Uncharted Waters: Selling a New Product Robustly

Kun Zhang[*]

November 18, 2022

Job Market Paper

Click here for the latest version

Abstract

A seller introduces a novel product to an unfamiliar market. The seller sets a price and chooses how much information to disclose about the product to a representative buyer, who may incur a search cost to discover an outside option. The buyer knows her outside option distribution, but the seller knows only its mean and an upper bound on its support, and evaluates any selling strategy by its guaranteed profit. The robustly optimal strategy balances the trade-off between demand and extraction: information design can boost demand by deterring the buyer's search, but this requires providing her with a high payoff via a low price. I find that full disclosure is optimal only when the search cost is high, and different kinds of partial disclosure policies are optimal for lower search costs. Perhaps surprisingly, the price is not monotone in the search cost. These results shed light on the large variations in information disclosure policies among new products, and suggest that improvements in information technology that reduce search costs may increase prices and make information provision noisier.

[*]Arizona State University, Email: kunzhang@asu.edu
I am indebted to Hector Chade for his continuing guidance, encouragement, and support, and I am grateful to Mark Whitmeyer for numerous stimulating conversations and insightful suggestions. I would like to thank Zach Bethune, Nina Bobkova, Jaroslav Borovička, Albin Erlanson, Yuhta Ishii, Navin Kartik, Kyungmin (Teddy) Kim, Andreas Kleiner, Tomas Larroucau, Fernando Leiva Bertran, Yue Li, Suraj Malladi, Alejandro Manelli, Eddie Schlee, Siyu Shi, Eran Shmaya, Ludvig Sinander, Galina Vereshchagina, Renkun Yang, Ziwei Wang, John Zhu, and the audience at ASU, Stony Brook, and the 2022 Stony Brook International Conference on Game Theory for their helpful comments and suggestions. All errors are my own.

1 Introduction

Rapid technological development has brought an increasing number of new products to us. Examples of such products include e-readers, electric vehicles, LED lights, and infrared cookers, to name a few. Because of its novelty, a seller of a new product often has considerable control over what a buyer can learn about it. In particular, the seller can provide information by offering, for example, free trials, product samples, and product descriptions. The buyer, however, can still check the pricing and features of related existing products, and the seller may know little about a buyer's knowledge of her alternatives. For instance, the seller might be uncertain about how a buyer acquires or processes information, what substitute products the buyer has in mind, what kind of stores a buyer has access to, and so on.

In this paper, I examine the following questions. In the setting described above, what is the optimal selling strategy if the seller can set a price and choose how much information to provide about the new product? Would the buyer be better off if learning about her alternatives becomes easier? Finally, how do the answers to the previous questions shed light on selling different kinds of new products?

In my model, a seller faces a buyer whose match value with the seller's product is either high or low. Although the match value is unknown to both parties, the (prior) probability of the match value being high is common knowledge. Along with a posted price, the seller chooses a disclosure policy to provide information about the match value. The buyer has access to an unknown outside option that could be interpreted as the buyer's best alternative for the seller's new product. The distribution of this outside option is known to the buyer. Observing the price and the realized signal from the disclosure policy, the buyer updates her beliefs about the match value, and then decides whether to discover the value of the outside option (which I term "search") at a cost or buy the seller's product directly.

To capture the seller's uncertainty regarding what the buyer knows about her alternatives, I assume that she knows only the mean of the buyer's outside option distribution and an upper bound on the value of the outside option and that she employs a selling strategy to maximize her revenue guarantee. That is, the seller maximizes the worst possible revenue generated by an outside option distribution that is consistent with her information; such a strategy is "robust" in that it performs reasonably well in all possible scenarios. This can be metaphorically interpreted as the seller facing an "adversary" who designs the outside option distribution to minimize her profit.

To solve for the seller's robustly optimal selling strategy, I take a two-step approach. First, for each price, I solve for the seller's optimal information disclosure policy for that price. This is obtained by solving for a saddle point of a zero-sum game played by the seller who chooses a disclosure policy to maximize her revenue and the adversary who chooses an outside option distribution to minimize it. Second, having identified the saddle point, the seller's revenue guarantee is determined by her choice of price, and hence the robustly optimal price can be solved using standard optimization techniques.

Two observations are particularly useful in understanding the seller's robustly optimal strategy. Note first that the buyer would like to incur the search cost and search for an alternative when the belief about the match value is below some threshold. Therefore, it may be helpful for the seller to pool the beliefs just above the threshold to deter the buyer from searching, thereby increasing the likelihood of purchase. I call a disclosure policy with such a feature a *deterrence policy*. A deterrence policy, however, need not be robust. In fact, the robustly optimal selling strategy entails a deterrence policy only if the price is lower than a threshold that is proportional to the search cost–this is the first key observation. Second, a disclosure policy that continuously and evenly spreads out beliefs "hedges well" against the adversary: for a fixed price, this disclosure policy keeps the probability of purchase the same under any outside option distribution with the same mean. Such a property generates the desired robustness.

The first observation above highlights that the main trade-off the seller faces is between demand and extraction. Although using a deterrence policy can increase demand, this strategy is effective only when the price is below the threshold and thus can hurt surplus extraction. When the search cost is small, the price threshold is also low. As argued above, in this case using a deterrence policy is unprofitable. The second observation above then suggests that a disclosure policy that generates continuously and evenly spread out impressions of the product is optimal. In particular, it hedges well against the adversary and allows the seller to charge a higher price and hence extract a greater surplus. Such a disclosure policy provides noisy information: it is likely that the buyer's impression is neither favorable nor unfavorable.

When the search cost is large, however, so is the price threshold and thus the seller can charge a higher price even if it must be below the threshold. In this case, full disclosure turns out to be optimal. This disclosure policy only produces two posterior beliefs: the match value is high with probability one, or the match value is low with probability one. The buyer buys without search whenever the match value is high, and never buys otherwise. Intuitively, by informing the buyer of the exact match value (high or low), the seller can identify those for whom the innovative features of the new product are especially attractive, and make sure that they buy without search with probability one. This strategy helps the seller secure a sizable demand while charging a higher price.

For intermediate search costs, it can be optimal for the seller to combine the above two disclosure policies. The resulting disclosure policy informs the buyer that the match value is high with some probability, and with complementary probability it spreads out beliefs evenly. In the former case, the buyer buys immediately; otherwise, unlike full disclosure, she may return to buy the seller's product after search. This is optimal only when the prior probability of a high match value is sufficiently large relative to the mean of the outside option distribution. Put differently, the market must have enough confidence in the seller's product vis-à-vis the outside option. Consequently, it is likely that the buyer's expected value for the seller's product is much higher than her realized outside option value. Thus, the seller has an incentive to attract the buyer to come back to buy if she goes to search.

To summarize, full disclosure is optimal if the search cost is sufficiently high, and different kinds of partial disclosure policies are optimal for lower search costs. Importantly, the results help explain the large variation in disclosure policies among new products. For example, the e-ink tablet reMarkable allows a full refund within 100 days, while many other new products offer only a 7-day free trial.

Furthermore, the results have concrete implications for the sale of different kinds of new products. Some of these products are *revolutionary*: for example, iPhone and 3D printer. Some of them are *evolutionary products*, namely, existing products made slightly better, like smart lamps and foods made with healthier ingredients. Others are *alternatives to existing products*, which are revolutionary in some aspects at the cost of losing some existing features. One may think of portable wireless speakers, which are much more convenient at the cost of sound quality. One way to interpret the search cost is a measure of the ease with which a consumer can figure out the best alternative in the market. For evolutionary products, the search cost is usually low, while it is likely to be higher for alternatives to existing products. This is because evolutionary products only differ from existing products in a certain aspect, but alternatives to existing products are in a "completely different direction" and thus it can be significantly harder for the buyer to figure out what is the best alternative.

Consequently, for evolutionary products, noisy information provision is optimal. This can be done by, for instance, offering a short trial period, limiting the number of features available in a free trial, or succinct product descriptions. For alternatives to existing products, however, it is optimal to provide full information, divide the consumers into "lovers" and "haters", and serve the former only. Examples of such disclosure policies include a long trial period, a money-back guarantee, or a no-hassle return with a long return window. The e-ink tablet mentioned above falls into this category. For a revolutionary product, the market usually has enough confidence in it. My results, therefore, predict that the optimal disclosure strategy would create some "die-hard fans", and the rest of the potential consumers obtain noisy signals. This feature matches what we observed on, for example, iPhone and Tesla.

The model produces some surprising comparative statics. The conventional wisdom in the search literature is that a higher search cost makes the buyer more likely to buy without search, which enables the seller to charge a higher price. In my model, however, although the robust price as a function of the search cost is increasing nearly everywhere, the function can "jump down" at one point (see Figure 6 below). This feature stems from the trade-off between demand and extraction. As discussed above, when the search cost is small, the seller does not use a deterrence policy and charges a higher price. As the search cost increases, the price threshold also increases, and hence a deterrence policy becomes more attractive. When the search cost is sufficiently large, the demand advantage makes it profitable for the seller to switch to a deterrence policy even if she must lower the price. Accordingly, the seller may charge a lower price in exchange for more effective search deterrence.

Another comparative statics result I derive is that, for a large range of parameters, as the search cost increases, the seller's information disclosure policy becomes more informative. As long as an increase in the search cost does not cross the point where the price jumps down, the robust price increases. In most cases, raising the price requires generating more favorable beliefs more frequently, which leads to a more informative disclosure policy. Moreover, if an increase in the search cost makes the seller switch to a deterrence policy, such a strategy pools beliefs at the top and hence is more likely to reveal that the match value is high, which increases informativeness. The insight that emerges from these comparative statics results is that although some technological advancements may reduce search costs, they can lead to higher prices and noisier information provision for certain new products.

To better understand how search frictions and robustness concerns drive the results, I investigate two variations of the main model. In the first variation, the buyer can discover the value of her outside option for free. In the absence of search frictions, the seller's incentive to pool beliefs disappears, and her robustness concerns render spreading out beliefs continuously and evenly optimal. In particular, the seller never provides full information. In the second variation, I eliminate the seller's robustness concerns by assuming that she knows the exact outside option distribution. The trade-off between demand and extraction persists, as well as the features caused by it, including the nonmonotonicity of the optimal price. Interestingly, full disclosure is always optimal. The intuition is that full disclosure maximally differentiates the seller's product from the buyer's outside option, allowing the seller to extract a greater surplus. This feature stands in stark contrast to the main model: when the seller seeks robustness, since full disclosure pools beliefs at the top (and hence it is a deterrence policy), it can only be optimal when the price is below a threshold. This threshold may limit the scope of surplus extraction even if there is ample differentiation, in which case partial disclosure can be optimal.

As an extension, I allow the seller to recognize whether the buyer is a first-time visitor or is coming back from search: she can either commit to an exploding offer, commit to a buy-now discount, or increase the posted price when the buyer comes back from search. I show that exploding offers and lack of commitment to the posted price lead to the same outcome, which is superior to the outcome in the baseline model, and buy-now discounts need not be useful. Furthermore, I argue that allowing for an extra fixed outside option that the buyer can consume without incurring a cost would not change the results qualitatively, and many important observations continue to hold if the match value is continuously distributed on an interval.

1.1 Related Literature

While a majority of the literature on selling a new product focuses on strategic pricing,[1] there are a few papers that consider the case where sellers can choose both the price and an information disclosure policy.[2] Heiman and Muller (1996) study how the length of demonstration affects the probability of purchasing different kinds of new prod-

[1]For a survey, see Chatterjee (2009). Many papers cited therein study the pricing dynamics of new products, an issue that I abstract away.

[2]In Milgrom and Roberts (1986), the seller of a new product chooses both a price and an advertisement spending level. Although the choice of the latter signals the quality of the product, it does not have any information content.

ucts. Fainmesser, Lauga, and Ofek (2021) consider a model in which the seller provides information about the new product's quality to first-generation consumers, and second-generation consumers learn about the quality through first-period buyer's product reviews.[3] Boleslavsky, Cotton, and Gurnani (2017) model competition between a firm selling a new product with an unknown match value and a firm selling an established alternative whose match value is known. The authors show that if the innovative firm sets the price first and then chooses the disclosure policy, partial disclosure is optimal; if the pricing decision has to be made after the choice of the disclosure policy, however, full disclosure is optimal. To the best of my knowledge, this paper is the first to study the roles of search frictions and the seller's robustness concerns in selling a new product.

On a higher level, this paper lies at the intersection of two strands of literature: robust monopoly pricing, and monopoly pricing with information disclosure and consumer search. Unlike papers studying robust monopoly pricing,[4] in my model the nonquantifiable uncertainty that the seller faces is not about the distribution over the buyer's valuation of the product she sells, but about the buyer's outside option distribution. Moreover, in my model, the seller also has control over how much information to disclose, which allows me to study the interaction between price and information.

There are some papers that connect monopoly pricing with information disclosure to consumer search. Anderson and Renault (2006), Wang (2017), and Lyu (2021) study the problem of pricing and information disclosure of a monopolist who is selling a search good; this paper focuses instead on experience goods. Furthermore, none of these papers consider a robustness-seeking seller. Among these papers, the most related one is Lyu (2021). In his model, the seller of a search good can disclose product information to the buyer with a private outside option. Upon seeing the signal realization, the buyer chooses whether to search (the true match value is revealed after search) or leave: the buyer cannot buy without search.

In my model, search frictions create a search deterrence motive for the seller. One strand of literature explores price-based deterrence tactics. Armstrong and Zhou (2016) study how a seller could use price tools, including buy-now discounts, exploding offers, and nonrefundable deposits, to deter the buyer from searching for products from com-

[3]Although these authors call it "quality", it is initially unobservable to the seller; in particular, it can be interpreted as the match value.

[4]Representative contributions in this literature include (this list is by no means exhaustive) Bergemann and Schlag (2008, 2011), Carrasco, Luz, Kos, Messner, Monteiro, and Moreira (2018), Du (2018), Hinnosaar and Kawai (2020), and Che and Zhong (2022). For a recent survey on robust contracting, see Carroll (2019).

peting sellers.[5] Another strand of literature allows a seller to strategically increase the search cost; such strategies are known as "search obfuscation".[6] Bar-Isaac, Caruana, and Cuñat (2010) consider a monopoly seller who can increase the search cost by choosing a marketing strategy that makes it harder for a buyer to learn her true valuation.

Unlike these papers, to dissuade the buyer from searching, the seller in my model uses a different channel: information disclosure. As far as I am aware, the only paper that studies a similar channel is Wang (2017). His model is similar to that studied in Bar-Isaac et al. (2010), but instead of choosing the search cost directly, the search cost is fixed, and the seller discloses information to prevent the buyer from searching for a finer signal about her product. In my model search deterrence is, like in Armstrong and Zhou (2016), with respect to an outside option.

As mentioned, for each price, my model can be thought of as a zero-sum game between the seller and an adversary. This feature connects this work, from a technical perspective, to the study of information design contests. In this literature, each sender provides information about her object by committing to a disclosure policy, and the sender with the most appealing signal realization wins. Boleslavsky and Cotton (2015, 2018) work on a setting where the receiver's prior is binary, and Hwang, Kim, and Boleslavsky (2019) conduct analysis with continuous priors; the sellers are also allowed to choose prices in the latter. He and Li (2021) and Au and Whitmeyer (2022) add search frictions to information design contests. However, the economic focus of my paper is different from the papers in this literature, which leads to distinct insights.

Finally, this work is related to information design under non-probabilistic uncertainty.[7] In Kosterina (2022) the sender faces non-probabilistic uncertainty over the receiver's prior: it may depart from a "reference prior" to some degree. Hu and Weng (2021) and Dworczak and Pavan (2022) assume that it is common knowledge that the seller and the buyer share the same prior, and the uncertainty concerns the receiver's additional signal. Most closely related is Sapiro-Gheiler (2021), who considers a setting in which the receiver takes the

[5] The theoretical predictions therein are experimentally tested in Brown, Viriyavipart, and Wang (2018) and Pan and Zhao (2020).

[6] For a review of this literature that covers both theory and empirics, see Ellison (2016).

[7] I list only the papers in which senders maximize their worst-case payoffs, as in this paper. Two other papers consider a different objective: the sender minimizes the difference between the worst-case scenario and the no-uncertainty benchmark. In Babichenko, Talgam-Cohen, Xu, and Zabarnyi (2021) the sender has non-probabilistic uncertainty over the receiver's utility function. Parakhonyak and Sobolev (2022) assume that the sender evaluates the worst-case among all possible joint distributions over the state, the prior, and the outside option of the receiver.

sender's preferred action if and only if the posterior mean of the state exceeds the receiver's outside option, and the seller faces non-probabilistic uncertainty over the outside option distribution.[8] Adding the pricing channel shifts the designer's objective from increasing the odds of certain actions being played to maximizing profit, which, together with search frictions, generates novel insights into the design of information disclosure policies.

2 The Model

A seller of a product (Seller) faces a risk-neutral buyer (Buyer) whose match value with the product is $x \in \{0, 1\}$, with a commonly known distribution such that $\mu = \mathbb{P}(x = 1) \in (0, 1)$. The probability μ can be interpreted as the (prior) mean match value. Initially, neither Seller nor Buyer knows the realization of x. Seller's production cost is assumed to be zero.[9] Seller chooses a price p and a *disclosure policy* (χ, S) consisting of a signal space S and a mapping $\chi : \{0, 1\} \to \Delta(S)$. It is well-known that a cumulative distribution function (cdf) H over posteriors $w \in [0, 1]$ can be induced by a disclosure policy if and only if it satisfies the constraint

$$\int_0^1 w \, \mathrm{d}H(w) = \mu; \tag{1}$$

that is, the expected posterior equals the prior.[10] Letting $\mathcal{M}(\mu)$ denote the set of all distributions over posteriors that satisfy (1), the analysis can be recast as Seller choosing $(p, H) \in [0, 1] \times \mathcal{M}(\mu)$ instead of $(p, (\chi, S))$. I call (p, H) a *selling strategy*, and refer to the choice of H as the choice of a disclosure policy. If the realized posterior is w, Buyer's *net value* of purchasing Seller's product is given by $w - p$.

I assume that Buyer has unit demand and that she has an outside option with an unknown value; to discover its value v, a search cost $s \geq 0$ must be incurred. One way to interpret this is a reduced form of Buyer's sequential search. Buyer knows that v is distributed according to cdf G with $\mathrm{supp}(G) \subseteq [0, 1]$.[11] Seller cannot observe v, and does not even know the distribution G; she only knows that the mean of G is ξ, and G is supported

[8] The problem studied by Sapiro-Gheiler (2021) is mathematically equivalent to an information design contest, and hence the results therein share many common features of the results in that literature.

[9] This is equivalent to assuming that the trade between Seller and Buyer is socially efficient.

[10] See, for example, Kamenica and Gentzkow (2011).

[11] For any cdf F, $\mathrm{supp}(F)$ denotes the support of F:

$$\mathrm{supp}(F) = \{w : F(w + \varepsilon) - F(w - \varepsilon) > 0 \text{ for all } \varepsilon > 0\}.$$

on a subset of $[0, 1]$. To focus on interesting cases in which Buyer prefers to search to buying nothing, I assume that the search cost satisfies

$$s < \xi. \tag{2}$$

I study Seller's problem of maximizing the revenue guarantee, which is the worst case (expected) revenue generated by an outside option distribution whose mean is ξ and the upper bound of the support is 1. Metaphorically, after Seller chooses (p, H), a "malevolent" adversary (henceforth Nature) chooses a distribution G with mean ξ to minimize Seller's payoff. I will use this metaphor in the analysis below, as it is helpful in solving Seller's optimization problem as well as interpreting the results.

Let

$$S_G(t) := \mathbb{E}_G[\max\{t, v\}] - t = \int_t^1 (v - t)\, \mathrm{d}G(v)$$

denote the expected benefit of search when Buyer's net value of purchasing Seller's product is t and the outside option is distributed according to some G with mean ξ. Let a be such that

$$S_G(a) = s. \tag{3}$$

Assumption (2) ensures that such a exists.[12] Buyer will purchase Seller's product without search whenever the expected benefit from search is no more than the search cost, that is, $S_G(w-p) \leq s$. Then since S_G is decreasing, (3) implies that this is equivalent to $w-p \geq a$.[13] Intuitively, a represents the net surplus the buyer needs to obtain from the seller to forgo search, or (in jargon) the *reservation value* of the outside option. It can be checked that $a \in [\xi - s, 1 - s/\xi]$, where the lower bound can be induced by δ_ξ, the degenerate distribution at $v = \xi$, and the upper bound is uniquely induced by the binary distribution with support $\{0, 1\}$ and mean ξ.[14] By taking convex combinations of these two distributions, any $a \in$

[12]Since G is chosen by Nature, a is dependent on G. For notational ease, I suppress the dependence and simply write a.

[13]Unless explicitly specified, I use "increasing" and "decreasing" in the weak sense; that is, "increasing" means "weakly increasing". "Strictly" would be added whenever needed–similarly with "positive", "negative", "above", "below", "more", and "less".

[14]Denoting its cdf by G_B,

$$S_{G_B}(a) = \mathbb{E}_{G_B}[\max\{v, a\}] - a = \xi \max\{1, a\} + (1 - \xi) \max\{0, a\} - a = \xi(1 - a),$$

where the last equality follows from the fact that $a \in [0, 1]$. Then by Equation (3), $S_{G_B}(a) = \xi(1 - a) = s$, and thus $a = 1 - s/\xi$.

In search problems, it is well known that (see, for example, Kohn and Shavell (1974)), *ceteris paribus*, a

$(\xi - s, 1 - s/\xi)$ can be achieved.

