Competitive Information Disclosure in Search Markets*

Simon Board† and Jay Lu‡

August 15, 2015

Abstract

Buyers often search across multiple retailers or websites to learn which product best fits their needs. We study how sellers manage these search incentives through their disclosure policies (e.g. advertisements, product trials and reviews), and ask how competition affects information provision. If sellers can observe the beliefs of buyers (e.g. they can track buyers via their cookies) then, in a broad range of environments, there is a unique equilibrium in which sellers provide the "monopoly level" of information. However, if buyers are anonymous, then there is an equilibrium in which sellers provide full information as search costs vanish. Tracking software thus enables sellers to implicitly collude, providing a motivation for regulation.

1 Introduction

The internet has led to a proliferation of products, and a proliferation of sellers providing any given product. For example, eBay has around 5 billion listings from 2 million sellers in 50,000 product categories. Sophisticated consumers search across sellers to learn the properties of different products to benefit from this ever-expanding variety. Starting with Stigler (1961), there has been a large literature examining the incentives of buyers to search for better prices. This paper studies the incentives for buyers to search for better information, and asks how sellers manage these incentives through their disclosure policies.

Sellers have a lot of flexibility in managing buyers' information through advertisements, product trials and advice. For example, when selling books online, a retailer chooses which products to suggest, what information to provide in the descriptions, whether to let the buyer "look inside",

*We are grateful to John Asker, Alessandro Bonatti, Moritz Meyer-ter-Vehn, Marek Pycia and Asher Wolinsky for comments, as well as seminar audiences at Berkeley, ESSET (Gerzensee), SED (Warsaw), SITE, UNSW, UT Austin, UT Sydney and U Queensland. Keywords: Information Disclosure, Advertising, Learning, Search, Targeting, Cookies. JEL: D40, D83, L13.

†Department of Economics, UCLA. http://www.econ.ucla.edu/sboard

‡Department of Economics, UCLA. http://www.econ.ucla.edu/jay

and how to present other customers' reviews. Such information can be increasingly tailored to buyers' characteristics and their shopping behavior at previous sellers. For example, Facebook's "Like" and Twitter's "Tweet" buttons carry code enabling the companies to track users' movements across the internet. Similarly, a slew of third-party tracking companies follow users from site to site, monitoring popular websites and piggybacking on data collected by other data brokers.[1]

In this paper we characterize a general model of sequential search for information in which the seller is free to choose any disclosure policy. We then ask whether competition forces sellers to provide all pertinent information, and show that the answer depends on the sellers' information about buyers. When sellers can observe buyers' beliefs (e.g. they can track buyers via their cookies) then, in a broad range of environments, there is a unique equilibrium in which all sellers choose the monopoly disclosure policy. That is, sellers manipulate their customers to purchase the most profitable, rather than the most suitable product. In contrast, when buyers are anonymous, there is an equilibrium in which sellers fully reveal all their information as search costs vanish.

We thus show that tracking software allows sellers to implicitly collude. Intuitively, tracking enables sellers to discriminate between new and old buyers. Since old buyers are already informed, it is optimal for a seller to provide them with less information; this lowers their outside option and undermines competition. This finding is important since, over the last decade, internet giants such as Google and Facebook and a plethora of data brokers have been collecting ever more information about their users, giving rise to concerns that such information makes consumers somehow "exploitable".[2] We provide a theoretical foundation for such claims and show that neither competition nor individual privacy measures – such as deleting one's cookies – can correct such distortions, providing a rationale for regulating tracking programs and the data sharing agreements that underlie them.

Beyond online retailing, our model is applicable to a wide variety of situations where a sequence of "sellers" try to convince an uninformed "buyer" to choose them. When a company wishes to file for an IPO, it will typically approach banks sequentially, obtaining advice about whether to file and how to price the stock, with the winning bank obtaining a 7% fee. When a consumer wishes to find information on the internet, they receive recommendations from one search engine and, if not satisfied, try another. And, when a Government wishes to find an expert to design an auction, it will enter discussions with a number of academics, choosing one when it is sufficiently certain of the right design. This paper examines the effectiveness of competition in such environments, and thus has implications for antitrust and the regulation of advisory industries (e.g. mortgage lending).

[1] To illustrate, *The Economist* reports that the 100 most widely used websites are monitored by more than 1,300 firms. See their "Special Report: Advertising and Technology" (13th September 2014).

[2] For example, see "Facebook Tries to Explain Its Privacy Settings but Advertising Still Rules," *New York Times* (13th November 2014) or "Google to Pay $17 Million to Settle Privacy Case," *New York Times* (18th November 2013).

In the model, we consider a prospective buyer ("he") who samples from sellers in order to learn which product best suits his needs. Each seller ("she") sells the same set of products and, when sampled, chooses how much information to disclose in order to encourage the buyer to purchase. A product is defined by the utility $u(s)$ to the agent in state s, and the profit $\tilde{\pi}$ to the seller. We suppose that the seller does not know the state (e.g. whether the buyer likes a book when they "look inside"), and model the seller's information disclosure as a distribution of signal realizations that may increase or decrease the buyer's assessment of the product, as in Kamenica and Gentzkow (2011). After updating his belief, the buyer chooses to buy a product, to exit, or to pay a search cost c and randomly sample another seller. We study how much information sellers disclose and how this depends on sellers' information about buyers.

In Section 2, we suppose that buyers' beliefs are public; for example, the seller may learn a buyer's beliefs via cookies on his computer. First, we show that the monopoly disclosure policy is always an equilibrium. Intuitively, if all sellers use the same disclosure policy no buyer has an incentive to continue to search and, since the policy maximizes profits, no seller has the incentive to defect. This monopoly policy may be very undesirable for customers: if there is a single product and the buyer is skeptical, preferring not to buy in the absence of information, then the monopoly policy leaves the buyer with zero utility.

Next, we derive conditions under which the monopoly policy is the unique equilibrium. This is the case if there is a single product. Intuitively, the search cost allows one seller to provide a little less information than the market, iteratively pushing the equilibrium towards the monopoly outcome. However, in other examples, e.g. two vertically differentiated products, non-monopolistic equilibria may also exist. Intuitively, a seller would like to provide a lot less information than the market, but the search cost only allows for local deviations which may not exist. We then show that if products are "dispersed" relative to the underlying state space, meaning that a buyer wishes to purchase at most one product at each belief, then these local deviations are enough to prove that the monopoly policy is the unique equilibrium.

In Section 3, we suppose that buyers' beliefs are private. That is, a seller knows a buyer's initial prior but does not know his current belief nor his past search behavior; the seller then chooses a signal that is independent of past signals. In equilibrium, the buyer purchases (or exits) in the first period. However, the possibility that the buyer can continue to search and mimic a new buyer forces sellers to provide more information than when beliefs are public. First, we show that as the search cost vanishes, there is a sequence of equilibria that converges to full information. Intuitively, if all other sellers provide (almost) full information and search costs are small, then any one seller must match this and also provide (almost) full information. Formally, when the search cost is small, the best-response correspondence maps the set of signals close to full information into itself. We then show this best-response correspondence has a fixed point to prove an equilibrium exists, and let the search cost vanish.

Next, we derive conditions under which full information is the unique limit equilibrium. As search costs become small, any equilibrium has the property that beliefs are "partitional", meaning that buyers either learn everything or nothing about which of two states occurs. If products are "sufficiently dispersed" in that a buyer prefers not to buy unless they receive a little information about each state, then full information is the only limit equilibrium.

Comparing the private and public cases, we find that anonymity is a powerful force to ensure sellers compete against one another. When beliefs are public, they discriminate between new and old buyers, lowering buyers' outside options, and making them less likely to search. When beliefs are private, sellers cannot discriminate and the option of going elsewhere forces sellers to provide all the pertinent information in the limit.

In Section 4, we consider two extensions. First, we ask what happens when buyers can choose to avoid being tracked at a small cost. For example, they can activate the "Do Not Track" function in their browser or delete their cookies. The analysis of the previous two sections implies that becoming anonymous exerts a positive externality on other buyers, inducing sellers to provide more information to everyone. Moreover, the temptation to free-ride is so strong that no matter how small the cost of becoming anonymous, there is no equilibrium where all buyers voluntarily do so. Second, we consider a case in between public and private beliefs where a seller can observe whether a buyer approached previous sellers, but cannot observe the outcome of the signals. As in the public beliefs case, a seller is able to discriminate between new and old buyers. Consequently, the monopoly policy is always an equilibrium and, in a single-product example, it is the only equilibrium.

1.1 Literature

There is a large literature in which sellers post prices and buyers pay a fixed cost for each search. In the benchmark model, Diamond (1971) shows that if all sellers sell a single identical good then, in equilibrium, all sellers charge the monopoly price. Intuitively, the search cost generates a holdup problem, allowing any one seller to raise her price slightly above the price set by others.[3] In our paper, sellers choose disclosure policies rather than prices. The logic in the case of public beliefs is analogous to Diamond, while the case of private beliefs contrasts sharply. In comparison to the Diamond model, when a seller provides information to a visiting buyer, this changes his beliefs and changes how he acts at subsequent sellers.

There is a growing literature on information disclosure with competition, based on the Kamenica-Gentzkow framework.[4] In a general model, Gentzkow and Kamenica (2012) consider two senders

[3]This result continues to hold if sellers make multiple offers to each buyer (Board and Pycia (2014)). However, it breaks down if sellers sell heterogeneous products (Wolinsky (1986)), multiple products (Zhou (2014); Rhodes (2014)), or if buyers pre-commit to a number of searches (Burdett and Judd (1983)).

[4]See also Rayo and Segal (2010) and Aumann and Maschler (1995).

who simultaneously choose a disclosure policy, assuming that senders' signals are "coordinated" in that a common random variable determines how states are mapped into signals. They show how each sender takes the information released by the other as given, choosing the optimal monopoly policy on the residual, implying that two senders always provide more information than one. Li and Norman (2014) reconsider this model when signals are independent, and show that adding a second sender may reduce the amount of information provided as the first sender tries to protect herself from the new information provided by the second. Finally, Hoffmann, Inderst, and Ottaviani (2014) consider a model with more structure, assuming heterogeneous sellers compete to sell to a single buyer, and that information revelation is a binary decision. When the seller's disclosure strategy is observed, an increase in competition increases the amount of information provided. Intuitively, in order to make a sale, the buyer's utility at a seller has to exceed a threshold; this threshold increases in the level of competition, implying that the seller wishes to raise the variance of the buyer's posterior. In contrast to these papers, we consider sequential competition, giving rise to the public vs. private dichotomy that is the centerpiece of our analysis.[5,6]

As in much of the communication literature, starting with Crawford and Sobel (1982), we assume payoffs to the parties are exogenous, rather than having the seller choose the prices or qualities. With the buyer-seller application, we think of this as markets where prices are not flexible (e.g. IPO fees, books with RPM), or where the seller chooses aggregate prices when faced with heterogeneous buyers, and then individually targets product recommendations. With other applications, prices may be out of the control of the "seller" (e.g. an intermediary like a search engine), or there may be no transfers (e.g. a government seeking advice). Anderson and Renault (2006) study a monopolist selling a single good who can choose information and prices, and 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. With more goods, such as two vertically differentiated goods, matters become much more complicated, and a monopolist must trade off the efficiency of the allocation and the rents acquired by the buyer.[7]

[5]There are other related papers. Gentzkow and Kamenica (2015) considers a very general model and ask when competition yields more information than collusion. Forand (2013) studies a model of information revelation with directed search.

[6]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"). For example, Milgrom and Roberts (1986) finds conditions under which competition leads to full revelation. More recently, Bhattacharya and Mukherjee (2013) consider a model where senders are stochastically informed, characterizing the receiver's preferences over the biases of experts. A second literature supposes information is unverifiable ("cheap talk"). Here, Krishna and Morgan (2001) consider sequential communication, while Battaglini (2002) analyzes a simultaneous game, both finding conditions under which full revelation is possible when the senders have opposed preferences.

[7]The seller can again achieve first-best if she can either make the price contingent on the state, or can use an up-front fee (e.g. Eső and Szentes (2007)).

2 Public Beliefs

We first consider the setting where buyers' beliefs are public. For example, sellers have cookies on a buyer's computer that track which signals he sees and his reaction to these signals.

We first describe the model and consider a simple single-product example. Our first main result is that the monopoly policy is always an equilibrium. In general, there may be multiple equilibria and we characterize these in terms of local deviations. We then show that if products are "dispersed", then the monopoly policy is the unique equilibrium.

2.1 Model

Basics. There is one (male) buyer and infinite (female) sellers who sell the same products. A buyer is uncertain about a finite set of states relevant for his decision that we call S. Let ΔS be the set of beliefs about S. The buyer starts with an initial prior that, without loss, places positive probability on every state.

Actions. When a buyer approaches a seller, she observes the buyer's prior $p \in \Delta S$, which may incorporate information from other sellers, and chooses an information disclosure policy that we describe below. After the buyer sees the seller's signal, he updates his belief to form a posterior $q \in \Delta S$ and chooses whether to (1) buy a product from the seller, (2) exit the market or (3) pay a search cost $c > 0$ and pick a new seller at random. In the latter case, the buyer arrives at the new seller, the seller observes the buyer's new prior, and the game proceeds as above.

Buyer's strategy. Denote $U \subset \mathbb{R}^S$ as a finite set of available products[8] including the exit option $0 \in U$. This means that product $u \in U$ delivers utility $u(s)$ in state $s \in S$ and expected utility $q \cdot u$ to the buyer under belief $q \in \Delta S$. Define the *acceptance set* Q as the set of beliefs for which the buyer decides to stop searching. If the buyer decides to accept at posterior q, let $u^*(q) \subset U$ denote the set of optimal choices.

Seller's strategy. A seller observes a buyer's prior p and chooses a distribution of posteriors such that the average posterior equals the prior. Formally, a *disclosure policy* K is a Markov kernel satisfying $\int_{\Delta S} q \; K(p, dq) = p$ for each prior p.

Let $\tilde{\pi}(u)$ be the seller's payoff given the buyer's choice $u \in U$. To avoid ties, we assume each choice u yields different profits, i.e. $\tilde{\pi}(u) \neq \tilde{\pi}(u')$ for $u \neq u'$, and normalize $\tilde{\pi}(0) = 0$. Assuming that ties in the buyer's actions are resolved in the seller's favor, the seller's payoff given the buyer's belief $q \in \Delta S$ is $\pi(q) := \max_{u \in u^*(q)} \tilde{\pi}(u)$. A leaving buyer never returns, so sending a buyer outside

[8] One could also interpret each $u \in U$ as a bundle of two products, three products, etc..

of his acceptance set Q is never optimal for the seller. Thus, a disclosure policy K is *optimal given* Q iff for all possible disclosure policies L

$$\int_Q \pi(q)\, K(p, dq) \geq \int_Q \pi(q)\, L(p, dq)\,. \tag{1}$$

Moreover, to break ties, assume that if it is weakly optimal for the seller to reveal no information, then she will do so. We call this the *signal tie-breaking rule.* This simplifies the analysis and has the interpretation that providing some information is a little costly, in a lexicographic sense.

The *full information policy* is a disclosure policy such that every posterior is degenerate on some state, $q = \delta_s$. A *monopoly policy* is a disclosure policy that is optimal given $Q = \Delta S$. The monopoly policy corresponds to the static optimal information disclosure policy in Kamenica and Gentzkow (2011).

Symmetric equilibrium. If all sellers use policy K, then a buyer with posterior q who continues to search has continuation utility

$$V_c(K, q) := -c + \int_{\Delta S} \max\left\{\max_{u \in U} r \cdot u,\ V_c(K, r)\right\} K(q, dr)$$

A disclosure policy and acceptance set (K, Q) form an *equilibrium* iff K is optimal given

$$Q_c(K) = \left\{ q \in \Delta S \;\middle|\; \max_{u \in U} q \cdot u \geq V_c(K, q) \right\}$$

In this case, we call K an *equilibrium policy.*

Remarks. We model information disclosure in the style of Kamenica and Gentzkow (2011). It is most natural to think of the buyer and seller sharing uncertainty about the buyer's preference. For example, when a buyer contemplates purchasing a book, Amazon can let the buyer "look inside", have temporary access, or try a basic version for free, but does not know how the buyer will interpret this information. An alternative interpretation is that the seller has private information about the buyer's taste, but commits to reveal information "truthfully" and does not try to copy another type of seller. For example, Amazon may have a policy of providing the most recent book reviews, even if it knows these reviews are negative.

We assume there are a large number of sellers, which ensures a buyer never returns to a seller. The analysis is the same if there are a finite number of sellers. The one-shot deviation principle means that the acceptance set Q is unaffected by any deviation. Moreover, when a buyer leaves a seller they return with probability less than one, so it is never optimal to induce a posterior outside the acceptance set, Q. Each seller thus chooses posteriors to maximize profits subject to $q \in Q$,

Figure 1: Single Product (Public Beliefs)

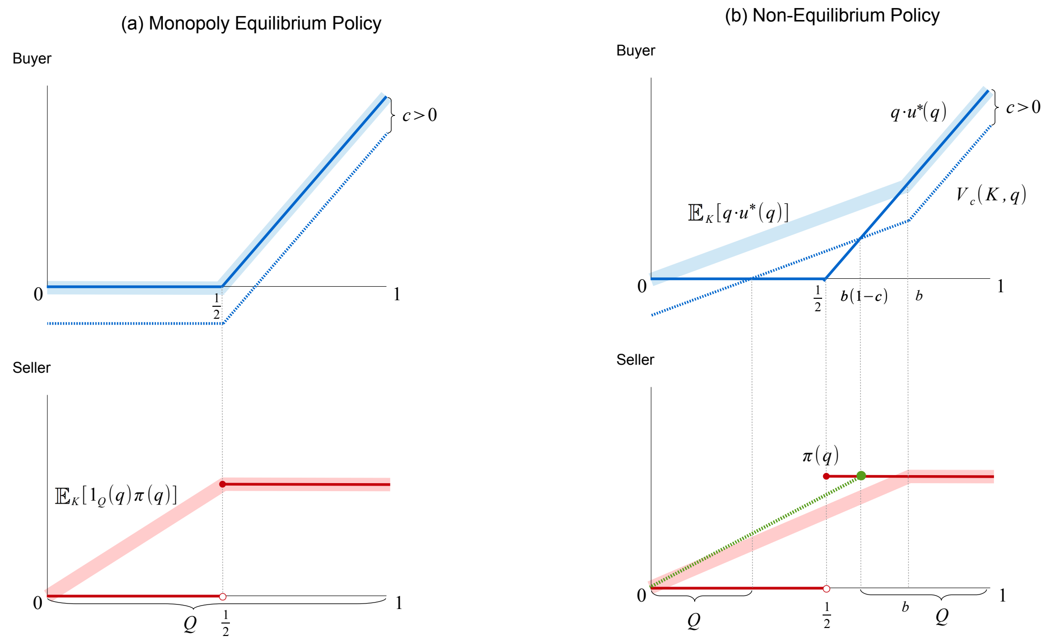

This figure illustrates Example 1 with public beliefs. The top panels show the buyer's utility: the bold line is the posterior utility function $q \cdot u^*(q)$, the shaded line is the expected utility from the policy $E_K[q \cdot u^*(q)]$, and the dotted line is the value of the policy $V_c(K,q) = E_K[q \cdot u^*(q)] - c$. The bottom panels show the seller's profit: the bold line is the posterior profit function $\pi^*(q)$ and the shaded line is the expected profit $\mathbb{E}_K[\mathbf{1}_Q \pi]$. **Figure 1(a)** demonstrates that the monopoly policy is an equilibrium. Since the dashed line always lies below the buyer's utility function, he always accepts. Given this, the seller has no incentive to defect. **Figure 1(b)** shows a more informative policy cannot be an equilibrium. Such a policy induces acceptance set Q, and the seller can raise his expected profits from the shaded line to the dotted line by providing a little less information, sending the buyer to $b(1-c)$ instead of b.

while Q is given by the buyer's expected continuation value, as in the infinite seller case.[9]

Finally, we make three assumptions about the nature of equilibria. First, we assume the seller chooses a Markovian disclosure policy that depends on the buyer's prior belief at that seller, but not on the order in which the buyer approaches sellers. Second, we assume that if "no information" is a best-response for the seller, then she chooses it. Third, we focus on symmetric equilibria, which is without loss if the seller's optimal policy is unique. In Appendix C1, we show that for our leading example, none of these assumptions are necessary. Indeed, any rationalizable equilibrium induces monopoly payoffs.

2.2 Motivating Example

Example 1 (Single product). A seller has a single product u to sell, there are two states $\{L, H\}$, and the buyer wishes to buy the product if state H is more likely. For example, a consumer considers buying a game theory book, but does not know whether he will find it useful. Suppose that $(u(L), u(H)) = (-1, 1)$, so the buyer prefers to buy if $p = \Pr(H) \geq \frac{1}{2}$, and that a sale generates profits $\tilde{\pi}(u) = 1$ for the seller. The resulting monopoly policy provides just enough information to persuade the buyer to buy,

$$K(p) = \begin{cases} (1-2p)\,\delta_{\{0\}} + 2p\;\delta_{\{\frac{1}{2}\}} & \text{if } p \in \left[0, \frac{1}{2}\right) \\ p & \text{if } p \in \left[\frac{1}{2}, 1\right] \end{cases}$$

as illustrated in Figure 1(a). That is, if the buyer starts with a prior $p \geq \frac{1}{2}$, the seller provides no information and the buyer buys; if $p < \frac{1}{2}$, the seller sends the buyer to posteriors 0 and $\frac{1}{2}$. For shorthand, we will sometimes write this "perfect bad news" policy as $p \to \{0, \frac{1}{2}\}$.

The monopoly policy is an equilibrium of the game. If all sellers choose such a policy a buyer learns nothing from a second seller after leaving a first and, given the search cost, buys immediately. Since the buyer purchases and the seller makes monopoly profits, she has no incentive to deviate.

