Buyer-Optimal Learning and Monopoly Pricing*

Anne-Katrin Roesler† and Balázs Szentes‡

January 20, 2017

Abstract

This paper analyzes a bilateral trade model where the buyer's valuation for the object is uncertain and she observes only a signal about her valuation. The seller gives a take-it-or-leave-it offer to the buyer. Our goal is to characterize those signal structures which maximize the buyer's expected payoff. We identify a buyer-optimal signal structure which generates (i) efficient trade and (ii) a unit-elastic demand. Furthermore, we show that every other buyer-optimal signal structure yields the same outcome as the one we identify, in particular, the same price.

1 Introduction

Rapid developments in information technology have given consumers access to new information sources that in many cases allow them to acquire product information prior to trading. For an example, consider Yelp, a mobile phone app that aggregates customer reviews and makes them available to new customers. Ceteris paribus, a single consumer is better off with such an app because she becomes more efficient at choosing the store selling her most preferred products. On the other hand, after the appearance of the app, stores will face customers with a different (and perhaps higher) distribution of willingness-to-pay. Therefore, stores are likely to respond by increasing their prices. As a consequence, a new information source - even when freely available - might worsen the buyer's outcome if being better informed results in higher prices. The question then becomes: How does an information source affect prices and consumer welfare? Our paper contributes to this problem by identifying information structures which are best for the buyers in a monopolistic market. Building on this result, we then characterize all the possible combinations of consumer and producer surplus which can arise from *some* learning.

*This paper supersedes the solo paper of the first author which was circulated with the title "Is Ignorance Bliss? Rational Inattention and Optimal Pricing." We have benefited from discussions with Dirk Bergemann, Ben Brooks, Daniele Condorelli, Eddie Dekel, Yingni Guo, Johannes Hörner, Emir Kamenica, Daniel Krähmer, Benny Moldovanu, Tymofiy Mylovanov.

†Department of Economics, University of Michigan, Ann Arbor, MI 48109. E-mail: akroe@umich.edu.

‡Department of Economics, London School of Economics, London, UK. E-mail: b.szentes@lse.ac.uk.

We consider a stylized model where the seller has full bargaining power and the buyer receives a costless signal about her true valuation. The seller sets a price knowing the joint distribution of the buyer's valuation and the signal but not their realizations. We assume that the buyer's valuation is always positive and the seller's outside option is zero, such that efficiency requires trade with probability one. The paper's main theorem identifies the least informative buyer-optimal signal distribution. This distribution generates a unit-elastic demand and the seller is indifferent between charging any price on its support. In equilibrium, the seller sets the lowest price, p^*, and, since this is the lowest possible value estimate of the buyer, trade always occurs. We show that all buyer-optimal signal structures yield the same outcome, namely that the price is p^* and trade always takes place.

An important observation is that the buyer's signal does not enhance the efficiency of her purchasing decision along the equilibrium path. The buyer always buys the good at the equilibrium price, irrespective of her value estimate and typically even if the price exceeds her true valuation. The buyer's optimal learning is therefore driven exclusively by the goal of generating a demand function which induces the lowest possible price subject to efficient trade.

Finally, for any given value distribution of the buyer, we characterize the combinations of consumer and producer surpluses which can arise as an equilibrium outcome if the buyer has access to *some* information source.[1] We show that any consumer-producer surplus pair can be implemented if and only if (i) consumer surplus is non-negative, (ii) producer surplus is weakly larger than p^*, and (iii) total surplus is weakly smaller than the first-best surplus.[2]

Our observation that, in a contracting environment, an imperfectly informed agent can be better off than a perfectly informed one is also noticed by Kessler (1998). The author considers a principal-agent model where, prior to contracting, an agent observes a payoff relevant state with a certain probability. Of course, the principal's contract depends on this probability. The author shows that the agent's payoff is not maximized when this probability is one, that is, an agent might benefit from having less than perfect information. We not only confirm this result, but also characterize buyer-optimal learning without imposing any constraints on possible signal structures.

