Lemonade from Lemons: Information Design and Adverse Selection*

Navin Kartik[†] Weijie Zhong[‡]

December 2, 2019

Abstract

A seller posts a price for a single object. The seller's and buyer's values may be interdependent. We characterize the set of payoff vectors across all information structures. Simple feasibility and individual-rationality constraints identify the payoff set. The buyer can obtain the entire surplus; often, other mechanisms cannot enlarge the payoff set. We also study payoffs when the buyer is more informed than the seller, and when the buyer is fully informed. All three payoff sets coincide (only) in notable special cases—in particular, when there is complete breakdown in a "lemons market" with an uninformed seller and fully informed buyer.

Incomplete draft. Please do not circulate without permission.

*We thank Laura Doval and Pietro Ortoleva for useful conversations and various audiences for helpful comments.

[†]Department of Economics, Columbia University, E-mail: nkartik@gmail.com

[‡]Cowles Foundation, Yale University, E-mail: wz2269@columbia.edu

1. Introduction

Motivation. Asymmetric information affects market outcomes, both in terms of efficiency and distribution. For example, adverse selection can generate dramatic market failure (Akerlof, 1970) or skew wages in labor markets (Greenwald, 1986), while consumers can secure information rents from a monopolist (Mussa and Rosen, 1978). Much existing work takes the market participants' information as given and studies properties of a particular market structure or mechanism, or tackles these properties across various mechanisms.

This paper instead asks: what is the scope for different market outcomes as the participants' information varies? We are motivated by the fact that in the digital age, the nature of information that sellers (e.g., Amazon) have about consumers is ever changing. Consumers and regulators do have some control over this information, of course. In some cases it is plausible that a seller's information is a subset of the consumer's. But in other cases, the seller may well know *more* about the consumer' value for a product, or at least have some information the consumer herself does not. This is especially relevant for products the consumer is not already familiar with. Indeed, numerous firms make tailored recommendations to consumers about products they carry. With social media and other sources of information diffusion, the possible correlation in information across two sides of a market seems truly limitless.

Our paper fixes a simple, canonical market mechanism and studies the possible market outcomes across a variety of information structures, including *all* of them. We model two parties, Buyer and Seller, who can trade a single object. Buyer's value for the object is a random $v \in [\underline{v}, \overline{v}] \subset \mathbb{R}$. Seller's cost of providing the object, or equivalently, her value from not trading, is $c(v) \leq v$. Thus, values may be interdependent, but trade is always efficient.[1] The environment, i.e., the function $c(v)$ and the distribution of v, is commonly known. Seller posts a price $p \in \mathbb{R}$, and Buyer decides whether to buy.

This stylized setting subsumes a variety of possibilities, depending on the shape of the cost function $c(\cdot)$ and the parties' information about the value v. With an informed Buyer and an uninformed Seller, there is adverse selection when $c(\cdot)$ is increasing while there is favorable or advantageous selection when $c(\cdot)$ is decreasing.[2] If, on the other hand, Seller is better informed than Buyer, signaling becomes relevant; the price can serve as credible signal if the two parties' information is suitably correlated (e.g., Bagwell and Riordan, 1991).

[1] We describe here our baseline model presented in Section 2. Section 4 discusses extensions, including cases when Buyer's value does not pin down Seller's cost and when trade is not always efficient.

[2] Jovanovic (1982) uses the term 'favorable selection'. Einav and Finkelstein (2011) use 'advantageous', and discuss both adverse and advantageous selection in the context of insurance markets, with empirical references to empirical evidence on both.

A constant $c(\cdot)$ captures an environment in which there no uncertainty about Seller's cost; this is the canonical monopoly pricing problem when Seller is uninformed about v, and third-degree price discrimination when Seller has some partial information while Buyer is better informed.

Summary of results. For any given environment (i.e., Seller's cost function and the distribution of Buyer's values), we seek to identify the possible market outcomes. Specifically, we are interested in the ex-ante expected payoffs that obtain, given sequentially rational behavior, in an equilibrium under *some* information structure.[3] We provide three results, each of which covers a different class of information structures. Our main theorems are Theorem 1/Theorem 1*, which impose no restrictions on information, and Theorem 2, which applies when Buyer is better informed than Seller in the sense of Blackwell (1953); in fact, Theorem 2 applies more broadly, as elaborated later. Theorem 3 concerns a fully-informed Buyer (i.e., who knows his value v). We view each of these three cases as intellectually salient and economically relevant. Plainly, these payoff sets must be ordered by set inclusion: Theorem 1's is the largest; Theorem 2's is intermediate; and Theorem 3's the smallest. Figure 1 below summarizes.