More surprising, the monopoly policy is the the unique equilibrium of the game. To gain some intuition, fix $b > \frac{1}{2}$, and suppose that the seller provides a perfect bad news signal $p \to \{0, b\}$, as illustrated in Figure 1(b). This induces an acceptance set $Q = [0, \frac{b}{2b-1}c] \cup [b(1-c), 1]$, as shown in the lower panel. Now, consider a seller who deviates and uses the policy $p \to \{0, b(1-c)\}$.[10] 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 policy, there is a profitable deviation. $\triangle$

2.3 Monopoly is an Equilibrium

We now return to the general model, with multiple states and multiple products. We first characterize a seller's optimal disclosure policy. From Kamenica and Gentzkow (2011) we know that the seller's optimal profits coincides with the convex hull of the profit function, $\pi(q)\,\mathbf{1}_Q(q)$. We will be particularly interested in the states where the seller does not release any further information. Given

[9] This assumes that a seller can choose a different policy each time a buyer approaches her. In contrast, if a seller commits to her policy at the start of the game, a deviation in the policy affects the acceptance set Q. Using Lemma 1, we can restrict sellers to policies where a returning buyer receives no more information, meaning the buyer must keep searching for a different seller. Thus, say, with two sellers, the search cost is essentially double that in the infinite-seller model.

[10] This assumes that $b(1-c) > \frac{1}{2}$, so the buyer does not always accept. Otherwise, a seller should deviate to the monopoly policy.

any disclosure policy K, define the set of *absorbing beliefs* as

$$A_K := \{p \in \Delta S \mid K(p) = \delta_p\}.$$

The following lemma establishes that sellers will only send a buyer to absorbing posteriors; that is, if the buyer were to continue searching, he would get no further information. Once matched, buyers therefore never search again. Define the support of K, *supp(K)*, to be the largest set for which every neighborhood of every point in the set has strictly positive $K(p)$-measure for some $p \in \Delta S$.

Lemma 1. *If K is optimal given Q, then $supp(K) = cl(A_K)$. Consequently,*

$$V_c(K,q) = -c + \int_{\Delta S} \left(\max_{u \in U} r \cdot u\right) K(q, dr)$$

Proof. See Lemma A1.2(c) and Lemma A2.1 in Appendix A. □

Lemma 1 says that the support of an optimal disclosure policy coincides with A_K. Intuitively, if a seller sends the buyer to posteriors outside of A_K then a subsequent seller would provide further information to the buyer. Given the tie-break rule, this must be strictly profitable for the second seller, so the first seller should mimic her and provide all the information at once. Since buyers only get information one time, Lemma 1 implies that their value function is the value from a single search.

We now turn to equilibrium. We first establish that any equilibrium policy at search cost c is also an equilibrium policy at any search cost less than c. Hence the set of equilibria increases as costs decrease.

Lemma 2. *Any equilibrium policy for $c > 0$ is also an equilibrium policy for all $c' \leq c$.*

Proof. Let (K, Q) be an equilibrium for some $c > 0$, and consider $c' \in (0, c]$. First, observe that the buyer will accept at any absorbing point, $A_K \subset Q_{c'}(K)$, and by Lemma 1, K sends the buyer to absorbing points, $K(p, A_K) = 1$. Hence a seller's profits from policy K is unaffected by the reduction in cost. Second, the reduction in cost means that a buyer's continuation value $V_c(K,q)$ increases pointwise, and the acceptance set shrinks, i.e. $Q_{c'} \subset Q_c$. Since the seller only chooses posteriors in $Q_{c'}$, the scope for deviations is smaller than before. We now put this together: Since K is optimal given Q_c, for any $L \in \mathcal{K}$

$$\int_{Q_{c'}} \pi(q) K(p, dq) = \int_{Q_c} \pi(q) K(p, dq) \geq \int_{Q_c} \pi(q) L(p, dq) \geq \int_{Q_{c'}} \pi(q) L(p, dq)$$

and K is optimal given $Q_{c'}$.[11] Hence, $(K, Q_{c'})$ is an equilibrium. □

[11] We also need to check that the tie-breaking rule holds. Suppose that $K(p)$ yields the same profit on $Q_{c'}$ as some

Intuitively, a decrease in costs does not change the profit from the previous equilibrium policy K, but does make it harder to deviate. Hence the policy remains an equilibrium.

We can now show that the monopoly policy is always an equilibrium; this also implies existence of an equilibrium.

Theorem 1. *The monopoly policy is an equilibrium policy for all $c > 0$.*

Proof. Let K be the monopoly policy and find some $\bar{c}$ large enough such that

$$\left\{ q \in \Delta S \;\middle|\; \max_{u \in U} q \cdot u \geq -\bar{c} + \int_{\Delta S} \left(\max_{u \in U} r \cdot u \right) K(q, dr) \right\} = \Delta S$$

Then K is an equilibrium policy for any $c \geq \bar{c}$. Conversely, if $c < \bar{c}$, then Lemma 2 implies K is also an equilibrium policy for c. Hence, the monopoly policy is an equilibrium policy for any $c > 0$. $\square$

Intuitively, if all sellers choose the monopoly policy, a buyer will purchase at the first seller; since all sellers are making maximal profits, none has an incentive to defect.

2.4 Uniqueness of the Monopoly Equilibrium

Theorem 1 establishes that the monopoly policy is always an equilibrium. In this section we first provide an example where there is a non-monopoly equilibrium. We then return to the general case, characterize the set of equilibria by looking at local deviations and show that when products are "dispersed", the monopoly policy is the unique equilibrium policy.

Example 2 (Vertical differentiation). Suppose there are two vertically differentiated products and two states $\{L, H\}$, with $p = \Pr(H)$. The first product is a cheap, low-quality book yielding utility $u_1 = \left(-\frac{1}{3}, \frac{2}{3}\right)$; the second is an expensive, high-quality book, yielding utility $u_2 = (-1, 1)$. Thus, the buyer prefers no book if $p \in \left[0, \frac{1}{3}\right)$, prefers the cheap book if $p \in \left[\frac{1}{3}, \frac{2}{3}\right)$ and prefers the expensive book if $p \in \left[\frac{2}{3}, 1\right]$. Assume the expensive book also has a higher profit for the seller, $\tilde{\pi}(u_1) = 1$ and $\tilde{\pi}(u_2) = \frac{3}{2}$. The monopoly policy here is as follows:

$$K(p) = \begin{cases} (1 - 3p)\, \delta_{\{0\}} + 3p\delta_{\left\{\frac{1}{3}\right\}} & \text{if } p \in \left[0, \frac{1}{3}\right) \\ (2 - 3p)\, \delta_{\left\{\frac{1}{3}\right\}} + (3p - 1)\, \delta_{\left\{\frac{2}{3}\right\}} & \text{if } p \in \left[\frac{1}{3}, \frac{2}{3}\right) \\ p & \text{if } p \in \left[\frac{2}{3}, 1\right] \end{cases}$$

By Theorem 1, the monopoly policy is an equilibrium; this is illustrated in Figure 2(a). One can interpret this policy as the seller recommending the high-quality book for high beliefs, recommending

$L(p) = \delta_p$ for some $p \in \Delta S$. Then it must be that K also yields the same profit as δ_p on Q_c. Hence, the tie-breaking property of K given Q_c implies the tie-breaking property of K given $Q_{c'}$.

Figure 2: Vertical and Horizontal Differentiation (Public Beliefs)

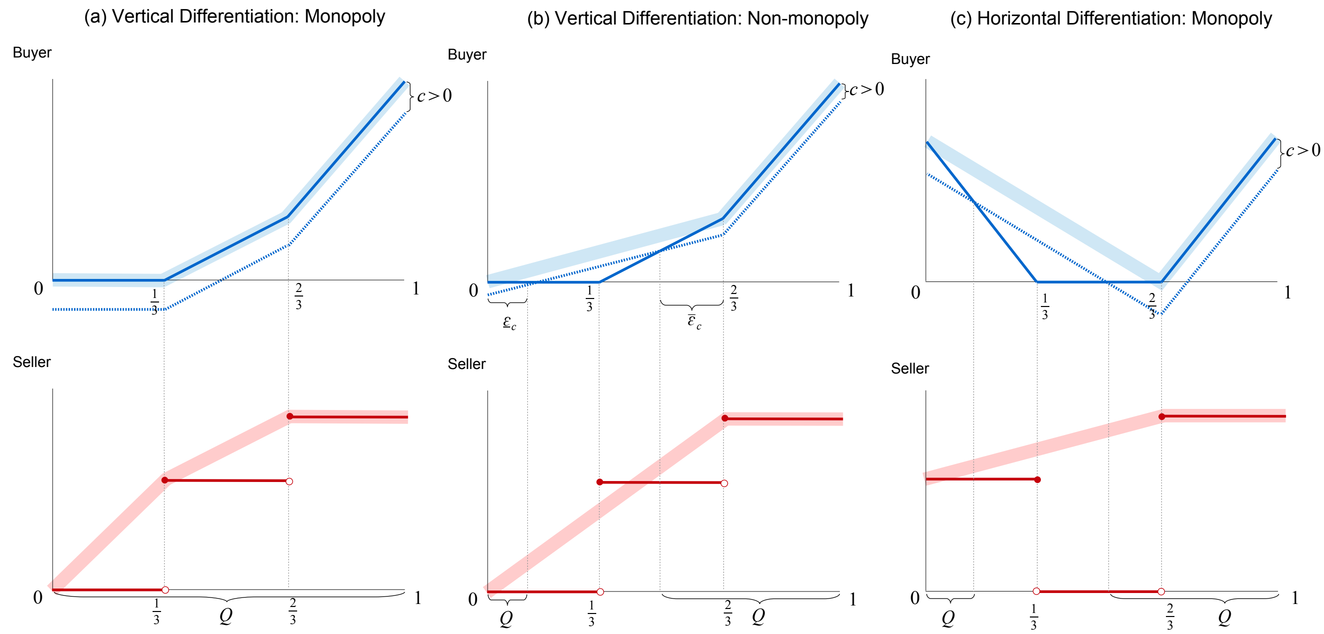

Figures 2(a) and (b) show the two equilibria that arise in the vertical differentiated example. **Figure 2(c)** shows the unique equilibrium that arises in the horizontally differentiated example; this coincides with the monopoly policy.

either the high- or low-quality book for middling beliefs, and recommending either the low-quality book or no book for low beliefs.

When $c > 0$ is small enough, there is also a second equilibrium where sellers provide more information,

$$K(p) = \begin{cases} \frac{2-3p}{2}\delta_{\{0\}} + \frac{3p}{2}\delta_{\left\{\frac{2}{3}\right\}} & \text{if } p \in \left[0, \frac{2}{3}\right) \\ p & \text{if } p \in \left[\frac{2}{3}, 1\right] \end{cases} \tag{2}$$

as illustrated in Figure 2(b). This gives rise to the acceptance set $Q = [0, \underline{\varepsilon}_c] \cup \left[\frac{2}{3} - \bar{\varepsilon}_c, 1\right]$ for some $\underline{\varepsilon}_c > 0$ and $\bar{\varepsilon}_c > 0$. Hence a seller's profits are

$$\pi(q)\,\mathbf{1}_Q(q) = \mathbf{1}_{\left[\frac{2}{3}-\bar{\varepsilon}_c, \frac{2}{3}\right)} + \frac{3}{2} \cdot \mathbf{1}_{\left[\frac{2}{3}, 1\right]}$$

This policy differs from the monopoly policy by never inducing the buyer to purchase the low-quality book. As shown in Figure 2(b), 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 one seller induces a buyer to have a posterior $q = \frac{1}{3}$, as under the monopoly policy, then the buyer would refuse to buy, correctly anticipating that other sellers will provide significantly more information.△

We now provide a characterization of equilibria. As illustrated by Examples 1-2, the search cost means that a seller can provide a little less information than the market. As a result, all equilibria can be characterized by a local optimality condition. We say a continuous disclosure policy K is *locally optimal* iff it generates more profits than any other policy that has support within an ε-ball of K's support. In other words, K is locally optimal if there exists some $\varepsilon > 0$ such that $\int_{\Delta S} \pi(q) K(p, dq) \geq \int_{\Delta S} \pi(q) L(p, dq)$ for all L with $L(p, B_\varepsilon(\text{supp}(K))) = 1$.[12] A monopolist seller who uses a locally optimal policy has no incentive to deviate if she were able to slightly perturb the support of her policy.

Proposition 1. *Any continuous equilibrium policy is locally optimal. Conversely, any locally optimal policy is an equilibrium policy for some $c > 0$.*

Proof. This follows from Propositions A2.3 and A2.5 in Appendix A. □

Local optimality is a necessary condition for equilibrium. Intuitively, a buyer with belief $q \in A_K$ gets no information since the belief is absorbing, and so purchases, $q \in Q$. By continuity of K, a buyer with any nearby belief $q \in B_\varepsilon(A_K)$ gets little information and also purchases, $q \in Q$. Since an equilibrium policy K must be better than any deviation in Q, it must be better than any deviation in $B_\varepsilon(A_K)$. Conversely, any locally optimal policy is an equilibrium when costs are sufficiently small. As $c \to 0$, the acceptance set converges to the absorbing beliefs, $Q \to A_K$. Hence for any ϵ, there is a sufficiently small c such that $Q \subset B_\varepsilon(A_K)$ and any locally optimal policy is better than any deviation in Q, making it an equilibrium.

As illustrated by the above examples, Proposition 1 is useful for characterizing the equilibria of the game. Since local optimality makes no reference to the buyer's acceptance set, we can check whether any given disclosure policy is an equilibrium policy simply by looking at local deviations. In Example 1, if K discloses more than the monopoly level, one can that see there is a local deviation. Conversely, in Example 2, one can verify that there are no local deviations from the proposed non-monopoly equilibrium policy, equation (2), meaning that it is an equilibrium when the search cost is small.

We now provide a sufficient condition for the monopoly policy to be the unique equilibrium. We say that products are *dispersed* if for any two products u and v, $q \cdot u \geq 0$ implies $q \cdot v \leq 0$. This condition says that the buyer is only willing to purchase one product at any given belief. It implies that no two products that are too alike and no products are "stuck in the middle" of the belief space, as in Example 2.

Theorem 2. *If products are dispersed then any equilibrium policy is a monopoly policy.*

[12] The continuity assumption is satisfied in most general cases. For example, if Q is a polytope or if products are dispersed, then the optimal disclosure policy is continuous. Intuitively, if Q is a polytope, then the concavified value function of the seller is the minimum of a finite number of hyperplanes. In this case, the best-response policy correspondence of the seller (as a function of buyer posteriors) is lower hemi-continuous. Hence, by the Michael selection theorem, a continuous optimal disclosure policy exists.

Proof. See Corollary A3.4 in Appendix A. □

The idea behind Theorem 2 is as follows. First observe that, if products are dispersed then the monopoly policy has a lexicographic perfect bad news structure. Suppose there are three states and three goods, as shown in Figure 3(a) with the most profitable good u_1 favored in state s_1, the second most profitable u_2 favored in state s_2, and the third most profitable u_3 favored in state s_3. The monopoly policy can then be viewed in two steps: first, it releases a perfect bad news signal about s_1, taking $p \to \{q_1, q_{23}\}$ where q_1 just persuades the buyer to buy u_1, and q_{23} is on the edge between s_2 and s_3; second, it releases a perfect bad news signal about s_2 such that $q_{23} \to \{q_2, \delta_{\{s_3\}}\}$, where q_2 just persuades the buyer to purchase good u_2.

Now, consider Figure 3(a), where the region H_i describes the beliefs for which the buyer prefers good u_i. We first show that any belief in H_1 must be absorbing under K. If this were not the case, the seller's optimal policy would place weight on posteriors in the interior of H_1 and the $\{s_2, s_3\}$ edge. She could then provide a little less information, enlarging the support of the policy within H_1 from $A_K \cap H_1$ to $B_\epsilon(A_K) \cap H_1$, analogous to Example 1. Given a prior like that in the figure, this deviation places more weight on beliefs in H_1 and less on the $\{s_2, s_3\}$ edge, yielding strictly higher profits for the seller. By induction, we next consider the $\{s_2, s_3\}$ edge and show that any belief in the shaded red line (for good u_2) and the red dot at s_3 (for good u_3) must also be absorbing. Hence, any belief in the support of the monopoly policy is absorbing and lies in the buyer's acceptance set, allowing the seller to obtain monopoly profits.

Theorem 2 assumes that products are dispersed. With two states, Example 2 implies that we may have a problem whenever there is a "middle" product; with three or more states, a similar issue may arise whenever a product is stuck in the middle along some dimension. For example, consider Figure 3(b) which is the three-dimensional extension of Example 2. In this case, the regions where goods u_1 and u_2 are optimal intersect and for sufficiently profitable u_2, the monopoly policy involves giving the buyer no information at p. However, as in Example 2, there is another equilibrium: if others sellers "skip over" product u_2 and send the buyer $p \to \{q_1, \delta_{s_3}\}$, it is optimal for the current seller to ignore u_2 as well. Such non-monopoly equilibria cannot arise if products are dispersed.

Theorem 2 has two useful corollaries. First, if there is a single product, then the only equilibrium is the monopoly policy. This result thus extends Example 1 to any finite number of states.

Corollary 1. *If there is a single product, then any equilibrium policy is a monopoly policy.*

Proof. Follows immediately from Theorem 2. □

Second, if there are two horizontally differentiated products then the only equilibrium is the monopoly policy. Broadly speaking, two products are horizontally differentiated if any state that is good for u_1 is bad for u_2. Formally, we say two products are *opposed* if $\delta_s \cdot u_1 \geq 0$ implies $\delta_s \cdot u_2 \leq 0$. Note that if two products are dispersed, then they must be opposed.

Figure 3: Monopoly Policy (Public Beliefs)

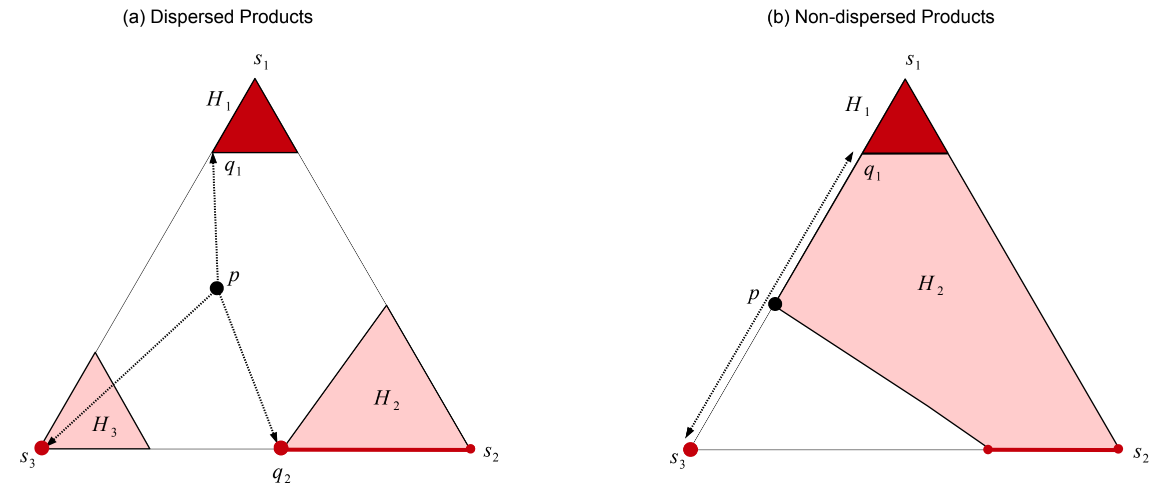

Figure 3(a) shows the monopoly policy if products are dispersed. **Figure 3(b)** considers an example where complications can arise when beliefs are not dispersed. In this case p is an absorbing belief under the monopoly policy, but may not be one under some other equilibrium policy, resulting in multiple equilibria.

Corollary 2. *If there are two opposed products, then any equilibrium policy is a monopoly policy.*

Proof. See Corollary A3.5 of Appendix A. □

This result is illustrated in Figure 2(c) which supposes there are two states. The first yields utility $u_1 = \left(-1, \frac{1}{2}\right)$, the second yields utility $u_2 = \left(\frac{1}{2}, -1\right)$ and profits are $\tilde{\pi}(u_1) > \tilde{\pi}(u_2) > 0$. When products are horizontally differentiated, we may have $u_1(q) > 0$ and $u_2(q) > 0$ for some q, meaning the condition is not implied by products being dispersed. Nevertheless, the condition that products are opposed means that no goods can be "stuck in the middle", and the optimal monopoly policy is a lexicographic perfect bad news policy. The rest of the proof is analogous to Theorem 2.

3 Private Beliefs

In this section, we suppose buyers' beliefs are private, so sellers cannot provide different information to buyers with different beliefs. More precisely, we suppose that a buyer starts with a known prior; if the buyer arrives at a seller, she observes neither the buyer's current belief nor his past search behavior; she then chooses to disclose an independent signal. On the equilibrium path, buyers will purchase at the first seller, but the possibility that an old buyer can mimic a new buyer forces sellers to provide more information than when beliefs are public.

We first set up the model and demonstrate how private beliefs affect the equilibrium in the single product example. Our first main result is that as search costs become small, there is is always a limit equilibrium that is payoff-equivalent to full-information. In general, there may be multiple equilibria and we characterize these limit equilibrium policies. We then show that if products are "sufficiently dispersed", then the full-information limit equilibrium is unique.

3.1 Model

Basics. There is one buyer and infinite sellers. All sellers know the buyers initial prior p which lies in the interior of ΔS. The timing is the same as before: when a buyer arrives at a seller, she chooses a signal structure, the buyer updates his prior and chooses whether to buy, exit or pay the search cost c and continue.

Seller's Strategy. A seller chooses a signal structure on some space Θ that is independent of all previous signals received by the buyer. Any potential buyer with prior r can also observe the signal and learn about the state. Let $q(\theta)$ be the posterior of a buyer with prior p after observing a signal $\theta \in \Theta$, and let $q_r(\theta)$ denote the posterior of a buyer with prior r who observes the same signal. The posteriors of the two buyers are related as follows

$$q_r(\theta) = \phi_r(q(\theta))$$

where $\phi_r : \Delta S \to \Delta S$ is a mapping from the posteriors of buyer p to the posteriors of buyer r satisfying

$$[\phi_r(q)](s) := \frac{q(s)\frac{r(s)}{p(s)}}{\sum_{s'} q(s')\frac{r(s')}{p(s')}}$$

Intuitively, for any two buyers, the proportionality of their likelihood ratios for any two states remains constant when we update via Bayes' rule. If buyer r starts off twice as optimistic as buyer p, i.e. r has double the likelihood ratio than p, then q will remain twice as optimistic as p after any signal realization.

