definition of λ_n^i , this implies that $\lambda_n^i \rightarrow 1$ and $q_n^i \rightarrow q^i$. Since p is in the interior of the convex hull of the average beliefs q^i , for large enough n , p is also in the convex hull of the beliefs q_n^i . Hence, for each n , we can define a strategy ν_n^* with finite support on each q_n^i such that $\nu_n^*(\{q_n^i\}) \rightarrow \nu^*(Q^i)$.

Finally, since $V_c(\mu_n, \cdot)$ is convex,

$$\begin{aligned} V_c(\mu_n, q_n^i) &\leq \lambda_n^i V_c(\mu_n, q^i) + (1 - \lambda_n^i) V_c(\mu_n, 1_s) \\ &= \lambda_n^i V_c(\mu_n, q^i) + (1 - \lambda_n^i)(-c) \\ &= \lambda_n^i(q^i \cdot u_i) \leq q^i \cdot u_i. \end{aligned}$$

Hence, each $q_n^i \in Q_c(\mu_n)$, so

$$\Pi_c(\mu_n, \nu_n^*) = \sum_i \pi_i, \nu_n^*(\{q_n^i\}) \rightarrow \sum_i \pi_i \nu^*(Q^i) = \Pi_c(\mu, \nu^*)$$

as desired. QED

In this last step, we use the sequence of strategies ν_n^* that converges to the optimal strategy ν^* to show that the limit strategy ν is also optimal. This proves that φ_c has a closed graph.

LEMMA 15. If $\varepsilon \leq \varepsilon^*$, then φ_c has a closed graph.

Proof. Let $\mu_n \rightarrow \mu$, $\nu_n \in \varphi_c(\mu_n)$, and $\nu_n \rightarrow \nu$. We want to prove that $\nu \in \varphi_c(\mu)$. First, we show that $\nu(Q_c(\mu)) = 1$. Note that $\nu_n(Q_c(\mu_n)) = 1$ for all n , as otherwise the seller can profitably deviate by providing information at any belief that is not stopping. In order to prove that $\nu(Q_c(\mu)) = 1$, we will show that the stopping set mapping Q_c is upper hemicontinuous and takes on nonempty, closed values. We can then use theorem 17.13 of Aliprantis and Border (2006), which says that in this case, $\nu_n(Q_c(\mu_n)) = 1$ for all n implies that $\nu(Q_c(\mu)) = 1$. Note that the non-emptiness and closedness of Q_c are trivial (the latter follows from the continuity of the value function from lemma 11). For upper hemicontinuity, suppose $(\mu_m, q_m) \rightarrow (\mu, q)$, where $q_m \in Q_c(\mu_m)$. Thus, $q_m \cdot u(q_m) \geq V_c(\mu_m, q_m)$, so by continuity again, $q \cdot u(q) \geq V_c(\mu, q)$, implying $q \in Q_c(\mu)$. Hence, Q_c is upper hemicontinuous, so applying the theorem yields $\nu(Q_c(\mu)) = 1$.

Since $\nu(Q_c(\mu)) = 1$, the seller profits are $\Pi_c(\mu, \nu) = \int_{\Delta S} \pi(q) \nu(dq)$. Since π is upper semicontinuous, the expected profit of the seller is also upper semicontinuous with respect to strategies (see theorem 15.5 of Aliprantis and Border [2006]). Using ν_n^* as defined from lemma 14, we thus have

$$\begin{aligned} \Pi_c(\mu, \nu) &= \int_{\Delta S} \pi(q) \nu(dq) \geq \limsup_n \int_{\Delta S} \pi(q) \nu_n(dq) = \limsup_n \Pi(\mu_n, \nu_n) \\ &\geq \lim_n \Pi_c(\mu_n, \nu_n^*) = \Pi_c(\mu, \nu^*) \geq \Pi_c(\mu, \nu). \end{aligned}$$

In other words, since the payoffs from the optimal strategies ν_n must be higher than those of the approximate strategies ν_n^* and the profits of the latter converge to that of the optimal strategy ν^* , so must the profits from ν_n . Thus $\Pi_c(\mu, \nu) = \Pi_c(\mu, \nu^*)$, so $\nu \in \varphi_c(\mu)$ as desired. QED

Finally, we put everything together. Consider a sequence of $\varepsilon_n \rightarrow 0$ such that $\varepsilon_n < \varepsilon^*$ and let c_n be such that φ_{c_n} is nonempty as guaranteed by lemma 12. Combined with lemmas 9, 10, and 15, the preconditions of KFG (corollary 17.55 of Aliprantis and Border [2006]) are satisfied, so there exists an equilibrium $\mu_n \in M_{c_n}$. Moreover, from lemma 12, we know that $|V_{c_n}(\mu_n, q) - W(\bar{\mu}, q)| < \varepsilon_n \rightarrow 0$, so the buyer's value function converges to his full-information payoff in the limit. We have thus constructed a full-information limit equilibrium as desired. QED