If Buyer, instead, prefers to investigate the outside option, that is, if $w - p < a$, she will return to buy Seller's product when the outside option turns out to be worse than Seller's offer, that is, when $w - p > v$.[15] I assume that Seller cannot recognize whether Buyer is a first-time visitor or not; consequently, she cannot change the price when Buyer comes back from search.[16]

The timing of the game is as follows:

1. Seller chooses (p, H);
2. Nature observes Seller's choice and chooses a distribution G.
3. Buyer's posterior w realizes according to H, and she also sees G; she buys immediately if $w - p \geq a$. Otherwise, she pays the search cost s and observes a realization of v from G.
4. Buyer returns to Seller to buy if $w - p > v$.

By choosing (p, H), for a given distribution over outside options G, Seller's expected revenue is given by

$$\Pi(p, H \mid G) := p\, \mathbb{E}_G[1 - H(p + \min\{a, v\})]. \tag{4}$$

Equation (4) is intuitive: Buyer *eventually* purchases Seller's product if and only if either $w - p \geq a$ or $w - p > v$. Consequently, $\mathbb{E}_G[1 - H(p + \min\{a, v\})]$ is the probability of eventual purchase, or the demand that Seller faces, under G.

Since Buyer's optimal behavior after any history is already embedded in the description of the game, the analysis reduces to Seller's revenue guarantee maximization problem, given by

$$\max_{(p,H)\in[0,1]\times\mathcal{M}(\mu)} \min_{G\in\mathcal{M}(\xi)} \Pi(p, H \mid G),$$

decision maker prefers a more dispersed distribution. In my setting, this arises because a higher incidence of very good outside options increases the value of searching, while the higher incidence of very bad outside options is not too detrimental because Buyer can always return to Seller and buy there. Now observe that the degenerate distribution is the most "concentrated" distribution with mean ξ, and the binary distribution is the most dispersed one.

[15] The tie-breaking assumption is implicitly embedded in the two inequalities above: Buyer does not search if she is indifferent between search or not, and she does not return to Seller if she is indifferent between Seller's offer and her outside option. This assumption is not only realistic, but also necessary for an equilibrium to exist.

[16] For the same reason, Seller cannot benefit from offering a menu of prices depending on Buyer's report of the outside option she discovered: no matter whether Buyer has searched or not, she would report the outside option that is associated with the lowest price. In Section 4.3 I discuss what could happen if Buyer identity is recognizable.

A solution to this problem is called a *robustly optimal selling strategy*; its components are called *robust price* and *robustly optimal disclosure policy*, respectively.

Equivalently, Seller solves

$$\max_{p\in[0,1]}\left\{\max_{H\in\mathcal{M}(\mu)}\min_{G\in\mathcal{M}(\xi)}\Pi(p,H\mid G)\right\}.$$

Let $\Phi(p) := \max_{H\in\mathcal{M}(\mu)}\min_{G\in\mathcal{M}(\xi)}\Pi(p,H\mid G)$. Seller's problem can be solved in two steps: first, Seller chooses $H\in\mathcal{M}(\mu)$ to maximize $\min_{G\in\mathcal{M}(\xi)}\Pi(p,H\mid G)$, namely Seller's revenue guarantee for a fixed price p; second, Seller chooses $p\in[0,1]$ to maximize $\Phi(p)$.

2.1 The Role of Information Disclosure

It can be readily seen from (4) that the "demand" of Seller's product for a fixed price p and a fixed outside option distribution G is determined by Seller's information disclosure policy, namely the choice of the distribution over posteriors H. Furthermore, (4) implies that Seller's expected payoff, taking the outside option distribution G as given, is

$$\begin{aligned}\Pi(p,H\mid G) &= p\left[\int_a^1 (1-H(p+a))\,\mathrm{d}G(v)+\int_0^a (1-H(p+v))\,\mathrm{d}G(v)\right]\\ &= p\left[1-H(p+a)+\int_0^a (H(p+a)-H(p+v))\,\mathrm{d}G(v)\right].\end{aligned}$$

For a fixed price $p\in[0,1]$, the first part in squared brackets, $1-H(p+a)$, is the probability that Buyer purchases without search when the distribution over posteriors is H. The second part in squared brackets,

$$\int_0^a (H(p+a)-H(p+v))\,\mathrm{d}G(v)=\mathbb{E}_G[\max\{H(p+a)-H(p+v),0\}],$$

is the probability that Buyer returns to buy Seller's product after searching: $H(p+a)$ is the probability that Buyer chooses to search, and $H(p+v)$ is the probability that Buyer chooses the outside option after searching (that is, the probability of $w-p\le v$) when the value of the outside option is v. I call the first term *buy-now demand* and the second term *buy-later demand* in the parlance of Armstrong and Zhou (2016).

Fix an outside option distribution G, and hence a is also fixed. If p is such that $p+a<1$, by pooling more beliefs at or above $p+a$, $1-H(p+a)$ increases and thus Buyer is more likely to buy without search. At the same time, it may also affect Buyer's buy-later demand. Therefore, one can say that information disclosure guides Buyer's decision regarding whether to search and return to Seller to buy or not after searching.

3 Main Results

As hinted above, I solve for Seller's robustly optimal selling strategy in two steps. Section 3.1 concerns the first step of finding the optimal disclosure policy for each fixed price; the interaction between price and information is also discussed. In Section 3.2 I proceed to find the robustly optimal selling strategy by solving for the robust price.

3.1 Optimal Disclosure Policy For a Fixed Price

In this subsection, I consider the first step of Seller's problem. For any fixed $p \in [0, 1]$, Seller's robustly optimal choice of disclosure policies is summarized in Proposition 1. Proof of all results in Section 3 are relegated to Appendix A.

Proposition 1. *Suppose $p > s/\xi$. If $\mu > (1 + p)/2$, the distribution over posteriors $U_{[2\mu-1,1]}$ is optimal.*[17] *If $\mu \leq (1 + p)/2$, there exists $\bar{w} \in [2\mu - p, 1]$ such that the optimal distribution is*[18]

$$H_{\bar{w}}(w \mid p) = \begin{cases} 1 - \frac{2\mu}{\bar{w}+p} & \text{if } w \in [0, p), \\ 1 - \frac{2\mu}{\bar{w}+p} + \frac{2\mu}{\bar{w}+p}\left(\frac{w-p}{\bar{w}-p}\right) & \text{if } w \in [p, \bar{w}), \\ 1 & \text{if } w \in [\bar{w}, 1]. \end{cases} \tag{5}$$

Now suppose $p \leq s/\xi$. If $\mu \geq p + 1 - s/\xi$, the degenerate distribution δ_μ is optimal. If $\mu < p + 1 - s/\xi$ and $p \geq (1 - 2\xi)(\xi - s)/(2\xi^2)$, the binary distribution with support on $\{0, p + 1 - s/\xi\}$ is optimal. Otherwise, there are two cases:

(i) if $p + (1 - s/\xi)/2 \leq \mu < p + 1 - s/\xi$, the optimal distribution is

$$H_u^h(w \mid p) = \begin{cases} 0 & \text{if } w \in [0, p), \\ 2\frac{\xi^2(p+1-\mu)-s\xi}{(\xi-s)^2}(w - p) & \text{if } w \in [p, p + 1 - s/\xi), \\ 1 & \text{if } w \in [p + 1 - s/\xi, 1]. \end{cases} \tag{6}$$

(ii) if $\mu < p + (1 - s/\xi)/2$, there exists $\bar{w} \in [2\mu - p, p + 1 - s/\xi)$ such that the optimal distribution is $H_{\bar{w}}(\cdot \mid p)$ defined in (5).

---

[17] $U_{[a,b]}$ is the cdf of the uniform distribution over $[a, b]$.

[18] See Appendix A.1.2 for the definition of $\bar{w}$.

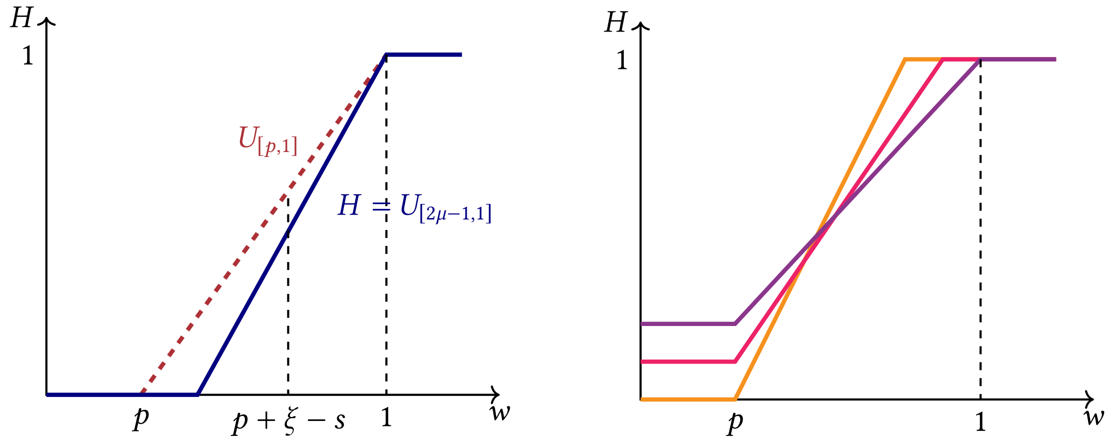

Figure 1: Distributions that can be optimal when $p > s/\xi$. The left panel corresponds to the case that $\mu > (1+p)/2$, and the right panel corresponds to $\mu \leq (1+p)/2$. In the right panel, the orange curve corresponds to an optimal distribution with $\bar{w} < 1$ and no mass point at $w = 0$, the pink curve corresponds to an optimal distribution with $\bar{w} < 1$ and a mass point at $w = 0$, and the violet curve corresponds to an optimal distribution with $\bar{w} = 1$ and a mass point at 0.

Observe that in all cases an optimal distribution has an affine component. This is because affinity of the distribution over posteriors H on $(p, \sup\{\text{supp}(H)\})$ hedges well against Nature. To see why, let $\tilde{w} := \sup\{\text{supp}(H)\}$, and note that the probability of eventual purchase satisfies (recall Equation (3))

$$\mathbb{E}_G[1 - H(p + \min\{a, v\})] = 1 - H(p + \mathbb{E}_G[\min\{a, v\}]) = 1 - H(p + \xi - s), \tag{7}$$

where the first equality holds by affinity, and the second equality follows from the definition of a in (3) and the fact that $\mathbb{E}_G[v] = \xi$:

$$\mathbb{E}_G[\min\{a, v\}] = \mathbb{E}_G[v + a - \max\{a, v\}] = \mathbb{E}_G[v] - S_G(a) = \xi - s.$$

In words, no matter what outside options distribution Nature chooses, the probability of eventual purchase is the same. Put differently, Nature is indifferent between spreading and contracting mass in designing the outside option distribution under affinity. This guarantees that there is not a single choice of the outside option distribution that Nature can take significant advantage of, which gives rise to the desired robustness. In the proof, I show that for every set of parameters there exists an optimal distribution that is affine on $(p, \tilde{w})$.

An important consequence of this observation is that there is no mass point in an optimal distribution on $(p, \tilde{w})$. However, Seller may have an incentive to pool mass at $\tilde{w}$. If Seller knows the distribution over outside options G and hence a, so long as the price is such that $p + a < 1$, she could benefit from deterring search by putting a mass point at $p + a$: when Buyer gets the posterior $w = p + a$, she would buy without search. When Seller takes a robust approach, however, she is only able to deter search if the sum of the price and the *maximum* reservation value is below one, that is, $p+1-s/\xi \leq 1$, or $p \leq s/\xi$. When $p > s/\xi$, if Seller attempts to deter search by setting an atom at any $w \in [p, 1]$, Nature can always frustrate it by choosing G such that $a = w - p + \varepsilon$ for some $\varepsilon > 0$ small enough: by doing this, $w - p = a - \varepsilon < a$ and hence Buyer would search for sure. Intuitively, when Seller's price is sufficiently high, she believes that Nature is always able to make the outside option attractive enough so that deterring search is not possible no matter how optimistic Buyer's prior is. When $p \leq s/\xi$, however, placing a mass point at $w = p+1-s/\xi$ can be helpful. This leads to the second observation: when $p > s/\xi$, there is no mass point in an optimal distribution on $[p, 1]$; and when $p \leq s/\xi$, the only possible mass point in an optimal distribution on $[p, 1]$ is at $w = p + 1 - s/\xi$.

To understand Proposition 1 in more detail, below I thoroughly discuss the two cases, $p > s/\xi$ and $p \leq s/\xi$.

The case of $p > s/\xi$. As discussed above, in this case, an optimal distribution over posteriors H is affine on $(p, \tilde{w})$ and has no mass point on $[p, 1]$. When $\mu \leq (1+p)/2$, if H is affine on $(p, 1)$, it must be that the posterior w is uniformly distributed on $[p, 1]$. Note that, however, the uniform distribution over $[p, 1]$, denoted by $U_{[p,1]}$, has mean $(1+p)/2 \geq \mu$. In fact, if $\mu = (1 + p)/2$, the mean of $U_{[p,1]}$ is exactly μ and is indeed optimal. If $\mu < (1+p)/2$, however, to maintain that (i) H is affine on $(p, \tilde{w})$ and (ii) the mean of H is μ, it must be that either (a) $\tilde{w}$ is strictly less than 1, or (b) H puts some mass on $[0, p)$ (more formally, $H(p) > 0$), or (c) both. Observing that the best way of putting mass on $[0, p)$ is to have a mass point at 0,[19] (a), (b), and (c) correspond to the orange, pink, and violet curves in the right panel of Figure 1. This gives rise to the formula of optimal distribution for this case in (5).

When $\mu > (1 + p)/2$, the power of Nature's adversarial choice neutralizes the advantage of a larger mean match value: the probability of eventual purchase can be at most

[19] Since the expected posterior must be the prior μ, by moving mass on $(0, p)$ to 0, Seller can put more mass on $[p, 1]$, which would increase the probability of eventual purchase.

as large as in the case of $\mu = (1+p)/2$ where the optimal distribution is $U_{[p,1]}$. To see this, consider the binary outside option distribution $\breve{G}$ whose probability mass function $\breve{g}$ is given by

$$\begin{array}{c|c|c} v & 0 & \frac{(1-p)\xi}{\xi-s} \\ \hline \breve{g}(v) & 1-\frac{\xi-s}{1-p} & \frac{\xi-s}{1-p} \end{array}. \tag{8}$$

It can be checked that $\mathbb{E}_G[v] = \xi$, and $a = 1-p$. Because $p > s/\xi$, $1-p < 1-s/\xi$ and thus $\breve{G}$ is feasible for Nature. Then the probability of eventual purchase is

$$\begin{aligned} &\mathbb{E}_G[1 - H(p + \min\{a, v\})] \\ &= 1 - \left(1 - \frac{\xi - s}{1-p}\right) H(p + \min\{0, a\}) - \frac{\xi - s}{1-p} H\left(p + \min\left\{\frac{(1-p)\xi}{\xi - s}, a\right\}\right) \\ &= 1 - \left(1 - \frac{\xi - s}{1-p}\right) H(p) - \frac{\xi - s}{1-p} H(1) \\ &\leq 1 - \frac{\xi - s}{1-p} \end{aligned} \tag{9}$$

where the second equality holds because $a = 1-p$, and the inequality follows from the facts that $H(p) \geq 0$ and $H(1) = 1$. Because $U_{[p,1]}$ is affine on $(p,1)$, (7) indicates that the probability of eventual purchase is exactly $1-(\xi-s)/(1-p)$. In particular, $H = U_{[2\mu-1,1]}$ has mean μ and attains this upper bound: taking $\breve{G}$ as given, inequality (9) holds as equality since $H(p) = 0$; and it can be shown that there is no outside option distribution outperforms this one for Nature. Consequently, $H = U_{[2\mu-1]}$ is optimal; this distribution and $U_{[p,1]}$ correspond to the solid blue curve and the dashed maroon curve in the left panel of Figure 1, respectively.

The case of $p \leq s/\xi$. In this case, Seller's price is low enough so that pooling mass at $w = p+1-s/\xi$ can be helpful in deterring search. If $\mu \geq p+1-s/\xi$, Buyer is sufficiently optimistic about the product match value of Seller's product, and thus disclosing no information is optimal for Seller because Buyer buys without search with probability 1. This corresponds to the degenerate distribution over posteriors δ_μ depicted in the lower panel of Figure 2. By providing nontrivial information, however, posteriors higher than the prior does not help Seller, and she is hurt by low posteriors as such realizations may make Buyer decide to search.

The more interesting case is $\mu - 1 + s/\xi < p \leq s/\xi$. In this region, hedging against Nature becomes relevant, which calls for the affinity of the distribution over posteriors

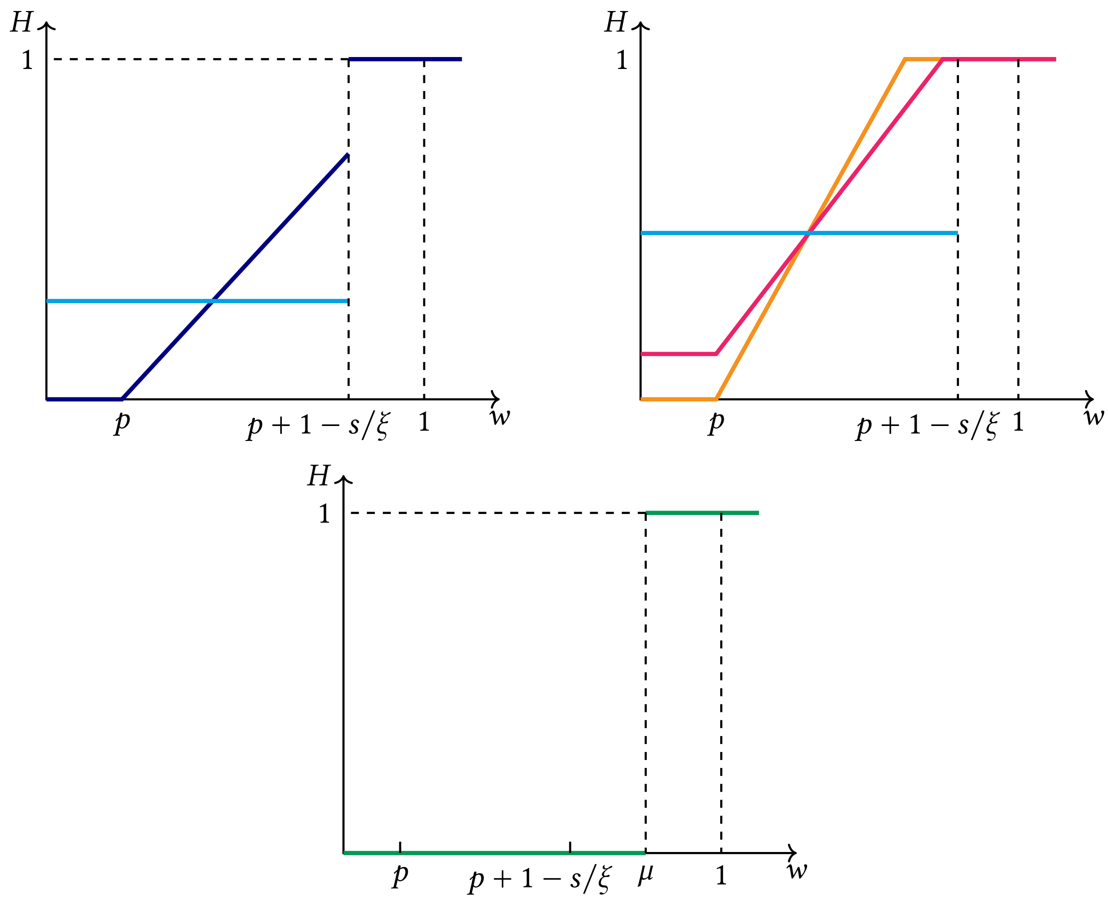

Figure 2: Distributions that can be optimal when $p \le s/\xi$. The upper left panel corresponds to the case that $p + 1 - s/\xi > \mu > p + (1 + s/\xi)/2$; the blue curve is H_u^h, and the light blue curve is the binary distribution. The upper right panel corresponds to $\mu \le p + (1 + s/\xi)/2$; again the light blue curve is the binary distribution, and the yellow and pink curves correspond to $H_{\tilde{w}}$ with and without a mass point at 0, respectively. The lower panel depicts the case of $\mu \ge p + 1 - s/\xi$.

on $(p, \tilde{w})$. When p is relatively large (that is, $p \ge (1 - 2\xi)(\xi - s)/(2\xi^2)$), to guarantee a decent chance of eventual purchase, Seller must make sure that her product is sufficiently differentiated from the outside option. A natural candidate for this purpose is the binary distribution, which not only generates the desired differentiation, but also allows Seller to take advantage of deterrence in the sense that Buyer would buy immediately if the high posterior realizes. Such a distribution is displayed in the upper panels of Figure 2.

When p is relatively small, intermediate posteriors can also be useful in persuading Buyer to buy. In fact, since the expected posterior must be the prior, high posteriors are "produced" at the cost of generating low posteriors that "push" Buyer to her outside option more often; in this sense, generating intermediate posteriors is more "cost-efficient". When the mean match value μ is relatively large for the price, it is optimal for Seller to pool beliefs at $p + 1 - s/\xi$ to create some "safe demand",[20] and also generate some intermediate posteriors since they are "less costly". This gives rise to H_u^h defined in (6), which corresponds to the navy blue curve in the upper left panel of Figure 2. When μ is relatively small, generating search-deterring posteriors is too cost-inefficient, and hence the optimal disclosure policy produces exclusively intermediate posterior beliefs. Then affinity implies that the robustly optimal distribution over posteriors is isomorphic to the $\mu \leq (1 + p)/2$ case when $p > s/\xi$. The pink and yellow curves in the upper right panel of Figure 2 illustrate the optimal distribution with or without a mass point at $w = 0$, respectively.