Since this mapping ϕ_r is independent of the signal structure, sellers can simply choose a *signal policy* μ which is a distribution of posteriors for a buyer with prior p and satisfies $\int_{\Delta S} q\ \mu(dq) = p$. If a buyer p observes a signal inducing posterior q, then buyer r has posterior $\phi_r(q)$; in addition, buyer r has a different distribution over states and therefore over signals. Putting this together, buyer r has posteriors distributed according to the Markov kernel $K_\mu(r) = \mu_r \circ \phi_r^{-1}$, where μ_r is the signal distribution on Θ from the perspective of buyer r. Let $\bar{\mu}$ denote the *full information signal policy* which has full support on all degenerate posteriors. In other words, $\bar{\mu}$ reveals the true state to the buyer.

Symmetric equilibrium. We consider symmetric equilibria such that, on the equilibrium path, a buyer leaves the market after his first match (with or without a purchase). Given the buyer's acceptance set Q, a policy μ is *optimal given* Q iff for all possible signal policies ν

$$\int_Q \pi(q)\,\mu(dq) \geq \int_Q \pi(q)\,\nu(dq).$$

and the usual tie-breaking rule applies.

A buyer with posterior q who continues to search obtains continuation utility

$$V_c(\mu, q) := -c + \int_{\Delta S} \max\left\{\max_{u\in U} r\cdot u,\ V_c(\mu, r)\right\} K_\mu(q, dr)$$

We also define $\bar{V}(q) := V_0(\bar{\mu}, q)$ as the buyer profit from using the full-information policy $\bar{\mu}$. A signal policy and acceptance set (μ, Q) form an *equilibrium* iff μ is optimal given

$$Q_c(\mu) := \left\{q \in \Delta S \;\middle|\; \max_{u\in U} q\cdot u \geq V_c(\mu, q)\right\}.$$

In this case, we call μ an *equilibrium policy.*

We also consider what happens to equilibria when search costs vanish. A signal policy μ is a *limit equilibrium policy* iff there are equilibrium policies μ_i for each $c_i > 0$ such that $\mu_i \to \mu$ as $c_i \to 0$.

Remarks. Since we are interested in large markets where coordination might be difficult, we assume signals are independent. As we will show, the crucial property is that a buyer learns all relevant information as the number of samples grows large; this suggests that any model of imperfect correlation will suffice. With this said, if sellers could perfectly coordinate their signals, then such a signal would be useful to a new buyer but not to an old one, and we would be back in the setting of Section 2.

We also assume that the seller chooses a single signal policy, rather than offering a menu to differentiate between new and old buyers. On the equilibrium path, buyers purchase from the first seller, so there are no old buyers and the single policy we study is a best-response in the more general model. Moreover, even if sellers did offer a second policy that were accepted by old buyers off the equilibrium path, this would only increase their utility, reinforcing Theorem 3.

3.2 Motivating Example

Example 1 (cont). As before there are two states $\{L, H\}$ and the buyer's initial prior is $p = \Pr(H)$. First, suppose that $p < \frac{1}{2}$, so that the monopoly policy provides a positive amount of information. Then as $c \to 0$, the only equilibrium policy is full revelation $\bar{\mu}$. To see this, first observe that the value function is convex and increasing in the posterior q. Since information has no value at the boundaries, the acceptance set is of the form $Q_c(\mu) = [0, a] \cup [b, 1]$, where $a \leq \frac{1}{2} \leq b$. As a result, the seller's optimal policy is a perfect bad news signal $p \to \{0, b\}$, as shown in Figure 4(a).

In such an equilibrium, the buyer with posterior b must prefer to accept rather than searching and obtaining one more signal. If the buyer with prior p receives a signal $p \to \{0, b\}$, a buyer with prior r will have posteriors

$$\begin{array}{ll} \phi_r(b) = \frac{b\frac{r}{p}}{b\frac{r}{p} + (1-b)\frac{1-r}{1-p}} & \text{with probability } \frac{r}{\phi_r(b)} \\ \phi_r(0) = 0 & \text{with probability } 1 - \frac{r}{\phi_r(b)} \end{array}$$

Hence buyer b prefers to purchase immediately rather than search again if

$$2b - 1 \geq \frac{b}{\phi_b(b)}(2\phi_b(b) - 1) - c$$

Rearranging, this becomes

$$-b^2 + b[1 + p - c(1-p)] - p \leq 0. \tag{3}$$

The LHS of the inequality (which has the same sign as the value of searching) equals $-c(1-p)$ at the endpoints $b \in \{p, 1\}$ and is larger in between. When c is small, the LHS is positive on $\left[\frac{1}{2}, b^*\right)$ and negative on $(b^*, 1]$, where b^* is the larger root of the quadratic. Since the seller wishes to provide as little information as possible, this means there is a unique equilibrium in which the seller uses the policy $p \to \{0, b^*\}$ which induces an acceptance set $Q = [0, a^*] \cup [b^*, 1]$, so b^* is indifferent between buying and searching. Moreover, as $c \to 0$, the value of searching increases and $b^* \to 1$, as shown in Figure 4(b). Economically, this means that the seller provides full information. Intuitively, since the monopoly policy provides some information then, as the cost vanishes, a buyer can obtain a large number signals at low cost and become almost fully informed. Hence the current seller must provide a lot of information in order to beat this outside option and make a sale.

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 c. If sellers provide no information, the buyer will not search and, since "no information" is the monopoly policy, no seller will defect. There are two further equilibria which correspond to the two roots of the quadratic equation (3), denoted by $\{b_L^*, b_H^*\} \subset \left[\frac{1}{2}, 1\right]$. As $c \to 0$, $b_L^* \to p$, which corresponds to no information, whereas $b_H^* \to 1$, which corresponds to full information. △

This example shows how there are equilibria that are much more informative than the monopoly policy. When beliefs are public, a buyer who receives information from one seller obtains no useful information if he were to go back to the market. However, when beliefs are private, the buyer can reject the first offer, pretend to be uninformed and receive more information from another seller in the market. This possibility of taking this outside option allows competition to work.

Figure 4: Single Product (Private Beliefs)

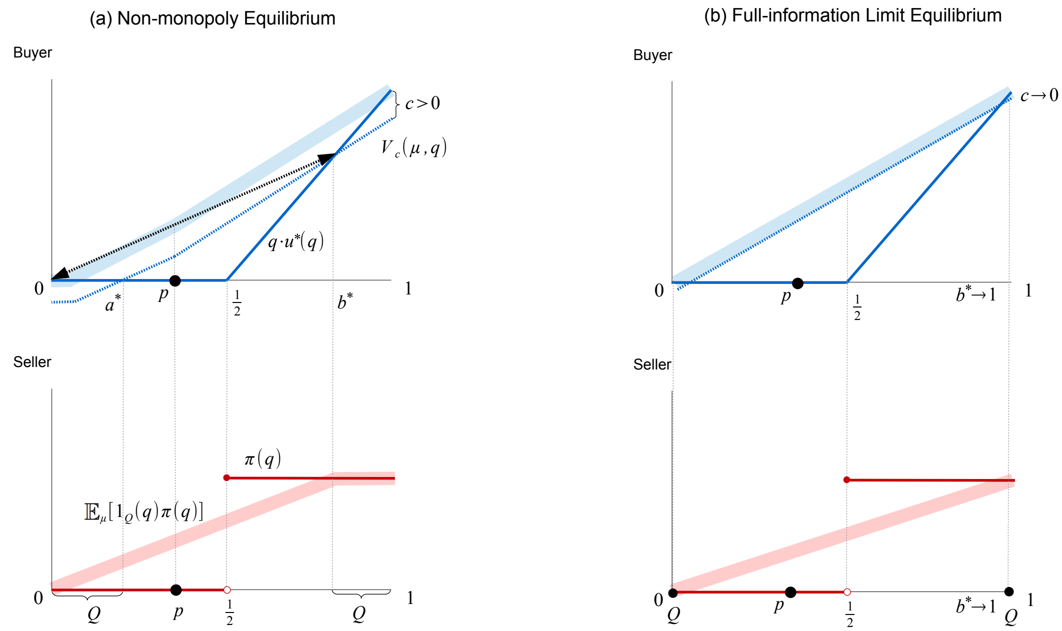

This figure illustrates Example 1 with private beliefs for a prior $p < \frac{1}{2}$. **Figure 4(a)** shows a non-monopoly equilibrium: A perfect bad news policy $p \to \{0, b^*\}$ that induces the acceptance set $Q = [0, a^*] \cup [b^*, 1]$. **Figure 4(b)** shows that when the search cost vanishes, $c \to 0$, this equilibrium converges to full-information, $b^* \to 1$.

3.3 Full Information is a Limit Equilibrium

We now return to the general model, with multiple states and multiple goods. First, we address the issue of equilibrium existence. Fix a search cost $c > 0$ and suppose all other sellers use policy μ. The other sellers' policies generate a value function $V_c(\mu, \cdot)$ for the buyer and an acceptance set $Q_c(\mu)$; the current seller then has a set of optimal signal policies $\varphi_c(\mu)$. We thus define the *best-response correspondence* $\mu \twoheadrightarrow \varphi_c(\mu)$ which describes the best-response policies given that all other sellers use policy μ. Note that the set of all signal policies is a convex compact space. If the best-response correspondence is nonempty, convex valued and has a closed graph, then the Kakutani-Fan-Glicksberg theorem implies that the correspondence has a fixed point, and an equilibrium exists.

Unfortunately, φ_c may not have a closed graph. To see this, let us return to the vertical differentiation case in Example 2 with initial prior $p = \frac{1}{3}$. Consider a sequence of policies μ_i and optimal policies $\nu_i \in \varphi_c(\mu_i)$ such that $(\mu_i, \nu_i) \to (\mu, \nu)$ and $Q_c(\mu_i) = \left[0, \frac{1}{3} - \varepsilon_i\right] \cup \left[\frac{2}{3}, 1\right]$, where $\varepsilon_i \to 0$. Note that for each $\varepsilon_i > 0$, the buyer will never purchase the low-quality book, so the seller's optimal policy ν_i is a perfect bad news signal $p \to \left\{0, \frac{2}{3}\right\}$. However, in the limit as $Q_c(\mu_i) \to \left[0, \frac{1}{3}\right] \cup \left[\frac{2}{3}, 1\right]$ the optimal policy $\nu^* \in \varphi_c(\mu)$ provides no information. Hence $\nu \notin \varphi_c(\mu)$

and the correspondence does not have a closed graph.

To address this issue, we assume that products are robust to full-information in that for every good $u \in U$, there is a state $s \in S$ such that the buyer purchases u if he knows the state is s for sure, as shown in Figure 5(a). Formally, let H_u denote the set of beliefs where the buyer chooses u,

$$H_u := \left\{ q \in \Delta S \;\middle|\; u = \text{argmax}_{u' \in u^*(q)} \tilde{\pi}\left(u'\right) \right\}$$

We say that the set of products U is *robust* iff for all non-zero $u \in U$, there exists some $s(u) \in S$ such that $\delta_{s(u)} \in H_u$. Note that if products are dispersed, then they are robust.

Proposition 2. *If products are robust then an equilibrium exists.*

Proof. See Corollary B2.5 in Appendix B. □

We now provide a sketch of the main argument. It is straightforward to show that the correspondence φ_c is nonempty and convex valued; the key step is to show that φ_c has a closed graph. Consider a sequence of policies μ_i and best-responses $\nu_i \in \varphi_c(\mu_i)$ such that $(\mu_i, \nu_i) \to (\mu, \nu)$. Let $\nu^* \in \varphi_c(\mu)$ be the optimal policy given μ, and suppose ν^* puts weight on some $q \in Q_c(\mu) \cap H_u$ for some $u \in U$. By the definition of robustness, we know that there is a state $s(u) \in S$ such that $\delta_{s(u)} \in H_u$. Since buyer value functions are convex and continuous, for each μ_i, we can find some $q_i \in Q_c(\mu_i) \cap H_u$ along the line segment from q to $\delta_{s(u)}$ such that $q_i \to q$. Hence, by shifting the weights from q onto q_i, we can construct policies $\nu_i^* \to \nu^*$ such that ν_i^* has support in $Q_c(\mu_i)$ and the profits coincide in the limit. Since ν_i is weakly more profitable than ν_i^* for every μ_i, it follows that $\nu = \lim_i \nu_i$ is weakly more profitable than $\nu^* = \lim_i \nu_i^*$, and hence $\nu \in \varphi_c(\mu)$, as required. We can thus apply Kakutani-Fan-Glicksberg to obtain the fixed point.

We now show that full information is always a limit equilibrium.

Theorem 3. *There exists a limit equilibrium μ such that $V_{c_i}(\mu_i, p) \to \bar{V}(p)$ as $c_i \to 0$.*

Proof. See Theorem B3.6 in Appendix B. □

Intuitively, as costs become small, a buyer either learns everything or nothing about a state. On the equilibrium path, a buyer only visits a single seller; however, for each search cost c_i, he always has the option of visiting $1/\sqrt{c_i}$ sellers for a total search cost of $\sqrt{c_i}$. As c_i approaches zero, so does the total search cost. Hence, if each seller provides some information about a state, the buyer will be fully-informed about that state in the limit. This means that if all sellers provide information about all states then a single seller must match them and also provide information about those states. Note that Theorem 3 says that the limit equilibrium policy must be payoff-equivalent to the full-information buyer payoff, but the seller may not provide all information. This occurs because a buyer who makes the same decision in state s_1 and s_2 will not pay a search cost to separate these states.

Figure 5: Equilibrium Existence (Private Beliefs)

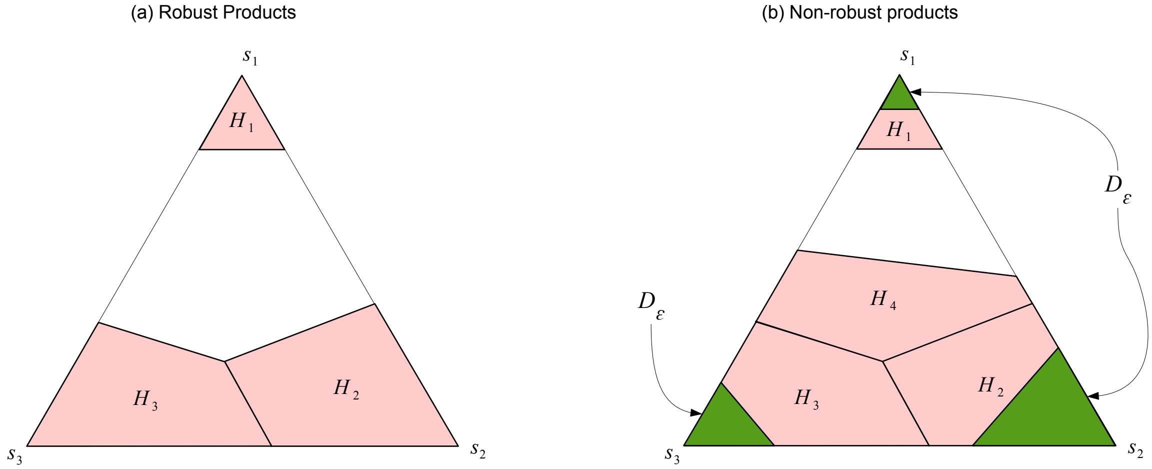

Figure 5(a) shows an example with robust products. **Figure 5(b)** shows how we prove that there always exists a full-information limit equilibrium.

To understand how the proof works, suppose that information about every state is valuable to the buyer as in Figure 5(b). We wish to construct a sequence of equilibrium policies μ_i that converges to full-information $\bar{\mu}$ as $c_i \to 0$. To do this, let M_ε be the set of signal policies with support contained in some ε-neighborhood of the vertices of ΔS, denoted D_ε. Suppose that all other sellers use a policy $\mu \in M_\varepsilon$; we claim that the best-response satisfies $\varphi_c(\mu) \subset M_\varepsilon$ for some small c. First observe that, since μ reveals something about every state, the buyer's value function $V_c(\mu, \cdot)$ converges to the full-information payoff $\bar{V}$ monotonically. This means we can find some small c such that $V_c(\mu, \cdot)$ is close to $\bar{V}$ and $Q_c(\mu) \subset D_\varepsilon$. Since any best-response policy $\nu \in \varphi_c(M_\varepsilon)$ puts full support in $Q_c(\mu)$, it follows that $\varphi_c(M_\varepsilon) \subset M_\varepsilon$ for small c, as required. Note that the products close to the vertices of ΔS are robust so Proposition 2 implies the existence of an equilibrium in M_ε.[13] Hence, for every $\varepsilon_i > 0$, we can find some small enough c_i such that there exists an equilibrium policy $\mu_i \in M_\epsilon$. As $\varepsilon_i \to 0$, $c_i \to 0$ and our sequence of equilibrium policies μ_i approaches the full-information policy $\bar{\mu}$ in the limit.

3.4 Uniqueness of the Full-Information Limit Equilibrium

Theorem 3 establishes that there always exists a limit equilibrium that is payoff equivalent to full-information. In some situations however, there are multiple limit equilibria where some provide less than full information. In this section, we characterize these limit equilibria as revealing partitional

[13] Products close to the vertices are robust if buyers have strict preferences at the vertices. The problematic case is when a buyer has $u^* = u_1$ at the vertex, and $u^* = u_2$ at all nearby points. In such a case, a seller strictly prefers the buyer to purchase u_1, and we define D_ϵ to coincide with the vertex.

Figure 6: Limit Equilibria (Private Beliefs)

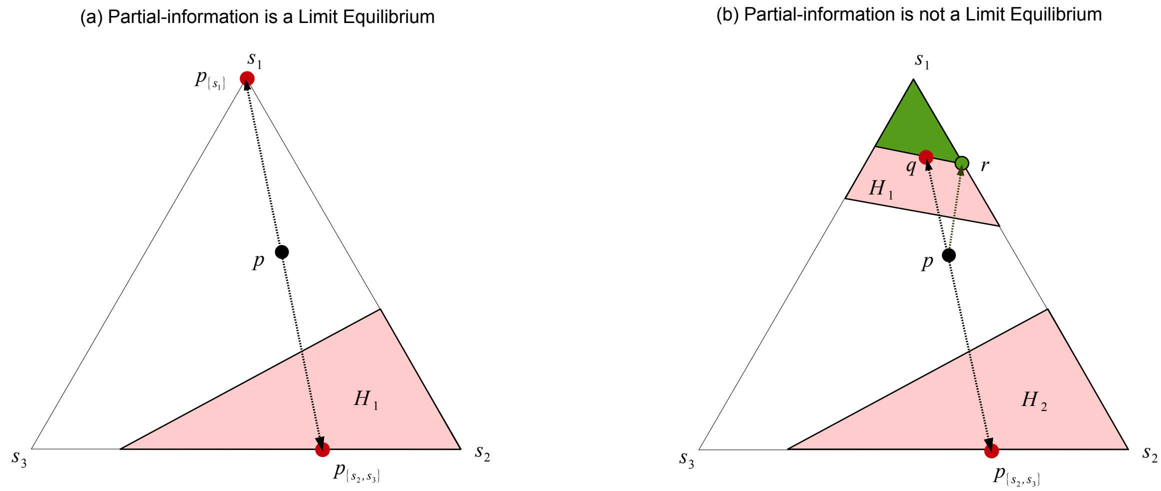

Figure 6(a) shows a partitional limit equilibrium where the buyer only learns $\mathcal{T} = \{\{s_1\}, \{s_2, s_3\}\}$. Note that $\mathcal{T}$ is partitionally optimal because at $p_{\{s_2,s_3\}}$, the buyer is already purchasing the most profitable (only) product. **Figure 6(b)** shows that not every partition that is partitionally optimal corresponds to a limit equilibrium. Even though all other sellers are providing information about s_1 only and $\mathcal{T} = \{\{s_1\}, \{s_2, s_3\}\}$ is partitionally optimal, each seller has an incentive to deviate and provide some information about s_2 versus s_3 (i.e. sending the buyer to posterior r instead of q).

information about the state space, and show that when products are "sufficiently dispersed", the full-information limit equilibrium payoff is unique.

We first present an example of a limit equilibrium that is not payoff equivalent to full-information. Suppose there are three states $S = \{s_1, s_2, s_3\}$ and a single product u_2 as in Figure 6(a). Suppose that all other sellers provide information about s_1 only, so a buyer never receives any information that allows him to distinguish between s_2 or s_3. The buyer's posterior beliefs thus lie on the line connecting s_1 with the belief $p_{\{s_2,s_3\}}$, the projection of p on the $\{s_2, s_3\}$ edge. From the current seller's perspective, a buyer near the belief $p_{\{s_2,s_3\}}$ purchases her product and the seller has no incentive to add dispersion in his beliefs. This means that given that all other sellers are providing no information about s_2 versus s_3, the current seller also has no incentive to provide information. As search costs go to zero, buyers know for sure whether s_1 is true or not, but they learn nothing about s_2 versus s_3. This partial-information policy is a limit equilibrium and gives the buyer a strictly lower payoff than full-information. Of course, Theorem 3 implies that full-information is also a limit equilibrium, so this example exhibits multiple limit equilibria.

We now provide a characterization of these limit equilibrium policies. Let $\mathcal{T}$ be some partition of the state space S. For any event $T \in \mathcal{T}$ and belief $p \in \Delta S$ such that $p(T) > 0$, let $p_T \in \Delta S$ be the conditional belief of p given T. In other words,

$$p_T(s) = \frac{p(s)}{p(T)}$$

for all $s \in T$, and $p_T(s) = 0$ otherwise. Geometrically, if ΔT is the sub-simplex corresponding to the event T, then p_T is the projection of p onto ΔT. For each partition $\mathcal{T}$, let $\mu^{\mathcal{T}}$ denote the signal policy corresponding to revealing $\mathcal{T}$ to the buyer. Finally, let $V_{\mathcal{T}}$ be the buyer's payoff from receiving information $\mathcal{T}$.[14]

We illustrate these ideas with two examples. First consider the case of no information release. $\mathcal{T}$ is then the trivial partition where $\mathcal{T} = \{S\}$ and $\mu^{\mathcal{T}} = \delta_p$. In this case, providing no information is a limit equilibrium if a monopolist provides no information, as seen in Example 1. Intuitively, if all other sellers provide no information, then a buyer will never search and a seller's problem is identical to the monopoly problem. Second, consider the case of full-information release. $\mathcal{T} = \bar{\mathcal{T}}$ is then the finest partition of S and $\mu^{\mathcal{T}} = \bar{\mu}$. In this case, providing full information is a limit equilibrium since all other sellers are already providing full information, as discussed in Theorem 3.