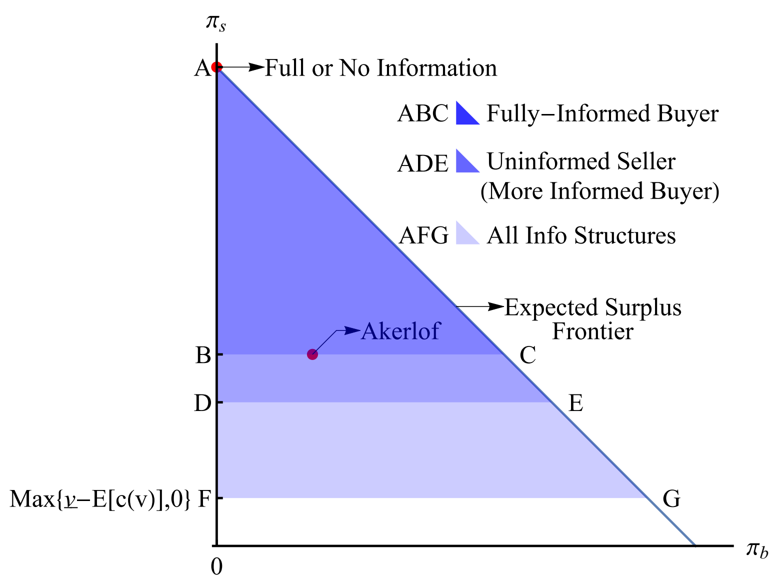

Figure 1: Outcome under different restrictions on information structures

In the figure's axes, π_b and π_s represent respectively Buyer's and Seller's ex-ante ex-

[3] As detailed in Section 2, an information structure specifies a joint distribution of private signals for each party conditional on the value v. This induces an extensive-form game of incomplete information. Our primary solution concept is weak Perfect Bayesian equilibrium; we also address refinements for our constructive arguments.

pected utilities or payoffs (for readability, we often drop the "expected" qualifier). The no-trade payoffs are normalized to zero. The three triangles, AFG, ADE, and ABC, depict Theorems 1–3 respectively. That payoffs must lie within AFG is straightforward: Buyer can guarantee himself a payoff of zero by not purchasing; Seller can guarantee herself not only a payoff of zero (by posting any price $p > \overline{v}$, which will not be accepted) but also $\underline{v} - \mathbb{E}[c(v)]$ (by pricing at or just below $\underline{v}$, which will be accepted); and the sum of payoffs cannot exceed the trading surplus $\mathbb{E}[v - c(v)]$. We refer to the first two constraints as *individual rationality* and the third as *feasibility*.

Theorem 1 says that every feasible and individually-rational payoff pair can be implemented, i.e., obtains in an equilibrium under some information structure. It is immediate that point A obtains whens both parties learn v (full information) or neither party has any information (no information). More interestingly, at the point G trade occurs with probability one and Buyer obtains the entire surplus, despite Seller posting the price. While perhaps surprising, this outcome obtains with sparse information structures. For simplicitly, suppose $\mathbb{E}[c(v)|v > \underline{v}] \leq \max\{\underline{v}, \mathbb{E}[c(v)]\}$. Then Buyer can be uninformed while Seller learns whether $v = \underline{v}$ or $v > \underline{v}$. In equilibrium, Seller prices at $p = \max\{\underline{v}, \mathbb{E}[c(v)]\}$ regardless of her signal and Buyer purchases. If Seller were to deviate to a higher price, Buyer would reject because he believes $v = \underline{v}$. Subsection 3.1 explains how a single information structure in fact implements every point in the triangle AFG. Theorem 1* there discusses how a richer information structure using imperfectly-correlated signals ensures implementation in Kreps and Wilson's (1982) sequential equilibrium in discretized versions of the model.