3.1.1 The Interaction Between Price and Information

Proposition 1 also allows us to understand how Seller's choice of distribution over posteriors changes in the price.

Corollary 1. *(1) If $p > s/\xi$, as p increases, the optimal disclosure policy becomes Blackwell more informative.*

(2) If $p \leq s/\xi$ there are two cases:

- *(i) if p_1, p_2 are such that $p_1 < p_2 < (1 - 2\xi)(\xi - s)/(2\xi^2)$ and $p_1, p_2 \in [(\mu - (1 - s/\xi))/2, \mu - (1 - s/\xi))$, the corresponding optimal disclosure policies cannot be Blackwell ranked;*
- *(ii) otherwise, as p increases, the optimal disclosure policy becomes Blackwell more informative.*

When $p > s/\xi$, Corollary 1 indicates that price and information are substitutes: if Seller would like to increase her price, it is optimal for her to provide more information. Intuitively, as price increases, the distribution over posterior generates higher posteriors more often so that the likelihood of eventual purchase does not fall too much. Then to ensure that the resulting distribution has the same mean, there must be a commensurate increase in the likelihood of lower posteriors. Thus, the new optimal distribution over posteriors must be a mean-preserving spread (MPS) of the previous one; this is illustrated

[20]It is "safe" because Buyer buys immediately whenever the realized posterior is $w = p + 1 - s/\xi$, and hence unaffected by Nature's choice of the outside option distribution.

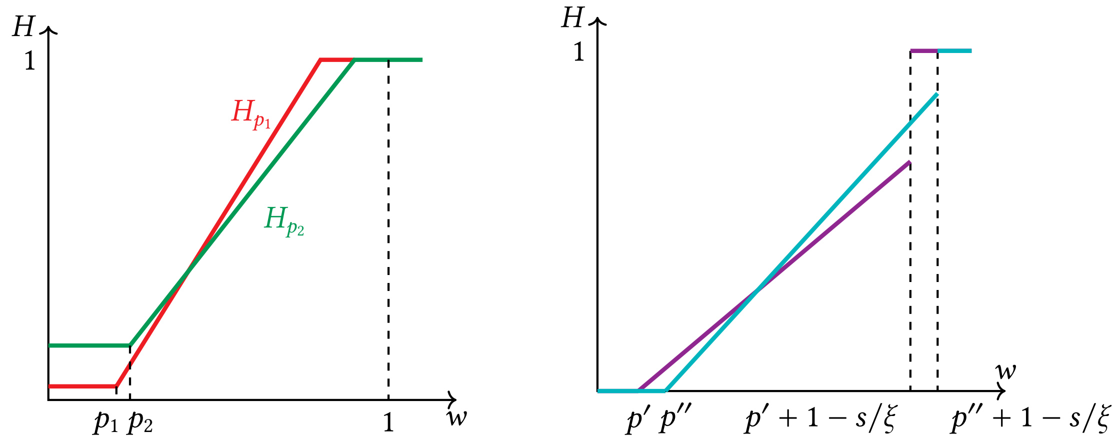

Figure 3: The left panel displays the optimal distributions corresponding to p_1 and p_2 with $s/\xi < p_1 < p_2$, denoted by H_{p_1} and H_{p_2}. It can be readily seen that the latter is a mean-preserving spread of the former. The right panel illustrates the case that the optimal distributions corresponding to p' and p'' with $p' < p'' \leq s/\xi$ cannot be Blackwell ordered.

in the left panel of Figure 3.[21] It is well-known that a disclosure policy is more Blackwell informative than another if and only if it is a MPS of the latter,[22] and hence the optimal distribution over posteriors associated with the higher price is more informative.

In the case of $p \leq s/\xi$, the idea of "higher price is paired with more precise information" vaguely persists, but not exactly. When $p \leq \mu - (1 - s/\xi)$, the price is so low that the highest reservation value that Nature could generate, namely $1 - s/\xi$, is below $\mu - p$; hence, by providing no information, the probability of (eventual) purchase is 1. In contrast, when the price is sufficiently high, that is, $p \geq (1 - 2\xi)(\xi - s)/(2\xi^2)$ (provided that $p > \mu-(1-s/\xi)$), a binary distribution with support on $\{0, p+1-s/\xi\}$ is optimal according to Proposition 1. Evidently, in this price range, as price increases, the binary distribution becomes more Blackwell informative: to maximize the chance of making a sale, a higher price must be accompanied by a binary distribution whose higher point of the support is larger. Finally, for an intermediate price, the optimal distribution over posteriors falls between the two extremes: it is not completely uninformative, but not as informative as any binary distribution that can be optimal.

The only exception is that when p_1, p_2 are such that $p_1 < p_2 < (1 - 2\xi)(\xi - s)/(2\xi^2)$ and $p_1, p_2 \in [(\mu-(1-s/\xi))/2, \mu-(1-s/\xi))$, the respective optimal distributions, H_{p_1} and

[21] Let F_1 and F_2 be two distributions defined on $[0, 1]$. F_1 is a mean-preserving spread of F_2 if $\int_0^x F_2(s)\mathrm{d}s \leq \int_0^x F_1(s)\mathrm{d}s$ for all $x \in [0, 1]$, where the inequality holds with equality at $x = 1$.

[22] This is formally established in Blackwell (1953).

H_{p_2}, cannot be ranked in the Blackwell order. This is because, as shown in the right panel of Figure 3,

$$\inf\left\{\text{supp}\left(H_{p_1}\right)\right\} < \inf\left\{\text{supp}\left(H_{p_2}\right)\right\} < \sup\left\{\text{supp}\left(H_{p_1}\right)\right\} < \sup\left\{\text{supp}\left(H_{p_2}\right)\right\}, \quad (10)$$

and hence not any of them is more "spread out" than another. This exception, however, has an intuitive explanation based on the forces of affinity and mass points described above. In this price range, the mean is high enough relative to prices; consequently, there is no need to have a mass point at $w = 0$. Then since $p \leq s/\xi$, the search deterrence motive gives rise to the mass point at $p + 1 - s/\xi$; and to hedge against Nature, the distribution is affine on $(p, p + 1 - s/\xi)$. Thus, for $p_1 < p_2$, (10) must hold: the joint effect of search frictions and robustness concerns renders the resulting distributions unrankable by Blackwell informativeness.

3.1.2 Proof Sketch for Proposition 1

Define Buyer's *effective outside option* as $z := \min\{v, a\}$, and let $\hat{G}$ denote its cumulative distribution function.[23] It can be shown that $z \in [0, 1-\xi/s]$, and $\mathbb{E}_{\hat{G}}[z] = \xi - s$. Observe that Nature's choice of outside option distribution only affects Seller and Nature's expected payoffs through the induced distribution over effective outside options:[24] using (4), Seller's expected revenue can be written as $\Psi\left(p, H \,\middle|\, \hat{G}\right) := \mathbb{E}_{\hat{G}}[1 - H(p + z)]$.[25] Consequently, for a fixed price, the seller's problem becomes

$$\max_{H \in \mathcal{M}(\mu)} \min_{\hat{G} \in \mathcal{M}(\xi - s)} \Psi\left(p, H \,\middle|\, \hat{G}\right). \quad (11)$$

It is well-known that the solution to this problem is a saddle point, or an equilibrium, of the zero-sum game in which Seller chooses H to maximize her revenue and Nature chooses an effective outside option distribution $\hat{G}$ to minimize it.[26]

It can be seen by inspecting problem (11) that observing Seller's choice of (p, H), Na-

[23] That is,

$$\hat{G}(z) := \begin{cases} G(z) & \text{if } z < a, \\ 1 & \text{if } z \geq a. \end{cases}$$

[24] This is first observed by Armstrong (2017) and Choi, Dai, and Kim (2018).

[25] Because in this step p is taken as given, I drop the multiplicative p from the objective functions to economize notation.

[26] See, for example, Proposition 22.2 (b) in Osborne and Rubinstein (1994).

ture's problem of choosing an effective outside option distribution is equivalent to

$$\max_{\hat{G}\in\mathcal{M}(\xi-s)} \int_0^{1-\frac{s}{\xi}} H(p+z)\,\mathrm{d}\hat{G}(z). \tag{12}$$

Define

$$G_p(w) := \begin{cases} 0 & \text{if } w < p, \\ \hat{G}(w-p) & \text{if } w \geq p; \end{cases}$$

then Seller's problem, taking Nature's choice as given, can be written as (after integration by parts)

$$\max_{H\in\mathcal{M}(\mu)} \int_0^1 G_p(w)\,\mathrm{d}H(w). \tag{13}$$

Consequently,

Lemma 1. *For a fixed p, $\left(H^*, \hat{G}^*\right)$ solves problem (11) if and only if*

$$H^* \in \operatorname*{arg\,max}_{H\in\mathcal{M}(\mu)} \int_0^1 G_p^*(w)\,\mathrm{d}H(w), \quad \textit{and} \quad \hat{G}^* \in \operatorname*{arg\,max}_{\hat{G}\in\mathcal{M}(\xi-s)} \int_0^{1-\frac{s}{\xi}} H^*(p+z)\,\mathrm{d}\hat{G}(z),$$

where

$$G_p^*(w) = \begin{cases} 0 & \textit{if } w < p, \\ \hat{G}^*(w-p) & \textit{if } w \geq p. \end{cases}$$

By Corollary 2 in Kamenica and Gentzkow (2011), the solution of problem (12) is identified by the concave hull of $H\big|_{[0,1-s/\xi]}(p+z)$;[27] and the value of problem (12) is just the concave hull evaluated at $\xi-s$, which I denote by $\widetilde{H}\big|_{[0,1-s/\xi]}(p+\xi-s)$. Similarly, the solution of problem (13) is identified by $\widetilde{G}_p$, and the value of problem (13) is given by $\widetilde{G}_p(\mu)$.

To find a worst-case effective outside option distribution, I first guess a candidate distribution over posteriors H^*. Next I find $\widetilde{H}^*\big|_{[0,1-s/\xi]}(p+z)$, which identifies the necessary and sufficient conditions that a worst-case distribution must satisfy: an effective outside option distribution $\hat{G}$ is a worst-case distribution if and only if $\hat{G}\in\mathcal{M}(\xi-s)$ and

$$\int_0^{1-\frac{s}{\xi}} H(p+z)\,\mathrm{d}\hat{G}(z) = \widetilde{H}\Big|_{[0,1-s/\xi]}(p+\xi-s).$$

[27] Let $f:\mathbb{R}\to\mathbb{R}$ be a function. The *concave hull* of f, denoted by $\tilde{f}$, is the smallest upper semicontinuous and concave function that majorizes f. Let $f\big|_{[c,d]}$ denote the restriction of f to $[c,d]\subset[0,1]$, and let $\tilde{f}\big|_{[c,d]}$ denote the concave hull of the restriction.

Then I find a distribution $\hat{G}^*$ that not only satisfy these conditions, but also makes H^* a solution to the problem $\max_{H \in \mathcal{M}(\mu)} \int_0^1 G_p^*(w)\,\mathrm{d}H(w)$; the latter can be checked by using $\widetilde{G}_p$. By Lemma 1, $(H^*, \hat{G}^*)$ solves problem (11). Finally, to show that H^* is indeed the robustly optimal distribution over posteriors, it only remains to find an outside option distribution G that induces $\hat{G}^*$. This step is nontrivial but mostly technical, and hence I do not explain it here.

3.2 Seller's Robustly Optimal Strategy

Now I am ready to solve for the optimal prices; together with the disclosure policies that are optimal for these prices, I will be able to identify Seller's robustly optimal selling strategy.

To state the main result, define

$$\underline{\mu} := \frac{2-\xi+s-\sqrt{2(\xi-s)-(\xi-s)^2}}{2\,(1-\xi+s)}, \text{ and } \bar{\mu} := 1-\frac{\sqrt{\xi-s}}{2};$$

algebra reveals that $\underline{\mu} < \bar{\mu}$ for all $\xi \in (0,1)$ and $s \in (0,\xi)$. Let p^* and p^{**} be given by

$$p^* := \begin{cases} \frac{1-\sqrt{2(\xi-s)-(\xi-s)^2}}{1-\xi+s} & \text{if } \mu \le \underline{\mu} \\ 2\mu-1 & \text{if } \underline{\mu} < \mu \le \bar{\mu} \\ 1-\sqrt{\xi-s} & \text{if } \mu > \bar{\mu} \end{cases}$$

and $p^{**} := s/\xi$, respectively.

Say that Seller uses *uniform disclosure* if $U_{[2\mu-1,1]}$ is used when $\mu > 1-\left(\sqrt{\xi-s}/2\right)$, and the distribution over posteriors is

$$H_h^*(w) = \begin{cases} 1-\frac{2\mu}{1+p^*} & \text{if } w \in [0, p^*) \\ 1-\frac{2\mu}{1-p^{*2}}(1-w) & \text{if } w \in [p^*, 1] \end{cases}$$

when $\mu \le 1-\left(\sqrt{\xi-s}/2\right)$.[28] Illustrated in panels (c) and (d) of Figure 4, uniform disclosure induces a distribution over posteriors that is uniform over $[p^*, 1]$, and may have a mass point at $w = 0$ when the mean match value μ is small. The information provision about the match value is "noisy" under uniform disclosure: every posterior belief $w \in [0,1]$ is

[28]This distribution is a convex combination of a point mass at $w = 0$ and a uniform distribution over $[p^*, 1]$.

assigned to a mix of high and low match values.

Say that Seller uses *full disclosure* if the distribution over posteriors is the binary distribution with support $\{0, 1\}$.[29] This distribution, illustrated in panel (b) of Figure 4, fully reveals the match value: if the posterior is 1 (0), the match value is high (low) with probability one.

Say that Seller uses *mixture disclosure* if the distribution over posteriors is

$$H_u^*(w) = \begin{cases} 0 & \text{if } w \in [0, p^{**}), \\ \frac{2\xi^2(1-\mu)}{(\xi-s)^2}(w - p^{**}) & \text{if } w \in [p^{**}, 1), \\ 1 & \text{if } w = 1. \end{cases} \tag{14}$$

As the name suggests, such a disclosure policy shares features of the previous two: as shown in panel (a) of Figure 4, it induces a distribution over posteriors that is uniform over $(p^{**}, 1)$ and has a mass point at $w = 1$. Like uniform disclosure, information provision is noisy; however, if posterior $w = 1$ realizes, the match value must be high.

For the last two disclosure policies, whenever $p \leq s/\xi$ and posterior $w = 1$ realizes, Buyer would buy without search because $w = 1 \geq p + 1 - s/\xi = p + \max a$. For this reason, I call them *deterrence policies*: no matter what outside option distribution Nature chooses, these disclosure policies guarantee that Buyer buys immediately with strictly positive probability.

Finally, let $B_1(\xi) = \xi(\xi-1)^2/(\xi^2+1)$, $B_2(\xi) = \xi(\xi-1)^2/(\xi+1)^2$, and $B_3(\xi) = \xi - 2\xi^2$. It can be checked that $B_1(\xi) > B_2(\xi)$ and $B_1(\xi) > B_3(\xi)$ for all $\xi \in (0, 1)$.

I am now ready to state the main result.

Theorem 1. *If $s \geq B_1(\xi)$, then full disclosure is optimal, and the robust price is $p^{**} = s/\xi$. If $s < B_2(\xi)$, then uniform disclosure is optimal, and the robust price is $p^* > s/\xi$. If $B_2(\xi) \leq s < B_1(\xi)$, there are two cases:*

(1) If $B_3(\xi) \leq s < B_1(\xi)$, then there exists $\hat{\mu} \in (0, 1)$ such that for $\mu < \hat{\mu}$, uniform disclosure is optimal, and the robust price is $p^* > s/\xi$; and for $\mu \geq \hat{\mu}$, full disclosure is optimal, and the robust price is $p^{**} = s/\xi$.

(2) If $B_2(\xi) \leq s < B_3(\xi)$, then there exists $\check{\mu} \in (0, 1)$ such that for $\mu < \check{\mu}$, uniform disclosure is optimal, and the robust price is $p^* > s/\xi$; and for $\mu \geq \check{\mu}$, mixture disclosure is optimal, and the robust price is $p^{**} = s/\xi$.

[29] See (23) for its cdf.

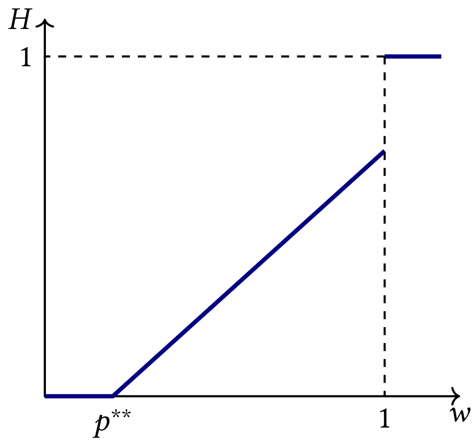

(a) Mixture disclosure

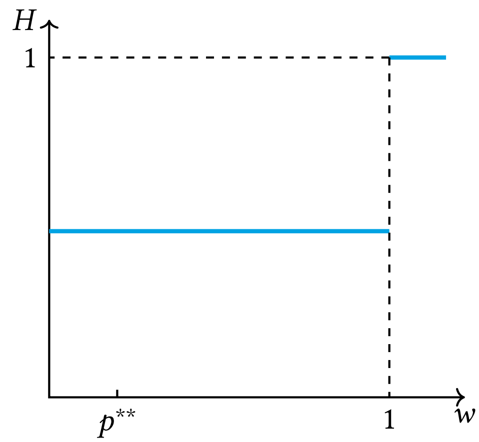

(b) Full disclosure

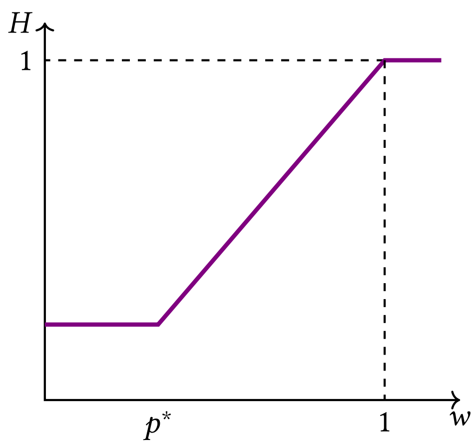

(c) Uniform disclosure when $\mu \leq 1 - \frac{\sqrt{\xi - s}}{2}$

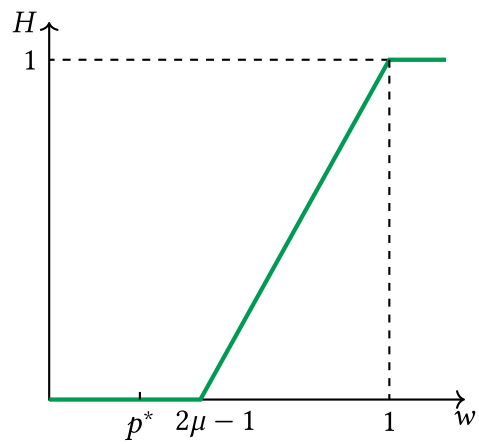

(d) Uniform disclosure when $\mu > 1 - \frac{\sqrt{\xi - s}}{2}$

Figure 4: Three kinds of information disclosure policies that can be optimal for Seller.

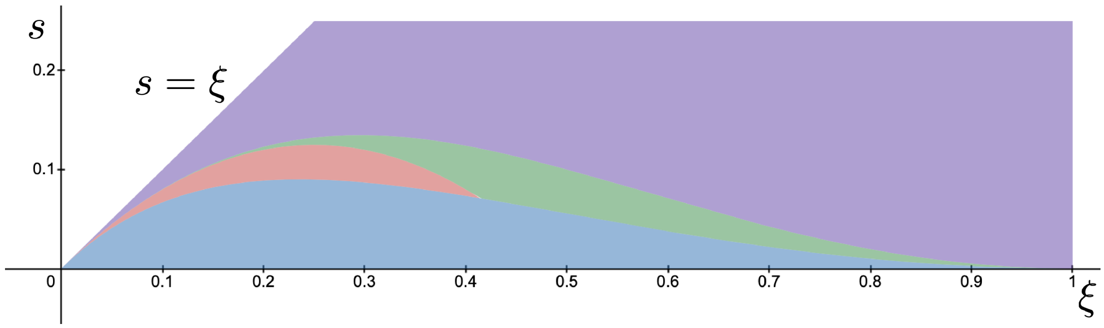

Figure 5: Illustration of the four regions in Theorem 1.

Although Theorem 1 might seem intricate, the economics behind it is intuitive. The main trade-off that Seller faces is between demand and extraction. Using a deterrence policy, namely either full disclosure or mixture disclosure, Seller can increase the demand she faces. This is because the top mass point makes Buyer more likely to purchase Seller's product without search, and this part of demand cannot be "messed up" by Nature's choice of the outside option distribution. However, this is only effective when the price is such that $p \le s/\xi$, and is hence bad for surplus extraction when s/ξ is small.

The details of Theorem 1 are illustrated in Figure 5. In Figure 5, the horizontal axis is the mean of the outside option distribution ξ, and the vertical axis is the search cost s; the 45-degree line reflects the assumption that $s < \xi$. In the blue region, namely when $s < B_2(\xi)$, the search cost is sufficiently small. Although putting a mass point "at the top" boosts the likelihood of eventual purchase, by doing that, the price has to be bounded above by s/ξ, which is too small in this case. Thus it is optimal for Seller to use uniform disclosure, which allows her to charge a higher price $p^* > s/\xi$. In the violet region, $s \ge B_1(\xi)$, the search cost is sufficiently large, and hence Seller can charge a higher price even under the restriction that $p \le s/\xi$. Put differently, the tension between demand and extraction is alleviated for s large enough. Full disclosure creates maximal differentiation between the new product and Buyer's outside option, and thus increases Buyer's willingness to pay whenever the match value is revealed to be high. Then since the upper bound on price is not too restrictive, full disclosure not only helps in extracting surplus, but also guarantees a demand of size μ even when the price becomes high.[30] These two properties, together, render full disclosure optimal.