We now show that any limit equilibrium policy must have a partitional informational structure corresponding to some $\mathcal{T}$. In other words, as search costs become arbitrarily small, a buyer learns everything about some states and nothing about others according to the partition $\mathcal{T}$. Moreover, it must also be true that $\mathcal{T}$ is *partitionally optimal* in that for all $T \in \mathcal{T}$,

$$\pi(p_T) \geq \int_{\Delta T} \pi(q)\, \nu(dq)$$

for all ν consistent with p_T. In other words, a monopolistic seller would prefer to provide no information to a buyer with belief p_T for every event $T \in \mathcal{T}$. Generically, this means that the buyer purchases the most profitable good that can be reached from p_T.[15]

Proposition 3. *If μ is a limit equilibrium, then there is a partitionally optimal $\mathcal{T}$ such that $V_{c_i}(\mu_i, p) \to V_{\mathcal{T}}(p)$ as $c_i \to 0$.*

Proof. See Proposition B3.7 in Appendix B. □

The intuition behind this result is that if sellers reveal a little information about state s versus s' then, as the search cost vanishes, a buyer can visit a large number of sellers and perfectly distinguish them. In order to make a sale, any one seller must therefore match the market, meaning that any limit information structure must be partitional. Now, suppose the market induces partition $\mathcal{T} = \{\{s_1\}, \{s_2, s_3\}\}$, as in Figure 6(a). It must be the case that, given belief $p_T = p_{\{s_2,s_3\}}$, the seller does not want to release any extra information about s_2 versus s_3. Otherwise, the buyer could

[14] Formally, $\mu^{\mathcal{T}} := \sum_{T \in \mathcal{T}} p(T)\, \delta_{p_T}$ and $V_{\mathcal{T}}(p) := \int_{\Delta S} (\max_{u \in U} q \cdot u)\, K_{\mu^{\mathcal{T}}}(p, dq)$.

[15] To illustrate, in Example 2, the monopolist provides no information if $p \in \left\{\frac{1}{3}\right\} \cup \left[\frac{2}{3}, 1\right]$. This means that, generically, the buyer purchases the high-quality book.

again visit many sellers and perfectly distinguish them. This means that any limit equilibrium must be partitionally optimal.

Even though every limit equilibrium corresponds to a partitionally optimal $\mathcal{T}$, not all partitionally optimal $\mathcal{T}$ correspond to limit equilibria. For example, consider Figure 6(b) in which there are two products u_1 and u_2, and let $\mathcal{T} = \{\{s_1\}, \{s_2, s_3\}\}$. In the limit, as the buyer learns either s_1 or $\{s_2, s_3\}$ perfectly, the seller has no incentive to provide information about s_2 versus s_3, since $\mathcal{T}$ is partitionally optimal. However, away from the limit, she does prefer to provide information about s_2 versus s_3. In particular, the seller would like to send the buyer to posterior r rather than q. Hence, there is no equilibrium that corresponds to $\mathcal{T}$.

We now derive conditions under which full information is the unique limit equilibrium. We say a set of products U is *sufficiently dispersed* iff it is dispersed and $u^*(p_T) = \{0\}$ for all events T that consist of at least two states. For example, products in Figure 3(a) are sufficiently dispersed since the buyer chooses to not purchase at each p_T for all T with at least two states. In other words, products are sufficiently spread out in belief space such that the seller needs to reveal some information about each state in order to induce the buyer to purchase a product.

Theorem 4. *If products are sufficiently dispersed, then for any limit equilibrium μ, $V_{c_i}(\mu_i, p) \rightarrow \bar{V}(p)$ as $c_i \rightarrow 0$.*

Proof. See Theorem B3.8 in Appendix B. □

Taken in combination, Theorems 2 and 4 imply that when products are sufficiently dispersed and search costs go to zero, buyers receive a higher payoff under private beliefs than public beliefs. When beliefs are public, buyers only attain the monopoly payoff, even when search costs are arbitrarily small. The ability to discriminate between new and old buyers allows sellers to implicitly collude. When beliefs are private, buyers can attain their full-information payoff in the limit. Even though buyers purchase from the first seller on-path, the option to mimic an uninformed buyer and receive more information forces sellers to compete against each other and provide much more information to buyers.

With two states however, the analysis is even simpler. This extends the analysis of Example 1 to multiple goods.

Corollary 3. *Suppose there are two states. If a monopolist provides some information, full information is the unique limit equilibrium. Otherwise, both full and no information are limit equilibria.*

Proof. Follows from Theorem 3 and Proposition 3. □

4 Extensions

In this Section we consider two extensions. In Section 4.1 we suppose a buyer can eliminate their cookies at a small cost and choose between the public and private cases. In Section 4.2 we consider a case in between the public belief model and the private belief model. In particular, we suppose a seller sees whom a buyer has visited in the past, but does not see the realizations of the signals. We argue the forces are broadly similar to the public belief model.

4.1 The Choice of "Do Not Track"

We now endogenize the decision to become anonymous, supposing that buyers can pay a small cost k to activate the "Do Not Track" mode on their browser. We first argue that buyers who become anonymous exert a positive externality on other buyers. We then observe that the free-riding is so severe that there is no equilibrium in which all buyers choose to become anonymous.

Consider a continuum of both buyers and sellers and suppose all buyers start with prior p. At time $t = 0$, a buyer can choose to pay cost k in order to become anonymous. After this decision, buyers can be described by two dimensions: their belief q when meeting a seller and their type $\{A, T\}$, which describes whether they are anonymous (A) or tracked (T). When revealing information, a seller must treat an anonymous buyer (A, q) in the same way as a new tracked buyer (T, p). The seller can, however, give different signals to old tracked buyers.

One can understand much of the economics by considering a simple 2×2 game where buyers choose whether or not to be anonymous. The following table provides an illustration showing the pure strategy payoffs for the single-product example with initial prior $p = \frac{1}{3}$ and $c = 0.05$.

		Other Buyers	
		Anonymous	Tracked
Buyer i	Anonymous	$0.3 - k$	$0.15 - k$
	Tracked	0.3	0

First, suppose all other buyers are tracked, so we are in the model from Section 2. If buyer i is also tracked then he receives a single "monopoly" signal from a seller and immediately buys since he knows a second seller will refuse to provide more information. In the single-product example, this means the buyer receives 0 utility. If buyer i becomes anonymous then he can get the initial "monopoly" signal, delete the seller's cookie and then receive "monopoly" signals from subsequent sellers, paying cost c for each search. In the example, this yields expected utility $0.15 - k$. Next, suppose all other buyers are anonymous, so we are in the model from Section 3. If buyer i is tracked then he receives a single "competitive" signal that is more informative than the monopoly signal. Since he is tracked, no sellers provide any further information to him. If buyer i becomes anonymous he has the option to receive a sequence of "competitive" signals. However, on the equilibrium path, the buyer will purchase at the first opportunity.

The first point to observe is that when other buyers become anonymous, this exerts a positive externality on buyer i. Intuitively, when faced with lots of anonymous buyers, competition is more effective and sellers provide more information in equilibrium. This benefits any buyer, independent of whether they are anonymous or tracked.

Second, observe that no matter how small the cost k, there is no equilibrium where everyone is anonymous. When all other buyers are anonymous then sellers provide enough information to prevent the searching; hence buyer i has no incentive to become anonymous himself. To illustrate, consider the above table. If $k > 0.15$ there is a unique equilibrium in which everyone is tracked. If $k \in (0, 0.15)$ then there is no pure strategy equilibrium.

If we wish to characterize the mixed strategy equilibrium played by the buyers, we have to move beyond the simple 2×2 game. Unfortunately, this is difficult to analyze since anonymous buyers will search multiple times before buying, meaning that a single seller faces a heterogeneous population of buyers. In Appendix C2, we heuristically discuss the single-product example and argue that, in the most informative equilibrium, an increase in the proportion of anonymous buyers raises the amount of information disclosed.

4.2 Intermediate Observability

In the previous sections we have considered two extreme cases wherein a seller either sees the belief of the buyer, or she knows nothing about the buyer's history. Instead, suppose that a seller sees which other sellers a buyer has previously visited, but that she does not know the outcome of the signal. The main message of this section is that this model variant is similar to the case of public beliefs (Section 2). The key force driving the monopoly equilibrium is that sellers treat old buyers differently from new buyers; to do this, sellers just need to know whether a buyer is new or old.

To be more precise, suppose that each seller chooses information independently. The n^{th} seller that a buyer approaches then chooses an information policy μ_n as a function of previous disclosure policies $(\mu_1, \ldots, \mu_{n-1})$. We first observe that the monopoly policy is a sequential equilibrium.

Proposition 4. *If sellers observe the number of past sellers visited and their policies, then the monopoly policy is a sequential equilibrium.*

Proof. We consider the following strategies: Seller 1 provides the monopoly level of information, and the buyer accepts this offer; if the buyer rejects, sellers $n \geq 2$ believe the posterior is drawn from the support of the monopoly posteriors μ^* and provide no information. To finish the description, we should describe what happens if seller 1 deviates, however this is rather complicated and is unnecessary to establish the result. Broadly, the subsequent sellers will see the deviation and may provide more information; the buyer can then accept seller 1 or reject, whichever gives them higher utility.

To show this is a sequential equilibrium, we have to verify no player wishes to defect. Seller 1 obtains monopoly profits, so cannot be better off. The buyer gains nothing by rejecting, since subsequent sellers provide no information. Now, suppose a buyer approaches seller $n \geq 2$, and previous sellers followed the equilibrium. She believes the buyer's belief q is in the support of the monopoly policy, and that subsequent sellers provide no information; hence providing no information is optimal for her. This is also the limit of mixed strategies: if the buyer approaches seller 2, then she can see who trembled and, if it is the buyer, will optimally provide no information. $\square$

The underlying intuition is that when a buyer approaches seller 2, she knows that seller 1 provides the optimal amount of information. Given that the two sellers are in essentially the same position, they have the same acceptance sets and it is optimal for seller 2 to provide no more information.

The next question is whether there are any other equilibria. In Example 1, the monopoly policy is the unique equilibrium.

Example 1 (cont). Since the buyer's belief is private information, his value function is convex and each seller n faces an acceptance set of the form $[0, \alpha_n] \cup [\beta_n, 1]$. Seller 1 knows the buyer has initial prior p and thus provides a perfect bad news signal $p \to \{0, \beta_1\}$. If the buyer approaches seller 2, then she knows the buyer's belief q lies in $\{0, \beta_1\}$. If $q = 0$, then seller 2 cannot affect the buyer's posterior; her optimal policy is thus to assume the buyer received a positive signal from the first seller and use a perfect bad news policy $\beta_1 \to \{0, \beta_2\}$. Notice, that seller 2 would use exactly the same policy if she could observe the belief of the buyer! The same logic applies to subsequent sellers, so the equilibria of the game coincide with the game with public beliefs. Given that the monopoly policy is the only rationalizable outcome, as shown in Appendix C1, it is also the equilibrium in this case with intermediate observability. $\triangle$

5 Conclusion

This paper studied a market where buyers search for better information about a product and sellers choose how much information to disclose. When beliefs are public, the monopoly policy is an equilibrium; when products are dispersed it is the only equilibrium. When beliefs are private, full revelation is a limit equilibrium; when products are sufficiently dispersed, it is the only equilibrium.

Our results have interesting policy implications. Much of the discussion of the increased individualization of the Internet has focused on the role of targeting in enhancing the level of personalization of firm policies. This paper demonstrates that tracking technology may be even more insidious, undermining the competitive mechanism and supporting collusive equilibria. More broadly, regulators should be wary of any technology that enables sellers to discriminate between new and old buyers and thereby discourage searching, including methods of correlating disclosure policies.

The model can be extended in a variety of ways. First, while we assumed that the buyer's prior and values are known by the seller – e.g. because sellers have demographic data about their customers – it is natural to consider heterogeneous buyers. We conjecture that the analysis is similar to our model if all types of buyers purchase at the first seller, which would happen if, say, the priors are distributed over a small interval, e.g. $U[p-\epsilon, p+\epsilon]$. More broadly, when buyers are heterogeneous, the seller might be able to use menus to discriminate between different types (e.g. Kolotilin et al. (2015)). Second, as in many of the traditional search papers, it is natural to consider match specific tastes, or let the sellers choose other dimensions of their strategy, such a price, product quality, and product range (e.g. Wolinsky (1986)). Third, while we have considered a model with two types of players, one would like to understand the incentives of third-party data-brokers to collect and sell information to sellers (e.g. Bergemann and Bonatti (2015)).

References

Aliprantis, C., and K. Border (2006): *Infinite Dimensional Analysis*. Springer.

Anderson, S. P., and R. Renault (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

Aumann, R. J., and M. Maschler (1995): *Repeated games with incomplete information*. MIT press.

Battaglini, M. (2002): "Multiple referrals and multidimensional cheap talk," *Econometrica*, 70(4), 1379–1401.

Bergemann, D., and A. Bonatti (2015): "Selling Cookies," *American Economic Journal: Microeconomics*, 7(3), 259–294.

Bhattacharya, S., and A. Mukherjee (2013): "Strategic Information Revelation when Experts Compete to Influence," *RAND Journal of Economics*, 44(3), 522–544.

Board, S., and M. Pycia (2014): "Outside options and the failure of the Coase conjecture," *The American Economic Review*, 104(2), 656–671.

Burdett, K., and K. L. Judd (1983): "Equilibrium price dispersion," *Econometrica: Journal of the Econometric Society*, pp. 955–969.

Çinlar, E. (2011): *Probability and Stochastics*. Springer.

Crawford, V. P., and J. Sobel (1982): "Strategic information transmission," *Econometrica*, 50(6), 1431–1451.

Diamond, P. A. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3(2), 156–168.

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

Forand, J. G. (2013): "Competing through Information Provision," *International Journal of Industrial Organization*, 31(5), 438–451.

Gentzkow, M., and E. Kamenica (2012): "Competition in Persuasion," Discussion paper, University of Chicago.

——— (2015): "Information Environments and the Impact of Competition on Information Provision," Discussion paper, University of Chicago.

Hoffmann, F., R. Inderst, and M. Ottaviani (2014): "Persuasion through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation," Discussion paper, Bocconi.

Kamenica, E., and M. Gentzkow (2011): “Bayesian Persuasion,” *American Economic Review*, 101(6), 2590–2615.

Kolotilin, A., M. Li, T. Mylovanov, and A. Zapechelnyuk (2015): “Persuasion of a Privately Informed Receiver,” Working paper.

Krishna, V., and J. Morgan (2001): “A Model of Expertise,” *Quarterly Journal of Economics*, 116(2), 747–775.

Li, F., and P. Norman (2014): “On Bayesian Persuasion with Multiple Senders,” Discussion paper, UNC.

Milgrom, P., and J. Roberts (1986): “Relying on the Information of Interested Parties,” *The RAND Journal of Economics*, pp. 18–32.

Rayo, L., and I. Segal (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118(5), 949–987.

Rhodes, A. (2014): “Multiproduct Retailing,” *Review of Economic Studies*, forthcoming.

Stigler, G. J. (1961): “The Economics of Information,” *Journal of Political Economy*, 69(3), 213–225.

Wolinsky, A. (1986): “True Monopolistic Competition as a Result of Imperfect Information,” *Quarterly Journal of Economics*, 101(3), 493–511.

Zhou, J. (2014): “Multiproduct search and the joint search effect,” *The American Economic Review*, 104(9), 2918–2939.

Appendix A: Public Beliefs

In this appendix, we present proofs for the case when beliefs are public. Appendix A1 characterizes different properties of absorbing beliefs, i.e., the set of beliefs where the seller provides no information. Appendix A2 covers results on local optimality and how it relates to equilibrium policies. Finally, Appendix A3 covers our main theorem on the uniqueness of the monopoly policy equilibrium with dispersed products.

First, we introduce some notation. Let $\mathcal{F}$ be the Borel algebra on ΔS and Π be the set of all Borel probability measures on ΔS. Formally, an *acceptance set* is a Borel set $Q \subset \Delta S$. We say a distribution over beliefs $\mu \in \Pi$ is *consistent* with the belief $p \in \Delta S$ iff it satisfies the martingale property $\int_{\Delta S} q \; \mu(dq) = p$. Hence, a *disclosure policy* is a Markov kernel $K : \Delta S \to \Pi$ such that $K(p)$ is consistent with p for all $p \in \Delta S$. Let $\mathcal{K}$ be the set of all disclosure policies. Recall that a disclosure policy K is *optimal given* some acceptance set Q if the firm's expected profit inside the set Q is greater under K than any other $L \in \mathcal{K}$ and it satisfies the *signal tie-breaking rule* (i.e. the seller provides no information when she is indifferent).

A1. Properties of Absorbing Beliefs

In this section, we present three useful lemmas describing properties of absorbing beliefs. Recall that given some disclosure policy K, the *absorbing set* A_K is the set of beliefs that provides no information under K, i.e. $K(p) = \delta_p$. Also define the *degenerate set* as the set of beliefs where the buyer knows the state for sure, i.e. $p(s) = 1$ for some $s \in S$.

Our first lemma shows that all degenerate beliefs are absorbing and the set of absorbing beliefs is Borel measurable.

Lemma (A1.1)**.** *For all $K \in \mathcal{K}$, $D \subset A_K \in \mathcal{F}$.*

Proof. We first show that A_K is a Borel set. Note that $K(p) = \delta_p$ iff $K(p, \{p\}) = 1$. Hence,

$$A_K = \{p \in \Delta S \ \mid K(p, \{p\}) = 1\} = \left\{ p \in \Delta S \;\middle|\; \int_{\Delta S} \mathbf{1}_{\{p\}}(q)\, K(p, dq) = 1 \right\}.$$

By Proposition I.6.9 of (Çinlar, 2011), A_K is $\mathcal{F}$-measurable as desired.

To show that $D \subset A_K$, let $p = \delta_s \in D$ for some $s \in S$. Since $K(p)$ is consistent with p,

$$\int_{\Delta S} q(s)\, K(p, dq) = p(s) = 1$$

Since $q(s) \leq 1$, this implies $q(s) = 1$ $K(p)$-a.s.. Hence, $p = q$ $K(p)$-a.s. so $K(p, \{p\}) = 1$ and $p \in A_K$. Hence, $D \subset A_K$ as desired. □

The next lemma list several useful properties of absorbing beliefs under optimal disclosure policies. First, optimal disclosure policies only send buyers to posterior beliefs that are absorbing. Second, if a policy is optimal given Q, then it is also optimal given any set that is sandwiched between its absorbing set and Q. Third, we can characterize the closure of any absorbing set by the support of its disclosure policy. Recall that the *support* of K (denoted by $\text{supp}\,(K)$) is the largest set $F \in \mathcal{F}$ such that for any $q \in F$ and $\varepsilon > 0$, there is some $p \in \Delta S$ such that $K\,(p, B_\varepsilon\,(q)) > 0$, where $B_\varepsilon\,(q)$ is the ε-ball around q.[16]

Lemma (A1.2)**.** *Let $K \in \mathcal{K}$ be optimal given $Q \in \mathcal{F}$. Then*

(a) $K\,(p, A_K) = 1$ *for all* $p \in \Delta S$.

(b) $A_K \subset F \subset Q$ *implies* K *is also optimal given* $F \in \mathcal{F}$.

(c) $\text{cl}\,(A_K) = \text{supp}\,(K)$

Proof. Let $K \in \mathcal{K}$ be optimal given $Q \in \mathcal{F}$. We prove the lemma in order.

(a) First, define $G \in \mathcal{F}$ to be the set of beliefs such that using K gives the seller a strictly higher profit than providing no information. Formally,

$$G := \left\{ q \in \Delta S \;\middle|\; \int_Q \pi\,(r)\,K\,(q, dr) > \int_Q \pi\,(r)\,\delta_q\,(dr) \right\}$$

Our first step is to show that K will never send the buyer to a posterior in G, i.e. $K\,(p, G) = 0$ for all $p \in \Delta S$. Once we have this, then the signal tie-breaking rule will ensure that K will always send the buyer to a posterior in A_K.

We prove this by contradiction. Suppose there is some $p \in \Delta S$ such that $K\,(p, G) > 0$. For notational convenience, set $\mu := K\,(p)$ and define the measure ν that coincides with μ outside of G but provides one more iteration of information according to K inside G. Formally,

$$\nu\,(F) := \mu\,(F \cap (\Delta S \backslash G)) + \int_G K\,(q, F)\,\mu\,(dq)$$

for all $F \in \mathcal{F}$. We show that using using ν is strictly more profitable for the seller so K is not optimal. First, we show that ν is consistent with p. In order to see this, note that

$$\begin{aligned}\int_{\Delta S} q\;\nu\,(dq) &= \int_{\Delta S \backslash G} q\;\mu\,(dq) + \int_G \left(\int_{\Delta S} r\;K\,(q, dr) \right) \mu\,(dq) \\ &= \int_{\Delta S \backslash G} q\;\mu\,(dq) + \int_G q\;\mu\,(dq) = \int_{\Delta S} q\;\mu\,(dq) = p\end{aligned}$$

where the second and last equalities all following from the fact that $K\,(q)$ is consistent with

[16] Formally, $B_\varepsilon\,(q) := \{p \in \Delta S \;|\; |p - q| \leq \varepsilon\}$.

q for all $q \in \Delta S$. However, the seller's profit for using ν is

$$\int_Q \pi(q)\,\nu(dq) = \int_{Q\cap(\Delta S\backslash G)} \pi(q)\,\mu(dq) + \int_G \left(\int_Q \pi(r)\,K(q,dr)\right)\mu(dq)$$
$$> \int_{Q\cap(\Delta S\backslash G)} \pi(q)\,\mu(dq) + \int_G \left(\int_Q \pi(r)\,\delta_q(dr)\right)\mu(dq) = \int_Q \pi(q)\,\mu(dq)$$

where the strict inequality follows from the definition of G and the fact that $K(p,G) > 0$. This contradicts the optimality of $K(p) = \mu$.