Several papers examine buyers' incentives to acquire costly information about their valuations before participating in auctions. The buyers' learning strategies depend on the selling mechanism announced by the seller. Persico (2000) shows that if the buyers' signals are affiliated then they acquire more information in a first-price auction than in a second-price one. Compte and Jehiel (2007) show that dynamic auctions tend to generate higher revenue than simultaneous ones. Shi (2011) also analyses models where it is costly for the buyers to learn about their valuations and identifies the revenue-maximizing auction in private-value environments. In all of these setups, the

[1] We are grateful to Ben Brooks for drawing our attention to this problem.

[2] In contrast, Bergemann et al. (2015) characterize the possible outcomes in a model where the buyer is fully informed and the seller observes a signal about the buyer's valuation.

seller is able to commit to a selling mechanism before the buyers decide how much information to acquire. In contrast, we characterize buyer-optimal learning in environments where the monopolist best-responds to the buyer's signal structure.[3]

Condorelli and Szentes (2015) also consider a bilateral trade model. In contrast to our setup, the distribution of the buyer's valuation is not given exogenously. Instead, the buyer chooses her value-distribution supported on a compact interval and perfectly observes its realization. The seller observes the buyer's distribution but not her valuation and sets a price. The authors show that, as in our model, the equilibrium distribution generates a unit-elastic demand and trade always occurs.

2 Model

There is a seller who has an object to sell to a single buyer. The buyer's valuation, v, is distributed according to the continuous CDF F supported on $[0,1]$. Let μ denote the expected valuation, that is, $\int_0^1 v dF(v) = \mu$. The buyer observes a signal s about v. The joint distribution of v and s is common knowledge. The seller then gives a take-it-or-leave-it price offer to the buyer, p.[4] Finally, the buyer trades if and only if her expected valuation conditional on her signal weakly exceeds p.[5] If trade occurs, the payoff of the seller is p and the payoff of the buyer is $v-p$; otherwise, both have a payoff of zero. Both the seller and the buyer are von Neumann-Morgenstern expected payoff maximizers. In what follows, we fix the CDF F and analyze those signal structures which maximize the buyer's expected payoff.

Since the buyer's trading decision only depends on $E(v|s)$, we may assume without loss of generality that each signal s provides the buyer with an unbiased estimate about her valuation, that is, $E(v|s) = s$. In what follows, we restrict attention to such signals and refer to them as *unbiased signals*.

3 Results

First, we argue that the payoffs of both the buyer and the seller are determined by the marginal distribution of the signal. To this end, let $D(G,p)$ denote the demand at price p if the signal's distribution is G, that is, $D(G,p)$ is the probability of trade at p. Note that $D(G,p) = 1 - G(p) + \Delta(G,p)$ where $\Delta(G,s)$ denotes the the probability of s according to the CDF G.[6] The seller's optimal price, p, solves $\max_s sD(G,s)$ and the buyer's payoff is $\int_p^1 (s-p)\, dG(s)$. Therefore, the

[3] Another strand of the literature analyzes the seller's incentives to reveal information about the buyers' valuations prior to participating in an auction, see for example, Ganuza (2004), Ganuza and Penalva (2010) and Bergemann and Pesendorfer (2007).

[4] It is well-known that any profit-maximizing mechanism can be implemented by a take-it-or-leave-it offer.

[5] Assuming that the buyer is non-strategic in the final stage of the game and buys the object whenever her value estimate weakly exceeds the price has no effect on our results but makes the analysis simpler.

[6] Formally, $\Delta(G,s) = G(s) - \sup_{x<s} G(x)$. If G has no atom at s then $\Delta(G,s) = 0$.

problem of designing a buyer-optimal signal structure can be reduced to identifying the marginal signal distribution which maximizes the buyer's expected payoff subject to monopoly pricing. Of course, not every CDF corresponds to a signal distribution. In what follows, we characterize the set of distributions that do.