Turning to Figure 1's triangle ADE, Theorem 2 establishes that the payoff pair in any equilibrium when Buyer is better informed than Seller arises in an equilibrium of an(other) information structure in which Seller is uninformed.[4] In other words, there is no loss of generality in studying an uninformed Seller so long as Buyer is better informed. When $c(\cdot)$ is increasing, such information generates a game with adverse selection; when $c(\cdot)$ is decreasing there is favorable selection. Seller's payoff along the line segment DE is the lowest payoff she can get in any information structure in which she is uninformed.[5] Theorem 2 further establishes that any point within the ADE triangle can be implemented with some such information structure by suitably varying Buyer's information. In fact, we show that higher slices of the triangle (i.e., those corresponding to larger Seller's payoff)

[4] We stipulate that a better-informed Buyer does not update his value from Seller's price (even off the equilibrium path), in line with the "no signaling what you don't know" requirement (Fudenberg and Tirole, 1991) that is standard in versions of Perfect Bayesian Equilibrium and implied by sequential equilibrium.

[5] It is because Seller cannot commit to the price as a function of her signal that she can be harmed (i.e., receive a payoff lower than that on the DE segment) with more information. However, Theorem 2 assures that Seller is not harmed so long as Buyer is better informed.

can always be implemented by reducing Buyer's information in the sense of Blackwell (1953). We also explain in Subsection 3.2 why the triangle ADE actually characterizes all payoffs that can obtain when Buyer does not update from Seller's price, even if Buyer is not better informed than Seller.

Finally, the ABC triangle in Figure 1 depicts Theorem 3, which characterizes all payoff pairs when Buyer is fully informed, i.e., learns v. We use the term "Akerlof" to describe a fully-informed Buyer and an uninformed Seller, as these information structures are standard in the adverse-selection literature; the corresponding payoff pair is marked as such in the figure. Depending on the environment's primitives, the Akerlof point can be anywhere on the segment BC, including at the extreme points. Any feasible payoff pair that satisfies Buyer's individual rationality and gives Seller at least her Akerlof payoff can be implemented with a fully-informed Buyer by suitably varying Seller's information.

An implication of Theorems 1–3 is that it is without loss, in terms of ex-ante equilibrium payoffs, to focus on information structures in which Buyer is fully informed if and only if Seller's Akerlof payoff coincides with her individual rationality constraint. This coincidence occurs if and only if the Akerlof market can have full trade (Seller prices at $p = \underline{v}$ and gets payoff $\underline{v} - \mathbb{E}[c(v)] \geq 0$) or no trade (the price is $p \geq \overline{v}$ and both parties' payoffs are 0). As detailed in Remark 4 of Subsection 3.3, in all other cases the point B in Figure 1 is distinct from the point D (and hence also F), which means that Buyer can obtain a higher payoff with less-than-full information, while keeping Seller uninformed. Furthermore, under a reasonable condition, if Seller's individual rationality constraint is zero (i.e., $\underline{v} \leq \mathbb{E}[c(v)]$), then point D is also distinct from F; see Remark 3 in Subsection 3.2. When D and F are distinct, maximizing Buyer's payoff, i.e., achieving point G, requires Seller to have some information Buyer does not and an equilibrium with price-dependent beliefs: after conditioning on his signal, Buyer must update about v from the price either on or off the equilibrium path.

Related literature. Our perspective and results are most closely related to Bergemann, Brooks, and Morris (2015), Roesler and Szentes (2017), and Makris and Renou (2018, Section 4). These papers—only the relevant section of the third paper—study the monopoly pricing problem in which there is uncertainty only about Buyer's valuation. This is the special case of our interdependent-values model with a constant function $c(v) \leq \underline{v}$.[6]

[6] The online appendix of Roesler and Szentes (2017) relaxes the assumption that there are always gains from trade, but maintains no uncertainty about Seller's cost. Related to Roesler and Szentes (2017) are also Du (2018) and Libgober and Mu (2019), who consider worst-case profit guarantees for Seller in static and dynamic environments respectively. Terstiege and Wasser (2019) qualify Roesler and Szentes (2017) by allowing

Bergemann et al. (2015) assume Buyer is fully informed, and hence can only vary Seller's information. Our Theorem 3, corresponding to triangle ABC in Figure 1, is a generalization of their main result to our environment; the key step in our proof methodology is to construct "incentive compatible distributions", which reduces to their "extreme markets" in the monopoly-pricing environment. Another contribution of our result is to ensure approximately unique implementation, in a sense explained in Subsection 3.3.

Roesler and Szentes (2017) assume Seller is uninformed and only vary Buyer's information. For the monopoly-pricing environment, they derive one part of our Theorem 2, viz., they identify the triangle ADE in Figure 1 as the implementable set given Seller is uninformed. Even for this result, our methodology is quite different from theirs because we do not stipulate a linear $c(\cdot)$ function; our methodology delivers new insights, including that noted in Remark 2 of Subsection 3.2. When we specialize to a linear $c(\cdot)$, we can obtain a sharper characterization of the point E, which extends Roesler and Szentes' characterization of Buyer-optimal information to an interdependent-values environment; see Proposition 2 in Subsection 4.3.