Hence, $K(p,G) = 0$ and since we know that the seller's profit from K cannot be strictly less than providing no information, it must be that for all $p \in \Delta S$,

$$K_p\left\{q \in \Delta S \;\middle|\; \int_Q \pi(r)\,K(q,dr) = \int_Q \pi(r)\,\delta_q(dr)\right\} = 1$$

Thus, K is always going to send the buyer to posteriors where the seller is indifferent between the policy and providing no information. By the signal tie-breaking rule, this implies $K(q) = \delta_q$ K_p-a.s. so $K(p, A_K) = 1$ for all $p \in \Delta S$.

(b) Suppose $A_K \subset F \subset Q$. We want to prove that K is also optimal given F. Since $K(p, A_K) = 1$ from part (a) and K is optimal given Q,

$$\int_F \pi(q)\,K(p,dq) = \int_{A_K} \pi(q)\,K(p,dq) = \int_Q \pi(q)\,K(p,dq)$$
$$\geq \int_Q \pi(q)\,L(p,dq) \geq \int_F \pi(q)\,L(p,dq)$$

for any $p \in \Delta S$ and $L \in \mathcal{K}$. We also need to check the tie-breaking rule. Suppose $\int_F \pi(q)\,K(p,dq) = \int_F \pi(q)\,\delta_p(dq)$ so from the above inequalities, we have

$$\int_Q \pi(q)\,K(p,dq) = \int_Q \pi(q)\,\delta_p(dq)$$

Since K is optimal given Q, $K(p) = \delta_p$ follows from the tie-breaking rule under Q. Hence, K is also optimal given F as desired.

(c) We first show that $\mathrm{cl}(A_K) \subset \mathrm{supp}(K)$. Let $q \in \mathrm{cl}(A_K)$ and consider some neighborhood $B_\varepsilon(q)$ for $\varepsilon > 0$. Since $q \in \mathrm{cl}(A_K)$, we can find some $p \in B_\varepsilon(q) \cap A_K$. By the definition of A_K, we know $K(p,\{p\}) = 1$ and since $p \in B_\varepsilon(q)$, $K(p, B_\varepsilon(q)) = 1 > 0$. Since this is true for any neighborhood around any $q \in \mathrm{cl}(A_K)$, $\mathrm{cl}(A_K) \subset \mathrm{supp}(K)$.

We now show that $\mathrm{supp}(K) \subset \mathrm{cl}(A_K)$. Consider $q \in \mathrm{supp}(K)$ and suppose $q \notin \mathrm{cl}(A_K)$. We show that this must imply a contradiction. Find some $\varepsilon > 0$ such that $B_\varepsilon(q) \cap \mathrm{cl}(A_K) = \emptyset$. Since $q \in \mathrm{supp}(K)$, by the definition of support, it must be that there is some $p \in \Delta S$ where $K(p, B_\varepsilon(q)) > 0$. However, from part (a), we know that $K(p, A_K) = 1$ for all $p \in \Delta S$

implying $K(p, B_\varepsilon(q)) = 0$ a contradiction. Hence, $\text{supp}(K) \subset \text{cl}(A_K)$ so $\text{supp}(K) = \text{cl}(A_K)$ as desired.

□

The last lemma of this section shows that if the disclosure policy is continuous, then its set of absorbing beliefs is exactly its support.

Lemma (A1.3)**.** If $K \in \mathcal{K}$ is continuous, then $A_K = \text{supp}(K)$.

Proof. Let $K \in \mathcal{K}$ be continuous. We show that A_K is closed so $\text{cl}(A_K) = A_K$ and the result follows immediately from Lemma A1.2(c). Consider a sequence $p_i \to p$ such that $p_i \in A_K$. If we can show that $p \in A_K$, then A_K is closed and we are done. Note that $K(p_i) = \delta_{p_i}$ for all p_i. We show that $K(p) = \delta_p$. Since K is continuous, $K(p_i) \to K(p)$ weakly so for any continuous function ξ,

$$\int_{\Delta S} \xi(q) K(p, dq) = \lim_i \int_{\Delta S} \xi(q) K(p_i, dq) = \lim_i \int_{\Delta S} \xi(q) \delta_{p_i}(dq) = \lim_i \xi(p_i) = \xi(p) \qquad (4)$$

If we can replace ξ with $\mathbf{1}_p$, i.e. the indicator function on $\{p\}$, then we are done. However, $\mathbf{1}_p$ is not continuous, so we approximate it using sequence of continuous functions $\xi^j \to \mathbf{1}_{\{p\}}$ such that $\xi^j(p) = 1$ for all j. By dominated convergence,

$$K(p, \{p\}) = \int_{\Delta S} \mathbf{1}_{\{p\}}(q) K(p, dq) = \lim_j \int_{\Delta S} \xi^j(q) K(p, dq)$$

Since ξ^j is continuous, we can apply equation (4) to get

$$K(p, \{p\}) = \lim_j \xi^j(p) = 1$$

Hence $p \in A_K$ so A_K is closed as desired. □

A2. Results for Local Optimality

We now present results showing how equilibrium policies are characterized by local optimality. First, we introduce a couple of technical lemmas. For notational convenience, define the single-search payoff for a buyer with belief q under policy K as

$$W(K, q) := \int_{\Delta S} \left(\max_{u \in U} r \cdot u \right) K(q, dr)$$

Since optimal disclosure policies will only send buyers to absorbing beliefs (see Lemma A1.2(a)), the first lemma below shows that the buyer only searches once in equilibrium. This means that we have a very simple expression for the buyer's value function.

Lemma (A2.1)**.** *Let $K \in \mathcal{K}$ be optimal given $Q \in \mathcal{F}$. Then for any $c > 0$,*

$$V_c(K, q) = -c + W(K, q)$$

Proof. Fix some $c > 0$ and consider some $r \in A_K$ so

$$V_c(K, r) = -c + \max\left\{\max_{u \in U} r \cdot u,\ V_c(K, r)\right\}$$

Note that if $V_c(K, r) > \max_{u \in U} r \cdot u$, then $V_c(K, r) = -c + V_c(K, r)$ contradicting $c > 0$. Thus, $\max_{u \in U} r \cdot u \geq V_c(K, r)$ for all $r \in A_K$. From Lemma A1.2(a), we know that $K(q, A_K) = 1$ for all $q \in \Delta S$ so

$$\begin{aligned} V_c(K, q) &= -c + \int_{A_K} \max\left\{\max_{u \in U} r \cdot u,\ V_c(K, r)\right\} K(q, dr) \\ &= -c + \int_{\Delta S} \left(\max_{u \in U} r \cdot u\right) K(q, dr) = -c + W(K, q) \end{aligned}$$

as desired. □

The following lemma says that in equilibrium, if the buyer's value function is continuous on some set of absorbing beliefs, then that set must be contained in the acceptance set. In other words, in any equilibrium, we can always find a neighborhood around that set that is still contained inside the acceptance set. Put differently, the seller always has some additional room to deviate locally but chooses not to. This captures the spirit of local optimality and will be very useful in proving subsequent results.

Lemma (A2.2)**.** *Let (K, Q) be an equilibrium and $V_c(K, \cdot)$ be continuous on some closed set $F \subset A_K$. Then $B_\varepsilon(F) \subset Q$ for some $\varepsilon > 0$.*

Proof. Let (K, Q) be an equilibrium where $Q := Q_c(K)$. Consider a closed set $F \subset A_K$ such that $V_c(K, \cdot)$ is continuous on F. Define

$$\psi_c(r) := \max_{u \in U} r \cdot u - V_c(K, r)$$

so ψ_c is also continuous on F. Since F is closed, by Corollary 3.31 of Aliprantis and Border (2006), ψ_c is uniformly continuous on F. Now, for every $q \in F \subset A_K$, $\psi_c(q) = c > 0$. Hence, by uniform continuity, we can find some $\varepsilon > 0$ such that $q \in B_\varepsilon(F)$ implies $\psi_c(q) \geq \frac{c}{2} > 0$. Hence, $q \in Q_c(K) = Q$ so $B_\varepsilon(F) \subset Q$ as desired. □

We are now ready for our first main proposition in this section. Recall that a continuous policy K is *locally optimal* iff it is optimal given $B_\varepsilon(\text{supp}(K))$ for some $\varepsilon > 0$. Note that from Lemma A1.3, we know that $\text{supp}(K) = A_K$.

Proposition (A2.3)**.** *Any continuous equilibrium policy is locally optimal.*

Proof. Let $K \in \mathcal{K}$ be an equilibrium policy and $Q := Q_c(K)$. Since K is continuous, we know that $A_K = \text{supp}(K)$ and is closed from Lemma A1.3. Moreover, since support functions are continuous, $W(K, \cdot)$ is continuous. Since $V_c(K, \cdot) = -c + W(K, \cdot)$ by Lemma A2.1, the buyer's value function is also continuous. Hence, from Lemma A2.2, we can set $F = A_K$ and there is some $\varepsilon > 0$ such that $B_\varepsilon(A_K) \subset Q$. Note that $A_K \subset Q$ follows easily from Lemma A2.1. Thus, $A_K \subset B_\varepsilon(A_K) \subset Q$ so from Lemma A1.2(b), K is optimal given $B_\varepsilon(A_K)$ proving local optimality. □

Before we proceed to the final proposition in this section, we first present a technical lemma that will be useful in the proof of the proposition. It says that if the buyer is indifferent between staying and searching under zero search cost, then the seller's profit is bounded below by her profit if the buyer stays.

Lemma (A2.4)**.** *Let K be locally optimal and $p \in \Delta S$ be such that $p \cdot u^*(p) = W(K, p)$. Then*

$$\int_{\Delta S} \pi(q) K(p, dq) \geq \pi(p)$$

Proof. Let K be locally optimal and $v := u^*(p)$. Since support functions are convex,

$$p \cdot v = W(K, p) = \int_{\Delta S} \left(\max_{u \in U} q \cdot u \right) K(p, dq)$$

implies that K_p must place all its mass on the affine part of the support function $\max_{u \in U} q \cdot u$. In other words,

$$K_p \left\{ q \in \Delta S \;\middle|\; q \cdot v = \max_{u \in U} q \cdot u \right\} = 1$$

Moreover, recall that $K_p(A_K) = 1$ from Lemma A1.2(a). Consider some $q \in A_K$ and $q \cdot v = \max_{u \in U} q \cdot u$. Since ties in the buyer's actions are resolved in the seller's favor and K is optimal given $B_\varepsilon(A_K)$, it must be that $\pi(q) \geq \pi(p)$. Hence

$$1 = K_p \left\{ q \in A_K \;\middle|\; q \cdot v = \max_{u \in U} q \cdot u \right\} \leq K_p \{ q \in A_K \mid \pi(q) \geq \pi(p) \}$$

The desired result then follows immediately. □

We now show the converse of the proposition above. Any locally optimal policy is an equilibrium policy for some small enough search cost.

Proposition (A2.5)**.** *Any locally optimal policy is an equilibrium policy for some $c > 0$.*

Proof. Suppose $K \in \mathcal{K}$ is locally optimal so there is some $\varepsilon > 0$ such that K is optimal given $B_\varepsilon(A_K)$. Note that K is also continuous by definition. Also, note that from Lemma A2.1, for any

$c > 0$,

$$Q_c(K) = \left\{ p \in \Delta S \;\middle|\; \max_{u \in U} p \cdot u \geq -c + W(K, p) \right\}$$

The main argument rests on showing that there is some $c > 0$, such that K is optimal given $B_\varepsilon(A_K) \cup Q_c(K)$. In other words, we can find some $c > 0$ small enough such that the seller will not have a strict incentive to send the buyer to a posterior somewhere in $B_\varepsilon(A_K) \cup Q_c(K)$ but outside A_K. Once we have this and since $A_K \subset Q_c(K) \subset B_\varepsilon(A_K) \cup Q_c(K)$, Lemma A1.2(b) implies that K is optimal given $Q_c(K)$ demonstrating that K is an equilibrium policy. To simplify notation, define $B := B_\varepsilon(A_K)$, $Q_c := Q_c(K)$ and

$$Q_0 := \left\{ p \in \Delta S \;\middle|\; \max_{u \in U} p \cdot u \geq W(K, p) \right\}$$

Finally, define $\hat{Q}_c := Q_c \backslash B$ for all $c \geq 0$.

We prove that K is optimal given $B \cup Q_c$ by contradiction. Consider some small $c > 0$ and $p \in \Delta S$ and suppose there is some measure μ consistent with p such that

$$\int_{B \cup Q_c} \pi(q)\, \mu(dq) > \int_{B \cup Q_c} \pi(q)\, K(p, dq)$$

We can assume $\mu(B \cup Q_c) = 1$ without loss of generality. Note that $\mu\left(\hat{Q}_c\right) = \mu(Q_c/B) > 0$ since otherwise, μ only sends the buyer to posteriors in B and the optimality of K given B would contradict the strict inequality. Now, consider the measure ν that coincides with μ in B but provides one more iteration of information according to K in $\hat{Q}_c$. Formally,

$$\nu(F) := \mu(F \cap B) + \int_{\hat{Q}_c} K(q, F)\, \mu(dq)$$

for all $F \in \mathcal{F}$. Note that since μ is consistent with p, ν is also consistent with p by the same argument as in Lemma A1.2(a). Define $Q_c^* := \hat{Q}_c \backslash \hat{Q}_0$ so we can partition $\hat{Q}_c = \hat{Q}_0 \cup Q_c^*$. We now have

$$\begin{aligned}
\int_{B \cup Q_c} \pi(q)\, \nu(dq) &= \int_B \pi(q)\, \mu(dq) + \int_{\hat{Q}_c} \left(\int_{B \cup Q_c} \pi(r)\, K(q, dr) \right) \mu(dq) \\
&= \int_B \pi(q)\, \nu(dq) + \int_{\hat{Q}_0} \left(\int_{\Delta S} \pi(r)\, K(q, dr) \right) \mu(dq) \\
&\quad + \int_{Q_c^*} \left(\int_{\Delta S} \pi(r)\, K(q, dr) \right) \mu(dq)
\end{aligned} \tag{5}$$

Now, $q \in \hat{Q}_0 = Q_0 \backslash B$ implies $\max_{u \in U} q \cdot u \geq W(K, q)$ so by Lemma A2.4, $\int_{\Delta S} \pi(r)\, K(q, dr) \geq \pi(q)$. Moreover, by the continuity of K, we can choose c small enough such that $\int_{\Delta S} \pi(r)\, K(q, dr) \geq$

$\pi(q)$ for all $q \in Q_c^*$. Substituting these two inequalities into equation (5) yields

$$\int_{B\cup Q_c} \pi(q)\,\nu(dq) \geq \int_B \pi(q)\,\mu(dq) + \int_{\hat{Q}_0} \pi(q)\,\nu(dq) + \int_{Q_c^*} \pi(q)\,\mu(dq)$$
$$\geq \int_B \pi(q)\,\mu(dq) + \int_{\hat{Q}_c} \pi(q)\,\mu(dq) = \int_{B\cup Q_c} \pi(q)\,\mu(dq)$$

However, since $\nu(B) = 1$, by the optimality of K given B, we have

$$\int_{B\cup Q_c} \pi(q)\,K(p,dq) = \int_B \pi(q)\,K(p,dq) \geq \int_B \pi(q)\,\nu(dq)$$
$$\geq \int_{B\cup Q_c} \pi(q)\,\nu(dq) \geq \int_{B\cup Q_c} \pi(q)\,\mu(dq)$$

yielding a contradiction. □

A3. Uniqueness of the Monopoly Equilibrium Policy

In this section, we show our main result that the monopoly policy is the unique equilibrium policy with dispersed products. First, we prove a useful lemma that says that if a monopoly policy's absorbing beliefs is contained in an equilibrium policy's absorbing beliefs, then the equilibrium policy must be a monopoly policy. The reason is that if the seller can send the buyer to all the posterior beliefs that the monopoly policy uses, then the seller can just use the monopoly to get her maximal profit.

Lemma (A3.1)**.** *Let K be an equilibrium policy and K^* be a monopoly policy. If $A_{K^*} \subset A_K$, then K is a monopoly policy.*

Proof. Let $K \in \mathcal{K}$ be an equilibrium policy and $Q := Q_c(K)$. Let K^* be a monopoly policy. Suppose $A_{K^*} \subset A_K \subset Q$ where the last set inclusion follows from Lemma A2.1. Hence, $A_{K^*} \subset Q \subset \Delta S$ so by Lemma A1.2(b), K^* is also optimal given Q. Thus,

$$\int_{\Delta S} \pi(q)\,K(p,dq) = \int_Q \pi(q)\,K(p,dq) = \int_Q \pi(q)\,K^*(p,dq) = \int_{\Delta S} \pi(q)\,K^*(p,dq)$$
$$\geq \int_{\Delta S} \pi(q)\,L(p,dq)$$

for any $L \in \mathcal{K}$. For the tie-breaking rule, suppose the above is satisfied with equality for $L(p) = \delta_p$. Hence, $K^*(p) = \delta_p$ so $K(p) = \delta_p$ by the tie-breaking rule from the optimality of K. Hence, K is a monopoly policy as desired. □

Recall that products are *dispersed* iff $q \cdot u \geq 0$ implies $q \cdot v \leq 0$ for any $\{u, v\} \subset U$. We first fully characterize the monopoly policy under dispersed products. We then prove that Lemma A3.1 applies so that any equilibrium policy must be an monopoly policy.

First, we introduce some notation. For each $u \in U$, let $\bar{H}_u \subset \Delta S$ be the set of beliefs where choosing u is optimal for the buyer (under no information). Formally,

$$\bar{H}_u := \{p \in \Delta S \mid p \cdot u \geq p \cdot v \text{ for all } v \in U\}$$

Also, let $H_u \subset \bar{H}_u$ is the set of beliefs where the buyer chooses u given our assumption that ties in the buyer's actions are resolved in the seller's favor. Formally,

$$H_u := \left\{ p \in \bar{H}_u \;\middle|\; u \in \arg\max_{v \in u^*(p)} \tilde{\pi}(v) \right\}$$

Note that each H_u is convex and $\bigcup_{u \in U} H_u$ forms a partition of the belief space ΔS.

Let us order the set of products $U = \{u_k, u_{k-1}, \ldots, u_1, 0\}$ such that $\tilde{\pi}(u_i) > \tilde{\pi}(u_{i-1})$ for all $i \in \{k, \ldots, 1\}$ and $u_0 = 0$. Let $T_i \subset S$ denote the set of states such that u_i is optimal. Formally,

$$T_i := \{s \in S \mid \delta_s \in H_{u_i}\}$$

For $i \in \{k, \ldots, 0\}$, let Δ_i^U denote the sub-simplex such that no product with profit greater than that of u_i is chosen. Formally,

$$\Delta_i^U := \Delta\left(S \backslash \left(\bigcup_{j>i} T_j\right)\right)$$

where ΔT denotes the probability simplex for event $T \subset S$. In other words, u_i is the most profitable product for the seller in each region Δ_i^U. Finally, let

$$E_i^U := \Delta_i^U \cap H_{u_i}$$

denote the region in Δ_i^U where u_i is chosen.

We say a set of products U is *regular* iff for any $p \in H_{u_j} \cap \left(\Delta_i^U \backslash \Delta_{i-1}^U\right)$ where $j < i$, there is some $p' \in H_{u_j}$ and $q \in H_{u_i}$ such that $p = aq + (1-a)p'$ for some $a \in (0,1)$. We will show that regularity is one of the sufficient conditions for the monopoly policy to be a unique equilibrium policy. First, we show that with regular products, the absorbing beliefs of any monopoly policy are contained in the E_i^U sets. This provides us with a characterization of the monopoly policy.

Lemma (A3.2). *Let U be regular and K^* be a monopoly policy. Then $A_{K^*} \subset \bigcup_i E_i^U$.*

Proof. We prove this by induction on the number of products. If $U = \{u_0\}$, then $E_0^U = \Delta S$ and the result follows trivially. Hence, assume the conclusion holds for at least $k-1$ products and consider a set of products $U = \{u_k, \ldots, u_1, 0\}$. Let $p \in A_{K^*}$ and we wish to prove that $p \in \bigcup_i E_i^U$. We consider three cases:

(i) Suppose $p \in H_{u_k}$. Then $\Delta_k^U = \Delta S$ and $p \in E_k^U$ by the definition of E_k^U.

(ii) Suppose $p \in \Delta_{k-1}^U$. Define the set of products $U' := U \backslash \{u_k\}$ so U' has $k-1$ products.

Note that since U and U' are the same on Δ_{k-1}^U, K^* is also a monopoly policy over Δ_{k-1}^U under U'. Hence, by the induction assumption, $p \in \bigcup_i E_i^{U'}$. Again, since U and U' coincide on Δ_{k-1}^U,

$$p \in \Delta_{k-1}^U \cap \left(\bigcup_i E_i^{U'} \right) = \Delta_{k-1}^U \cap \left(\bigcup_i E_i^U \right)$$

so $p \in \bigcup_i E_i^U$ as desired.

(iii) Finally, suppose $p \notin H_{u_k}$ and $p \notin \Delta_{k-1}^U$. Since $(H_u)_{u \in U}$ forms a partition of ΔS, let $p \in H_u$ for some $u \neq u_k$. Since $p \in A_{K^*}$, $K^*(p) = \delta_p$ so

$$\int_{\Delta S} \pi(q) K^*(p, dq) = \pi(p) = \tilde{\pi}(u)$$

We now use regularity. Since $p \in \Delta_k^U \backslash \Delta_{k-1}^U$ and U is regular, we can find some $q \in H_{u_k}$ and $p' \in H_u$ such that $p = aq + (1-a)p'$ for some $a \in (0,1)$. Suppose the seller releases some additional information that puts a mass on posterior q and $1-a$ mass on posterior p'. Formally, if we let $\nu := a\delta_q + (1-a)\delta_r$, then

$$\int_{\Delta S} \pi(q) \nu(dq) = a\tilde{\pi}(u_k) + (1-a)\tilde{\pi}(u) > \tilde{\pi}(u)$$

Note that ν is consistent with p but provides the seller with a strictly higher profit than no information, thus contradicting the optimality of K^*. Hence, this third case cannot happen so $A_{K^*} \subset \bigcup_i E_i^U$ as desired. □

Theorem (A3.3)**.** *Let U be regular and K be an equilibrium policy such that $p \in H_u$ implies $K(p, H_u) > 0$. Then K is a monopoly policy.*

Proof. Let (K, Q) be an equilibrium. The key is to prove that $\bigcup_i E_i^U \subset A_K$. Once we have this, then Lemma A3.2 gives us that

$$A_{K^*} \subset \bigcup_i E_i^U \subset A_K$$

We can then use Lemma A3.1 to prove that K is a monopoly policy as desired.