The area between the blue and violet regions is shaped by the trade-off between demand and extraction. As stated in (1) and (2) in Theorem 1, the optimal selling strategy

[30]Under full disclosure, posterior $w = 1$ realizes with probability μ.

in the green and maroon regions exhibits a "cutoff" feature in μ, and this stems from the interaction of two effects. One is a "price effect", as Seller can charge a higher price if she chooses uniform disclosure. Another one is a "demand effect", namely as μ increases, the demand generated by a distribution with a mass point "at the top" grows faster than the demand from uniform disclosure. For a small μ, the price effect makes it optimal to use uniform disclosure and charge a higher price; and the demand effect dominates when μ is large, which favors a deterrence policy. Thus, the cutoff structure results.

In the maroon region, the mean of the outside option distribution is relatively small. In this region, when the mean match value is above some cutoff $\check{\mu}$, it is optimal for Seller to use mixture disclosure. Intuitively, a relatively small ξ indicates that conditional on checking the outside option, Buyer is more likely to get an unsatisfactory draw. Therefore, there is an incentive for Seller to attract Buyer to come back to buy after searching. In particular, for a sufficiently large μ, the mean of the distribution over posteriors is high enough for Seller to put a mass point "at the top" and spread out the remaining mass evenly on the support of the distribution; thus it creates room for deterring search and attracting Buyer who searches and gets an unsatisfactory draw to come back at the same time. In the green region, the mean of the outside option distribution is relatively large. This means that Buyer is more likely to get a good draw upon checking the outside option, which in turn indicates that Seller faces fierce competition. To soften competition, Seller should maximally differentiate her product from the outside option; this incentive renders full disclosure optimal when the mean match value is above some cutoff $\hat{\mu}$.

Interestingly, as Figure 5 illustrates, the cutoffs in s as a function of ξ are all hump-shaped. This is because there are two countervailing forces shaping the cutoffs as ξ increases. One is that a larger ξ makes the outside option more attractive ex ante, and hence a higher search cost is needed to make a deterrence policy profitable, as otherwise the "search deterrence price" $p^{**} = s/\xi$ would be too small. Another is that a larger ξ indicates that Buyer is more unlikely to come back to buy if she goes to search, which strengthens Seller's deterrence motive. When ξ is relatively small, the first force dominates, and the second dominates when ξ is relatively large.

It may seem paradoxical that, except for boundary cases,[31] the mean match value μ does not enter Seller's robust price. This is because, when there is search deterrence, the revenue guarantee is increasing in price and hence the optimal price coincides with the threshold, which does not depend on the mean match value. When Seller chooses uniform

[31] See Appendix A.4 for a formal definition of "boundary cases".

disclosure, this feature stems from Seller's robustness concerns. If μ is small, it is optimal for Seller to evenly spread out the posterior beliefs over the entire interval. Thus, the support of the robustly optimal distribution over posteriors has nothing to do with the mean match value. Consequently, it only enters the revenue guarantee as a multiplicative term, and hence does not enter the expression of the robust price. If instead μ is large, as discussed after Proposition 1, there exists a choice of Nature that completely dissipates the effect of the mean match value.

Another interesting feature of my results is that no disclosure is *never* optimal. Fixing price p, the degenerate distribution over posteriors δ_μ is optimal only if $p \leq \mu + s/\xi - 1$. Although such a price makes Buyer buy without search with probability 1, Seller's revenue guarantee is minimal since the price is too low. In particular, no disclosure is strictly dominated by full disclosure paired with $p^{**} = s/\xi$.

As explained in the introduction, the robustly optimal selling strategies identified in Theorem 1 have sharp implications for selling new products.

3.3 Comparative Statics

I now derive some comparative statics on the robust price, the robustly optimal disclosure policy, and the revenue guarantee.

Theorem 2 (Comparative statics)**.**

(i) The robust price p_r is non-monotone in the search cost s: holding μ and ξ fixed, there exist $\hat{s}$ such that p_r is increasing on $[0, \hat{s})$ and $(\hat{s}, \xi)$, but $p_r(\hat{s}-) > p_s(\hat{s}+)$.

(ii) For any $s_1 < s_2$, the robustly optimal disclosure policy corresponding to s_2 is more informative than the one corresponding to s_1 unless $s_1, s_2 \in (B_2(\xi), B_3(\xi))$ and μ is sufficiently large.

(iii) Seller's revenue guarantee is strictly increasing in s, strictly decreasing in ξ, and increasing in μ.

Part (i) of Theorem 2 states that the robust price is not always increasing in the search cost; this is illustrated in Figure 6. This is a consequence of the trade-off between demand and extraction. The conventional wisdom in the search literature says that an increase in the search cost decreases the value of search, and hence makes Buyer more likely to buy without search; consequently, Seller can extract more surplus from Buyer by simply charging a higher price. In fact, the conventional wisdom does work in the sense that

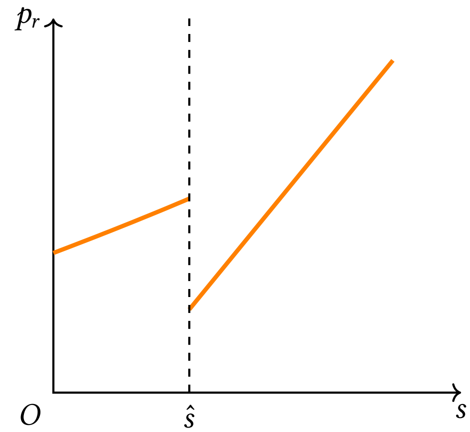

Figure 6: The robust price is non-monotone in the search cost. When $s < \hat{s}$, uniform disclosure is optimal; and when $\hat{s} \leq s < \xi$, it is optimal for Seller to put a "top mass" in the distribution over posteriors.

as s increases, both p^* and p^{**} are increasing. However, as s increases, p^{**} grows faster than p^*, and hence the gap between the two prices shrinks. Moreover, the probability of eventual purchase when Seller uses uniform disclosure is always bounded above by the counterpart when a distribution with a top mass is used: loosely speaking, charging a price $p > s/\xi$ gives Nature more power on "messing up" Seller's demand. Therefore, as s gets larger, Seller would eventually find that using a deterrence policy is more profitable albeit the corresponding price $p^{**} = s/\xi$ is smaller than p^*.[32]

Part (ii) asserts that except for one parameter region, as the search cost increases, the robustly optimal disclosure policy gets more informative. An increase in the search cost has two effects: it makes a deterrence policy more attractive and also affects the price. So long as the search cost does not cross the "jump down point" $\hat{s}$, an increase in the search cost also makes the price higher; then as the discussion in Section 3.1.1 indicates, in most of the parameter regions the information also gets more precise. If the increase makes Seller adopt a deterrence policy instead, such a strategy comes with more information because it reveals that the match value is surely high with positive probability. In most circumstances, the two effects work hand in hand to make the robustly optimal disclosure policy more informative.

[32] Choi et al. (2018) find that in an oligopolistic market with sequential consumer search, the equilibrium price is *decreasing* in the search cost under some conditions. This stems from their assumption that the match values are hidden but the prices are posted and hence can be used to direct search, which is completely different from the driving force behind my result.

When the prior is sufficiently large and the mean of the outside option distribution is sufficiently small, the informativeness of the robustly optimal disclosure policy is decreasing in the search cost when mixture disclosure is optimal. In this region, the robust price is $p^{**} = s/\xi$, and the robustly optimal disclosure policy features a mass point at $w = 1$, and is otherwise uniformly distributed on $(p^{**}, 1)$. Then as the search cost increases, the high match value is less likely to be fully revealed, which renders the disclosure policy less informative.

Part (iii) has a simple economic intuition: a higher search cost implies a rise in the market power, as it is less likely for Buyer to check her outside option; and a higher mean of the outside option distribution indicates more competition in the market. Moreover, a higher mean match value makes it easier to generate higher posteriors, and hence increases the probability that Buyer eventually buys.

4 Variations and Extensions

The two key features of the main model are search frictions and Seller's robustness concerns. To better understand the role of these features in the main results, I consider two variations of the main model; in each of them, only one of the key features is present. In Section 4.1, I shut down the search frictions, and Seller is still taking a robust approach. In Section 4.2, Seller's robustness concerns are absent in the sense that the outside option distribution is known to her, and hence the effect of search frictions can be isolated. As an extension, I investigate what happens if Seller is able to recognize the identity of the buyer in Section 4.3.

4.1 Zero Search Cost

When $s = 0$, Buyer always checks the outside option since doing that is costless. Furthermore, having a "mass at the top" is never optimal because $s/\xi = 0$. Then since the hedging motive persists, the robustly optimal selling strategy is isomorphic to the "low search cost" case in the main model.

To state the result formally, define

$$\underline{\mu}^0 := \frac{2 - \xi - \sqrt{2\xi - \xi^2}}{2(1-\xi)}, \text{ and } \bar{\mu}^0 := 1 - \frac{\sqrt{\xi}}{2}.$$

Proposition 2. *Suppose $s = 0$. The robust price is*

$$p_0^* = \begin{cases} \frac{1-\sqrt{2\xi-\xi^2}}{1-\xi} & \text{if } \mu \leq \underline{\mu}^0, \\ 2\mu - 1 & \text{if } \underline{\mu}^0 < \mu \leq \bar{\mu}^0, \\ 1-\sqrt{\xi} & \text{if } \mu > \bar{\mu}^0. \end{cases}$$

If $\mu > \bar{\mu}^0$, an robustly optimal distribution over posteriors is $U_{[2\mu-1,1]}$; and if $\mu \leq \bar{\mu}^0$, the robustly optimal distribution over posteriors is

$$H(w) = \begin{cases} 1 - \frac{2\mu}{1+p_0^*} & \text{if } w \in [0, p_0^*), \\ 1 - \frac{2\mu}{1-p_0^{*2}}(1-w) & \text{if } w \in [p_0^*, 1]. \end{cases}$$

Proof of this result, and all other results in Section 4, can be found in Appendix B.

It can be seen from Proposition 2 that when search frictions are absent, the affinity of the distribution over posteriors is still the optimal way to address robustness concerns, but the trade-off between demand and extraction disappears. In particular, full disclosure is never optimal: only when she wants to take advantage of the search frictions is Seller willing to provide full information about the match value.

4.2 Known Outside Option Distribution

The only difference between the model considered in this subsection and the main model is that Buyer's outside option distribution G is assumed to be *known* to Seller in the former. For simplicity, I assume that G has full support, and admits a log-concave density g. The optimal selling strategy for this problem is strikingly simple.

Proposition 3. *The optimal selling strategy consists of a disclosure policy that fully discloses the match value and an optimal price p^o, where*

$$p^o = \begin{cases} 1-a & \text{if } 1-a \geq p_h G(1-p_h), \\ p_h & \text{if } 1-a < p_h G(1-p_h), \end{cases}$$

where a is defined by Equation (3), and p_h is the unique solution of the equation[33]

$$p = \frac{G(1-p)}{g(1-p)}. \tag{15}$$

To understand Proposition 3, observe that when $p = 1-a$ is paired with full disclosure, Buyer buys without search if the match value is high, and would not come back for sure if the match value is low. Consequently, this strategy fully deters search. The resulting revenue is the product of the price and the (prior) probability that the match value is high, namely $\mu(1-a)$. Alternatively, if Seller does not deter search and charges price p, full disclosure makes Buyer comes back with probability $G(1-p)$ when the match value is high, and when the match value is low she never comes back. Hence, Seller's payoff from setting price p is $p\mu G(1-p)$, and the profit-maximizing price is precisely p_h. As a consequence, Seller's profit is given by $\max\{\mu(1-a), \mu p_h G(1-p_h)\}$: when the former is larger, Seller charges $p^o = 1-a$ to fully deter search; otherwise, Seller charges $p^o = p_h$ and lets Buyer search.

Importantly, regardless of whether Seller deters search or not, full disclosure is always optimal. This can be shown by noticing that for any selling strategy that does not feature full disclosure, Seller's profits can be improved by either increasing the price or providing more information, or both. By providing full information, Seller maximally differentiates her product from Buyer's outside option. This strategy softens the competition brought by the outside option, and thus allows Seller to maximally extract surplus.[34]

In contrast, full disclosure is not always optimal in the main model. When Seller seeks robustness, full disclosure can only be optimal if $p \leq s/\xi$. When using full disclosure, this upper bound on price may limit the extent to which Seller can extract surplus from Buyer even if the latter highly values the innovative features of the new product. Consequently, full disclosure can be suboptimal when s is relatively small; in particular, uniform disclosure allows Seller to charge a higher price, and mixture disclosure may generate higher demand.

[33] Because g is log-concave, the right-hand side (RHS) of the equality above is decreasing in p. Since G has full support, the left-hand side (LHS) is strictly less than the RHS when $p = 0$, and is strictly greater than the RHS when $p = 1$. Consequently, the solution to Equation (15) must be unique.

[34] More precisely, by providing more information, it is also more likely for Buyer to realize that her match value is low, and hence increases the likelihood that Buyer prefers the outside option or opts out without search. However, the benefit from extracting more surplus from "likers" by jointly charging a higher price and providing more information dominates this loss. In this sense, price and information are "complementary" when the outside option distribution is known to Seller.

Roughly, one can think of Seller having "a lot of uncertainty" about the outside option distribution in the main model; in the model studied in this subsection, however, she has no uncertainty at all. Interestingly, the information that Seller provides is noisier when she has much more uncertainty.[35]

Because $p^o = p_h$ only when $1 - a < p_h G(1 - p_h)$, when p_h is the optimal price, it must be that $p_h > 1 - a$. This highlights the trade-off between demand and extraction similar to the main model. In fact, for a known outside option distribution G, the optimal selling strategy is completely dictated by the magnitude of the search cost.

Corollary 2. *For every outside option distribution G, there exists $\hat{s}_G \in (0, \xi)$ such that $p^o = p_h$ for every $s < \hat{s}_G$, and $p^o = 1 - a$ for every $s \geq \hat{s}_G$. Furthermore, at $s = \hat{s}_G$, the optimal price drops from p_h to $1 - a(\hat{s}_G)$.*

Corollary 2 is intuitive. In this model, when deterring search Buyer eventually buys with probability μ, and otherwise Buyer eventually buys with probability $\mu G(1 - p_h)$. Therefore, deterring search increases the "demand" Seller faces. To deter search, however, the maximal price that Seller can charge is capped at $1-a$. When s is small, so is $1-a$, and hence an increased chance of eventual purchase does not justify search deterrence since the price has to be very low; instead, charging p_h and letting Buyer search is optimal. Analogous to the main model, as s gets sufficiently large, deterring search becomes more profitable. Consequently, the tradeoff between demand and extraction remains, and the nonmonotonicity of the optimal price in the search cost also holds here for the same reason as in the main model.

Another implication of Corollary 2 is that the middle sliver in the robust case disappears: there is not a region that whether to deter search depends on μ. This is because the optimal disclosure policy is always full disclosure in this model, and hence regardless of whether she deters search or not, μ enters Seller's revenue in the same (multiplicative) way.

Corollary 3 summarizes some comparative statics.

Corollary 3. *(i) The optimal price does not depend on μ, and Seller's profit is strictly increasing in μ.*

(ii) For $s < \hat{s}_G$, both the optimal price and Seller's profit do not depend on s. For $s \geq \hat{s}_G$, both the optimal price and Seller's profit are strictly increasing in s.

[35] It would be interesting to understand how the optimal disclosure policy changes when Seller's uncertainty gradually reduces, but that is beyond the scope of this paper.

To summarize, many insights persist when Seller knows the outside option distribution, but Seller's robustness concerns in the main model beget the feature of continuously and evenly spread out impressions. Consequently, this model does not generate as clear-cut implications for selling new products as the main model.

4.3 Recognizable Buyer Identity

In this subsection, I allow Seller to recognize whether Buyer is a first-time visitor or came back from search.

4.3.1 Exploding Offers and Renegotiation

One way that Seller can take advantage of this is to make an exploding offer: she commits not to sell to Buyer if she does not buy during her first visit. In this case, Buyer buys without search if and only if $w - p \geq \xi - s$, namely when her value of Seller's product is no less than the expected value of the unknown outside option net of the search cost; and if she goes to search, she would never come back. The probability of this event is $1 - H((p + \xi - s)^-)$, and hence Seller's revenue from an exploding offer is $p[1 - H((p + \xi - s)^-)]$.[36] One striking feature of exploding offers is that Nature's choice of outside option does not play any role in Seller's problem: it is outcome equivalent to that the outside option distribution is δ_ξ, the degenerate distribution at ξ, and Buyer must incur a cost s to consume the outside option.

Another possibility is that Seller posts a price first, and then if Buyer comes back she may have an incentive to increase the price. As noted in Armstrong and Zhou (2016), in the current framework very little can be said; but if Buyer must incur an exogenous cost $r > 0$ to return to buy Seller's product after search, no matter how small r is, a Diamond paradox style argument shows that once Buyer goes to search, she would never come back. Hence, the equilibrium outcome is the same as Seller committing to exploding offers.

Proposition 4 summarizes the findings when Buyer's identity is recognizable.

Proposition 4. *Suppose that Seller can recognize whether Buyer is a first-time visitor. Then*

(i) if Seller can commit to an exploding offer, it is optimal to offer $p = 1 - \xi + s$ with full disclosure;

[36] For a function $f : \mathbb{R} \to \mathbb{R}$, $f(s^-) = \lim_{x \searrow s} f(x)$ whenever this limit exists.

(ii) for all $\mu, \xi \in (0, 1)$ and $0 \leq s < \xi$, Seller earns strictly higher profits than the case that she cannot distinguish between first-time visitors and searchers.

(iii) if Seller cannot commit to the price, and there is a cost of returning to Seller $r > 0$, then the equilibrium outcome is the same as Seller committing to exploding offers.

The optimality of full disclosure stems from the fact that full disclosure is "efficient" in the sense that it leaves Buyer with no uncertainty on whether she should choose Seller's product or her outside option, and that when the outside option distribution is δ_ξ Seller can appropriate all the surplus simply by pricing at $1 - \xi + s$. Recognizable Buyer identity helps Seller because δ_ξ is always suboptimal for Nature when her choice matters.

4.3.2 Price Discrimination

Now I assume that Seller can deviate from the robustly optimal selling strategy in the sense that while the disclosure policy cannot be changed, she can commit to a price path (p_1, p_2) with $p_1 < p_2$ such that p_1 and p_2 are the prices charged if Buyer buys immediately or after search, respectively.[37] In particular, I allow Seller to deviate by either charging a higher price in the second period, or offering a "buy-now discount": a lower price is offered if Buyer purchases without search, but if she comes back from search she has to pay the equilibrium price.

Proposition 5. *Suppose that Seller can recognize whether Buyer is a first-time visitor. Let (p_r, H^*) be a robustly optimal selling strategy identified in Theorem 1, and let G^* be the corresponding worst-case outside option distribution. If Seller deviates by committing to a pair of prices (p_1, p_2), where either $p_1 = p_r$ or $p_2 = p_r$, then*

(i) If Nature cannot detect this deviation and hence the outside option distribution is still G^*, Seller can benefit from such a deviation unless H^* corresponds to full disclosure;

(ii) If Nature can detect this deviation and optimally responds to it by choosing a new outside option distribution, Seller cannot benefit from such a deviation.

When Nature cannot detect Seller's deviation and hence the outside option distribution is fixed at G^*, a celebrated result in Armstrong and Zhou (2016) can be used off-the-shelf. They show that whenever the buy-now demand is strictly more elastic than

[37] This form of price discrimination is studied in, for example, Nocke, Peitz, and Rosar (2011) and Armstrong and Zhou (2016).

the buy-later demand,[38] charging a higher "buy-later price" benefits Seller because by doing that she either extracts more surplus if Buyer buys after search, or makes Buyer more likely to buy without search. Unless H^* corresponds to full disclosure, in which case there is no buy-later demand and hence such selling tactics would not be useful, the buy-now demand is always strictly more elastic than the buy-later demand.

When Nature can detect Seller's deviation, she would choose $G = \delta_\xi$ to counter it. For any fixed H^*, δ_ξ creates the largest buy-now demand among all outside option distributions. The affinity of H implies that the "total demand", namely the probability of eventual purchase, is the same for all outside option distributions with the same mean (recall Equation (7)). Consequently, by "transforming" buy-later demand to buy-now demand, δ_ξ minimizes the effect of the price discrimination strategy since it only helps when it is optimal for Buyer to search under $((p_r, H^*), G^*)$. In fact, δ_ξ makes Seller's expected revenue under (p_1, p_2) the same as that she cannot price discriminate and the posted price is p_1. Thus, Seller is at most as well off as in the main model.

Fully solving the problem of Seller choosing $((p_1, p_2), H)$ under robustness concerns is beyond the scope of this paper, and is left for future research.

5 Discussion

I conclude by discussing a few of the assumptions.

Deterministic price. In the main model, I assume that the seller posts a single price. Because there is no information asymmetry at the time of contracting, and the buyer's identity is not recognizable, screening via a menu does not help. A natural question is whether randomizing over prices could help the seller. For the current setting, the answer is no, because the adversary observes the price before she chooses the outside option distribution.[39]

[38] A sufficient condition is that $1-F(p)$ is strictly log-concave, where F is the distribution of Buyer's value of Seller's product. In this model, F is identical to the distribution over posteriors.

[39] One may be tempted to ask what if the seller is allowed to choose a joint distribution over prices and signals, and this is the only thing that is observable to the adversary. Such an assumption can be a bit tricky for this model: because the buyer's identity is not recognizable, and going back to the seller is costless, the buyer may prefer to "sample" many times until she is sufficiently informed about the match value and/or encounters an attractive price.