While our main interest is in interdependent values, our results provide new insights even for monopoly pricing. Theorem 2 implies that the Roesler and Szentes (2017) bounds are without loss so long as Buyer is better informed than Seller; or, more generally, in equilibria in which Buyer's belief is price independent (after conditioning on his own signal). On the other hand, Theorem 1 establishes that any feasible and individually rational payoff pair can be implemented absent these restrictions: in particular, Buyer may even get all the surplus. This latter point has a parallel with Makris and Renou (2018). As an application of their general results on "revelation principles" for information design in multi-stage games, Makris and Renou's (2018) Proposition 1 deduces an analog of our Theorem 1 for the (independent values) monopoly pricing problem. We share with Makris and Renou an emphasis on sequential rationality;[7] we differ in showing that a single information structure implements all payoffs in the relevant triangle, and also in establishing in Theorem 1* off-the-equilibrium-path belief consistency in the sense of sequential equilibrium (Kreps and Wilson, 1982) in discretizations.

Other authors have studied different aspects of more specific changes of information in adverse-selection settings, maintaining that one side of the market is better informed than

Seller to supply Buyer with additional information, although Seller cannot have any private information of her own.

[7] Makris and Renou use an apparatus of "sequential Bayes correlated equilibrium", which we do not. They draw a contrast with Bergemann and Morris's (2016) Bayes correlated equilibrium. In our approach, note that Seller's individual rationality constraint described earlier hinges, when $\mathbb{E}[c(v)] < \underline{v}$, on Buyer's behavior being sequentially rational even off the equilibrium path.

the other. Levin (2001) identifies conditions under which the volume of trade decreases when one party is kept uninformed and the other's information become more effective in the sense of Lehmann (1988); see also Kessler (2001). Assuming a linear payoff structure, Bar-Isaac, Jewitt, and Leaver (2018) consider how certain changes in Gaussian information affect the volume of trade, surplus, and a certain quantification of adverse selection.

Finally, Garcia, Teper, and Tsur (2018) solve for socially optimal information provision—rather than characterizing the range of market outcomes—in an insurance setting with adverse selection; owing to a cross-subsidization motive, full information disclosure is typically not optimal.

The rest of the paper proceeds as follows. We introduce our model, equilibrium concept(s), and certain classes of information structures in Section 2. Section 3 presents the main results: implementable payoffs when the information structure is arbitrary or varies within canonical classes. Section 4 contains discussion and extensions. Proofs and some additional material are contained in the Appendices.

2. Model

2.1. Primitives

Seller may sell an indivisible good to Buyer. Buyer's value for the good is $v \in V \subset \mathbb{R}$, where V is a compact (finite or infinite) set with $\underline{v} \equiv \min V < \max V \equiv \bar{v}$. The value v is drawn from a probability measure μ with support V. Seller's cost of production is given by a function $c(v)$.[8] We assume $c : V \to \mathbb{R}$ is continuous, $v - c(v) \geq 0$ for all v, and $\mathbb{E}[v - c(v)] > 0$. So trading surplus is nonnegative for all v and positive for a positive measure of v. (Throughout, expectations are with respect to the prior measure μ unless indicated otherwise; 'positive' means 'strictly positive' and similarly elsewhere.) Note that function $c(v)$ need not be monotonic. We call $\Gamma = (c, \mu)$ an *environment*. We refer to an environment with a constant $c(\cdot)$ function as that of *monopoly pricing*.

An information structure consists of signal spaces for each party and a joint signal distribution. (We abuse notation and sometimes write 'distribution' even though 'measure' would be more precise.) Formally, there is a probability space $(\Omega, \mathcal{F}, P)$, complete and separable metric spaces T_s and T_b (with measurable structure given by their Borel sigma algebras), and an integrable function $X : \Omega \to T_s \times T_b \times V$. We hereafter suppress the probability space and define, with an abuse of notation, $P(D) = P(X^{-1}(D))$ for any $D \in T_s \times T_b \times V$ measurable. Each realization of random variable X is a triplet (t_b, t_s, v), where $t_b \in T_b$

[8] Subsection 4.1 discusses Seller's cost being stochastic even conditional on v.

is Buyer's signal and $t_s \in T_s$ is Seller's signal. For $i \in \{s, b, v\}$, let P_i denote the corresponding marginal distribution of P on dimension of T_i with the convention $T_v \equiv V$. We require $P_v = \mu$; this is the iterated expectation or Bayes plausibility requirement. Denote an information structure by τ.