For each unbiased signal structure, v can be expressed as $s+\varepsilon$ for a random variable ε with $E(\varepsilon|s)=0$. This means that G is the distribution of *some* unbiased signal about v if and only if F is a mean-preserving spread of G (see Definition 6.D.2 of Mas-Colell et al., 1995). Let $\mathcal{G}_F$ denote the set of CDFs of which F is a mean-preserving spread. By Proposition 6.D.2 of Mas-Colell et al. (1995) this set can be defined as follows

$$\mathcal{G}_F=\left\{G\in\mathcal{G}:\int_0^x F(v)\,dv\geq\int_0^x G(s)\,ds \text{ for all } x\in[0,1],\ \int_0^1 s dG(s)=\mu\right\}. \tag{1}$$

The problem of designing a buyer-optimal signal structure can be stated as follows

$$\begin{aligned} &\max_{G\in\mathcal{G}_F}\int_p^1 (s-p)\,dG(s)\\ \text{s.t. } p \quad &\in \quad \arg\max_s sD(G,s). \end{aligned}$$

In what follows, we call a pair (G,p) an *outcome* if $G\in\mathcal{G}_F$ and $p\in\arg\max_s sD(G,s)$. In other words, the pair (G,p) is an outcome if there exists an unbiased signal about v which is distributed according to the CDF G and it induces the seller to set price p.

Next, we define a set of distributions and prove that a buyer-optimal signal distribution lies in this set. For each $q\in(0,1]$ and $B\in[q,1]$ let the CDF G_q^B be defined as follows:

$$G_q^B(s)=\begin{cases} 0 & \text{if } s\in[0,q),\\ 1-\frac{q}{s} & \text{if } s\in[q,B),\\ 1 & \text{if } s\in[B,1]. \end{cases}$$

Observe that the support of G_q^B is $[q,B]$ and it specifies an atom of size q/B at B. An important attribute of each CDF in this class is that the seller is indifferent between charging any price on its support. To see this, note that G_q^B generates a unit-elastic demand on its support, $D\left(G_q^B,p\right)=q/p$, so that the seller's profit is q irrespective of the price on $[q,B]$. Of course, the seller is strictly worse off by setting a price outside of the support. Notice that G_q^q is a degenerate distribution which specifies an atom of size one at q.

The next lemma states that for each outcome, (G,p), there is another outcome $\left(G_q^B,q\right)$ which makes the buyer weakly better off while generating the same profit to the seller as G.

Lemma 1 *Suppose that (G,p) is an outcome and let π denote $pD(G,p)$. Then there exists a unique $B\in[\pi,1]$ such that*

(i) G is a mean-preserving spread of G_π^B,

(ii) $\left(G_\pi^B,\pi\right)$ is an outcome and

(iii) $\int_\pi^1(s-\pi)\,dG_\pi^B(s)\geq\int_p^1(s-p)\,dG(s)$ and the inequality is strict if $D(G,p)<1$.

Part (i) states that the signal distributed according to G_π^B is weakly less informative than the one distributed according to G. Parts (ii) and (iii) imply that this less informative signal generates the same payoff to the seller and a higher payoff to the buyer.

Proof. We first argue that there exists a unique $B \in [\pi, 1]$ such that G_π^B generates an expected value of μ. Since $p \in \arg\max_s sD(G,s)$, $sD(G,s) \leq pD(G,p)$ for all $s \in [0,1]$. Using the definitions of $D(G,p)$ and π, this inequality can be rewritten as

$$1 - \frac{\pi}{s} + \Delta(G,s) \leq G(s).$$

Since $\Delta(G,s) \in [0,1]$, the previous inequality implies

$$G_\pi^1(s) \leq G(s), \tag{2}$$

that is, G_π^1 first-order stochastically dominates G. This implies that

$$\int_0^1 sdG_\pi^1(s) \geq \int_0^1 sdG(s) = \mu \geq pD(G,p) = \pi = \int_0^1 sdG_\pi^\pi(s),$$

where the first inequality follows from first-order stochastic dominance, the first equality from $G \in \mathcal{G}_F$ and the last equality from G_π^π specifying an atom of size one at $s = \pi$. Since $\int_0^1 sdG_\pi^B(s)$ is continuous and strictly increasing in B, the Intermediate Value Theorem implies that there exists a unique $B \in [\pi, 1]$ such that