3. Proof of Lemma 4

Suppose $\alpha > 0$ and $\beta > 0$. First, suppose $\beta = 1$, so there is no coordination. Thus, we can just work with payoff beliefs and associate a strategy σ with its posterior distribution $\mu = K_{p_0}^\sigma$ as in Sections B.1 and B.2. Consider a sequence of decreasing search costs $c_n \rightarrow 0$ and associated equilibria μ_n such that the buyer value functions $V_{c_n}(\mu_n, \cdot)$ converge to the limit function V^* . Since the set of simple strategies are compact, we can assume $\mu_n \rightarrow \mu$ for some simple strategy μ . Let \mathcal{E} be the partition induced by μ and $\mu_\mathcal{E}$ be the simple strategy corresponding to revealing \mathcal{E} . We will show that $V^*(p_0) = W(\mu_\mathcal{E}, p_0)$.

Let t_n be the largest integer less than $c_n^{-1/2}$, so the buyer's value function is at least his payoff from visiting t_n sellers, that is,

$$V_{c_n}(\mu_n, q) \geq -t_n c_n + W^{t_n}(\mu_n, q) \geq -c_n^{1/2} + W^{t_n}(\mu_n, q).$$

Since $t_n \rightarrow \infty$ as $c_n \rightarrow 0$, we can assume $t_n > T$ for some T without loss of generality. Thus,

$$V_{c_n}(\mu_n, q) \geq -c_n^{1/2} + W^T(\mu_n, q)$$

as t_n rounds of information is more informative than T rounds of information. Since $\mu_n \rightarrow \mu$ and W^T is continuous, by taking limits, we have $V^*(q) \geq W^T(\mu, q)$. Since this is true for all T , we can take $T \rightarrow \infty$, so $W^T(\mu, q) \rightarrow W(\mu_\mathcal{E}, q)$ by lemma 7. Thus, $V^*(q) \geq W(\mu_\mathcal{E}, q)$. To see why $V^*(p_0)$ is weakly less than $W(\mu_\mathcal{E}, p_0)$, note that since μ_n is an equilibrium, by lemma 5,

$$V_{c_n}(\mu_n, p_0) = -c_n + \int_{\Delta S} q \cdot u(q) \mu_n(dq) = -c_n + W(\mu_n, p_0).$$

Taking limits, we have $V^*(p_0) = W(\mu, p_0) \leq W(\mu_\mathcal{E}, p_0)$, where the inequality follows from the fact that $\mu_\mathcal{E}$ is more informative than μ . This proves that $V^*(p_0) = W(\mu_\mathcal{E}, p_0)$, so all limit equilibria are partitional.

Now, suppose $\beta < 1$. Consider a sequence of equilibria $\sigma_n = (\sigma_n^g, \sigma_n^i)$ and let μ_n be the corresponding posterior distributions on ΔS for σ_n^i . As above, let μ be the limit of μ_n and \mathcal{E} be its induced partition. We will now show that $V^*(p_0) \geq W(\mu_\mathcal{E}, p_0)$. Let t_n be the largest integer less than $c_n^{-1/2}$, and for any $s \leq t_n$, let $P^s(\kappa)$ denote the probability that the buyer receives at most κ draws from the independent policy σ_n^i out of a total of t_n total draws. Hence, the buyer's value function is at least his independent payoffs from visiting t_n sellers, that is,

$$\begin{aligned} V_{c_n}(\sigma_n, \mathbf{q}) &\geq -t_n c_n + P^{t_n}(\kappa) \cdot 0 + [1 - P^{t_n}(\kappa)] W^{\kappa}(\mu_n, q) \\ &\geq -c_n^{1/2} + [1 - P^{t_n}(\kappa)] W^{\kappa}(\mu_n, q). \end{aligned}$$

Recall that β is the mean of this binomial distribution and suppose $\kappa \leq \beta t_n$. By Hoeffding's inequality,

$$P^{t_n}(\kappa) \leq e^{-2(t_n \beta - \kappa)^2 / t_n}$$

and

$$V_{c_n}(\sigma_n, \mathbf{q}) \geq -c_n^{1/2} + \left[1 - e^{-2(t_n \beta - \kappa)^2 / t_n} \right] W^{\kappa}(\mu_n, q).$$