As in the proof of Lemma A3.2, we prove that $\bigcup_i E_i^U \subset A_K$ by induction on the number of products. If $U = \{u_0\}$, then $E_0^U = \Delta S$ and the result follows trivially. Hence, assume the conclusion holds for at least $k-1$ products and consider a set of products $U = \{u_k, \ldots, u_1, 0\}$. Let $p \in \bigcup_i E_i^U$ and we wish to prove that $p \in A_K$.

First, suppose $p \in E_i^U$ for some $i < k$. Define the set of products $U' := U \backslash \{u_k\}$ so U' has $k-1$ products. Note that since U and U' are the same on Δ_{k-1}^U, K is also an equilibrium policy over Δ_{k-1}^U under U'. Since $p \in E_i^U \subset \Delta_{k-1}^U$ and E_i^U coincides with $E_i^{U'}$ on Δ_{k-1}^U, it must be that $p \in E_i^{U'}$. By the induction assumption, $p \in A_K$ as desired.

Now, suppose $p \in E_k^U = H_{u_k}$. Let $a := K_p(H_{u_k}) = K_p(A_K \cap H_{u_k})$ as $K_p(A_K) = 1$ from

Lemma A1.2(a). By the premise, $a > 0$. We will show that $a = 1$. In order to do this, we will make use of Lemma A2.2 which requires that $V_c(K, \cdot)$ be continuous on $A_K \cap H_{u_k}$. We now proceed to prove this latter claim. Let $q \in A_K \cap H_{u_k}$ and consider a sequence $q_i \to q$. Since K is optimal given Q, by the Theorem of the Maximum, $\int_Q \pi(r) K(\cdot, dr)$ is continuous. Hence, as $q_i \to q$,

$$\int_Q \pi(r) K(q_i, dr) \to \int_Q \pi(r) K(q, dr) = \tilde{\pi}(u_k)$$

Hence, $K(q_i, H_{u_k}) \to 1$. In other words, $K(q_i)$ will eventually put all its mass in H_{u_k} in the limit. From Lemma A2.1, we know that $V_c(K, q) = -c + W(K, q)$. Now,

$$W(K, q_i) = \int_{\Delta S} \left(\sup_{u \in U} r \cdot u \right) K(q_i, dr) \to \int_{H_{u_k}} \left(\sup_{u \in U} r \cdot u \right) K(q, dr) = q \cdot u_k$$

so $V_c(K, q_i) \to -c + q \cdot u_k = V_c(K, q)$. Hence, $V_c(K, \cdot)$ be continuous on $A_K \cap H_{u_k}$. Note that by a similar argument as in the proof of Lemma A1.3, the set $A_K \cap H_{u_k}$ is also closed.

The argument above implies that we can use Lemma A2.2 so that we can find some $\varepsilon > 0$ where $B_\varepsilon(A_K \cap H_{u_k}) \subset Q$. In order to show that $a = 1$, we will use this fact to show that whenever $a < 1$, the seller has a strict incentive to deviate from the equilibrium policy. We construct this deviation as follows. Since $a > 0$, we can split $K(p)$ into two measures μ_1 and μ_2 on H_{u_k} and $\Delta S \backslash H_{u_k}$ respectively. Formally,

$$K(p) = a\mu_1 + (1-a)\mu_2$$

where $\mu_1(H_{u_k}) = 1$ and $\mu_2(H_{u_k}) = 0$. Now, consider the policy that takes all beliefs in the support of μ_1 and moves them all slightly closer to p. If we choose our weights correctly, then the support of this "perturbed" policy μ_ε will be completely contained in $B_\varepsilon(A_K \cap H_{u_k}) \subset Q$. Formally, define the mapping $\gamma_\varepsilon(q) := \varepsilon p + (1-\varepsilon) q$, $a_\varepsilon := \frac{a}{1-(1-a)\varepsilon} > a$ and

$$\mu_\varepsilon := a_\varepsilon \left(\mu_1 \circ \gamma_\varepsilon^{-1} \right) + (1 - a_\varepsilon) \mu_2$$

First, we check that μ_ε is consistent with p. Note that

$$\begin{aligned}
\int_{\Delta S} q \ \mu_\varepsilon(dq) &= a_\varepsilon \int_{\Delta S} \gamma_\varepsilon(q) \mu_1(dq) + (1-a_\varepsilon) \int_{\Delta S} q \ \mu_2(dq) \\
&= a_\varepsilon \left(\varepsilon p + (1-\varepsilon) \int_{\Delta S} q \ \mu_1(dq) \right) + (1-a_\varepsilon) \int_{\Delta S} q \ \mu_2(dq) \\
&= a_\varepsilon \varepsilon p + \frac{(1-\varepsilon)}{1-(1-a)\varepsilon} p = p
\end{aligned}$$

so μ_ε is consistent with p as desired. However, since $\text{supp}(\mu_\varepsilon) \subset B_\varepsilon(A_K \cap H_{u_k}) \subset Q$ and $a_\varepsilon > a$

$$\begin{aligned}
\int_Q \pi(q)\,\nu_\varepsilon(dq) &= \int_{\Delta S} \pi(q)\,\nu_\varepsilon(dq) \\
&= a_\varepsilon \mu_1\left(\gamma_\varepsilon^{-1}(H_{u_k})\right)\tilde{\pi}(u_k) + (1-a_\varepsilon)\int_{\Delta S \backslash H_{u_k}} \pi(q)\,\mu_2(dq) \\
&> a\mu_1(H_{u_k})\,\tilde{\pi}(u_k) + (1-a)\int_{\Delta S \backslash H_{u_k}} \pi(q)\,\mu_2(dq) \\
&> \int_{\Delta S} \pi(q)\,K(p,dq) = \int_Q \pi(q)\,K(p,dq)
\end{aligned}$$

contradicting the optimality of K. Hence, it must be that $a = 1$.

For the final step, since $a = K_p(H_{u_k}) = 1$, from Lemma A2.1,

$$\begin{aligned}
V_c(K,p) &= -c + \int_{\Delta S} \max_{u \in U}(q \cdot u)\,K(p,dq) \\
&= -c + p \cdot u_k < p \cdot u_k
\end{aligned}$$

implying $p \in Q$. Hence, by the signal tie-breaking rule, $K(p) = \delta_p$ so $p \in A_K$ as desired. This concludes the proof. □

The next two corollaries are straightforward applications of Theorem A3.3. They involve checking that the assumptions of the theorem hold so that any equilibrium policy is a monopoly policy when products are dispersed or opposed. Recall that two products are *opposed* iff $u(s) \geq 0$ implies $v(s) \leq 0$ for any $\{u, v\} \subset U$.

Corollary (A3.4)**.** *If products are dispersed, then any equilibrium policy is a monopoly policy.*

Proof. Let U be dispersed. We check that the two assumptions of Theorem A3.3 hold. We first check regularity. Let $p \in H_{u_j} \backslash \Delta_{i-1}^U$ where $j > i$. Since U is dispersed, $\bar{H}_{u_j} = \{q \in \Delta S \mid q \cdot u_j \geq 0\}$. Consider $s \in T_i$ so $\delta_s \notin H_{u_i}$. By a standard separating hyperplane argument, we can find some $p' \in H_{u_j}$ such that $p = a\delta_s + (1-a)p'$ where $a \in (0,1)$. Hence, U is regular.

We now check the second condition. Let $p \in H_u$ and suppose $K(p, H_u) = 0$. Since $\Delta S \backslash H_u$ is convex, this implies $p \notin H_u$ yielding a contradiction. Hence, both conditions of Theorem A3.3 hold and the result follows. □

Corollary (A3.5)**.** *If products are opposed, then any equilibrium policy is a monopoly policy.*

Proof. Let $U = \{u_2, u_1, 0\}$ be opposed. We first check for regularity. Consider $p \in H_{u_1}$ and let $\delta_s \in H_{u_2}$ where $u_2(s) \geq 0$ and $u_1(s) \leq 0$. We can then find some $p' \in \Delta S \backslash \Delta S_1^U$ close to p such that $p = a\delta_s + (1-a)p'$ for some $a \in (0,1)$. Since $\delta_s \notin H_{u_1}$ and $p \in H_{u_1}$, it must be that $p' \in H_{u_1}$ so U is regular. Note that since there are only two products, $K(p, H_{u_i}) > 0$ for all $p \in H_{u_i}$ and $i \in \{1,2\}$. Hence, both conditions of Theorem A3.3 hold and the result follows. □

Appendix B: Private Beliefs

In this appendix, we present proofs for the case when beliefs are private. Appendix B1 provides some preliminary results on the buyer's value function. Appendix B2 characterizes conditions under which an equilibrium exists. Finally, Appendix B3 covers our main results on limit equilibria as search costs vanish.

First, we introduce some notation. Recall that $p \in \text{int}(\Delta S)$ is the initial prior and Π is the set of all Borel probability measures on ΔS. Let $M \subset \Pi$ be the set of all *signal policies*, i.e., the set of all signal policies $\mu \in M$ consistent with p. Recall that $\phi_q(r)$ is the posterior of a buyer with prior q given that a buyer with prior p has posterior r. More formally, given any signal policy $\mu \in M$, $\phi_q(r)$ must satisfy

$$\mu_q(dr)\left[\phi_q(r)\right](s) = \mu(dr)\, q(s)\, \frac{r(s)}{p(s)} \tag{6}$$

where μ_q be the distribution of signals from the perspective of a buyer with prior q.[17] The expression on the left of equation (6) is the joint distribution over signals and states. Since all buyers observe the same independent signal structure, the expression on the right is the same joint distribution but from the perspective of buyer p after applying Bayes' rule. Recall that K_μ is the Markov kernel induced by μ where $K_\mu(q) = \mu_q \circ \phi_q^{-1}$ is the distribution of posteriors for buyer q. Also recall that $V_c(\mu, \cdot)$ is the buyer's continuation utility. Finally, as in the public beliefs section, define the single-search buyer payoff as

$$W(\mu, q) := \int_{\Delta S} \left(\max_{u \in U} r \cdot u \right) K_\mu(q, dr)$$

B1. Preliminary Results

In this section, we introduce some preliminary results that will be useful in the subsequent analysis. The first lemma shows that the induced Markov kernel K_μ is in fact a continuous disclosure policy as defined in the section on public beliefs. The remaining two lemmas in this section demonstrate properties of the buyer's value function.

Lemma (B1.1)**.** *$K_\mu \in \mathcal{K}$ and is continuous for all $\mu \in M$.*

Proof. We first show that $K_\mu(q) = \mu_q \circ \phi_q^{-1}$ is continuous. First, note that by summing up both

[17] Note that this is well-defined since $p \in \text{int}(\Delta S)$.

sides of equation (6) over S, we have the following expression for μ_q

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

Note that this expression allows us to link the signal policy μ to the distribution of signals from the perspective of buyer q. To show that K_μ is continuous, we test weak convergence using a bounded continuous function $\xi : \Delta S \to \mathbb{R}$. Substituting the above expression for μ_q into the definition of K_μ, we have

$$\begin{aligned}\int_{\Delta S} \xi(r) K_\mu(q, dr) &= \int_{\Delta S} \xi(\phi_q(r)) \mu_q(dr) \\ &= \int_{\Delta S} \xi(\phi_q(r)) \sum_s q(s) \frac{r(s)}{p(s)} \mu(dr) \qquad (7)\end{aligned}$$

Since ϕ_q is continuous, this is clearly continuous in q. Hence, K_μ is continuous.

We now show that K_μ is a disclosure policy. Following equation (7) above and setting $\xi(r) = r(s)$, we have that

$$\begin{aligned}\int_{\Delta S} r(s) K_\mu(q, dr) &= \int_{\Delta S} [\phi_q(r)](s) \sum_s q(s) \frac{r(s)}{p(s)} \mu(dr) \\ &= \int_{\Delta S} q(s) \frac{r(s)}{p(s)} \mu(dr) = q(s)\end{aligned}$$

where the last equality follows from the fact that μ is consistent with p. Hence, $K_\mu(q)$ is consistent with $q \in \Delta S$ so $K_\mu \in \mathcal{K}$. □

The following quick lemma provides a useful explicit expression for the single-search buyer payoff W. The proof simply involves using the expression for μ_q in equation (7) of the proof of Lemma B1.1.

Lemma (B1.2)**.** *For any $\mu \in M$,*

$$W(\mu, q) = \int_{\Delta S} \max_{u \in U} \left(\sum_s q(s) \frac{r(s)}{p(s)} u(s) \right) \mu(dr)$$

Proof. Making use of equation (7) from Lemma B1.1, we have

$$\begin{aligned}W(\mu, q) &= \int_{\Delta S} \left(\max_{u \in U} \phi_q(r) \cdot u \right) \left(\sum_s q(s) \frac{r(s)}{p(s)} \right) \mu(dr) \\ &= \int_{\Delta S} \left(\max_{u \in U} \left(\left(\sum_s q(s) \frac{r(s)}{p(s)} \right) \phi_q(r) \right) \cdot u \right) \mu(dr) \\ &= \int_{\Delta S} \left(\max_{u \in U} \sum_s q(s) \frac{r(s)}{p(s)} u(s) \right) \mu(dr)\end{aligned}$$

The second equality follows from the fact that support functions are homogeneous of degree one while the last equality follows from the definition of ϕ_q. □

The final lemma in this section shows that the buyer's value function is jointly continuous in both signal policies and posterior beliefs and convex in his posterior belief. This drives many of our convergence results when we study limit equilibria.

Lemma (B1.3)**.** *$V_c(\mu, q)$ is continuous in (μ, q) and convex in q.*

Proof. First, fix some $c > 0$ and $\mu \in M$. First, consider a buyer who searches once. Hence, his value function is given by

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

Since support functions are continuous and ΔS is bounded, by Corollary 15.7 of Aliprantis and Border (2006), W is continuous in (μ, q). Hence, V_c^1 is continuous. Using the expression for W from Lemma B1.2 and the fact that support functions are convex, we get that W is convex in q. Hence, V_c^1 is convex in q as well. We can now define a sequence of finite-period value functions iteratively where for $i \in \mathbb{N}$

$$V_c^{i+1}(\mu, q) := -c + \int_{\Delta S} \max\left\{\max_{u \in U}(r \cdot u),\ V_c^i(\mu, r)\right\} K_\mu(q, dr)$$

We can now apply the same argument for V_c^1 inductively to conclude that V_c^i is continuous in (μ, q) and convex in q. Note that V_c^i converges uniformly to the value function V_c. Hence, V_c is also continuous in (μ, q) and convex in q. □

B2. Results for Equilibrium Existence

In this section, we will prove results regarding equilibrium existence. We will be using the Kakutani-Fan-Glicksberg (KFG) fixed point theorem to prove existence. However, the pre-conditions necessary for KFG are not always satisfied, so we will concentrate on a subset of signal policies defined below.

Given any $\varepsilon \geq 0$, let $D_\varepsilon \subset \Delta S$ be the set of beliefs such that buyer's acceptance payoff is at least ε-away from his full-information payoff. Recall that $\bar{V}$ is the buyer's payoff from the full information policy $\bar{\mu}$. Formally,

$$D_\varepsilon := \left\{q \in \Delta S \;\middle|\; \max_{u \in U} q \cdot u \geq \bar{V}(q) - \varepsilon\right\}$$

We will consider policies that only put weight in D_ε. In other words, we consider policies in the subset

$$M_\varepsilon := \{\mu \in M \mid \mu(D_\varepsilon) = 1\}$$

Note that we can set ε large enough such that $D_\varepsilon = \Delta S$ and $M_\varepsilon = M$. However, we will show that we can always find a small enough ε such that KFG applies on M_ε. Recall the *best-response correspondence* $\varphi_c : M \to 2^M$ such that $\nu \in \varphi_c(\mu)$ iff $\nu \in M$ is optimal given $Q_c(\mu)$, i.e., the buyer's acceptance set given signal policy $\mu \in M$. We will show that under certain conditions, we can apply the KFG on the mapping φ_c given the domain of policies M_ε.

First, we prove a couple of technical lemmas that are needed for KFG.

Lemma (B2.1)**.** *M_ε is a non-empty, compact and convex metric space.*

Proof. Let $\varepsilon > 0$. We first show that M_ε is non-empty. Consider some full information policy $\bar{\mu}$ such that $\bar{\mu}(D) = 1$ where D is the set of degenerate beliefs. Since $D \subset D_\varepsilon$, $\bar{\mu}(D_\varepsilon) = 1$ so $\bar{\mu} \in M_\varepsilon$ proving that M_ε is non-empty. Note that the convexity of M_ε follows trivially.

We now prove that M_ε is compact. Note that since ΔS is a compact metric space, Π is also a compact metric space by Theorem 15.11 of Aliprantis and Border (2006) and hence sequentially compact (see Theorem 3.28 of Aliprantis and Border (2006)). We now show that M_ε is also sequentially compact. Consider the sequence $\mu_i \in M_\varepsilon \subset \Pi$ so it has a convergent subsequence $\mu_k \to \mu \in \Pi$. Since $\mu_k \in M_\varepsilon$, we have

$$\int_{\Delta S} q \; \mu(dq) = \lim_k \int_{\Delta S} q \; \mu_k(dq) = p$$

so μ is also consistent with p. Moreover, since $\mu_k(D_\varepsilon) = 1$ and D_ε is closed, by Theorem 15.3 of Aliprantis and Border (2006),

$$\mu(D_\varepsilon) \geq \limsup_k \mu_k(D_\varepsilon) = 1$$

so $\mu \in M_\varepsilon$. Thus, M_ε is sequentially compact and therefore compact. □

Lemma (B2.2)**.** *$\varphi_c(\mu)$ is convex for all $\mu \in M$.*

Proof. Fix some $c > 0$. Let $\{\nu_1, \nu_2\} \subset \varphi_c(\mu)$ for some $\mu \in M$ and consider $\nu := a\nu_1 + (1-a)\nu_2$ for some $a \in [0, 1]$. Note that ν is also consistent with p so $\nu \in M$. Let $Q := Q_c(\mu)$ so for any $\nu' \in M$

$$\int_Q \pi(q)\,\nu(dq) = a\int_Q \pi(q)\,\nu_1(dq) + (1-a)\int_Q \pi(q)\,\nu_2(dq) \geq \int_{Q_c(\mu)} \pi(q)\,\nu'(dq)$$

so ν is optimal given $Q = Q_c(\mu)$ as well. Hence, $\nu \in \varphi_c(\mu)$ as desired. □

Lemmas B2.1 and B2.2 deliver nearly all the necessary conditions for applying KFG. The remaining condition is that the correspondence φ_c needs to have a closed graph. In general, this may not be true. As we result, we introduce sufficient conditions under which this condition holds. Note that these sufficient conditions are slightly more general than what we need right now so that they can be used in the next section when we characterize limit equilibria.

Recall that a set of products U is *robust* iff for all non-zero $u \in U$, there is some state $s \in U$ such that $\delta_s \in H_u$. We now generalize this definition of robustness. We say that the set of products

U is *ε-robust* iff for all non-zero $u \in U$ such that $H_u \cap D_\varepsilon \neq \emptyset$, there is some $s \in S$ such that $\delta_s \in H_u$. Note that if ε is large enough such that $D_\varepsilon = \Delta S$, then ε-robustness and robustness are the same. The following technical lemma shows that ε-robustness allows us to construct policies with just enough continuity such that we can eventually obtain closed-graphness for φ_c.

Lemma (B2.3)**.** *Let U be ε-robust and $\mu_i \in M_\varepsilon$ be such that $\mu_i \to \mu \in M_\varepsilon$. Then for any $\nu \in M$, there are $\hat{\nu}_i \in M$ such that*

$$\liminf_i \int_{Q_c(\mu_i)\cap D_\varepsilon} \pi(q)\,\hat{\nu}_i(dq) \geq \int_{Q_c(\mu)\cap D_\varepsilon} \pi(q)\,\nu(dq)$$

Proof. Let U be ε-robust and consider $\mu_i \in M_\varepsilon$ such that $\mu_i \to \mu \in M_\varepsilon$. Let $\nu \in M$. For notation convenience, define $Q_\varepsilon := Q_c(\mu) \cap D_\varepsilon$ and $Q^i_\varepsilon := Q_c(\mu_i) \cap D_\varepsilon$. We first prove an intermediary step. We show that for any non-zero $u \in U$ and $q \in Q_\varepsilon \cap H_u$, we can find $q_i \in Q^i_\varepsilon \cap H_u$ such that $q_i \to q$. This allows us construct the sequence of policies $\hat{\nu}_i \in M$ by shifting the weight that the ν puts on each q to q_i in order to obtain the desired result.

We now prove the intermediary step. Let $q \in Q_\varepsilon \cap H_u$ for some non-zero $u \in U$. We construct a sequence $q_i \in Q^i_\varepsilon \cap H_u$ as follows. If $q \in Q_c(\mu_i)$, then $q \in Q^i_\varepsilon \cap H_u$ so we can just set $q_i = q$. Now consider the case when $q \notin Q_c(\mu_i)$. Note that $H_u \cap D_\varepsilon \neq \emptyset$, so by ε-robustness, there is some $s \in S$ such that $\delta_s \in H_u$. We now choose a belief q_i along the line segment connecting the posterior q to the degenerate belief δ_s such that $q_i \in Q_c(\mu_i)$. Note that since $q \notin Q_c(\mu_i)$ and $q \in H_u$, $q \cdot u < V_c(\mu_i, q)$. We can now define

$$a_i := \frac{c}{c + V_c(\mu_i, q) - q \cdot u} < 1$$

and set $q_i := a_i q + (1 - a_i)\,\delta_s$. We show that $q_i \in Q_c(\mu_i)$. Since $V_c(\mu_i, \cdot)$ is convex by Lemma B1.3,

$$\begin{aligned} V_c(\mu_i, q_i) &\leq a_i V_c(\mu_i, q) + (1 - a_i)\, V_c(\mu_i, \delta_s) \\ &\leq a_i V_c(\mu_i, q) + (1 - a_i)(-c) = a_i q \cdot u < q \cdot u \end{aligned}$$

so $q_i \in Q_c(\mu_i)$. Note that $\{q, \delta_s\} \subset H_u \cap D_\varepsilon$ and since H_u and D_ε are convex, $H_u \cap D_\varepsilon$ is also convex so $q_i = a_i q + (1 - a_i)\,\delta_s \in H_u \cap D_\varepsilon$. Hence, $q_i \in Q^i_\varepsilon \cap H_u$. Finally, to prove that $q_i \to q$, note that by the continuity of V_c from Lemma B1.3, $V_c(\mu_i, q) \to V_c(\mu, q) \geq q \cdot u$ as $q \in Q_c(\mu)$. Hence $a_i \to 1$ so $q_i \to q$ as desired.