$$\int_0^1 sdG_\pi^B(s) = \mu. \tag{3}$$

We are ready to prove part (i). If $x \leq B$, then (2) and $G_\pi^B(s) = G_\pi^1(s)$ on $[0,B)$ imply that

$$\int_0^x G(s)\,ds \geq \int_0^x G_\pi^B(s)\,ds. \tag{4}$$

If $x \geq B$ then

$$\int_0^x G(s)\,ds = 1-\mu-\int_x^1 G(s)\,ds \geq 1-\mu-(1-x) = 1-\mu-\int_x^1 G_\pi^B(s)\,ds = \int_0^x G_\pi^B(s)\,ds, \tag{5}$$

where the inequality follows from $G(s) \leq 1$ and the second equality from $G_\pi^B(s) = 1$ if $s \geq B$. Equations (3)-(5) imply that G is a mean-preserving spread of G_π^B.

To show part (ii), note that since F is a mean-preserving spread of G and G is a mean-preserving spread of G_π^B, F is also a mean-preserving spread of G_π^B. Therefore, $G_\pi^B \in \mathcal{G}_F$. We have already argued that $[\pi, B] = \arg\max_s sD\left(G_\pi^B, s\right)$, so $\left(G_\pi^B, \pi\right)$ is indeed an outcome.

To prove part (iii), observe that

$$\int_\pi^1 (s-\pi)\,dG_\pi^B(s) = \mu - \pi \geq \int_p^1 (s-p)\,dG(s),$$

where the equality follows from (3) and the observation that the lower bound of the support of G_π^B is π. The inequality follows from the fact that the buyer's payoff cannot exceed the first-best total surplus (μ) minus the seller's profit (π). This inequality is strict whenever the total surplus in the outcome (G,p) is strictly less than μ, that is, $D(G,p)<1$. ■

Applying Lemma 1 to a buyer-optimal outcome (G,p) yields the result that there is also a buyer-optimal signal distribution in the set $\left\{G_\pi^B\right\}_{\pi,B}$. If the distribution of the buyer's signal is G_π^B, it is optimal for the seller to set price π which, in turn, provides the buyer with a payoff of $\mu-\pi$. A buyer-optimal signal distribution is therefore defined by the lowest π for which $G_\pi^B\in\mathcal{G}_F$. Let us define p^* as this lowest value, that is,

$$p^*=\min\left\{\pi:\exists\ B\in[\pi,1]\ \text{ s.t. }\ G_\pi^B\in\mathcal{G}_F\right\}, \tag{6}$$

and let B^* denote the unique value for which $G_{p^*}^{B^*}\in\mathcal{G}_F$.

We are ready to state our main result.

Theorem 1 *The outcome $\left(G_{p^*}^{B^*},p^*\right)$ maximizes the buyer's payoff across all outcomes. If the outcome (G,p) also maximizes the buyer's payoff then*

(i) $p=p^*$,

(ii) $D(G,p)=1$ and

(iii) G is a mean-preserving spread of $G_{p^*}^{B^*}$.

Parts (i) and (ii) of this theorem imply that trade occurs at price p^* with certainty in any buyer-optimal outcome. A consequence of this is that the buyer's payoff is $\mu-p^*$ whenever the signal is buyer-optimal. In general, the buyer's optimal CDF is not determined uniquely because there might be many CDFs which second-order stochastically dominate F and induce the seller to set price p^*. However, according to part (iii), the CDF $G_{p^*}^{B^*}$ second-order stochastically dominates any other buyer-optimal CDF. This means that $G_{p^*}^{B^*}$ provides the buyer with less information than any other buyer-optimal CDF. If $F=G_{p^*}^{B^*}$, the optimal information structure is unique and it specifies perfect learning, that is, $s=v$.