No "safe" outside option. An implicit assumption of the model is that the buyer does not have a "safe" outside option that she can consume without incurring the search cost. Alternatively, her "safe" outside option u_0 is less than or equal to zero. This assumption is only made to ease exposition. If $u_0 > 0$, the seller's expected revenue under $((p, H), G)$ is given by

$$p\,\mathbb{E}_G\left[1 - H\left(p + \max\left\{\min\{a, v\}, u_0\right\}\right)\right].$$

In solving the seller's problem in the baseline model, I define the buyer's effective outside option to be $z := \min\{a, v\}$. Then I solve the auxiliary problem in which the adversary chooses the effective outside option distribution, and show that the adversary's optimal choice can be induced by an outside option distribution.

To solve this new problem, it only suffices to define $y := \max\{z, u_0\}$ and work with the distribution of this new variable instead. It can be shown that Nature's optimal distribution of this new variable can be induced by an outside option distribution, and adding this (relevant) "safe" outside option does not change the qualitative features of the main results.

Binary match value. A key assumption in the model is that the buyer's match value is binary. If instead it can take a continuum of values, then the problem becomes much more complex. To see why, recall that in the first stage a continuum of information design problems needs to be solved, one for each price. The tractability provided by binary match values allows me to classify the solutions to this continuum of problems into a small number of "groups" such that all problems in the same group can be solved at once. With a continuum of match values, however, even if a strong assumption is imposed on the (prior) match value distribution (for example, the distribution admits a single-peaked density), there are way too many groups to consider. Furthermore, characterizing the robust price is another daunting task; in particular, unlike the binary match value case, a closed-form robust price cannot be obtained.

Despite the nontrivial additional complexity, some important observations remain to hold when the match value is continuously distributed. In particular, the tension between demand and extraction persists. Moreover, the optimal distribution is affine on a subset of its support in many cases.

The assumptions on the seller's knowledge about the outside option distribution. In the model I assume that the seller knows the mean of the outside option distribution as well as an upper bound on its support. For a new product, it should be reasonable to assume that if it is a good match for a buyer, this product should provide more value, when the price is zero, to the buyer than any alternatives in the market. Then since the (gross) value of the new product to "likers" is normalized to one, I assume that the value of the buyer's "best alternative" cannot exceed one.

One reasonable alternative assumption on the seller's knowledge about the outside option distribution is that she knows the mean and an upper bound on a higher moment (for example, variance). To solve for the optimal information for any fixed price, consider a relaxed problem without the higher moment constraint first. Equivalently, in this problem, the seller knows the mean and that the upper bound is infinity; as a consequence, the seller does not benefit from a top mass. If the adversary's choice of outside option distribution in this auxiliary problem satisfies the higher moment constraint, a solution to the original problem is found; otherwise, the higher moment constraint must bind. It is only in the latter case that the seller may benefit from deterrence, and the seller's optimal distribution may be strictly convex instead of affine. The qualitative insights, however, are not affected: deterring search imposes an upper bound on price, which begets the main trade-off; and there are large variations in disclosure policies.

Appendices

A Proofs and Omitted Details for Section 3

A.1 Proof of Proposition 1

A.1.1 Preliminaries

I proceed as follows: as outlined in Section 3.1.2, for every candidate robustly optimal distribution H^* in Proposition 1, I find a corresponding worst-case effective outside option distribution $\hat{G}^*$, and then use Lemma 1 to show that $(H^*, \hat{G}^*)$ solves problem (11). Next, I show that the worst-case effective outside option distribution $\hat{G}^*$ can be induced by an outside option distribution (or just "induced" for simplicity), which then implies that H^* is indeed a robustly optimal distribution.

To check that an effective outside option distribution can be induced, in some cases I use the following result due to Au and Whitmeyer (2022).

Lemma A.1 (Au and Whitmeyer, 2022)**.** *An effective outside option distribution can be induced if and only if there exists* $\nabla \in \Delta([\xi - s, 1 - s/\xi])$ *where for each* $a \in \operatorname{supp}(\nabla)$ *there is a distribution of outside options* $G_a \in \mathcal{M}(\xi)$ *such that for each* $z \in [0, 1 - s/\xi]$,

$$\hat{G}(z) = \nabla(z) + \int_{\operatorname{supp}(\nabla)\cap(z,1-s/\xi]} G_a(z)\, \mathrm{d}\nabla(a). \tag{16}$$

Observe also that any price $p > 1 - (\xi - s)$ is weakly dominated by $p = 0$: for any such price, by setting the outside option distribution to be the degenerate distribution at ξ, Nature is able to make Buyer not to buy from Seller for sure. This is because the highest posterior that Seller can generate is 1, and hence the highest net value is $1 - p$, but $1 - p < \xi - s$. Consequently, Seller does not sell at all and makes zero profit, which is the same as setting $p = 0$. Therefore, when optimizing over p, it suffices to consider $p \in [0, 1 - (\xi - s)]$.

A.1.2 The case of $p > s/\xi$

The following two claims establish the results for this case.

Claim A.1. *Suppose* $p > s/\xi$ *and* $\mu > (1 + p)/2$*. Then any distribution over posteriors* $H \in \mathcal{M}(\mu)$ *that first-order stochastically dominates* $U_{[p,1]}$ *is robustly optimal. Furthermore, Seller's revenue is* $\Phi(p) = p[1 - (\xi - s)/(1 - p)]$.

Proof of Claim A.1. Consider any distribution H^* that first-order stochastically dominates $U_{[p,1]}$. Then the concave hull of $H\Big|_{[0,1-s/\xi]}(p+z)$ coincides with $U_{[p,1]}$, and Nature can attain the value of problem (12) by using a binary distribution

$$\hat{G}^*(z) = \begin{cases} 1 - \frac{\xi - s}{1-p} & \text{if } 0 \le z < 1 - p, \\ 1 & \text{if } z \ge 1 - p. \end{cases}$$

Now $\mu > (1 + p)/2$ implies that $p < 2\mu - 1 < \mu$, and hence Seller's value is $\widetilde{G}_p(\mu) = 1 - (\xi - s)/(1 - p)$ for any choice of distribution over posteriors. Thus, $(H^*, \hat{G}^*)$ is a saddle point; and to show that H^* is robustly optimal, it only suffices to show that $\hat{G}^*$ can be induced by some outside option distribution G^*. Consider

$$G^*(v) = \begin{cases} 1 - \frac{\xi - s}{1-p} & \text{if } 0 \le v < \frac{\xi(1-p)}{\xi - s}, \\ 1 & \text{if } v \ge \frac{\xi(1-p)}{\xi - s}; \end{cases}$$

by (3), $a = 1 - p$. Thus, by definition of the effective outside option distribution, the effective value distribution induced by G^* is

$$\hat{G}(z) = \begin{cases} G^*(z) & \text{if } 0 \le z < 1-p, \\ 1 & \text{if } z \ge 1-p, \end{cases}$$

which is exactly $\hat{G}^*$. This completes the proof. ■

In particular, since $\mu > (1+p)/2$, $U_{[2\mu-1,1]}$ is well-defined, has mean μ, and first-order stochastically dominates $U_{[p,1]}$. Consequently, $U_{[2\mu-1,1]}$ is optimal.

Claim A.2. *If $\mu \le (1+p)/2$, there exists $\bar{w} \in [2\mu - p, 1]$ where*

$$\bar{w} = \begin{cases} 2\mu - p & \text{if } \mu > p \text{ and } \xi - s \le \frac{2(\mu-p)^2}{2\mu-p}, \\ \sqrt{(\xi-s)(\xi-s+2p)} + (\xi - s + p) & \text{if } \mu > p \text{ and } \frac{2(\mu-p)^2}{2\mu-p} < \xi - s \le \frac{(1-p)^2}{2}, \\ & \text{ or } \mu \le p \text{ and } \xi - s \le \frac{(1-p)^2}{2}, \\ 1 & \text{if } \xi - s \ge \frac{(1-p)^2}{2}, \end{cases}$$

such that the robustly optimal distribution is

$$H_{\bar{w}}(w) = \begin{cases} 1 - \frac{2\mu}{\bar{w}+p} & w \in [0,p), \\ 1 - \frac{2\mu}{\bar{w}+p} + \frac{2\mu}{\bar{w}+p}\left(\frac{w-p}{\bar{w}-p}\right) & w \in [p, \bar{w}), \\ 1 & w \in [\bar{w}, 1]. \end{cases}$$

And Seller's revenue is

$$\Phi(p) = \begin{cases} p\left[1 - \frac{\xi-s}{2(\mu-p)}\right] & \text{if } \mu > p \text{ and } \xi - s \le \frac{2(\mu-p)^2}{2\mu-p}, \\ \frac{p\mu}{\sqrt{(\xi-s)(\xi-s+2p)}+(\xi-s+p)} & \text{if } \mu > p \text{ and } \frac{2(\mu-p)^2}{2\mu-p} < \xi - s \le \frac{(1-p)^2}{2}, \\ & \text{ or } \mu \le p \text{ and } \xi - s \le \frac{(1-p)^2}{2}, \\ \frac{2\mu}{1+p}\left(1 - \frac{\xi-s}{1-p}\right) & \text{if } \xi - s \ge \frac{(1-p)^2}{2}. \end{cases}$$

Proof of Claim A.2. I consider the case of $\bar{w} = 2\mu - p$ first, where the distribution in the statement of the claim becomes

$$H_{2\mu-p}(w) = \begin{cases} 0 & w \in [0,p), \\ \frac{w-p}{2(\mu-p)} & w \in [p, 2\mu-p), \\ 1 & w \in [2\mu-p, 1]. \end{cases}$$

Observe that the concave hull of $H_{2\mu-p}\Big|_{[0,1-s/\xi]}(p+z)$ coincides with $H_{2\mu-p}(w)$ on $[p,1]$, and hence Nature's value is $H_{2\mu-p}(p+\xi-s)$, which can be obtained by effective outside option distribution with mean $\xi - s$ supported on a subset of $[0, 2(\mu-p)]$. Now consider effective outside option distribution

$$\hat{G}^{2\mu-p}(z) = \begin{cases} 1 - \frac{\xi-s}{\mu-p} + \frac{\xi-s}{2(\mu-p)^2} z & \text{if } z \in [0, 2(\mu-p)), \\ 1 & \text{if } z \in [2(\mu-p), 1], \end{cases}$$

which induces

$$\hat{G}_p^{2\mu-p}(w) = \begin{cases} 0 & \text{if } w \in [0, p) \\ 1 - \frac{\xi-s}{\mu-p} + \frac{\xi-s}{2(\mu-p)^2}(w-p) & \text{if } w \in [p, 2\mu-p), \\ 1 & \text{if } w \in [2\mu-p, 1]. \end{cases}$$

Then the concave hull of $\hat{G}_p^{2\mu-p}(w)$ coincide with the function itself on $[p,1]$; and since $\mu > p$, Seller's value is $\hat{G}_p^{2\mu-p}(\mu) = 1 - (\xi-s)/(2(\mu-p))$. It can be obtained by any distribution with mean μ supported on a subset of $[p, 2\mu-p]$, and this condition is satisfied by $H_{2\mu-p}$. Thus, $(H_{2\mu-p}, \hat{G}^{2\mu-p})$ is a saddle point; and to show that $H_{2\mu-p}$ is robustly optimal, it only remains to show that $\hat{G}^{2\mu-p}$ can be induced. Consider

$$\nabla(a) = \frac{a}{\mu-p} - 1 \quad \text{on } [\mu-p, 2(\mu-p)];$$

and let $G_a(v)$ be a binary distribution with support on $\{2(\mu-p)-a, a+2s(\mu-p)/(\xi-s)\}$. It is not difficult to show that $G_a(v) \in \mathcal{M}(\xi)$ for all $a \in [\mu-p, 2(\mu-p)]$; then by Lemma A.1, $\hat{G}^{2\mu-p}$ can be induced.

Next I consider the case of $\bar{w} = \sqrt{(\xi-s)(\xi-s+2p)} + (\xi-s+p) < 1$. The concave hull of $H_{\bar{w}}\Big|_{[0,1-s/\xi]}(p+z)$ coincides with $H_{\bar{w}}(w)$ on $[p,1]$, and hence Nature's value is $H_{\bar{w}}(p+\xi-s)$, which can be obtained by effective outside option distribution with mean $\xi-s$ and supported on a subset of $[0, \bar{w}-p]$. Now consider effective outside option distribution

$$\hat{G}^{\bar{w}}(z) = \begin{cases} (z+p)/\bar{w} & \text{if } z \in [0, \bar{w}-p), \\ 1 & \text{if } z \in [\bar{w}-p, 1], \end{cases}$$

which induces

$$\hat{G}_p^{\bar{w}}(w) = \begin{cases} 0 & \text{if } w \in [0,p) \\ w/\bar{w} & \text{if } w \in [p,\bar{w}), \\ 1 & \text{if } w \in [\bar{w},1]; \end{cases}$$

and the concave hull of $\hat{G}_p^{\bar{w}}$ is

$$\widetilde{G}_p^{\bar{w}}(w) = \begin{cases} w/\bar{w} & \text{if } w \in [0,\bar{w}), \\ 1 & \text{if } w \in [\bar{w},1]. \end{cases}$$

Consequently, Seller's value is $\widetilde{G}_p^{\bar{w}}(\mu) = \mu/\bar{w}$, which can be attained by any distribution over posteriors with mean μ supported on a subset of $\{0\} \cup [p,\bar{w}]$, and this condition is satisfied by $H_{\bar{w}}$. Thus, $(H_{\bar{w}}, \hat{G}^{\bar{w}})$ is a saddle point; and to show that $H_{\bar{w}}$ is indeed robustly optimal, it only remains to show that $\hat{G}^{\bar{w}}$ can be induced. Consider

$$\nabla(a) = \frac{2a}{\bar{w}-p} - 1 \quad \text{on } [(\bar{w}-p)/2, \bar{w}-p];$$

and let $G_a(v)$ be a ternary distribution with pmf $g_a(v)$ given by

v	0	$\bar{w}-p-a$	$\frac{2\bar{w}\xi}{\bar{w}-p} - (\bar{w}-p) + a$
$g_a(v)$	$p/\bar{w}$	$(1-p/\bar{w})/2$	$(1-p/\bar{w})/2$

It is not difficult to show that $G_a(v) \in \mathcal{M}(\xi)$ for all $a \in [(\bar{w}-p)/2, \bar{w}-p]$; then by Lemma A.1, $\hat{G}^{\bar{w}}$ can be induced.

Finally, I consider the case of $\bar{w} = 1$, where the proposed distribution becomes

$$H_1(w) = \begin{cases} 1 - \frac{2\mu}{1+p} & \text{if } w \in [0,p), \\ 1 - \frac{2\mu}{1+p} + \frac{2\mu}{1+p}\left(\frac{w-p}{1-p}\right) & \text{if } w \in [p,1]. \end{cases}$$

The concave hull of $H_1\Big|_{[0,1-s/\xi]}(p+z)$ coincides with $H_1(w)$ on $[p,1]$, and hence Nature's value is $H_1(p+\xi-s)$, which can be obtained by any effective outside option distribution with mean $\xi - s$ and supported on a subset of $[0,1-p]$. Now consider effective outside option distribution

$$\hat{G}^1(z) = \begin{cases} \frac{2(p+z)}{1+p}\left(1 - \frac{\xi-s}{1-p}\right) & \text{if } z \in [0,1-p), \\ 1 & \text{if } z \in [1-p,1], \end{cases}$$

which induces

$$\hat{G}_p^1(w) = \begin{cases} 0 & \text{if } w \in [0, p) \\ \frac{2w}{1+p}\left(1 - \frac{\xi - s}{1-p}\right) & \text{if } w \in [p, 1]; \end{cases}$$

and the concave hull of $\hat{G}_p^1$ is

$$\widetilde{G}_p^1(w) = \frac{2w}{1+p}\left(1 - \frac{\xi - s}{1-p}\right) \quad \text{on } [0, 1].$$

Consequently, Seller's value is $\widetilde{G}_p^1(\mu)$, which can be attained by any distribution over posteriors with mean μ supported on a subset of $\{0\} \cup [p, 1]$, and this condition is satisfied by H_1. Thus, $(H_1, \hat{G}^1)$ is a saddle point; and to show that H_1 is indeed robustly optimal, it only remains to show that $\hat{G}^1$ can be induced. Consider

$$\nabla(a) = \frac{a - (\xi - s)}{1 - p - (\xi - s)} \quad \text{on } [\xi - s, 1 - p];$$

and let $G_a(v)$ be a quarternary distribution with pmf $g_a(v)$ given by

v	0	$1-p-a$	$\phi - \frac{[1-p-(\xi-s)](1-p-a)}{\xi-s}$	1
$g_a(v)$	$\frac{2p[1-p-(\xi-s)]}{(1-p)(1+p)}$	$\frac{2[1-p-(\xi-s)]^2}{(1-p)(1+p)}$	$\frac{2(\xi-s)[1-p-(\xi-s)]}{(1-p)(1+p)}$	$\frac{2(\xi-s)-(1-p)^2}{(1-p)(1+p)}$

where

$$\phi = \frac{\xi\left(1 - p^2\right) + (1-p)^2}{2(\xi - s)} - 1 - \frac{[1 - p - (\xi - s)](1 - p)}{\xi - s}, \quad \text{and} \quad \tau = \frac{1 - p - (\xi - s)}{\xi - s}.$$

It is not difficult to show that $G_a(v) \in \mathcal{M}(\xi)$ for all $a \in [\xi - s, 1 - p]$; then by Lemma A.1, $\hat{G}^1$ can be induced. This completes the proof. ∎

A.1.3 The case of $p \le s/\xi$

The following four claims establish the results for this case.

Claim A.3. *If $p \le s/\xi$ and $\mu \ge p + 1 - s/\xi$, the degenerate distribution δ_μ is optimal. Furthermore, Seller's revenue is $\Phi(p) = p$.*

Proof of Claim A.3. By using distribution over posteriors δ_μ, because $\mu - p \ge 1 - s/\xi$, the highest possible effective outside option is below Buyer's net value from purchasing

Seller's product. Therefore, Buyer buys without search with probability one, and thus δ_μ must be optimal, and Seller's revenue equals her price p. ■

Claim A.4. *If $p \leq s/\xi$, $\mu < p+1-s/\xi$, and $p \geq (1-2\xi)(\xi-s)/(2\xi^2)$, the binary distribution with support on $\{0, p+1-s/\xi\}$, is robustly optimal. Its cdf is*

$$H_b(w) = \begin{cases} 1 - \frac{\mu\xi}{\xi(1+p)-s} & \textit{if } w \in [0, p+1-s/\xi), \\ 1 & \textit{if } w \in [p+1-s/\xi, 1]. \end{cases} \tag{17}$$

Proof of Claim A.4. Evidently, the concave hull of $H_b\Big|_{[0,1-s/\xi]}(p+z)$ is constant on $[p, p+1-s/\xi]$, and hence any effective outside option distribution with mean $\xi - s$ supported on a subset of $[0, 1-s/\xi]$ is optimal for Nature. When $(1-2\xi)(\xi-s)/(2\xi^2) \leq p < (1-2\xi)(\xi-s)/\xi^2$, consider effective outside option distribution

$$\hat{G}^0(z) = \begin{cases} \left(\frac{2}{\xi-s} - \left(\frac{1-\xi}{\xi}\right)\frac{1+\xi}{\xi(1+p)-s}\right)z + \frac{(1-2\xi)(\xi-s)-\xi^2 p}{\xi^2(1+p)-s\xi} & \text{if } z \in [0, \xi-s), \\ \frac{\xi}{\xi(1+p)-s}(z+p) & \text{if } z \in [\xi-s, 1-s/\xi]; \end{cases}$$

the restriction on p guarantees that $\hat{G}^0$ is a cdf, and it induces

$$\hat{G}^0_p(w) = \begin{cases} 0 & \text{if } w \in [0, p), \\ \left(\frac{2}{\xi-s} - \left(\frac{1-\xi}{\xi}\right)\frac{1+\xi}{\xi(1+p)-s}\right)(w-p) - \frac{(1-2\xi)(\xi-s)-\xi^2 p}{\xi^2(1+p)-s\xi} & \text{if } w \in [p, p+\xi-s), \\ \frac{\xi}{\xi(1+p)-s}w & \text{if } w \in [p+\xi-s, p+1-s/\xi), \\ 1 & \text{if } w \in [p+1-s/\xi, 1]. \end{cases}$$

Inspection shows that the concave hull of $\hat{G}^0_p(w)$ is

$$\widetilde{G}^0_p(w) = \begin{cases} \frac{\xi}{\xi(1+p)-s}w & \text{if } w \in [0, p+1-s/\xi], \\ 1 & \text{if } w \in [p+1-s/\xi, 1]. \end{cases}$$

Consequently, any distribution over posteriors with mean μ and with support on a subset of $\{0\} \cup [p+\xi-s, p+1-s/\xi]$ is optimal for Seller, and this condition is satisfied by the binary distribution H^b. Thus, $(H^b, \hat{G}^0)$ is a saddle point; to show that H^b is indeed robustly optimal, it only remains to show that $\hat{G}^0$ can be induced. Consider

$$\nabla(a) = \frac{a}{\xi(1+p)-s} + 1 - \frac{\xi - s}{\xi^2(1+p) - s\xi} \quad \text{on } [\xi - s, 1 - s/\xi];$$

and let $G_a^0(v)$ be a ternary distribution with pmf $g_a^0(v)$ given by

v	0	$\beta - \gamma a$	$a + s/\xi$
$g_a^0(v)$	$\frac{(1-2\xi)(\xi-s)-\xi^2 p}{\xi^2(1+p)-s\xi}$	p_L	ξ

where

$$p_L = 1 - \xi - \frac{(1-2\xi)(\xi - s) - \xi^2 p}{\xi^2(1+p) - s\xi},$$

$\gamma = \xi/p_L$, and $\beta = (\xi - s)\gamma/\xi$. It can be checked that $G_a^0(v) \in \mathcal{M}(\xi)$ for all $a \in [\xi - s, 1 - s/\xi]$; then by Lemma A.1, $\hat{G}^0$ can be induced.