We now show how to use the intermediary step to construct $\hat{\nu}_i$. First, define the following set of beliefs

$$Q^* := \bigcup_{u \neq 0} (Q_\varepsilon \cap H_u)$$

By the intermediary step, for each $q \in Q^*$, we can find some $q_i \in Q_c(\mu_i)$ that converges to q. We will construct $\hat{\nu}_i$ by shifting the weight that ν places on each $q \in Q^*$ to q_i. Since we need to ensure that each $\hat{\nu}_i$ is consistent with p, we will need to add a "correction" term for each $\hat{\nu}_i$. This term may decrease the seller's profit, but as $q_i \to q$ the difference will vanish in the limit. Formally, define the transition kernel L_i such that

$$L_i(q) := \lambda_i \left(\mathbf{1}_{Q^*}(q)\, \delta_{q_i} + \mathbf{1}_{\Delta S \setminus Q^*}(q)\, \delta_q \right) + (1 - \lambda_i)\, \delta_{p'_i}$$

where the "correction" terms p'_i are chosen so that consistency is maintained. We can now set

$$\hat{\nu}_i(dr) := \nu(dq)\, L_i(q, dr)$$

Note that

$$\begin{aligned} \hat{\nu}_i \left(Q^i_\varepsilon \cap H_u \right) &= \int_{\Delta S} L_i \left(q,\ Q^i_\varepsilon \cap H_u \right) \nu(dq) \\ &\geq \int_{Q_\varepsilon \cap H_u} L_i \left(q,\ Q^i_\varepsilon \cap H_u \right) \nu(dq) \geq \int_{Q_\varepsilon \cap H_u} \lambda_i \nu(dq) \end{aligned}$$

where the last inequality follows from the fact that $q_i \in Q^i_\varepsilon \cap H_u$. Hence, $\hat{\nu}_i \left(Q^i_\varepsilon \cap H_u \right) \geq \lambda_i \nu(Q_\varepsilon \cap H_u)$. Since $q_i \to q$, $\lambda_i \to 1$ and we have $\liminf_i \hat{\nu}_i \left(Q^i_\varepsilon \cap H_u \right) \geq \nu(Q_\varepsilon \cap H_u)$. Hence

$$\begin{aligned} \liminf_i \int_{Q^i_\varepsilon} \pi(q)\, \hat{\nu}_i(dq) &\geq \liminf_i \sum_u \pi(u)\, \hat{\nu}_i \left(Q^i_\varepsilon \cap H_u \right) \\ &\geq \sum_u \pi(u)\, \nu(Q_\varepsilon \cap H_u) = \int_{Q_\varepsilon} \pi(q)\, \nu(dq) \end{aligned}$$

as desired. □

We now use the lemma above to show that ε-robustness is sufficient for φ_c to have a closed graph. In other words, we can now use KFG to obtain the existence of an equilibrium.

Theorem (B2.4)**.** *Suppose U is ε-robust and $\varphi_c(M_\varepsilon) \subset M_\varepsilon$. Then there exists an equilibrium policy $\mu \in M_\varepsilon$.*

Proof. As a first step, we first show that the mapping $Q_c : M \to \mathcal{F}$ is upper hemi-continuous and takes on non-empty, closed values. Note that for all $\mu \in M$, $\delta_s \in Q_c(\mu)$ so $Q_c(\mu)$ is non-empty. Since V_c is continuous from Lemma B1.3, $Q_c(\mu)$ is also closed. We now demonstrate upper hemi-continuity. Let $(\mu_i, q_i) \to (\mu, q)$ where $q_i \in Q_c(\mu_i)$. Thus, $\max_u q_i \cdot u \geq V_c(\mu_i, q_i)$ so $\max_u q \cdot u \geq V_c(\mu, q)$ as V_c is continuous again by Lemma B1.3. Thus, $q \in Q_c(\mu_i)$ so Q_c is upper hemi-continuous as desired.

We now prove that the mapping $\varphi_c : M_\varepsilon \to 2^{M_\varepsilon}$ has a closed graph where $\varphi_c(M_\varepsilon) \subset M_\varepsilon$ follows from assumption. Let $(\mu_i, \nu_i) \to (\mu, \nu)$ where $\nu_i \in \varphi_c(\mu_i)$. Let $\nu^* \in \varphi_c(\mu)$ so from Lemma B2.3, we can find $\hat{\nu}_i \in M$ such that

$$\begin{aligned}\int_{Q_c(\mu)\cap D_\varepsilon} \pi(q)\,\nu^*(dq) &\leq \liminf_i \int_{Q_c(\mu_i)\cap D_\varepsilon} \pi(q)\,\hat{\nu}_i(dq) \\ &\leq \limsup_i \int_{Q_c(\mu_i)} \pi(q)\,\hat{\nu}_i(dq) \leq \limsup_i \int_{Q_c(\mu_i)} \pi(q)\,\nu_i(dq) \\ &\leq \limsup_i \int_{\Delta S} \pi(q)\,\nu_i(dq) \leq \int_{\Delta S} \pi(q)\,\nu(dq)\end{aligned} \tag{8}$$

where the second line follows from the fact that $\nu_i \in \varphi_c(\mu_i)$ and the last line follows from Theorem 15.5 of Aliprantis and Border (2006) since π is upper semicontinuous. Since $\nu^* \in M_\varepsilon$, $\nu^*(D_\varepsilon) = 1$ so

$$\int_{Q_c(\mu)} \pi(q)\,\nu^*(dq) = \int_{Q_c(\mu)\cap D_\varepsilon} \pi(q)\,\nu^*(dq) \tag{9}$$

Since $\nu_i(Q_c(\mu_i)) = 1$ and Q_c is upper hemi-continuous and takes on nonempty closed values from above, by Theorem 17.13 of Aliprantis and Border (2006), $\nu(Q_c(\mu)) = 1$. Combining this along with equations (8) and (9), we have

$$\int_{Q_c(\mu)} \pi(q)\,\nu^*(dq) \leq \int_{\Delta S} \pi(q)\,\nu(dq) = \int_{Q_c(\mu)} \pi(q)\,\nu(dq)$$

so $\nu \in \varphi_c(\mu)$ as desired. Combined with Lemmas B2.1 and B2.2, KFG (Corollary 17.55 of Aliprantis and Border (2006)) implies there exists some $\mu \in M_\varepsilon$ such that $\mu = \varphi_c(\mu)$ as desired. □

Finally, we apply this to the special case when products are robust.

Corollary (B2.5)**.** *If products are robust, then an equilibrium exists.*

Proof. Fix $c > 0$ and choose ε large enough such that $D_\varepsilon = \Delta S$. Hence, $M_\varepsilon = M$ and U is ε-robust by definition. Since $\varphi_c(M) \subset M$, Theorem B2.4 delivers the desired result. □

B3. Results for Limit Equilibria

This section covers results on limit equilibria. We will present three main results. First, we use Theorem B2.4 to construct a sequence of equilibria that converges to full information as search costs vanish. Hence, full information is always a limit equilibria. Second, we show that all limit equilibria have a partitional structure. Third, we show that when products are sufficiently dispersed, then full information is the unique limit equilibrium.

First, we prove a few technical lemma about convergence. Given $K \in \mathcal{K}$, define the Markov kernel K^j such that

$$K^j(q, dq_j) := K(q, dq_1)\,K(q_1, dq_2)\cdots K(q_{j-1}, dq_j)$$

Thus, K^j corresponds to disclosing j-iterations of information according to the policy K. Our first lemma uses the martingale convergence theorem to prove that K^j always converges to some disclosure policy K^* that has support in A_K. In other words, a limit policy always exists and it is at most as informative as the original policy K.

Lemma (B3.1)**.** *For any continuous $K \in \mathcal{K}$, there exists a $K^* \in \mathcal{K}$ such that $K^j \to K^*$ and $K^*(q, A_K) = 1$.*

Proof. We first set up a formal probability space such that the buyer's posterior beliefs follow a martingale. Fix $q_0 \in \Delta S$ and let $\Omega := \Delta S^{\mathbb{N}}$. By Ionescu-Tulcea's theorem (Theorem IV.4.7 of Çinlar (2011)), we can find a probability measure $\mathbb{P}$ on Ω such that for all j

$$q_j = \mathbb{E}_j[q_{j+1}] = \int_{\Delta S} q_{j+1} K(q_j, dq_{j+1})$$

In other words, beliefs follow a martingale. Since they are bounded, by the martingale convergence theorem (Theorem V.4.1 of Çinlar (2011)), we have $\mathbb{P}$-a.s.

$$q^* = \lim_j q_j$$

Hence, we can define the limit policy K^* such that $K^*(q_0) := \lim_j K^j(q_0)$, i.e. the distribution of the limit posterior q^*.

We now show that $K^*(q, A_K) = 1$. For $\varepsilon > 0$, define the random variable $X_j^\varepsilon := \mathbf{1}_{B_\varepsilon(q^*)}(q_{j+1})$. Hence, X_j^ε is 1 iff q_{j+1} is in an ε-neighborhood around the limit posterior q^*. Note that $\mathbb{P}$-a.s.

$$\lim_j X_j^\varepsilon = \lim_j \mathbf{1}_{B_\varepsilon(q^*)}(q_{j+1}) = 1$$

Since K is continuous, by dominated convergence, we have $\mathbb{P}$-a.s.

$$\begin{aligned} K(q^*, B_\varepsilon(q^*)) &= \lim_j K(q_j, B_\varepsilon(q^*)) = \lim_j \mathbb{E}_j\left[\mathbf{1}_{B_\varepsilon(q^*)}(q_{j+1})\right] \\ &= \lim_j \mathbb{E}_j\left[X_j^\varepsilon\right] = 1 \end{aligned}$$

Since this is true for all $\varepsilon > 0$, we must have $K(q^*, \{q^*\}) = 1$ so $q^* \in A_K$ $\mathbb{P}$-a.s. or $K^*(q_0, A_K) = 1$ for all $q_0 \in \Delta S$.

Finally, we show that $K^*(q_0)$ is consistent with q_0 for all $q_0 \in \Delta S$. Note that by bounded convergence and iterated expectations,

$$\begin{aligned} \int_{\Delta S} q^* K^*(q_0, dq^*) = \mathbb{E}[q^*] &= \mathbb{E}\left[\lim_j q_j\right] \\ &= \lim_j \mathbb{E}[q_j] = \lim_j \mathbb{E}[q_1] = q_0 \end{aligned}$$

as desired. Hence, $K^* \in \mathcal{K}$. □

We now introduce some notation regarding partitions that will be useful in the subsequent proofs. Let $\mathcal{T}$ be some partition of S and recall that $\mu^{\mathcal{T}}$, $K_{\mathcal{T}}$ and $V_{\mathcal{T}}$ are the signal policy, Markov kernel and buyer payoffs corresponding to $\mathcal{T}$ respectively. Also recall that q_T is the conditional belief of $q \in \Delta S$ given any event $T \in \mathcal{T}$. Since $\mu^{\mathcal{T}}$ sends buyer p to posterior beliefs $(p_T)_{T\in\mathcal{T}}$, it must be that $K_{\mathcal{T}}(q)$ must send buyer q to posterior beliefs $(q_T)_{T\in\mathcal{T}}$. Formally, this means that

$$K_{\mathcal{T}}(q) = \sum_{T\in\mathcal{T}} q(T)\,\delta_{q_T}$$

where δ_{q_T} is the degenerate measure on the conditional belief q_T.

Given a partition $\mathcal{T}$, define the conditional simplex

$$\Delta\mathcal{T} := \text{conv}\left(\bigcup_{T\in\mathcal{T}} p_T\right)$$

Hence, $\Delta\mathcal{T}$ is the $|\mathcal{T}|$-dimensional simplex with $(p_T)_{T\in\mathcal{T}}$ as extreme points. The first lemma below shows that $\Delta\mathcal{T}$ consists of exactly all the beliefs that have the same conditional beliefs as the initial prior p. In other words, $\Delta\mathcal{T}$ is the set of all possible posterior beliefs where the buyer only learns about the partition $\mathcal{T}$. This is a useful characterization of $\Delta\mathcal{T}$ that we will use later.

Lemma (B3.2)**.** *For any partition $\mathcal{T}$, $q \in \Delta\mathcal{T}$ iff $q_T = p_T$ for all $T \in \mathcal{T}$.*

Proof. Fix a partition $\mathcal{T}$. First, let $q \in \Delta\mathcal{T}$. Since q is in $\Delta\mathcal{T}$, we can express $q = \sum_{T'\in\mathcal{T}} q(T')\,p_{T'}$ as a convex combination of the conditional beliefs of p given $\mathcal{T}$. Note that $q(T)$ is the probability that q puts on the event $T \subset S$. Now, for any $s \in T \in \mathcal{T}$, we can calculate the conditional probability

$$q_T(s) = \frac{\sum_{T'\in\mathcal{T}} q(T')\,p_{T'}(s)}{\sum_{T'\in\mathcal{T}} q(T')\,p_{T'}(T)} = p_T(s)$$

so $q_T = p_T$ for all $T \in \mathcal{T}$ as desired.

Now, suppose $q_T = p_T$ for all $T \in \mathcal{T}$. For each $s \in T \in \mathcal{T}$,

$$\sum_{T\in\mathcal{T}} q(T)\,p_T(s) = q(T)\,p_T(s) = q(s)$$

so $\sum_{T\in\mathcal{T}} q(T)\,p_T = q$. Since q can be expressed as a convex combination of conditional beliefs of p given $\mathcal{T}$, $q \in \Delta_{\mathcal{T}}$ as desired. □

The next lemma shows that if μ only sends the buyer to posteriors in $\Delta\mathcal{T}$, then the buyer cannot learn more than $\mathcal{T}$ in the limit.

Lemma (B3.3)**.** *If $\mu(\Delta\mathcal{T}) = 1$ and $K_\mu^j \to K^*$, then $K^*(p, \Delta\mathcal{T}) = 1$.*

Proof. Let $\mu(\Delta\mathcal{T}) = 1$ and $K_\mu^j \to K^*$. First, we will show that if $q \in \Delta\mathcal{T}$, then $r \in \Delta\mathcal{T}$ implies $\phi_q(r) \in \Delta\mathcal{T}$. In other words, if a buyer with prior p only learns about $\mathcal{T}$, then a buyer with prior

q in $\Delta\mathcal{T}$ can also only learn about $\mathcal{T}$. Let $r \in \Delta\mathcal{T}$ so from Lemma B3.2, $r_T = p_T$ for every $T \in \mathcal{T}$. Since $q \in \Delta\mathcal{T}$, we also have $q_T = p_T$ for all $T \in \mathcal{T}$. Hence, for all $s \in T \in \mathcal{T}$,

$$[\phi_q(r)]_T(s) = \frac{q(s)\frac{r(s)}{p(s)}}{q(T)\frac{r(T)}{p(T)}} = \frac{q(s)}{q(T)} = q_T(s) = p_T(s)$$

so $\phi_q(r) \in \Delta\mathcal{T}$ again from Lemma B3.2. Hence, $q \in \Delta\mathcal{T}$ implies $\Delta\mathcal{T} \subset \phi_q^{-1}(\Delta\mathcal{T})$ where

$$\phi_q^{-1}(\Delta\mathcal{T}) := \{r \in \Delta S \mid \phi_q(r) \in \Delta\mathcal{T}\}$$

is the set of beliefs such that ϕ_q maps into $\Delta\mathcal{T}$.

Since $\mu(\Delta\mathcal{T}) = 1$, $\Delta\mathcal{T} \subset \phi_q^{-1}(\Delta\mathcal{T})$ implies $\mu\left(\phi_q^{-1}(\Delta\mathcal{T})\right) = 1$ for all $q \in \Delta\mathcal{T}$. From the definition of $K_\mu(q)$ (see equation (7) in the proof of Lemma B1.1), since μ has support in $\phi_q^{-1}(\Delta\mathcal{T})$,

$$\begin{aligned} K_\mu(q, \Delta\mathcal{T}) &= \int_{\phi_q^{-1}(\Delta\mathcal{T})} \sum_s q(s) \frac{r(s)}{p(s)} \mu(dr) = \int_{\Delta S} \sum_s q(s) \frac{r(s)}{p(s)} \mu(dr) \\ &= \sum_s q(s) \frac{1}{p(s)} \int_{\Delta S} r(s)\, \mu(dr) = 1 \end{aligned}$$

where the last equality follows from the fact that μ is consistent with p. In other words, as long as q is in $\Delta\mathcal{T}$, $K_\mu(q)$ will only send the buyer to beliefs in $\Delta\mathcal{T}$. Hence, for any j, we have

$$K_\mu^j(q, \Delta\mathcal{T}) = \int_{\Delta S} \cdots \int_{\Delta S} K_\mu(q_{j-1}, \Delta\mathcal{T})\, K_\mu(q_{j-2}, dq_{j-1}) \cdots K_\mu(q, dq_1) = 1$$

Since $K_\mu^j(q) \to K^*(q)$ and $\Delta\mathcal{T}$ is closed, by Theorem 15.3 of Aliprantis and Border (2006),

$$K^*(q, \Delta\mathcal{T}) \geq \limsup_j K_\mu^j(q, \Delta_{\mathcal{T}}) = 1$$

as desired. □

Finally, we can now use the lemmas above to show that every policy converges to some partitional policy if the buyer were to receive infinite draws of the signal.

Lemma (B3.4)**.** *For any $\mu \in M$, $K_\mu^j \to K_{\mathcal{T}}$ for some partition $\mathcal{T}$.*

Proof. Note that by Lemma B1.1, $K_\mu \in \mathcal{K}$ and is continuous. Hence, by Lemma B3.1, we can define $K^* \in \mathcal{K}$ such that $K_\mu^j \to K^*$ and $K^*\left(q, A_{K_\mu}\right) = 1$ for all $q \in \Delta S$. Recall that $q \in A_{K_\mu}$ iff $K_\mu(q) = \delta_q$. In other words, K^* has support only on the set of posterior beliefs such that K_μ provides no information.

Since $K_\mu(q) = \mu_q \circ \phi_q^{-1}$, we have $q \in A_{K_\mu}$ iff

$$1 = K_\mu(q, \{q\}) = \mu_q\left(\phi_q^{-1}\{q\}\right) \tag{10}$$

Now $\phi_q^{-1}\{q\}$ is the set of beliefs $r \in \Delta S$ such that $\phi_q(r) = q$. In other words, these are the posterior beliefs of buyer p such that a buyer with prior q does not update. Clearly, this depends on q. For example, if $q = \delta_s$ is degenerate, then regardless of what posterior buyer p obtains, buyer q will still have posterior $q = \delta_s$. In other words, $\phi_q^{-1}\{q\}$ is the entire simplex. Define $T_q := \{s \in S \mid q(s) > 0\}$ as the set of states where q puts strictly positive probability and define the partition

$$\mathcal{T}_q := \{T_q\} \cup \bigcup_{s \notin T_q} \{s\}$$

In other words, $\mathcal{T}_q$ partitions S into the event T_q where q has strictly positive probability and all other states where $q(s) = 0$. We will now show that

$$\phi_q^{-1}\{q\} = \Delta\mathcal{T}_q$$

In other words, $\phi_q^{-1}\{q\}$ is exactly the sub-simplex with p_{T_q} and every degenerate belief not in T_q as vertices. In order to prove this, note that $r \in \phi_q^{-1}\{q\}$ iff for all $s \in S$

$$q(s)\frac{r(s)}{p(s)} = \left(\sum_{s'} q(s')\frac{r(s')}{p(s')}\right) q(s) \tag{11}$$

Now, for any $\hat{s} \in T_q$, since $q(\hat{s}) > 0$, we can divide equation (11) by $q(\hat{s})$ to get

$$r(\hat{s}) = \left(\sum_{s'} q(s')\frac{r(s')}{p(s')}\right) p(\hat{s})$$

This implies that $r_{T_q} = p_{T_q}$. Moreover, for every $s \notin T_q$, $r_{\{s\}} = \delta_s = p_s$ trivially. Hence, we have $r_T = p_T$ for all $T \in \mathcal{T}_q$. By Lemma B3.2, $r \in \Delta\mathcal{T}_q$. Thus, $\phi_q^{-1}\{q\} \subset \Delta\mathcal{T}_q$. By similar reasoning, $\Delta\mathcal{T}_q \subset \phi_q^{-1}\{q\}$ so $\phi_q^{-1}\{q\} = \Delta\mathcal{T}_q$ as desired.

Returning to equation (11), we now have that $q \in A_{K_\mu}$ implies $\mu_q(\Delta\mathcal{T}_q) = 1$. We show that this implies $\mu(\Delta\mathcal{T}_q) = 1$. Suppose on the contrary that $\mu(\Delta S \backslash \Delta\mathcal{T}_q) > 0$. From the expression for μ_q in equation (6), we have

$$\begin{aligned} 0 = \mu_q(\Delta S \backslash \Delta\mathcal{T}_q) &= \int_{\Delta S \backslash \Delta\mathcal{T}_q} \sum_{s \in S} q(s)\frac{r(s)}{p(s)}\mu(dr) \\ &= \sum_{s \in T_q} q(s)\frac{1}{p(s)} \int_{\Delta S \backslash \Delta\mathcal{T}_q} r(s)\mu(dr) \end{aligned}$$

Now, in order for this equality to hold, it must be that for any $s \in T_q$, $\int_{\Delta S \backslash \Delta\mathcal{T}_q} r(s)\mu(dr) = 0$. Since $\mu(\Delta S \backslash \Delta\mathcal{T}_q) > 0$, this implies $r(s) = 0$ μ-a.s. for all $s \in T_q$. This implies μ only puts mass on posteriors such that $r(T_q) = 0$. Since μ is consistent with p and $p \in \operatorname{int}(\Delta S)$, this yields a contradiction. Hence, $\mu_q(\Delta\mathcal{T}_q) = 1$ implies $\mu(\Delta\mathcal{T}_q) = 1$.