Proof. Let (G,p) be a buyer-optimal outcome and let π denote $pD(G,p)$. By Lemma 1, there exists a $B\in[\pi,1]$ such that $\left(G_\pi^B,\pi\right)$ is an outcome. Then,

$$\int_p^1(s-p)\,dG(s)\le\int_\pi^1(s-\pi)\,dG_\pi^B(s)=\mu-\pi\le\mu-p^*=\int_{p^*}^1(s-p^*)\,dG_{p^*}^{B^*}(s),$$

where the first inequality follows from part (iii) of Lemma 1 and the second inequality follows from (6). Since (G,p) is buyer-optimal, the outcome $\left(G_{p^*}^{B^*},p^*\right)$ is also buyer-optimal and both inequalities are equalities. Part (iii) of Lemma 1 implies that the first inequality is strict unless $D(G,p)=1$, which proves part (ii). The second inequality is strict unless $\pi=p^*$, which proves part (i). Finally, given that $\pi=p^*$, part (i) of Lemma 1 implies part (iii). ■

Next, we solve for an optimal information structure in the case where the buyer's valuation is uniformly distributed, and compare the outcome to that realised in the full-information environment.

Example. Suppose that the buyer's valuation is distributed uniformly on $[0,1]$, that is, $F(v) = v$ and $\mu = 1/2$. The expectation of G_p^B is

$$\int_0^1 s dG_p^B(s) = \int_p^B \frac{p}{s} ds + p = p \log s|_p^B + p = p \log B - p \log p + p.$$

Therefore, in order to guarantee that the expected value generated by G_p^B is one half, it must be that

$$\log B = \frac{\frac{1}{2} + p \log p - p}{p}. \tag{7}$$

The CDF F is a mean-preserving spread of G_p^B if and only if, for all $x \in [p, B]$

$$\int_p^x \left(1 - \frac{p}{s}\right) ds = x - p - \int_p^x \frac{p}{s} ds = x - p - p \log x + p \log p \leq \int_0^x s ds = \frac{x^2}{2}. \tag{8}$$

The smallest p which satisfies these inequalities is approximately 0.2037. That is, $p^* \approx 0.2$ and, by (7), $B^* \approx 0.87$. In other words, by designing the signal structure optimally, the buyer is able to trade at a price just above 0.2. The buyer's payoff is $\mu - p^* \approx 0.3$ and the seller's payoff is $p^* \approx 0.2$.

It is insightful to compare the outcome resulting from a buyer-optimal signal structure to that realized in the full information case, i.e. $s = v$. If the buyer observes her valuation prior to trade then she is willing to trade at price p with probability $1 - p$. Hence, the seller's optimal price solves $\max_p p(1-p)$, so the equilibrium price is 0.5. In turn, the seller's payoff is 0.25 since the buyer trades with probability one half at price 0.5. The seller is therefore better off if the buyer receives a perfectly informative signal than if she receives a buyer-optimal one. The buyer's payoff is $\int_{0.5}^1 v - 0.5 dv = 0.125$ which is less than half of her payoff under the optimal signal structure. The deadweight loss due to the buyer's perfectly informative signal is also 0.125.

In the previous example, the seller's profit induced by the buyer-optimal signal structure is smaller than his profit in the full-information environment. We now generalize this observation and show that a buyer-optimal signal structure is in fact a signal structure that minimizes the seller's profit.

Corollary 1 *For each outcome (G,p), $pD(G,p) \geq p^*$.*

This corollary states that the seller's profit from an outcome is at least p^*. Moreover, by Theorem 1, the seller's profit induced by any buyer-optimal signal structure is p^*. The seller's minimum profit is therefore p^*.

Proof. Let π denote $pD(G,p)$. By Lemma 1, there exists a $B \in [\pi, 1]$ such that $\left(G_\pi^B, \pi\right)$ is an outcome and, in particular, $G_\pi^B \in \mathcal{G}_F$. Hence, by (6), $\pi \geq p^*$. ■

This corollary can be used to characterize the combinations of those consumer and producer surplus which can arise as an equilibrium outcome for *some* signal s. We argue that the set of these payoff profiles is the convex hull of the following three points: $(0, p^*)$, $(0, \mu)$ and $(\mu - p^*, p^*)$