And for $p \geq (1-2\xi)(\xi - s)/\xi^2$, consider effective outside option distribution

$$\hat{G}^{\xi-s}(z) = \begin{cases} \frac{(1-\xi)^2}{\xi^2(1+p)-s\xi} z & \text{if } z \in [0, \xi - s), \\ \frac{\xi}{\xi(1+p)-s}(z+p) & \text{if } z \in [\xi - s, 1 - s/\xi]; \end{cases}$$

the restriction on p guarantees that $\hat{G}^{\xi-s}$ is a cdf, and it induces

$$\hat{G}_p^{\xi-s}(w) = \begin{cases} 0 & w \in [0, p), \\ \frac{(1-\xi)^2}{\xi^2(1+p)-s\xi}(w-p) & \text{if } w \in [p, p+\xi-s), \\ \frac{\xi}{\xi(1+p)-s} w & \text{if } w \in [p+\xi-s, p+1-s/\xi), \\ 1 & \text{if } w \in [p+1-s/\xi, 1]. \end{cases}$$

Inspection shows that the concave hull of $\hat{G}_p^{\xi-s}(w)$ is

$$\widetilde{G}_p^{\xi-s}(w) = \begin{cases} \frac{\xi}{\xi(1+p)-s} w & \text{if } w \in [0, p+1-s/\xi], \\ 1 & \text{if } w \in [p+1-s/\xi, 1]. \end{cases}$$

Consequently, any distribution over posteriors with mean μ and with support on a subset of $\{0\} \cup [p+\xi-s, p+1-s/\xi]$ is optimal for Seller, and this condition is satisfied by the binary distribution H^b. Thus, $(H^b, \hat{G}^{\xi-s})$ is a saddle point; to show that H^b is indeed robustly optimal, it only remains to show that $\hat{G}^{\xi-s}$ can be induced. Consider

$$\nabla(a) = \frac{a}{\xi(1+p)-s} + 1 - \frac{\xi - s}{\xi^2(1+p) - s\xi} \quad \text{on } [\xi - s, 1 - s/\xi];$$

and let $G_a^{\xi-s}(v)$ be a binary distribution with pmf $g_a^{\xi-s}(v)$ given by

v	$(\xi(1-a)-s)/(1-\xi)$	$a + s/\xi$
$g_a^{\xi-s}(v)$	$1-\xi$	ξ

It can be checked that $G_a^0(v) \in \mathcal{M}(\xi)$ for all $a \in [\xi - s, 1 - s/\xi]$; then by Lemma A.1, $\hat{G}^{\xi-s}$ can be induced. This completes the proof. ■

Claim A.5. *If $p \le s/\xi$, $p \le \mu - (1 - s/\xi)/2$, and $p < (1-2\xi)(\xi - s)/(2\xi^2)$, the distribution*

$$H^h(w) = \begin{cases} 0 & \text{if } w \in [0, p), \\ 2\frac{\xi^2(p+1-\mu)-s\xi}{(\xi-s)^2}(w-p) & \text{if } w \in [p, p+1-s/\xi), \\ 1 & \text{if } w \in [p+1-s/\xi, 1] \end{cases}$$

is robustly optimal. Furthermore, Seller's revenue is

$$\Phi(p) = p\left[1 - 2\xi + 2\xi^2\frac{\mu - p}{\xi - s}\right].$$

Proof of Claim A.5. Observe that the concave hull of $H^h\Big|_{[0,1-s/\xi]}(p+z)$ coincides with $H^h(w)$ on $[p, p+1-s/\xi]$, which is affine. Consequently, any effective outside option distribution with mean $\xi - s$ supported on a subset of $[0, 1-s/\xi]$ is optimal for Nature. Now consider effective outside option distribution

$$\hat{G}^h(z) = 1 - 2\xi + \frac{2\xi^2}{\xi - s}z \quad \text{on } [0, 1 - s/\xi];$$

note that $\hat{G}^h(0) \ge 0$ because $p < (1-2\xi)(\xi - s)/(2\xi^2)$ implies $\xi < 1/2$. Then

$$\hat{G}_p^h(w) = \begin{cases} 0 & \text{if } w \in [0, p), \\ 1 - 2\xi + \frac{2\xi^2}{\xi-s}(w-p) & \text{if } w \in [p, p+1-s/\xi]; \end{cases}$$

and its concave hull is

$$\widetilde{G}_p^h(w) = 1 - 2\xi + \frac{2\xi^2}{\xi - s}(w - p) \quad \text{on } [0, p + 1 - s/\xi].$$

Therefore, any distribution over posteriors with mean μ and with support on a subset of $[0, p + 1 - s/\xi]$ is optimal for Seller, and this condition is satisfied by H^h. Furthermore, Seller's value is $\widetilde{G}_p^h(\mu)$. Thus, $(H^h, \hat{G}^h)$ is a saddle point; and to show that H^h is indeed robustly optimal, it only remains to show that $\hat{G}^h$ can be induced. Consider

$$\nabla(a) = \frac{2\xi}{\xi - s}a - 1 \quad \text{on } [(\xi - s)/(2\xi), 1 - s/\xi];$$

and let $G_a^h(v)$ be a ternary distribution with pmf $g_a^h(v)$ given by

v	0	$1 - s/\xi - a$	$a + s/\xi$
$g_a^h(v)$	$1 - 2\xi$	ξ	ξ

It can be checked that $G_a^h(v) \in \mathcal{M}(\xi)$ for all $a \in [(\xi-s)/(2\xi), 1-s/\xi]$; then by Lemma A.1, $\hat{G}^h$ can be induced by a mixture of outside option distributions. This completes the proof. ■

Claim A.6. *If $p \leq s/\xi$, $\mu - (1 - s/\xi) \geq p > \mu - (1 - s/\xi)/2$, and $p < (1 - 2\xi)(\xi - s)/(2\xi^2)$, there exists $\bar{w} \in [2\mu - p, 1]$ where*

$$\bar{w} = \begin{cases} 2\mu - p & \text{if } \mu > p \text{ and } \xi - s \leq \frac{2(\mu-p)^2}{2\mu-p}, \\ \sqrt{(\xi - s)^2 + 2p(\xi - s)} + (\xi - s + p) & \text{otherwise,} \end{cases}$$

such that the robustly optimal distribution is

$$H_{\bar{w}}(w) = \begin{cases} 1 - \frac{2\mu}{\bar{w}+p} & \text{if } w \in [0, p), \\ 1 - \frac{2\mu}{\bar{w}+p} + \frac{2\mu}{\bar{w}+p}\left(\frac{w-p}{\bar{w}-p}\right) & \text{if } w \in [p, \bar{w}), \\ 1 & \text{if } w \in [\bar{w}, 1]. \end{cases}$$

And Seller's revenue is

$$\Phi(p) = \begin{cases} p\left[1 - \frac{\xi-s}{2(\mu-p)}\right] & \text{if } \mu > p \text{ and } \xi - s \leq \frac{2(\mu-p)^2}{2\mu-p}, \\ \frac{p\mu}{\sqrt{(\xi-s)(\xi-s+2p)}+(\xi-s+p)} & \text{otherwise.} \end{cases} \tag{18}$$

Proof of Claim A.6. In this case,

$$\sqrt{(\xi - s)^2 + 2p(\xi - s)} + (\xi - s + p) < p + 1 - s/\xi \leq 1,$$

where the first inequality holds because $p < (1-2\xi)(\xi - s)/(2\xi^2)$, and the second follows from the assumption that $p \leq s/\xi$. Rearranging the above inequalities, $\xi - s < (1-p)^2/2$. The rest of the proof is analogous to the proof of Claim A.2, and hence omitted.[40] ■

The results in Appendix A.1.2 and Appendix A.1.3 together establish Proposition 1.

A.2 Proof of Corollary 1

Proof of Part (1). Fix μ, ξ, s such that $\mu, \xi \in (0,1)$ and $0 < s < \xi$. Take any $p_1, p_2 \in (0,1)$ such that $p_1 < p_2$. Denote the optimal distributions over posteriors corresponding to p_1 and p_2 by H_{p_1} and H_{p_2}, respectively. By Theorem 7 in Blackwell (1953), the optimal disclosure policy for p_2 is more informative than the one for p_1 if and only if H_{p_2} is a MPS of H_{p_1}. If $p_1 < p_2 \leq 2\mu - 1$, by Proposition 1, $H_{p_1} = H_{p_2}$, and hence H_{p_2} is a MPS of H_{p_1}. Suppose instead that $p_2 > 2\mu - 1$. By Proposition 1, it can be seen that H_{p_1} must cross H_{p_2} only once, and from below. Then since they must have the same mean (namely μ), by Theorem 3.A.44 in Shaked and Shanthikumar (2007), H_{p_2} is a MPS of H_{p_1}. This completes the proof. ■

Proof of Part (2). Let $p_1, p_2 \in (0, s/\xi)$ be such that $p_1 < p_2$. It can be seen from Proposition 1 that, if there exists $i \in \{1,2\}$ such that $p_i \notin [0, (1-2\xi)(\xi - s)/(2\xi^2)) \cap [(\mu - (1 - s/\xi))/2, \mu - (1 - s/\xi))$, by Proposition 1 there must exist $\alpha \in [0,1]$ such that $H_{p_1}(w) \leq H_{p_2}(w)$ for all $w \leq \alpha$, and $H_{p_1}(w) \geq H_{p_2}(w)$ for all $w > \alpha$. Again by Theorem 3.A.44 in Shaked and Shanthikumar (2007), H_{p_2} is a MPS of H_{p_1}, which establishes Part (2)(ii).

For Part (2)(i), take any p_1, p_2 are such that $p_1 < p_2 < (1-2\xi)(\xi - s)/(2\xi^2)$ and $p_1, p_2 \in [(\mu - (1 - s/\xi))/2, \mu - (1 - s/\xi))$. Proposition 1 indicates that, there exists $\varepsilon > 0$ sufficiently small such that $H_{p_1}(w) > H_{p_2}(w)$ for all $w \in (p_1, p_1 + \varepsilon)$. Thus,

$$\int_0^{p_1+\varepsilon} H_{p_1}(w)\, \mathrm{d}w > \int_0^{p_1+\varepsilon} H_{p_2}(w)\, \mathrm{d}w. \tag{19}$$

[40] There is one subtle point, though: for the case of $\bar{w} = 2\mu - p$, to make sure that my choice of the saddle point in the proof of Claim A.2 works, it has to be that $2(\mu - p) \leq 1 - s/\xi$. This must hold by the assumption that $p > \mu - (1 - s/\xi)/2$.

One can also find $\delta > 0$ small enough such that $H_{p_1}(w) > H_{p_2}(w)$ for all $w \in (p_2 + 1 - s/\xi - \delta, p_2 + 1 - s/\xi)$. Consequently, $\int_{p_2+1-s/\xi-\delta}^{1} H_{p_1}(w)\,\mathrm{d}w > \int_{p_2+1-s/\xi-\delta}^{1} H_{p_2}(w)\,\mathrm{d}w$. Because H_{p_1} and H_{p_2} have the same means, $\int_0^1 H_{p_1}(w)\,\mathrm{d}w = \int_0^1 H_{p_2}(w)\,\mathrm{d}w$. Therefore,

$$\int_0^{p_2+1-s/\xi-\delta} H_{p_1}(w)\,\mathrm{d}w > \int_0^{p_2+1-s/\xi-\delta} H_{p_2}(w)\,\mathrm{d}w. \tag{20}$$

But (19) and (20) together imply that H_{p_1} is not a MPS of H_{p_2}, and H_{p_2} is not a MPS of H_{p_1}. Thus, the corresponding optimal disclosure policies are not Blackwell ranked. This completes the proof. ■

A.3 Proof of Theorem 1

I first establish two preliminary results, Claim A.7 and Claim A.8, that concern the cases of $p > s/\xi$ and $p \leq s/\xi$, respectively.

Claim A.7. *Assume $p > s/\xi$. The robust price is*

$$p^* = \begin{cases} \frac{1-\sqrt{2(\xi-s)-(\xi-s)^2}}{1-\xi+s} & \mu \leq \underline{\mu}, \\ 2\mu - 1 & \underline{\mu} < \mu \leq \bar{\mu}, \\ 1 - \sqrt{\xi - s} & \mu > \bar{\mu}. \end{cases} \tag{21}$$

And if $\mu > 1 - (\sqrt{\xi - s}/2)$, any distribution over posteriors $H \in \mathcal{M}(\mu)$ such that $H(w) \leq U[1 - \sqrt{\xi - s}, 1]$ is optimal; if $\mu \leq 1 - (\sqrt{\xi - s}/2)$, the optimal distribution over posteriors is

$$H(w) = \begin{cases} 1 - \frac{2\mu}{1+p^*} & w \in [0, p^*), \\ 1 - \frac{2\mu}{1-p^{*2}}(1-w) & w \in [p^*, 1]. \end{cases}$$

The seller's revenue guarantee is

$$\Pi_h = \begin{cases} \mu\left(1 - \sqrt{2(\xi - s) - (\xi - s)^2}\right) & \mu \leq \underline{\mu}, \\ (2\mu - 1)\left(1 - \frac{\xi - s}{2 - 2\mu}\right) & \underline{\mu} < \mu \leq \bar{\mu}, \\ \left(1 - \sqrt{\xi - s}\right)^2 & \mu > \bar{\mu}. \end{cases} \tag{22}$$

Proof of Claim A.7. Consider any price p with $p > s/\xi$. By Claim A.1 and Claim A.2, after rearranging some inequalities, Seller's expected revenue can be written as a function of

p:

$$\Phi(p) = \begin{cases} p\left(1-\frac{\xi-s}{1-p}\right) & \text{if } p < 2\mu-1, \\ p\left[1-\frac{\xi-s}{2(\mu-p)}\right] & \text{if } \mu-t > p \geq 2\mu-1, \\ \frac{p\mu}{\sqrt{(\xi-s)(\xi-s+2p)}+(\xi-s+p)} & \text{if } 1-\sqrt{2(\xi-s)} > p \geq \max\{\mu-t, 2\mu-1\}, \\ \frac{2\mu}{1+p}\left(1-\frac{\xi-s}{1-p}\right) & \text{if } p \geq 1-\sqrt{2(\xi-s)} \text{ and } p \geq 2\mu-1, \end{cases}$$

where

$$t := \xi - s + \frac{\sqrt{(\xi-s)(\xi-s+8\mu)}}{4} > 0.$$

It can be checked that $\Phi(p)$ is concave and piecewise continuously differentiable. Then note that if $\mu - t > p > 2\mu - 1$, or $1-\sqrt{2(\xi-s)} > p > \max\{\mu-t, 2\mu-1\}$, $\Phi'(p)$ is strictly positive, which implies that these prices yield strictly less revenue than $p = 1-\sqrt{2(\xi-s)}$. Consequently, an optimal price cannot lie in these regions; thus, optimal prices can be solved by maximizing

$$\widetilde{\Phi}(p) = \begin{cases} p\left(1-\frac{\xi-s}{1-p}\right) & \text{if } p < 2\mu-1, \\ \frac{2\mu p}{1+p}\left(1-\frac{\xi-s}{1-p}\right) & \text{if } p \geq 2\mu-1; \end{cases}$$

note that $\widetilde{\Phi}$ is continuous at $p = 2\mu - 1$. To solve this problem, I first solve

$$\max_p \; p\left(1-\frac{\xi-s}{1-p}\right) \quad \text{subject to } \; 0 \leq p \leq 2\mu-1;$$

standard Lagrangian approach shows that the solution is

$$p_1 = \begin{cases} 1-\sqrt{\xi-s} & \text{if } \mu \geq \bar{\mu}, \\ 2\mu-1 & \text{if } \mu < \bar{\mu}. \end{cases}$$

Similarly, the solution to the problem

$$\max_p \; \frac{2\mu p}{1+p}\left(1-\frac{\xi-s}{1-p}\right) \quad \text{subject to } \; 2\mu-1 \leq p \leq 1,$$

is

$$p_2 = \begin{cases} \frac{1-\sqrt{2(\xi-s)-(\xi-s)^2}}{1-\xi+s} & \text{if } \mu \leq \underline{\mu}, \\ 2\mu-1 & \text{if } \mu > \underline{\mu}. \end{cases}$$

Thus, the optimal robust price when $p > s/\xi$, p^*, is given by (21). The other statements in Claim A.7 then follow from Claim A.1, Claim A.2, and (21). ∎

Claim A.8. *Assume $p \leq s/\xi$. Then the robust price is always $p^{**} = s/\xi$. Furthermore, if $s \geq \xi - 2\xi^2$, the optimal distribution over posteriors is the binary distribution*

$$H_b^*(w) = \begin{cases} 1-\mu & w \in [0,1), \\ 1 & w = 1, \end{cases} \tag{23}$$

and the revenue guarantee is $\Pi = \mu s/\xi$.

Suppose instead $s < \xi - 2\xi^2$. If $\mu > p^{**} + (1 - s/\xi)/2 = (1 + s/\xi)/2$, the optimal distribution over posteriors is

$$H_h^*(w) = \begin{cases} 0 & w \in [0, p^{**}), \\ \frac{2\xi^2(1-\mu)}{(\xi-s)^2}(w - p^{**}) & w \in [p^{**}, 1), \\ 1 & w = 1, \end{cases} \tag{24}$$

and the revenue guarantee is

$$\Pi = \frac{s}{\xi} - \frac{2s\xi\,(1-\mu)}{\xi - s}. \tag{25}$$

If $\mu \leq (1 + s/\xi)/2$, the optimal distribution over posteriors is

$$H_{\bar{w}}(w \mid p^{**}) = \begin{cases} 1 - \frac{2\mu}{\bar{w}+p^{**}} & w \in [0, p^{**}), \\ 1 - \frac{2\mu}{\bar{w}+p^{**}} + \frac{2\mu}{\bar{w}+p^{**}}\left(\frac{w-p^{**}}{\bar{w}-p^{**}}\right) & w \in [p^{**}, \bar{w}), \\ 1 & w \in [\bar{w}, 1]; \end{cases} \tag{26}$$

and Seller's revenue guarantee is

$$\Pi = \begin{cases} p^{**}\left[1 - \frac{\xi-s}{2(\mu-p^{**})}\right] & \text{if } \mu > p \text{ and } \xi - s \leq \frac{2(\mu-p^{**})^2}{2\mu-p^{**}}, \\ \frac{p^{**}\mu}{\sqrt{(\xi-s)(\xi-s+2p^{**})}+(\xi-s+p^{**})} & \text{otherwise.} \end{cases} \tag{27}$$

Proof of Claim A.8. If $p \geq (1-2\xi)(\xi - s)/(2\xi^2)$, by Claim A.4, the binary distribution H_b is always optimal whenever $p > \mu + s/\xi - 1$. Then for such a price p, Seller's revenue is

$$p\,[1 - H_b(p + \xi - s)] = \frac{p\mu\xi}{\xi(1+p) - s},$$

which is strictly increasing in p, so the optimal price is just $p^* = s/\xi$, and the resulting profit is $\mu s/\xi$. If instead $p \leq \mu + s/\xi - 1$, Proposition 1 implies that the degenerate distribution δ_μ is optimal, and clearly it is optimal to charge the highest possible price consistent with this case, which makes Buyer buys without search with probability 1. Thus, $p = \Pi = \mu + s/\xi - 1$. Now observe that

$$\left(\mu + \frac{s}{\xi} - 1\right) - \frac{\mu s}{\xi} = \left(1 - \frac{s}{\xi}\right)(\mu - 1) < 0,$$

where the inequality holds since $s < \xi$ and $\mu < 1$. Therefore, the strategy with the degenerate distribution is dominated by the strategy with the binary distribution, and thus one can say that the former is never a part of a revenue guarantee maximizing strategy, and there is no need to worry about this case henceforth.

When $(1 - 2\xi)(\xi - s)/(2\xi^2) > p \geq \mu - (1 - s/\xi)/2$, again by Proposition 1, $H_{\bar{w}}$ defined in (26) is optimal;[41] and a similar argument as in the proof of Claim A.7 shows that Seller's revenue is strictly increasing in p. When $p < \min\{\mu - (1 - s/\xi)/2, (1 - 2\xi)(\xi - s)/(2\xi^2)\}$, Seller chooses price to solve

$$\max_p \; p\left[1 - 2\xi \frac{\xi(p + 1 - \mu) - s}{\xi - s}\right].$$

The objective is strictly concave in p, so it suffices to look at the first-order condition (FOC). The FOC yields

$$p^o = \frac{(1 - 2\xi)(\xi - s)}{4\xi^2} + \frac{\mu}{2};$$

but it can be shown that $p^o \geq \min\{\mu - (1 - s/\xi)/2, (1 - 2\xi)(\xi - s)/(2\xi^2)\}$;[42] then since Seller's objective is strictly concave, in this case it is also strictly increasing in p. Then because the objective function is strictly increasing in p in all three cases above, when $p \leq s/\xi$ it is always optimal to choose $p^{**} = s/\xi$; and the optimal distribution over posteriors and revenue guarantee is determined by which of the three cases $p^{**} = s/\xi$ falls in.

Finally, simple algebra reveals that

$$p^{**} = \frac{s}{\xi} \geq \frac{(1 - 2\xi)(\xi - s)}{2\xi^2}$$

[41] If $(1 - 2\xi)(\xi - s)/(2\xi^2) < \mu - (1 - s/\xi)/2$, then $H_{\bar{w}}(\cdot \mid p^{**})$ is never optimal.