To summarize, we now have that $q \in A_{K_\mu}$ implies $\mu(\Delta\mathcal{T}_q) = 1$. By Lemma B3.3, this implies that $K^*(p, \Delta\mathcal{T}_q) = 1$ for every $q \in A_{K_\mu}$. Since $K^*\left(p, A_{K_\mu}\right) = 1$, this further implies that $K^*\left(p, \Delta\mathcal{T}_q \cap A_{K_\mu}\right) = 1$. We now show that K^* must have support on conditional beliefs of p, i.e., posterior beliefs of the form p_T where $T \subset S$. First, let $q \in A_{K_\mu}$ and suppose $T_q = S$. In this case, $\mathcal{T}_q = \{S\}$ i.e. the no information partition and $\Delta\mathcal{T}_q = \{p\}$. Hence, $K^*(p, \{p\}) = 1$ as desired. Next, suppose $T_q = S\backslash\{s'\}$. In this case, $\mathcal{T}_q = \{T_q, \{s'\}\}$ and $\Delta\mathcal{T}_q = \text{conv}\left\{p_{T_q}, \delta_{s'}\right\}$. Hence,

$$K^*\left(p,\ \text{conv}\left\{p_{T_q}, \delta_{s'}\right\} \cap A_{K_\mu}\right) = 1$$

In other words, K_p^* sends the buyer only to absorbing beliefs in the line segment connecting p_{T_q} and $\delta_{s'}$. Suppose there is an interior belief $r \in \text{int}\left(\text{conv}\left\{p_{T_q}, \delta_{s'}\right\}\right)$ and $r \in A_{K_\mu}$. This implies $r \in \text{int}(\Delta S)$ so $T_r = S$. By the step before, this implies $p = r$. Hence, K^* only puts weight on p_{T_q}, $\delta_{s'}$ or p. By iterating this reasoning over the cardinality of T_q for any $q \in A_{K_\mu}$, we obtain that K^* can only put weights on conditional beliefs of p. Hence, $K^* = K_{\mathcal{T}}$ for some partition $\mathcal{T}$. □

We are now ready to address our first main result showing that a full information limit equilibrium always exists (we will be using the notation from Appendix B2). The main idea is to choose a sequence of increasingly small ε and apply Theorem B2.4 on the set of signals M_ε. In order to use the theorem, we need one more lemma in order to ensure that the the correspondence φ_c only maps signals in M_ε into M_ε.

Lemma (B3.5)**.** *There exists some $\bar{\varepsilon} > 0$, such that for any $\varepsilon \leq \bar{\varepsilon}$, we can find some $c > 0$ where $\varphi_c(M_\varepsilon) \subset M_\varepsilon$.*

Proof. Set $\bar{\varepsilon} > 0$ small enough such that for all $\varepsilon \leq \bar{\varepsilon}$, D_ε is close enough to D_0 so that they share the same conditional beliefs of p. In other words, $p_T \in D_\varepsilon$ implies $p_T \in D_0$ for all $T \subset S$. Fix some $\varepsilon \leq \bar{\varepsilon}$ and note that by Lemma B3.4, for any $\mu \in M_\varepsilon$, $K_\mu^j \to K_{\mathcal{T}}$ for some partition $\mathcal{T}$.

First, we show that $p_T \in D_0$ for all $T \in \mathcal{T}$. We prove this by contradiction. Let $T \in \mathcal{T}$ and suppose $p_T \notin D_0$. By our choice of ε, this implies that $p_T \notin D_\varepsilon$. Note that by Lemma B3.1, $K_{\mathcal{T}}\left(p_T, A_{K_\mu}\right) = 1$. Hence, $p_T \in A_{K_\mu}$. By the same argument as in the proof of Lemma B3.4, this implies that $\mu(\Delta\mathcal{T}_q) = 1$ where $q = p_T$. Now, since $\mu \in M_\varepsilon$, $\mu(D_\varepsilon) = 1$ so $\mu(\Delta\mathcal{T}_q \cap D_\varepsilon) = 1$. However, $q \notin D_\varepsilon$ so μ puts no mass around $p_T = q$ and all its mass in $\Delta\mathcal{T}_q \cap D_\varepsilon$. This contradicts the fact that $p \in \text{int}(\Delta S)$ and μ is consistent with p. Hence, it must be that $p_T \in D_0$ for all $T \in \mathcal{T}$.

We now prove that as $c \to 0$, $V_c \to \bar{V}$ uniformly on $M_\varepsilon \times \Delta S$. Let $c_i \to 0$. Fix $\mu \in M_\varepsilon$ and $q \in \Delta S$. First, we first show that $V_{c_i}(\mu, q) \to \bar{V}(q)$ pointwise. Set n_i as the largest integer less than $c_i^{-\frac{1}{2}}$ so $n_i c_i \leq c^{\frac{1}{2}}$. Note that the buyer can always visit n_i sellers before choosing to go with the n_i-th seller. Hence, his payoff is at least as great as searching n_i times. If we let μ^{n_i} be the distribution of posteriors for a buyer who visits n_i sellers, then

$$V_{c_i}(\mu, q) \geq -n_i c_i + W(\mu^{n_i}, q) \geq -c^{\frac{1}{2}} + W(\mu^{n_i}, q) \tag{12}$$

As $c_i \to 0$, $n_i \to \infty$. Hence, by Lemma B3.4, $K_{\mu^{n_i}} = K_\mu^{n_i} \to K_\mathcal{T}$ for some partition $\mathcal{T}$. Moreover, we know that $p_T \in D_0$ for all $T \in \mathcal{T}$. Recall that

$$D_0 = \left\{ q \in \Delta S \;\middle|\; \max_{u \in U} q \cdot u \geq \bar{V}(q) \right\}$$

Since $\bar{V}$ is the full information payoff, it is linear in q. Hence, $p_T \in D_0$ for all $T \in \mathcal{T}$ implies $q_T \in D_0$ for all $T \in \mathcal{T}$. This implies that $K_\mathcal{T}(q, D_0) = 1$ for all $q \in \Delta S$. In other words, $K_\mathcal{T}(q)$ has support in D_0 for all $q \in \Delta S$ so we have

$$\begin{aligned} W(\mu^{n_i}, q) \to \int_{\Delta S} \left(\max_{u \in U} r \cdot u \right) K_\mathcal{T}(q, dr) = \int_{D_0} \bar{V}(r) K_\mathcal{T}(q, dr) &= \int_{\Delta S} \bar{V}(r) K_\mathcal{T}(q, dr) \\ =& \bar{V}\left(\int_{\Delta S} r \; K_\mathcal{T}(q, dr) \right) = \bar{V}(q) \end{aligned}$$

where the penultimate equality follows again from the fact that $\bar{V}$ is linear in q. Since $W(\mu^{n_i}, q) \to \bar{V}(q)$ it must be that $V_{c_i}(\mu, q) \to \bar{V}(q)$ pointwise. Note that $M_\varepsilon \times \Delta S$ is compact by Lemma B2.1 and V_c is continuous by Lemma B1.3. Since V_{c_i} is decreasing monotonically, by Dini's Theorem, (Theorem 2.66 of Aliprantis and Border (2006)), $V_{c_i} \to \bar{V}$ uniformly on $M_\varepsilon \times \Delta S$.

Now, by uniform continuity, we can find some $c > 0$ such that for all $(\mu, q) \in M_\varepsilon \times \Delta S$

$$\left| V_c(\mu, q) - \bar{V}(q) \right| < \varepsilon$$

Hence, for all $\mu \in M_\varepsilon$,

$$\begin{aligned} Q_c(\mu) &= \left\{ q \in \Delta S \;\middle|\; \max_{u \in U} q \cdot u \geq V_c(\mu, q) \right\} \\ &\subset \left\{ q \in \Delta S \;\middle|\; \max_{u \in U} q \cdot u \geq \bar{V}(q) - \varepsilon \right\} = D_\varepsilon \end{aligned}$$

Consider a best-response policy $\nu \in \varphi_c(\mu)$. Since we know that ν must send the buyer inside the acceptance set, $\nu(Q_c(\mu)) = 1$. Hence,

$$\nu(D_\varepsilon) \geq \nu(Q_c(\mu)) = 1$$

so $\nu \in M_\varepsilon$. Hence, $\varphi_c(\mu) \subset M_\varepsilon$ for every $\mu \in M_\varepsilon$ so $\varphi_c(M_\varepsilon) \subset M_\varepsilon$ as desired. □

We are now ready to state our first main result on the existence of full information limit equilibria. Recall that a *limit equilibrium* policy μ is a policy where there exists a sequence of equilibrium policies μ_i for each $c_i > 0$ such that $\mu_i \to \mu$ as $c_i \to 0$.

Theorem (B3.6)**.** *There exists a limit equilibrium μ such that $V_{c_i}(\mu_i, q) \to \bar{V}(q)$ as $c_i \to 0$.*

Proof. First, suppose U has strict preferences in that $u(s) \neq v(s)$ for all products u and v and $s \in S$. We first prove this for strict preferences. Consider a decreasing sequence $\varepsilon_i \to 0$. Now, for each ε_i, from Lemma B3.5, we can find some $c_i > 0$ such that $\varphi_{c_i}(M_{\varepsilon_i}) \subset M_{\varepsilon_i}$. We will now use Theorem B2.4 to show existence of an equilibrium. In order to do this, we need to show that U is ε_i-robust. Since U has strict preferences, we can alway set ε_i such that D_{ε_i} is small enough such that it only intersects those H_u that also touch at least one of the degenerate vertices. This is because we can always set ε_i such that $\varepsilon_i < |u(s) - v(s)|$ for all $\{u, v\} \subset U$ and $s \in S$. In other words, for every non-zero $u \in U$, $H_u \cap D_{\varepsilon_i} \neq \emptyset$ implies there is some $s \in S$ such that $\delta_s \in H_u$. Note that this is exactly the definition of ε-robustness, so by Theorem B2.4, there exists an equilibrium policy $\mu_i \in M_{\varepsilon_i}$. Note that $V_{c_i} \to \bar{V}$ follows directly from the proof of Lemma B3.5.

Now, suppose U does not have strict preferences. In this case, we can redefine the D_ε sets such that they only include the indifferent points at the vertices. Note that if we define D_ε in this manner, since ties in the buyer's actions are resolved in the seller's favor, sellers will never want to send any buyer to posterior beliefs outside of these vertices. Hence Lemma B3.5 that $\varphi_c(M_\varepsilon) \subset M_\varepsilon$ is satisfied and the argument then follows exactly as above. $\square$

Recall that a partition $\mathcal{T}$ is *partitionally optimal* iff a monopolistic seller would prefer to provide no information to a buyer with belief p_T for every event $T \in \mathcal{T}$. Formally, this means that for all $T \in \mathcal{T}$,

$$\pi(p_T) \geq \int_{\Delta T} \pi(q)\, \nu(dq)$$

for all ν consistent with p_T where ΔT is the sub-simplex corresponding to the event T. We now show that any limit equilibrium must be partitionally optimal.

Proposition (B3.7)**.** *If μ is a limit equilibrium, then there is a partitionally optimal $\mathcal{T}$ such that $V_{c_i}(\mu_i, q) \to V_{\mathcal{T}}(q)$ as $c_i \to 0$.*

Proof. Let μ be a limit equilibrium so $\mu_i \to \mu$ as $c_i \to 0$ where μ_i are all equilibrium policies. As in the proof of Lemma B3.5, the buyer can always search n_i times where n_i is the largest integer less than $c_i^{\frac{1}{2}}$ so $n_i c_i \leq c_i^{\frac{1}{2}}$. Hence, by the same argument for equation (12) in Lemma B3.5, we obtain

$$V_{c_i}(\mu_i, q) \geq -n_i c_i + W(\mu_i^{n_i}, q) \geq -c^{\frac{1}{2}} + W(\mu_i^{n_i}, q)$$

where $\mu_i^{n_i}$ is the distribution of posteriors for a buyer who visits n_i sellers under signal policy μ_i. Now, by Lemma B3.4, $\mu^{n_i} = K_\mu^{n_i}(p) \to K_{\mathcal{T}}(p)$ for some partition $\mathcal{T}$. Since $\mu = \lim_i \mu_i$ and W is continuous in μ from Lemma B1.2, we have

$$\begin{aligned} \lim_i W(\mu_i^{n_i}, q) &= \lim_i W(\mu^{n_i}, q) = W\left(\lim_i \mu^{n_i}, q\right) \\ &= W(K_{\mathcal{T}}(p), q) = V_{\mathcal{T}}(q) \end{aligned}$$

Hence, $\lim_i V_{c_i}(\mu_i, q) \geq V_{\mathcal{T}}(q)$. Moreover, since $\mu_i^{n_i}$ converges to $K_{\mathcal{T}}(p)$, it cannot be strictly more informative than $K_{\mathcal{T}}(p)$ so the inequality is an equality. Hence, $V_{c_i}(\mu_i, q) \to V_{\mathcal{T}}(q)$.

We now show that $\mathcal{T}$ is partitionally optimal. We prove this by contradiction. Suppose $\mathcal{T}$ is not partitionally optimal so we can find some $T \in \mathcal{T}$ and ν consistent with p_T such that

$$\pi(p_T) < \int_{\Delta T} \pi(q)\,\nu(dq)$$

Since μ_i is an equilibrium, we have that

$$\int_{Q_c(\mu_i)} \pi(q)\,\mu_i(dq) \geq \int_{Q_c(\mu_i)} \pi(q)\,\nu'(dq)$$

for all $\mu' \in M$. Since $V_{c_i}(\mu_i, q) \to V_{\mathcal{T}}(q)$, we can find c_i small enough such that $p_T \in Q_c(\mu_i)$. However, since $\pi(p_T) < \int_{\Delta T} \pi(q)\,\nu(dq)$ and μ_i is optimal given $Q_c(\mu_i)$, it must be that μ_i never sends the buyer to posteriors near p_T. In other words, we can find some $B_\varepsilon(p_T)$ such that $\mu_i(B_\varepsilon(p_T)) = 0$. This contradicts the fact that $V_{c_i}(\mu_i, q) \to V_{\mathcal{T}}(q)$ so $\mathcal{T}$ must be partitionally optimal. □

Recall that a set of products U is *sufficiently dispersed* iff it is dispersed and $u^*(p_T) = \{0\}$ for all event $T \subset S$ where T consists of at least two states. We now show that sufficiently dispersed is enough to guarantee that full information is the unique limit equilibrium.

Theorem (B3.8)**.** *If products are sufficiently dispersed, then for any limit equilibrium* μ*,* $V_{c_i}(\mu_i, q) \to \bar{V}(q)$ *as* $c_i \to 0$*.*

Proof. Suppose U is sufficiently dispersed and let μ be a limit equilibrium where $\mu_i \to \mu$ as $c_i \to 0$. From Proposition B3.7, we know that $V_{c_i}(\mu_i, q) \to V_{\mathcal{T}}(q)$ for some partitionally optimal $\mathcal{T}$. We now show that $V_{\mathcal{T}} = \bar{V}$, i.e., $\mathcal{T}$ is payoff-equivalent to full information for the buyer. Note that $K_{\mathcal{T}}$ only puts weight on p_T for all $T \in \mathcal{T}$. Hence, it is sufficient to just check that $V_{\mathcal{T}}(p_T) = \bar{V}(p_T)$ for all $T \in \mathcal{T}$

We prove this by contradiction. Suppose there is some $T \in \mathcal{T}$ such that $V_{\mathcal{T}}(p_T) < \bar{V}(p_T)$. Note that if $T = \{s\}$ is a singleton, $V_{\mathcal{T}}(\delta_s) = \bar{V}(\delta_s)$ so T must contain at least two states. Since U is sufficiently dispersed, we know that $u^*(p_T) = \{0\}$. Hence if we let H_0 denote the set of beliefs such that the buyer chooses 0, then $p_T \in \operatorname{int}(H_0)$. Hence, we can create a neighborhood $B_\varepsilon(p_T)$ such that $\mu_i(B_\varepsilon(p_T)) = 0$. Since $\mu_i \to \mu$, this implies $\mu(B_\varepsilon(p_T)) = 0$. In other words, μ will never send a buyer to any belief near p_T. This contradicts the fact that $K_{\mathcal{T}}$ is the limit policy as $V_{c_i}(\mu_i, q) \to V_{\mathcal{T}}(q)$. Hence, $V_{\mathcal{T}}(p_T) = \bar{V}(p_T)$ for all $T \in \mathcal{T}$ so $V_{\mathcal{T}} = \bar{V}$ as desired. □

C. Additional Analysis

C1. Non-Markovian Equilibria

In the main text, we restrict equilibria in three ways. First, we consider Markovian equilibria, assuming sellers only condition their disclosure strategies on the buyer's current belief (which is the payoff relevant state variable).[18] Second, we assume a signal tie-breaking rule that if "no information" is a best-response for a seller, then she will use it. Third, we focus on symmetric equilibria. Here we reconsider Example 1 and show that the monopoly strategy is the unique rationalizable outcome without these restrictions on strategies.

Example 1 (cont). Consider a single seller i and suppose that all other sellers pursue arbitrary policies. No matter what $-i$ do, even if they provide full information, a buyer would purchase from seller i if the state falls into the set

$$Q_1 := \{q \in [0,1] : \max\{2q-1, 0\} \geq V_0(q)\} = [0, a_1] \cup [b_1, 1] = [0, c] \cup [1-c, 1]$$

where the full-information value function $\bar{V}(q) = V_0(q) := -c + q$ provides an upper bound on the buyer's value if he searches.

Given that seller i faces $Q \supset Q_1$, then the only strategies that are not strictly dominated are (payoff equivalent to) perfect bad news policies that just persuade the buyer to purchase the product. If $b_1 \leq \frac{1}{2}$, seller i will use the monopoly policy, and we are done. Otherwise, she will use one of a set of perfect bad news policies; of these, the most information such a signal could generate for a buyer is $p \to \{0, b_1\}$. The same logic applies to all sellers, so a buyer's value from searching is bounded above by $V_1(q) = q[2 - \frac{1}{b_1}] - c$ which means that the acceptance set must contain

$$\begin{aligned} Q_2 :&= \{q \in [0,1] : \max\{2q-1, 0\} \geq V_1(q)\} = [0, a_2] \cup [b_2, 1] \\ &= \left[0, \frac{b_1}{2b_1 - 1} c\right] \cup [b_1(1-c), 1] = \left[0, \frac{(1-c)}{2(1-c)-1} c\right] \cup \left[(1-c)^2, 1\right] \end{aligned}$$

Continuing this process, $b_n = (1-c)^n$ which means that with $n \geq -\log(2)/\log(1-c)$ rounds, then $Q_n \supset [\frac{1}{2}, 1]$ and the seller chooses the monopoly policy. Note that in this case sellers may use different strategies, but all are payoff-equivalent to the monopoly policy. △

[18] Of course, if a seller cannot observe the order in which the buyer approaches sellers, then this is without loss of generality.

C2. Mixed Strategies in the Tracking Game - Single Product Example

We now consider mixed strategies in the tracking game discussed in Section 4.1. Consider Example 1, with $p < 1/2$, so that a buyer does not buy in the absence of information.[19] Suppose proportion α of buyers are initially anonymous while $1-\alpha$ are tracked. If a seller knows she is facing a tracked buyer, then she provides the monopoly level of information, as in Section 2. If a seller faces a buyer without tracking information, this may mean that the buyer is a new tracked buyer with prior p, or an anonymous buyer who has received any number of signals.

Consider an equilibrium in which sellers use perfect bad news signals such that $p \to \{0, q\}$. Using Bayes' rule define $\{p, q, q^{(2)}, q^{(3)}, \ldots\}$ as the sequence of beliefs that arises from repeatedly observing the good signal, so one more signal means $q^{(n)} \to \{0, q^{(n+1)}\}$. Observe that we must have $q \geq \frac{1}{2}$. If this were not the case, a seller could provide "two signals" taking $q^{(n)} \to \{0, q^{(n+2)}\}$, raising the seller's sales from anonymous buyers, and possibly leading to sales from tracked buyers. Tracked buyers thus purchase at the first seller; anonymous buyers have some threshold belief q_A, and purchase after receiving $N := \min\{n : q^{(n)} \geq q_A\}$ positive signals, where $N \geq 2$ to justify the cost of becoming anonymous.

We now verify there is an equilibrium of this form. A seller faces $1-\alpha$ tracked buyers with prior p, α anonymous buyers with prior p, $\alpha\frac{p}{q}$ anonymous buyers with belief q, …, and $\alpha\frac{p}{q^{(N-1)}}$ anonymous buyers with prior $q^{(N-1)}$. Since the seller's profit function is a mixture of step functions, it follows that her optimal disclosure policy will be at a corner, meaning one type of buyer is indifferent between purchasing and not. A symmetric equilibrium can thus be of two kinds: (a) the seller ignores the anonymous buyers and sets $q = \frac{1}{2}$, or (b) the seller chooses $q > \frac{1}{2}$ so that a tracked buyer with belief $q^{(N)}$ is indifferent between buying and continuing.

For $q = \frac{1}{2}$ to be an equilibrium the seller must not want to deviate by providing "two signals" and selling to two cohorts of informed buyers,[20]

$$(1-\alpha)\frac{p}{q} + \alpha\frac{p}{q^{(N)}} \geq (1-\alpha)\frac{p}{q^{(2)}} + \alpha\frac{p}{q^{(N)}} + \alpha\frac{p}{q^{(N+1)}},$$

where N is derived from the anonymous buyer's stopping problem.

For $q > \frac{1}{2}$ to be an equilibrium, we require that the seller does not want to deviate by providing the monopoly level of information, giving up on the anonymous buyers,[21]

$$(1-\alpha)\frac{p}{1/2} \leq (1-\alpha)\frac{p}{q} + \alpha\frac{p}{q^{(N)}}.$$

As one would expect, an increase in the number of anonymous buyers, α, slackens this constraint

[19] Without this assumption, no information is always an equilibrium (see Corollary 3).

[20] This corresponds to the local upwards (IC) constraint. Unfortunately, this does not imply other (IC) constraints bind.

[21] This corresponds to the (only) downward (IC) constraint. In the most informative equilibrium this is the relevant constraint.

and allows us to support a higher q. This means that, in the most informative equilibrium, sellers disclose more as α rises, and any one buyer is better off when other buyers are anonymous. $\triangle$