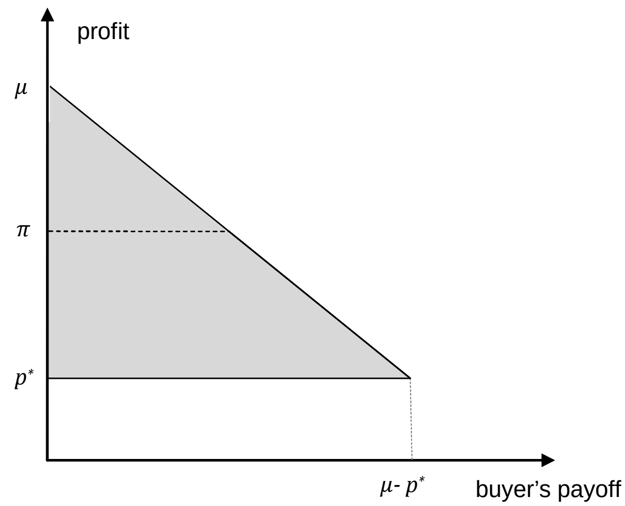

Figure 1: Outcome Triangle

which is represented by the shaded triangle in Figure 1. It is not hard to show that for each $\pi \in [p^*, \mu]$, there exists a $B \in [\pi, 1]$, such that $G_\pi^B \in \mathcal{G}_F$. If the buyer's demand is generated by G_π^B, the seller's profit is π and he is indifferent between any price on $[\pi, B]$. Depending on which of these prices the seller sets, the buyer's surplus can be anything on $[0, \mu - \pi]$, see the horizontal dashed line-segment on Figure 1.[7] Hence, any point in the triangle can be implemented. A payoff profile outside of the triangle cannot be an equilibrium outcome because the seller's profit must be at least p^* (see Corollary 1) and the total surplus cannot exceed μ.

4 Conclusion

The goal of this paper was to analyse a buyer's optimal learning when facing a monopolist. We characterized the buyer-optimal signal distribution which involves minimal learning. The demand induced by this signal is unit-elastic and makes the seller indifferent between setting any price on its support. We also proved that all buyer-optimal signal structures generate the same price and efficient trade.

It is not hard to extend our characterization of the least informative buyer-optimal signal structure to environments where the buyer's valuation can be negative and the seller has a positive valuation for the object.[8] There, the optimal signal structure may provide the buyer with a single value estimate which is smaller than the equilibrium price. Conditional on not getting this signal,

[7] The characterization of Bergemann et al. (2015) is also based on constructing information structures which make the seller indifferent on large sets of prices.

[8] For further details, see http://personal.lse.ac.uk/szentes/docs/extensions.pdf.

the optimal signal structure generates a demand which makes the seller indifferent between any price on its support, just like in our model. The buyer sometimes purchases the object even if her valuation is smaller than that of the seller, so trade is typically inefficient.

References

[1] Dirk Bergemann, Benjamin Brooks, and Stephen Morris, 2015, "The Limits of Price Discrimination," *American Economic Review*, *105*: 921–957.

[2] Dirk Bergemann and Martin Pesendorfer, 2007, "Information Structures in Optimal Auctions," *Journal of Economic Theory*, *137*(1): 580–609.

[3] Olivier Compte and Philippe Jehiel, 2007, "Auctions and Information Acquisition: Sealed-Bid or Dynamic Formats?" *RAND Journal of Economics*, *38*(2): 355–372.

[4] Daniele Condorelli and Balazs Szentes, 2016, "Buyer-optimal Demand and Monopoly Pricing," *Mimeo.*

[5] Juan Jose Ganuza, 2004, "Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations", *RAND Journal of Economics*, *35*(3): 583–598.

[6] Juan Jose Ganuza and Jose S. Penalva, 2010, "Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions," *Econometrica*, *78*(3): 1007–1030.

[7] Anke Kessler, 1998, "The Value of Ignorance," *RAND Journal of Economics*, *29*: 339–354.

[8] Andreu Mas-Colell, Michael D. Whinston, and Jerry Green, 1995, *Microeconomic Theory*, Oxford University Press.

[9] Nicola Persico, 2000, "Information Acquisition in Auctions", Econometrica, *68*: 135–148.

[10] Xianwen Shi, 2012, "Optimal Auctions with Information Acquisition," *Games and Economic Behavior*, *158*: 666–686.