[42] I prove this fact in Claim C.1 in Appendix C.

if and only if $s \geq \xi - 2\xi^2$. This completes the proof. ■

Notice that (26) is just (5), one candidate for optimal distribution over posteriors in the case of $p > s/\xi$, evaluated at p^{**}; and the revenue guarantee in (27) is $\Phi(p)$ defined in (18) evaluated at p^{**}. When $\mu \leq (1 + s/\xi)/2$, one can check that $\Phi'(p) > 0$ even for $p > s/\xi$, and hence the price-information pair $(p^{**}, H_{\bar{w}}(\cdot \mid p^{**}))$ can never be optimal. Consequently, the candidates for optimal selling strategies with $p \leq s/\xi$ are just (p^{**}, H_b^*) and (p^{**}, H_u^*).

Next I am going to compare Seller's revenue guarantee from the above two cases, namely $p > s/\xi$ and $p \leq s/\xi$. To that end, for fixed $\xi \in (0, 1)$ and $s \in (0, \xi)$, define

$$D = \{\mu \in [0,1] \,:\, \mu s/\xi \geq \Pi_h\},$$

and

$$N = \left\{\mu \geq \frac{1 + s/\xi}{2} \,:\, \frac{s}{\xi} - \frac{2s\xi\,(1-\mu)}{\xi - s} \geq \Pi_h\right\},$$

where Π_h is the revenue guarantee for the case of $p > s/\xi$ defined in (22). D and N are the sets of priors that the revenue guarantees from H_b^* and H_h^*, respectively, exceed Π_h.

Note that D is empty if and only if $\mu s/\xi\Big|_{\mu=1} < \Pi_h\Big|_{\mu=1}$, namely $s/\xi \; < (1 - \sqrt{\xi - s})^2$; rearrange,

$$s < \frac{\xi\,(\xi - 1)^2}{(\xi + 1)^2} = B_2(\xi). \tag{28}$$

Moreover, algebra reveals that N is empty if and only if D is empty. Thus, focusing on buy-later demand is optimal if (28) holds. If instead (28) does not hold, define

$$\hat{\mu} = \inf B, \text{ and } \check{\mu} = \inf N.$$

Because $\mu s/\xi$ is linear in μ, and Π_h is also linear for low μ, $\hat{\mu} = 0$ if and only if the slope of the linear part of Π_h is less than or equal to s/ξ, namely

$$\frac{s}{\xi} \geq 1 - \sqrt{2(\xi - s) - (\xi - s)^2}, \tag{29}$$

or equivalently[43]

$$s \geq \frac{\xi(\xi-1)^2}{\xi^2+1} = B_1(\xi). \tag{30}$$

If (30) holds, (p^{**}, H_b^*) is optimal for all $\mu \in [0,1]$; in other words, maximally deterring search is optimal. If both (28) and (30) do not hold, it must be that $B_2(\xi) \leq s < B_1(\xi)$, and $\hat{\mu}, \check{\mu} \in (0,1]$. In this case, Claim A.8 suggests that (p^{**}, H_u^*) dominates (p^{**}, H_b^*) if and only if $s < \xi - 2\xi^2 = B_3(\xi)$. Thus, if further $B_3(\xi) \leq s < B_1(\xi)$, there exists a cutoff $\hat{\mu}$ such that focusing on buy-later demand is optimal whenever $\mu \leq \hat{\mu}$, and maximally deter search is optimal otherwise. If instead $B_2(\xi) < s < B_3(\xi)$, there exists $\check{\mu} \in (0,1)$ such that for $\mu < \check{\mu}$, Seller focuses on buy-later demand; and for $\mu \geq \check{\mu}$, Seller optimally balances between buy-now and buy-later demand.

It only remains to verify that, whenever p^* is the robust price, it indeed satisfies $p^* > s/\xi$. To see this, observe that for (p^*, H) in Claim A.7 to be optimal for some $\mu \in (0,1)$ and $0 < s < \xi < 1$, it must be that

$$\frac{s}{\xi} < 1 - \sqrt{2(\xi-s)-(\xi-s)^2} < \frac{1-\sqrt{2(\xi-s)-(\xi-s)^2}}{1-\xi+s} < 1-\sqrt{\xi-s};$$

where the first inequality holds because a necessary condition for p^* to be optimal is that (29) does not hold, as otherwise p^{**} is always optimal; the second inequality follows because $\xi - s < 1$, and the third follows by algebra. Consequently, both $p^* = (1-\sqrt{2(\xi-s)-(\xi-s)^2})/(1-\xi+s)$ and $p^* = 1-\sqrt{\xi-s}$ satisfy $p^* > s/\xi$. Furthermore, $p^* = 2\mu - 1$ is optimal only if it is strictly above $(1-\sqrt{2(\xi-s)-(\xi-s)^2})/(1-\xi+s)$. Therefore, whenever p^* is optimal, it must be that $p^* > s/\xi$. This completes the proof.

A.4 The "Boundary Cases"

The only candidate robust price that depends on the prior μ is $p = 2\mu - 1$, and by Theorem 1, it is indeed the robust price if and only if

i. either $s < B_2(\xi)$, and $\underline{\mu} \leq \mu < \bar{\mu}$; or

ii. $B_2(\xi) \leq s < B_1(\xi)$, and either

[43] Observe that

$$B_2(\xi) = \frac{\xi(\xi-1)^2}{(\xi+1)^2} < \frac{\xi(\xi-1)^2}{\xi^2+1} = B_1(\xi)$$

for all $\xi \in (0,1)$.

(a) $B_3(\xi) \leq s < B_1(\xi)$, and $\underline{\mu} \leq \mu < \min\{\bar{\mu}, \hat{\mu}\}$, or

(b) $B_2(\xi) \leq s < B_3(\xi)$, and $\underline{\mu} \leq \mu < \min\{\bar{\mu}, \check{\mu}\}$.

These conditions are stringent: the distance between $\underline{\mu}$ and $\bar{\mu}$ is usually very small, and for ii. to be relevant, both $\underline{\mu} \leq \mu < \bar{\mu}$ and $\hat{\mu} \geq \underline{\mu}$ (or $\check{\mu} \geq \underline{\mu}$) must hold.

A.5 Proof of Theorem 2

A.5.1 Proof of Part (i)

For fixed μ and ξ, Theorem 1 indicates that $\hat{s}$, the level of the search cost at which Seller switches to a deterrence policy, must satisfy $\hat{s} = B_i(\xi)$ for some $i = 1, 2, 3$. It is easy to see, from the expressions of p^* and p^{**}, that the robust price

$$p_r = \begin{cases} p^* & \text{if } s \in [0, \hat{s}) \\ p^{**} & \text{if } s \in [\hat{s}, \xi) \end{cases}$$

is increasing in s on $(0, \hat{s})$ and $(\hat{s}, \xi)$. Now it only suffices to show that $p_r(\hat{s}-) > p_s(\hat{s}+)$; that is, $p^*(\hat{s}) > p^{**}(\hat{s})$. After some algebra, one can show that for all $\xi \in (0, 1)$, so long as $s < B_1(\xi)$,

$$\frac{1 - \sqrt{2(\xi - s) - (\xi - s)^2}}{1 - \xi + s} > \frac{s}{\xi}.$$

Then because $B_1(\xi) > \max\{B_2(\xi), B_3(\xi)\}$, it must be that $p^*(\hat{s}) > p^{**}(\hat{s})$ no matter what value does $\hat{s}$ take. This completes the proof.

A.5.2 Proof of Part (ii)

Suppose $s_1 < s_2$. Fix $\xi \in (0, 1)$; for any $s \in (0, \xi)$, let $\hat{\mu}(s)$ and $\check{\mu}(s)$ denote the cutoffs in the statement of Theorem 1 when the search cost is s and the mean of the outside option distribution is ξ. It can be shown that both $\hat{\mu}$ and $\check{\mu}$ are decreasing in s; consequently, $\hat{\mu}(s_1) \geq \hat{\mu}(s_2)$, and $\check{\mu}(s_1) \geq \check{\mu}(s_2)$. Note also that $B_3(\xi) \leq B_2(\xi)$ if and only if $\xi \geq \sqrt{2} - 1$, and hence whenever $\xi \geq \sqrt{2} - 1$, it cannot be that $B_2(\xi) \leq s < B_3(\xi)$. Consequently, it is convenient the two cases, $\xi < \sqrt{2} - 1$ and $\xi \geq \sqrt{2} - 1$, separately.

Case 1: $\xi \geq \sqrt{2} - 1$.

(1-1) $s_1 < s_2 < B_2(\xi)$.

In this case, both s_1 and s_2 correspond to some uniform disclosure policies; let H_{s_1} and H_{s_2} denote the corresponding distribution over posteriors. It suffices to show that H_{s_2} is a MPS of H_{s_1}. Since $s_1 < s_2 < B_2(\xi)$, by Theorem 1, the corresponding robust prices are $p^*(s_1)$ and $p^*(s_2)$, respectively. In particular, $p^*(s_1) > s_1/\xi$, and $p^*(s_2) > s_2/\xi$. Moreover, since $s_1 < s_2 < \hat{s}$, Part (i) implies that p^* is increasing in s on $(0, \hat{s})$. In this region, the search cost only enters the robustly optimal distribution over posteriors via the robust price; then by Corollary 1, H_{s_2} is a MPS of H_{s_1}.

(1-2) $s_1 < B_2(\xi) \leq s_2$.

If $s_2 \geq B_1(\xi)$ then it is obvious that the robustly optimal disclosure policy that corresponds to s_2 is more Blackwell informative since it is full disclosure. Otherwise, there are two cases: $\hat{\mu}(s_2) \leq \mu$ or $\hat{\mu}(s_2) > \mu$. In the first case, again the robustly optimal disclosure policy that corresponds to s_2 is full disclosure, and hence must be more Blackwell informative. In the second, since $\hat{\mu}(s_1) \geq \hat{\mu}(s_2)$, both s_1 and s_2 correspond to some uniform disclosure policies, and thus the same argument as in Case (1-1) would do the work.

(1-3) $B_2(\xi) \leq s_1 < s_2$. There are two sub-cases:

(a) $\mu < \hat{\mu}(s_2)$. Then there are three possibilities: $s_2 < B_1(\xi)$, and hence both s_1 and s_2 correspond to uniform disclosure; $s_1 < B_1(\xi) \leq s_2$, and hence s_1 corresponds to uniform disclosure and s_2 corresponds to full disclosure; $s_1 \geq B_1(\xi)$, and hence both s_1 and s_2 correspond to full disclosure. In all these possibilities, the robustly optimal disclosure policy that corresponds to s_2 is more Blackwell informative than the one that corresponds to s_1.

(b) $\mu \geq \hat{\mu}(s_2)$. Then it must be that s_2 corresponds to full disclosure, and hence the desired statement follows.

Case 2: $\xi < \sqrt{2} - 1$**.**

(2-1) $s_1 < s_2 < B_2(\xi)$.

The same argument as in Case (1-1) proves the desired conclusion.

(2-2) $s_1 < B_2(\xi) \leq s_2$.

If $s_2 \geq B_1(\xi)$ then it is obvious that the robustly optimal disclosure policy that corresponds to s_2 is more Blackwell informative since it is full disclosure. If $B_3(\xi) \leq$

$s_2 < B_1(\xi)$, the same argument as in Case (1-2) proves the desired conclusion. If $B_2(\xi) \leq s_2 < B_3(\xi)$, there are two cases: $\check{\mu}(s_2) > \mu$ or $\check{\mu}(s_2) \leq \mu$. In the first, again the same argument as in Case (1-1) shows that the robustly optimal disclosure policy that corresponds to s_2 is more Blackwell informative. In the second, however, it can be that the disclosure policies correspond to s_1 and s_2 cannot be Blackwell ranked.

(2-3) $B_2(\xi) \leq s_1 < s_2$.

If $s_1 \geq B_3(\xi)$, an analogous argument as in Case (1-3) proves the desired conclusion. If $s_2 \geq B_3(\xi) > s_1$, the definition of D and N in the proof of Theorem 1 implies that $\hat{\mu}(s_2) < \check{\mu}(s_1)$, and hence there are three possibilities: $\hat{\mu}(s_2) \geq \mu$, $\hat{\mu}(s_2) < \mu < \check{\mu}(s_1)$, and $\check{\mu}(s_1) \geq \mu$. The first one is the same as in Case (1-1), and in the other two the robustly optimal disclosure policy that corresponds to s_2 is full disclosure, and hence must be more Blackwell informative.

If instead $B_2(\xi) \leq s_1 < s_2 < B_3(\xi)$. There are three possibilities: $\check{\mu}(s_2) \geq \mu$, $\check{\mu}(s_2) < \mu < \check{\mu}(s_1)$, and $\check{\mu}(s_1) \geq \mu$. In the second, everything is the same as the last possibility considered in Case (2-2); and in the third, the same argument as in Case (1-1) shows that the disclosure policy associated with s_2 is more informative. Now suppose $\check{\mu}(s_2) \geq \mu$, and hence for both s_1 and s_2 the robustly optimal distribution over posteriors is H_u^* defined in (14). Since $s_1 < s_2$, $p^{**}(s_1) < p^{**}(s_2)$, and the slope of the affine segment is strictly higher for s_2. Therefore, H_{u,s_2}^* crosses H_{u,s_1}^* only once and from below; then by Theorem 3.A.44 in Shaked and Shanthikumar (2007), the robustly optimal disclosure policy corresponding to s_2 is *less* Blackwell informative.

Summarizing, for any $s_1 < s_2$, the robustly optimal disclosure policy that corresponds to s_2 is more Blackwell informative than the one that corresponds to s_1 unless $s_2 \in (B_2(\xi), B_3(\xi))$ and $\check{\mu}(s_1) \geq \mu$. This completes the proof.

A.5.3 Proof of Part (iii)

Claim A.9, which is a corollary of Theorem 1, Claim A.7, and Claim A.8, summarizes Seller's revenue guarantee for different parameter values.

Claim A.9. *If $s \geq \xi(\xi-1)^2/(\xi^2+1)$, Seller's revenue guarantee is $\Pi = \mu s/\xi$; and if $s < \xi(\xi-1)^2/(\xi+1)^2$, Seller's revenue guarantee is given by (22). If $\xi(\xi-1)^2/(\xi+1)^2 \leq s <$*

$\xi(\xi-1)^2/(\xi^2+1)$, *there are two cases:*

(1) *If either* $\xi \geq \sqrt{2}-1$, *or* $\xi < \sqrt{2}-1$ *and* $\xi - 2\xi^2 < s < \xi(\xi-1)^2/(\xi^2+1)$, *there exists* $\hat{\mu} \in (0,1)$ *such that for* $\mu < \hat{\mu}$, *Seller's revenue guarantee is given by* (22); *and for* $\mu \geq \hat{\mu}$, *Seller's revenue guarantee is* $\Pi = \mu s/\xi$.

(2) *If instead* $\xi < \sqrt{2}-1$ *and* $\xi(\xi-1)^2/(\xi+1)^2 \leq s < \xi - 2\xi^2$, *there exists* $\check{\mu} \in (0,1)$ *such that for* $\mu < \check{\mu}$, *Seller's revenue guarantee is given by* (22); *and for* $\mu \geq \check{\mu}$, *Seller's revenue guarantee is given by* (25).

Based on Claim A.9, the conclusions made in Part (iv) can be obtained from routine algebraic exercises.

B Proofs for Section 4

B.1 Proof of Proposition 2

When $s = 0$, (3) indicates that $a = 1$. Consequently, Seller's payoff when the distribution over outside options is G can be simplified to $p\,\mathbb{E}_G[1 - H(p+v)]$. Therefore, I can work with the outside option distribution directly; in particular, there is no need to find an effective outside option distribution first and then find an outside option distribution that generates it.

For each pair of distribution over posterior means and effective outside option distribution $(H, \hat{G})$ that is a saddle point in the proof of Claim A.1 and Claim A.2, by setting $s = 0$, the resulting pair (H', G') of distribution over posterior means and outside option distribution is a saddle point for the zero search cost problem. Therefore, the analogs of Claim A.1 and Claim A.2 can be obtained. Using these results, an argument similar to the proof of Claim A.7 establishes the proposition.

B.2 Proof of Proposition 3

For a fixed posterior w, the probability of Buyer buying from Seller is given by $G_p^a(w)$, where when $p + a \leq 1$,

$$G_p^a(w) := \begin{cases} 0 & \text{if } w < p, \\ G(w-p) & \text{if } p \leq w < p+a, \\ 1 & \text{if } w \geq p+a; \end{cases}$$

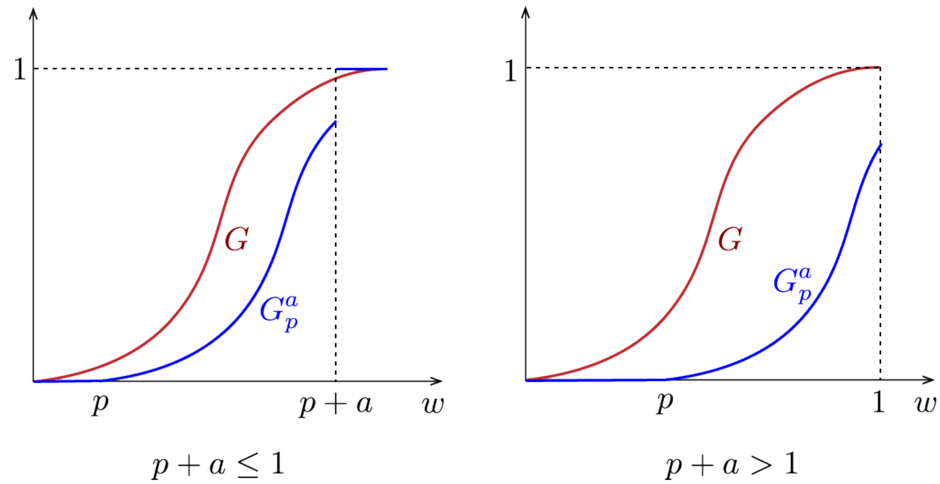

Figure 7: Constructing G_p^a from G. The left panel shows the case of $p+a \le 1$, and the case of $p+a > 1$ is displayed in the right panel.

and

$$G_p^a(w) := \begin{cases} 0 & \text{if } w < p, \\ G(w-p) & \text{if } w \ge p, \end{cases}$$

when $p+a > 1$. See Figure 7 for an illustration of this function. Therefore, if Seller chooses a distribution H, her payoff can be written as $p \int_0^1 G_p^a(w)\,\mathrm{d}H(w)$. Therefore, Seller solves

$$\max_{p \in [0,1]} \left\{ \max_{H \in \mathcal{M}(\mu)} p \int_0^1 G_p^a(w)\,\mathrm{d}H(w) \right\}.$$

To prove Proposition 3, I solve for the optimal information disclosure policy for a fixed price first. For a fixed $p \in [0,1]$, Seller's problem of choosing a distribution over posteriors is

$$\max_{H \in \mathcal{M}(\mu)} \int_0^1 G_p^a(w)\,\mathrm{d}H(w).$$

This problem is identical to the information design problem studied in Kamenica and Gentzkow (2011),[44] where Seller's (who plays the role of Sender in their framework) value function is exactly $G_p^a(w)$. Consequently, Seller's optimal distribution can be identified by finding the concave hull of $G_p^a(w)$.

To identify the concave hull of $G_p^a(w)$, starting from $v = 0$, I try to find a line segment

[44] See the problem on page 2596 in Kamenica and Gentzkow (2011).

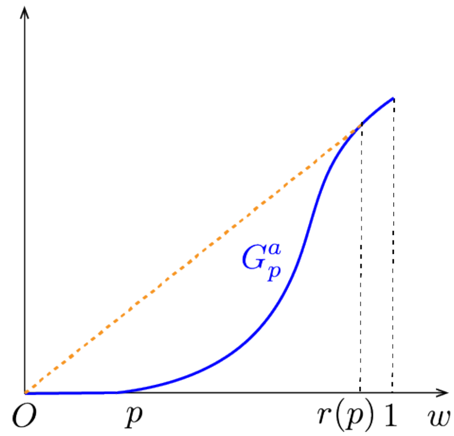

Figure 8: The concave hull of G_p^a when $p > 1 - a$, which is identified by the orange dashed curve.

that tangents to G_p^a at some $r(p)$: $r(p)$ solves

$$g(r(p) - p)r(p) = G(r(p) - p); \tag{31}$$

if a solution to the above equation does not exist for some p, set $r(p) = 1$. Note that when $r(p) < 1$, it must be that $g'(r(p) - p) < 0$. If $p > 1 - a$, it is not hard to see that G_p^a is a convex-concave function on $[0, 1]$, and it is concave on $[r(p), 1]$. Thus, the concave hull of G_p^a, as illustrated in Figure 8, is affine on $[0, r(p)]$, and identical to G_p^a on $[r(p), 1]$.

The case of $p \le 1 - a$ is more complicated. If $g(r(p) - p) \le 1/(p + a)$, that is, the slope of G_p^a at $r(p)$ is less than or equal to the slope of the line segment that connects $(0, 0)$ and $(p + a, 1)$, which I denote by ℓ_p, then the concave hull of G_p^a is essentially identified by ℓ_p. If instead $g(r(p) - p) > 1/(p + a)$, then draw another line segment from $(p + a, 1)$ that tangents to G_p^a at some $t(p)$: $t(p)$ solves

$$g(t(p) - p)(p + a - t(p)) = 1 - G(t(p) - p). \tag{32}$$

And in this case the concave hull is identified by $r(p)$ and $t(p)$. The left and right panel of Figure 9 illustrate the above two cases, respectively.

Being able to identify the concave hull of G_p^a, the optimal disclosure policy for a fixed price is immediate.