If we let $\kappa_n = \lambda t_n$ for $\lambda < \beta$, then

$$V_{c_n}(\sigma_n, \mathbf{q}) \geq -c_n^{1/2} + \left[1 - e^{-2(\beta-\lambda)^2 t_n}\right] W^{\lambda t_n}(\mu_n, q).$$

Since $t_n \rightarrow \infty$ as $c_n \rightarrow 0$, we can assume $t_n > T$ for some T without loss of generality. Thus,

$$V_{c_n}(\sigma_n, \mathbf{q}) \geq -c_n^{1/2} + \left[1 - e^{-2(\beta-\lambda)^2 t_n}\right] W^T(\mu_n, q)$$

as t_n rounds of information is more informative than T rounds of information. Since $\mu_n \rightarrow \mu$ and W^T is continuous, taking limits we have $V^*(\mathbf{q}) \geq W^T(\mu, q)$. Since this is true for all T , we can take $T \rightarrow \infty$, so $W^T(\mu, q) \rightarrow W(\mu_\varepsilon, q)$ by lemma 7. Thus, $V^*(\mathbf{q}) \geq W(\mu_\varepsilon, q)$ and in particular $V^*(\mathbf{p}_0) \geq W(\mu_\varepsilon, p_0)$ as desired. QED

4. Proof of Theorem 4

Suppose $\alpha > 0$ and $\beta > 0$. Let \mathcal{E} be the partition induced by the limit μ of the independent policies μ_n as in the proof of lemma 4 in Section B.3. We will prove that all events in \mathcal{E} are trivial by contradiction. Suppose otherwise and let $\mathcal{F} \subset \mathcal{E}$ denote the set of nontrivial events in the partition. By the premise, for every $E \in \mathcal{F}$, p_E is fully improvable, that is, $\pi(p_E) < \bar{\Pi}(p_E)$, where $\bar{\Pi}$ is the full-information profit. Define the strategy ν that is essentially the same as the partitional strategy μ_ε but provides another round of information according to full information at p_E for every $E \in \mathcal{F}$. As a result,

$$\int_{\Delta S} \pi(q) \mu_\varepsilon(dq) = \sum_{E \in \mathcal{E}} \pi(p_E) p(E) < \int_{\Delta S} \pi(q) \nu(dq). \tag{A5}$$

Now, since ν has support only on either trivial or degenerate beliefs, $\nu(\mathbf{Q}_{c_n}(\sigma_n)) = 1$. Given that σ_n is an equilibrium, μ_n must be optimal given σ_n , so

$$\int_{\Delta S} \pi(q) \mu(dq) \geq \limsup_n \int_{\Delta S} \pi(q) \mu_n(dq) \geq \int_{\Delta S} \pi(q) \nu(dq). \tag{A6}$$

We now show that the monopoly profits from μ and μ_ε are the same. First, note that by lemma 5, we have that $\mu_n(\mathbf{Q}_{c_n}(\sigma_n)) = 1$ for all n . Since $V_{c_n}(\sigma_n, \cdot) \rightarrow V^*$ and $\mu_n \rightarrow \mu$, we can use theorem 17.13 of Aliprantis and Border (2006) as in the proof of lemma 15 to obtain that

$$\mu\{q \in \Delta S | q \cdot u(q) \geq V^*(q)\} = 1.$$

Moreover, since $V^*(\mathbf{q}) \geq W(\mu_\varepsilon, q)$ from the proof of lemma 4 above, this implies that

$$\mu\{q \in \Delta S | q \cdot u(q) \geq W(\mu_\varepsilon, q)\} = 1,$$

so $q \cdot u(q) = W(\mu_\varepsilon, q)$ μ -a.s. Thus, if we let J_q^ε denote the posterior distribution corresponding to revealing \mathcal{E} to a buyer with belief q , then we have μ -a.s.

$$q \cdot u(q) = \int_{\Delta S} r \cdot u(r) J_q^\varepsilon(dr) = \sum_i \int_{H_i} r \cdot u_i J_q^\varepsilon(dr) = \sum_i J_q^\varepsilon(H_i) q_i \cdot u_i,$$

where

$$q_i := \frac{1}{J(q, H_i)} \int_{H_i} \eta J_q^\varepsilon(dr)$$

for $J_q^\varepsilon(H_i) > 0$. If we let $q \in H_i$, then this implies that $J_q^\varepsilon(H_i) = 1$. Thus,

$$\int_{\Delta S} \pi(q) \mu_\varepsilon(dq) = \int_{\Delta S} \left[\int_{\Delta S} \pi(r) J_q^\varepsilon(dr) \right] \mu(dq) = \int_{\Delta S} \pi(q) \mu(dq).$$