Proposition B.1. *Suppose that $p \le 1-a$. Then if $g(r(p)-p) \le 1/(p+a)$, when $\mu \in (0, p+a)$,*

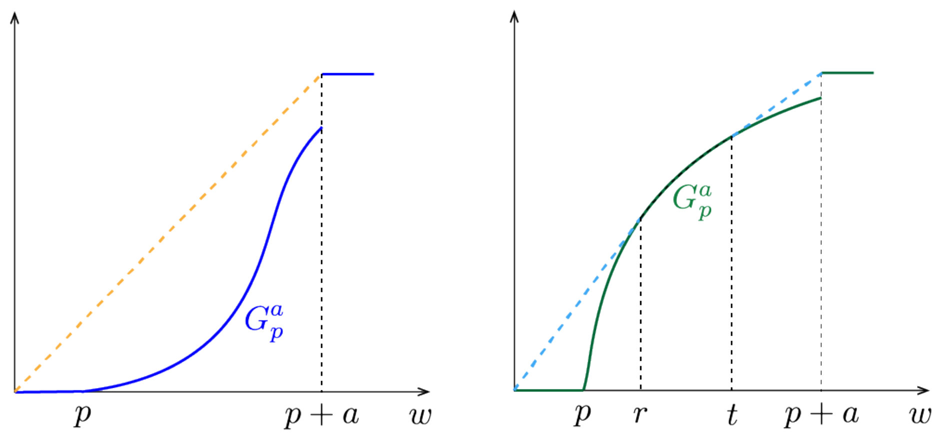

Figure 9: The concave hull of G_p^a when $p \le 1-a$. The left panel depicts the case of $g(r(p)-p) \le 1/(p+a)$, and the concave hull of G_p^a is identified by the orange dashed curve. The right panel illustrates the case of $g(r(p)-p) > 1/(p+a)$, and the concave hull of G_p^a is identified by the cyan dashed curve.

$\{0, p+a\}$ is optimal;[45] *and when $\mu \in [p+a, 1]$, $\{\mu\}$ is optimal. If $g(r(p)-p) > 1/(p+a)$,*

- *when $\mu \in (0, r(p))$, $\{0, r(p)\}$ is optimal;*
- *when $\mu \in [r(p), t(p)]$, $\{\mu\}$ is optimal;*
- *when $\mu \in (t(p), p+a)$, $\{t(p), p+a\}$ is optimal; and*
- *when $\mu \in [p+a, 1)$, $\{\mu\}$ is optimal.*

Suppose instead that $p > 1-a$. Then when $\mu \in (0, r(p))$, $\{0, r(p)\}$ is optimal; and when $\mu \in [r(p), 1)$, $\{\mu\}$ is optimal.

Now I am ready to prove Proposition 3.

Proof of Proposition 3. Suppose first that $p > 1-a$. Let

$$\hat{p} = \sup\{p \in [0,1] \, : \, \text{there exists } r(p) \in [0,1] \text{ that solves (31)}\}.$$

If G is concave, let $r(0) = 0$; and otherwise let $r(0)$ solve $g(r(0))r(0) = G(r(0))$ if a solution exists, or else let $r(0) = 1$. By the implicit function theorem,

$$r'(p) = r(p) - \frac{g(r(p)-p)}{g'(r(p)-p)}. \tag{33}$$

[45]Since every optimal distribution over posteriors is either degenerate or binary, I identify such a distribution by its support.

By Proposition B.1, the optimal disclosure policy depends on the location of μ: if $\mu \in (0, r(p))$, $\{0, r(p)\}$ is optimal; and if $\mu \in [r(p), 1)$, $\{\mu\}$ is optimal. Suppose $\mu \in (0, r(p))$, Seller's payoff by setting $p \in (1-a, \hat{p})$ is given by

$$p\frac{\mu}{r(p)}G(r(p)-p) = p\mu g(r(p)-p),$$

where the equality follows from (31) in the text; and by setting $p \in [\hat{p}, 1)$, Seller's payoff is $p\mu G(1-p)$. It can be checked that, using (33), when $p \in (1-a, \hat{p})$ Seller's payoff is strictly increasing in p, and hence Seller fully discloses information by using distribution $\{0, 1\}$, with price being p_h defined by the solution of Equation (15).

Now suppose Seller uses no disclosure strategy $\{\mu\}$; her payoff is given by $pG(\mu - p)$, which is maximized at

$$p_\mu = \frac{G(\mu - p_\mu)}{g(\mu - p_\mu)}. \tag{34}$$

Note that by definition of $r(p)$, $r(p) > p$ unless possibly at $p = 0$, which is never optimal. Then

$$\frac{G(r(p_\mu) - p_\mu)}{g(r(p_\mu) - p_\mu)} = r(p_\mu) > p_\mu = \frac{G(\mu - p_\mu)}{g(\mu - p_\mu)},$$

where the first equality follows from (31), and the second equality holds by (34). By log-concavity of g, it must be that $\mu < r(p_\mu)$. Then the optimal distribution is in fact $\{0, r(p_\mu)\}$, which implies that no disclosure is never optimal. To summarize, when $p > 1-a$, full disclosure is optimal, and Seller's optimal price and profits are given by p_h and $p_h\mu G(p_h - \mu)$, respectively, where p_h is the solution to (32).

Now suppose that $p \leq 1-a$. Again by the implicit function theorem,

$$t'(p) = 1; \tag{35}$$

consequently, both $r(p)$ and $t(p)$ are increasing in p, and by (35),

$$t(p) = t(0) + p. \tag{36}$$

If $g(r(0)) \leq 1/a$, it can be checked that for all $p \leq 1-a$, $g(r(p)-p) \leq 1/(p+a)$. Then by Proposition B.1, the distribution $\{0, p+a\}$ is optimal; and by using this distribution, Buyer buys if and only if she buys without search, which happens with probability $\mu/(p+a)$. Consequently, Seller's expected payoff is $p\mu/(p+a)$, which is strictly increasing in p. It is maximized at $p = 1-a$, with profit $(1-a)\mu$; the associated distribution is $\{0, 1\}$, namely full disclosure. To summarize, when $g(r(0)) \leq 1/a$, full disclosure is optimal; $p = 1 - a$ is the optimal price, and Seller's expected payoff is $\mu(1 - a)$.

Consider next the case that $g(r(0)) > 1/a$. Again by Proposition B.1, define

$$\bar{p} = \sup \left\{ p \in [0, 1 - a] \, : \, g(r(p) - p) > \frac{1}{p + a} \right\};$$

by log-concavity of g, $\bar{p}$ is unique. Now for a fixed p, optimal information disclosure policy again depends on the location of μ: if $\bar{p} \leq p \leq 1 - a$, $\{0, p + a\}$ is optimal; and for $p < \bar{p}$,

- if $\mu \in (0, r(p))$, $\{0, r(p)\}$ is optimal,
- if $\mu \in [r(p), t(p)]$, $\{\mu\}$ is optimal,
- if $\mu \in (t(p), p + a)$, $\{t(p), p + a\}$ is optimal,
- if $\mu \in [p + a, 1)$, $\{\mu\}$ is optimal.

An analogous argument like the case of $p > 1-a$ shows that no disclosure is never optimal; and $p < \bar{p}$ implies that $r(p) < p + a$, but then Seller's expected payoff is strictly increasing in p, which implies that it is strictly better for Seller to price at $\bar{p}$ and disclose according to $\{0, p + a\}$ instead. Consequently, it only remains to consider the third bulletpoint.

To this end, suppose Seller discloses according to $\{t(p), p + a\}$. Seller's problem of finding the optimal price is

$$p \left[G(t(p) - p) \frac{p + a - \mu}{p + a - t(p)} + \frac{\mu - t(p)}{p + a - t(p)} \right];$$

by (32), it can be written as

$$p \left[G(t(0)) \frac{p + a - \mu}{a - t(0)} + \frac{\mu - p - t(0)}{a - t(0)} \right];$$

because $G(t(0)) < 1$, it is strictly concave in p. Then the optimal price is given by

$$p_t = \frac{aG(t(0)) - t(0)}{2[1 - G(t(0))]} + \frac{\mu}{2}.$$

For this price to be indeed optimal, it has to be that $p_t < \bar{p}$. Observe that at $\bar{p}$, $r(\bar{p}) =$

$t(\bar{p}) = t(0) + \bar{p}$. Then by definition of $\bar{p}$,

$$g(r(\bar{p}) - p) = \frac{G(r(\bar{p}) - \bar{p})}{r(\bar{p})} = \frac{1}{\bar{p} + a},$$

and hence

$$\bar{p} = \frac{aG(t(0)) - t(0)}{1 - G(t(0))}.$$

But then $p_t < \bar{p}$ implies that $\mu < t(0) + p_t = t(p_t)$, which in turn implies that $\{t(p_t), p_t + a\}$ is not optimal at p_t. Therefore, the only candidate for optimal selling strategy is $(p, \{0, p+a\})$. Consequently, similar to the case of $g(r(0)) \leq 1/a$, full disclosure is optimal, $p = 1 - a$ is the optimal price, and Seller's expected payoff is $\mu(1 - a)$. Summarizing, when $p \leq 1 - a$, the selling strategy described above is optimal.

To conclude, full disclosure is always optimal, the choice of the optimal selling strategy boils down to comparing $1 - a$ and $p_h G(1 - p_h)$. Then $p = 1 - a$ is optimal if and only if $p_h G(1 - p_h) \leq 1 - a$, and otherwise p_h is optimal. This yields the statement in the proposition, and hence concludes the proof. ■

B.3 Proof of Corollary 2

Because the density of the outside option distribution g is strictly positive, (3) indicates that a is strictly decreasing in s. Because p_h does not depend on s, there exists a unique a^* that solves $1 - a = p_h G(1 - p_h)$; let the search cost that correspond to a^* by $\hat{s}_G$. Then the statement follows from Proposition 3.

B.4 Proof of Corollary 3

The price part of Part (i) follows directly from the expression of p^o in Proposition 3, and the profit part holds since

$$\pi^o = \begin{cases} \mu(1 - a) & \text{if } 1 - a \geq p_h G(1 - p_h), \\ \mu p_h G(1 - p_h) & \text{if } 1 - a < p_h G(1 - p_h). \end{cases}$$

For Part (ii), Corollary 2 and the fact that p_h does not depend on s together implies the first assertion. And when $s \geq \hat{s}_G$, by Corollary 2, the optimal price is $1 - a$, and Seller's profit is $\mu(1 - a)$. Then since a is strictly decreasing in s, the second assertion follows.

B.5 Proof of Proposition 4

I prove part (i) first. For a fixed p, Seller minimizes $H((p+\xi-s)^-)$. There are two cases: $p+\xi-s \leq \mu$ and $p+\xi-s > \mu$. If $p+\xi-s \leq \mu$, the optimal distribution over posteriors is the degenerate distribution at μ, which correspond to no disclosure. Consequently, Buyer buys with probability 1, and hence Seller's revenue is exactly p. Thus, it is optimal for Seller to set $p = \mu - \xi + s$ provided the right-hand side is nonnegative, and her profit is $\mu - \xi + s$.

Another case is $p+\xi-s > \mu$. In this case, to minimize $H((p+\xi-s)^-)$, it is optimal to put as much mass as possible at $p+\xi-s$, and put the rest of the mass at 0 so that $\mathbb{E}_H[w] = \mu$. Thus, the optimal distribution for a fixed p is the binary distribution with support on $\{0, p+\xi-s\}$. Consequently, Seller's revenue is $p\mu/(p+\xi-s)$, and one can show that this expression is strictly increasing in p. Therefore, the optimal price is $1-\xi+s$, the optimal distribution is the binary distribution with support on $\{0,1\}$, and Seller's profit is $\mu(1-\xi+s)$.

Finally, note that $\mu(1-\xi+s) \geq \mu-\xi+s$ and the inequality is strict for all $\mu \in (0,1)$. Hence full disclosure, namely the binary distribution with support on $\{0,1\}$, and price $1-\xi+s$ is the optimal selling strategy when Seller can commit to an exploding offer.

Part (ii) is obtained by comparing Seller's profits in this case, namely $\mu(1-\xi+s)$, with her revenue guarantee I solved in the proof of Theorem 1. Finally, Part (iii) follows from a similar argument as the proof of Proposition 4 in Armstrong and Zhou (2016).

B.6 Proof of Proposition 5

Proof of part (i). In this case the outside option distribution is fixed at G^*. If H^* corresponds to full disclosure, that is, it is the binary distribution with mean μ and with support on $\{0,1\}$. Under full disclosure, there is no buy-later demand and hence charging different prices for "buy-now" and "buy-later" does not increase Seller's payoff.

If H^* is not the binary distribution, then $1-H^*(p)$ is strictly log-concave on $(p_r, \sup\{\operatorname{supp}(H^*)\})$. Then by Proposition 1 in Armstrong and Zhou (2016), Seller benefits from such a deviation. This concludes the proof. ∎

Proof of part (ii). Suppose Seller offers prices (p_1, p_2) with $p_1 < p_2$, and upon observing this Nature can choose an outside option distribution different from G^*. Seller's total demand is given by

$$1-\mathbb{E}_G\left[H^*\left(p_2+\min\left\{v,S_G^{-1}(s+p_2-p_1)\right\}\right)\right],$$

and the buy-now demand is $1-H^*\left(p_2+S_G^{-1}(s+p_2-p_1)\right)$. Thus, Seller's buy-later demand is given by

$$\begin{aligned}&1-\mathbb{E}_G\left[H^*\left(p_2+\min\left\{v,S_G^{-1}(s+p_2-p_1)\right\}\right)\right]-\left[1-H^*\left(p_2+S_G^{-1}(s+p_2-p_1)\right)\right]\\=&H^*\left(p_2+S_G^{-1}(s+p_2-p_1)\right)-\mathbb{E}_G\left[H^*\left(p_2+\min\left\{v,S_G^{-1}(s+p_2-p_1)\right\}\right)\right].\end{aligned}$$

Consequently, Seller's expected payoff is

$$\begin{aligned}&p_1\left[1-H^*\left(p_2+S_G^{-1}(s+p_2-p_1)\right)\right]+\\&p_2\left[H^*\left(p_2+S_G^{-1}(s+p_2-p_1)\right)-\mathbb{E}_G\left[H^*\left(p_2+\min\left\{v,S_G^{-1}(s+p_2-p_1)\right\}\right)\right]\right],\end{aligned}$$

which is equivalent to

$$p_2\mathbb{E}_G\left[1-H^*\left(p_2+\min\left\{v,S_G^{-1}\left(s+p_2-p_1\right)\right\}\right)\right]-(p_2-p_1)\left[1-H^*\left(p_2+S_G^{-1}\left(s+p_2-p_1\right)\right)\right].$$

But then since H is affine on $(p_r,1)$, the first term is constant in the choice of G, and the second is minimized by choosing $G=\delta_\xi$. Then the above expression becomes

$$p_2\left[1-H^*\left(p_2+\xi-s-p_2+p_1\right)\right]-(p_2-p_1)\left[1-H^*\left(p_2+\xi-s-p_2+p_1\right)\right]=p_1\left[1-H^*\left(p_1+\xi-s\right)\right].$$

But this implies that Seller cannot benefit from setting different buy-now and buy-later prices. ■

C Supplementary Results

Claim C.1. *For any $\xi<1/2$,*

$$p^o=\frac{(1-2\xi)(\xi-s)}{4\xi^2}+\frac{\mu}{2}\geq\min\left\{\frac{(1-2\xi)\left(\xi-s\right)}{2\xi^2},\mu-\frac{1-s/\xi}{2}\right\}.$$

Proof. The desired inequality is equivalent to that of one of the two inequalities below holds:

$$\frac{(1-2\xi)(\xi-s)}{4\xi^2}+\frac{\mu}{2}\geq\frac{(1-2\xi)\,(\xi-s)}{2\xi^2},$$
$$\frac{(1-2\xi)(\xi-s)}{4\xi^2}+\frac{\mu}{2}\geq\mu-\frac{\xi-s}{2\xi}.$$

And the two inequalities are equivalent to

$$\mu\geq\frac{(1-2\xi)\,(\xi-s)}{2\xi^2} \tag{37}$$

and

$$\mu\leq\frac{\xi-s}{2\xi^2}, \tag{38}$$

respectively. Now observe that for every $\mu\in[0,1]$ such that (37) fails,

$$\mu<(1-2\xi)\,\frac{\xi-s}{2\xi^2}\leq\frac{\xi-s}{2\xi^2},$$

where the last inequality holds since $1-2\xi\in[0,1]$ for all $\xi\in(0,1/2)$, and $s<\xi$ by assumption. Therefore, (38) must hold whenever (37) fails. ■

References

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

Armstrong, M. (2017): "Ordered Consumer Search," *Journal of the European Economic Association*, 15, 989–1024.

Armstrong, M. and J. Zhou (2016): "Search Deterrence," *The Review of Economic Studies*, 83, 26–57.

Au, P. H. and M. Whitmeyer (2022): "Attraction Versus Persuasion," *Journal of Political Economy*, forthcoming.

Babichenko, Y., I. Talgam-Cohen, H. Xu, and K. Zabarnyi (2021): "Regret-Minimizing Bayesian Persuasion," Available at arXiv:2105.13870.

Bar-Isaac, H., G. Caruana, and V. Cuñat (2010): "Information Gathering and Marketing," *Journal of Economics & Management Strategy*, 19, 375–401.

Bergemann, D. and K. Schlag (2011): "Robust Monopoly Pricing," *Journal of Economic Theory*, 146, 2527–2543.

Bergemann, D. and K. H. Schlag (2008): "Pricing without Priors," *Journal of the European Economic Association*, 6, 560–569.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, 24, 265–272.

Boleslavsky, R. and C. Cotton (2015): "Grading Standards and Education Quality," *American Economic Journal: Microeconomics*, 7, 248–279.

——— (2018): "Limited Capacity in Project Selection: Competition Through Evidence Production," *Economic Theory*, 65, 385–421.

Boleslavsky, R., C. S. Cotton, and H. Gurnani (2017): "Demonstrations and Price Competition in New Product Release," *Management Science*, 63, 2016–2026.

Brown, A. L., A. Viriyavipart, and X. Wang (2018): "Search Deterrence in Experimental Consumer Goods Markets," *European Economic Review*, 104, 167–184.

Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira (2018): "Optimal Selling Mechanisms under Moment Conditions," *Journal of Economic Theory*, 177, 245–279.

Carroll, G. (2019): "Robustness in Mechanism Design and Contracting," *Annual Review of Economics*, 11, 139–166.

Chatterjee, R. (2009): "Strategic Pricing of New Products and Services," in *Handbook of Pricing Research in Marketing*, ed. by V. R. Rao, Cheltenham: Edward Elgar Publishing, 169–215.

Che, Y.-K. and W. Zhong (2022): "Robustly Optimal Mechanisms for Selling Multiple Goods," Available at arXiv:2105.02828.

Choi, M., A. Y. Dai, and K. Kim (2018): "Consumer Search and Price Competition," *Econometrica*, 86, 1257–1281.

Du, S. (2018): "Robust Mechanisms Under Common Valuation," *Econometrica*, 86, 1569–1588.

Dworczak, P. and A. Pavan (2022): "Preparing for the Worst But Hoping for the Best: Robust (Bayesian) Persuasion," *Econometrica*, 90, 2017–2051.

Ellison, S. F. (2016): "Price Search and Obfuscation: An Overview of the Theory and Empirics," in *Handbook on the Economics of Retailing and Distribution*, ed. by E. Basker, Cheltenham: Edward Elgar Publishing, 287–305.

Fainmesser, I. P., D. O. Lauga, and E. Ofek (2021): "Ratings, Reviews, and the Marketing of New Products," *Management Science*, 67, 7023–7045.

He, W. and J. Li (2021): "Competitive Information Disclosure in Random Search Markets," Chinese University of Hong Kong working paper.

Heiman, A. and E. Muller (1996): "Using Demonstration to Increase New Product Acceptance: Controlling Demonstration Time," *Journal of Marketing Research*, 33, 422–430.

Hinnosaar, T. and K. Kawai (2020): "Robust Pricing with Refunds," *The RAND Journal of Economics*, 51, 1014–1036.

Hu, J. and X. Weng (2021): "Robust Persuasion of A Privately Informed Receiver," *Economic Theory*, 72, 909–953.

Hwang, I., K. Kim, and R. Boleslavsky (2019): "Competitive Advertising and Pricing," Seoul National University working paper.

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

Kohn, M. G. and S. Shavell (1974): "The Theory of Search," *Journal of Economic Theory*, 9, 93–123.

Kosterina, S. (2022): "Persuasion with Unknown Beliefs," *Theoretical Economics*, 17, 1075–1107.

Lyu, C. (2021): "Information Design for Selling Search Goods and the Effect of Competition," University of Wisconsin working paper.

Milgrom, P. and J. Roberts (1986): "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94, 796–821.

Nocke, V., M. Peitz, and F. Rosar (2011): "Advance-purchase discounts as a price discrimination device," *Journal of Economic Theory*, 146, 141–162.

Osborne, M. J. and A. Rubinstein (1994): *A Course in Game Theory*, Cambridge, MA: MIT Press.

Pan, S. and X. Zhao (2020): "Commitment and Cheap Talk in Search Deterrence," *The RAND Journal of Economics*, forthcoming.

Parakhonyak, A. and A. Sobolev (2022): "Persuasion without Priors," University of Oxford working paper.

Sapiro-Gheiler, E. (2021): "Persuasion with Ambiguous Receiver Preferences," Available at arXiv:2109.11536v2.

Shaked, M. and J. G. Shanthikumar (2007): *Stochastic Orders*, Berlin: Springer.

Wang, C. (2017): "Advertising as A Search Deterrent," *The RAND Journal of Economics*, 48, 949–971.