However, this contradicts inequalities (A5) and (A6) above. Hence, all events in \mathcal{E} are trivial, so by lemma 4, all limit equilibria must be full-information as desired. QED

References

- Aliprantis, C., and K. Border. 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. 3rd ed. Berlin: Springer.
- Alonso, R., and O. Camara. 2016. "Bayesian Persuasion with Heterogeneous Priors." *J. Econ. Theory* 165:672–706.
- Anderson, S. P., and R. Renault. 2006. "Advertising Content." *A.E.R.* 96 (1): 93–113.
- Au, P. H. 2015. "Dynamic Information Disclosure." *RAND J. Econ.* 46 (4): 791–823.
- Aumann, R. J., and M. Maschler. 1995. *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press.
- Battaglini, M. 2002. "Multiple Referrals and Multidimensional Cheap Talk." *Econometrica* 70 (4): 1379–1401.
- Best, J., and D. Quigley. 2016. "Persuasion for the Long Run." Working paper, Oxford Univ.
- Bhattacharya, S., and A. Mukherjee. 2013. "Strategic Information Revelation When Experts Compete to Influence." *RAND J. Econ.* 44 (3): 522–44.
- Board, S., and M. Pycia. 2014. "Outside Options and the Failure of the Coase Conjecture." *A.E.R.* 104 (2): 656–71.
- Burdett, K., and K. L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–69.
- Clemens, J., and J. Gottlieb. 2014. "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *A.E.R.* 104 (4): 1320–49.
- Crawford, V. P., and J. Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.
- Diamond, P. A. 1971. "A Model of Price Adjustment." *J. Econ. Theory* 3 (2): 156–68.
- Ely, J. C. 2017. "Beeps." *A.E.R.* 107 (1): 31–53.
- Forand, J. G. 2013. "Competing through Information Provision." *Internat. J. Indust. Org.* 31 (5): 438–51.
- Gentzkow, M., and E. Kamenica. 2017a. "Bayesian Persuasion with Multiple Senders and Rich Signal Spaces." *Games and Econ. Behavior* 104:411–29.
- . 2017b. "Competition in Persuasion." *Rev. Econ. Studies* 84 (1): 300–322.
- . 2017c. "Disclosure of Endogenous Information." *Econ. Theory Bull.* 5 (1): 47–56.
- Green, J., and N. Stokey. 1978. "Two Representations of Information Structures and Their Comparisons." Discussion paper, Stanford Univ.
- Hoffmann, F., R. Inderst, and M. Ottaviani. 2014. "Persuasion through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation." Discussion paper, Univ. Bocconi.

- Inderst, R., and M. Ottaviani. 2012. "Competition through Commissions and Kickbacks." *A.E.R.* 102 (2): 780–809.
- Johnson, J., and D. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *A.E.R.* 96 (3): 756–84.
- Kamenica, E. 2008. "Contextual Inference in Markets: On the Informational Content of Product Lines." *A.E.R.* 98 (5): 2127–49.
- Kamenica, E., and M. Gentzkow. 2011. "Bayesian Persuasion." *A.E.R.* 101 (6): 2590–2615.
- Lewis, T. R., and D. E. Sappington. 1994. "Supplying Information to Facilitate Price Discrimination." *Internat. Econ. Rev.* 35 (2): 309–27.
- Li, F., and P. Norman. 2018a. "On Bayesian Persuasion with Multiple Senders." *Econ. Letters* 170 (September): 66–70.
- . 2018b. "Sequential Persuasion." Manuscript, Univ. North Carolina.
- Milgrom, P., and J. Roberts. 1986. "Relying on the Information of Interested Parties." *RAND J. Econ.* 17 (1): 18–32.
- Moscarini, G., and L. Smith. 2001. "The Optimal Level of Experimentation." *Econometrica* 69:1629–44.
- Mullainathan, S., M. Noeth, and A. Schoar. 2012. "The Market for Advice: An Audit Study." Working Paper no. 17929, NBER, Cambridge, MA.
- Rayo, L., and I. Segal. 2010. "Optimal Information Disclosure." *J.P.E.* 118 (5): 949–87.
- Rhodes, A. 2015. "Multiproduct Retailing." *Rev. Econ. Studies* 82 (1): 360–90.
- Stigler, G. J. 1961. "The Economics of Information." *J.P.E.* 69 (3): 213–25.
- Wald, A. 1947. *Sequential Analysis*. New York: Wiley.
- Wolinsky, A. 1986. "True Monopolistic Competition as a Result of Imperfect Information." *Q.J.E.* 101 (3): 493–511.
- Zhou, J. 2014. "Multiproduct Search and the Joint Search Effect." *A.E.R.* 104 (9): 2918–39.