The environment Γ and information structure τ define the following *game*:

1. The random variables (t_b, t_s, v) are realized. Signal t_b is privately observed by Buyer and signal t_s privately observed by Seller. Neither party observes v.

2. Seller posts a price $p \in \mathbb{R}$.

3. Buyer accepts or rejects the price. If Buyer accepts, his von-Neumann Morgenstern payoff is $v - p$ and Seller's is $p - c(v)$. If Buyer rejects, both parties' payoffs are normalized to 0.

Note that because the signal spaces T_b and T_s are abstract and the two parties' signals can be arbitrarily correlated conditional on v, there is no loss of generality in assuming that each party privately observes their own signal. For example, public information can be captured by perfectly correlating (components of) t_b and t_s.

2.2. Strategies and Equilibria

In the game defined by (Γ, τ), denote Seller's strategy by σ and Buyer's by α. Following Milgrom and Weber (1985), we define σ as a distributional strategy: σ is a joint distribution on $\mathbb{R} \times T_s$ whose marginal distribution on T_s must be the Seller's signal distribution. So $\sigma(\cdot|t_s)$ is Seller's price distribution given her signal t_s.[9] Buyer's strategy $\alpha : \mathbb{R} \times T_b \to [0, 1]$ maps each price-signal pair (p, t_b) into a trading probability. A strategy profile (σ, α) induces expected utilities for Buyer and Seller (π_b, π_s) in the natural way:

$$\pi_b = \int (v - p)\alpha(t_b, p)\sigma(\mathrm{d}p|t_s)P(\mathrm{d}t_s, \mathrm{d}t_b, \mathrm{d}v),$$
$$\pi_s = \int (p - c(v))\alpha(t_b, p)\sigma(\mathrm{d}p|t_s)P(\mathrm{d}t_s, \mathrm{d}t_b, \mathrm{d}v).$$

Our baseline equilibrium concept is weak Perfect Bayesian equilibrium. Since Seller's action is not preceded by Buyer's we can dispense with specifying beliefs for Seller. For Buyer, it suffices to focus on his belief about the value v given his signal and the price; we denote this distribution by $\nu(v|p, t_b)$.

[9] Here $\sigma(\cdot|t_s)$ is the regular conditional distribution, which exists and is unique almost everywhere because T_s is a standard Borel space (Durrett, 1995, pp. 229–230). Similarly for subsequent such notation; we drop "almost everywhere" qualifiers unless essential.

Definition 1. A strategy profile (σ, α) and a belief $\nu(v|p, t_b)$ is a *weak perfect Bayesian equilibrium* (wPBE) of game (Γ, τ) if:

1. Buyer plays optimally at every information set given his belief:

$$\alpha(p, t_b) = \begin{cases} 1 & \text{if } \mathbb{E}_{\nu(v|p,t_b)}[v] > p \\ 0 & \text{if } \mathbb{E}_{\nu(v|p,t_b)}[v] < p; \end{cases}$$

2. Seller plays optimally:

$$\sigma \in \arg\max_{\hat{\sigma}} \int (p - c(v))\alpha(p, t_b)\hat{\sigma}(\mathrm{d}p|t_s)P(\mathrm{d}t_s, \mathrm{d}t_b, \mathrm{d}v);$$

3. Beliefs satisfy Bayes rule on path: for every measurable $D \subseteq \mathbb{R} \times T_s \times T_b \times V$,

$$\int_D \nu(\mathrm{d}v|p, t_b)\sigma(\mathrm{d}p|t_s)P(\mathrm{d}t_s, \mathrm{d}t_b, V) = \int_D \sigma(\mathrm{d}p|t_s)P(\mathrm{d}t_s, \mathrm{d}t_b, \mathrm{d}v).$$

Notice that we have formulated Seller's optimality requirement ex ante, but Buyer's at each information set. The latter is needed to capture sequential rationality. The former is for (notational) convenience; this choice is inconsequential because Seller moves before Buyer.

Hereafter, "equilibrium" without qualification refers to a wPBE. As is well understood, wPBE permits significant latitude in beliefs off the equilibrium path. We will discuss refinements when appropriate.

2.3. Implementable Payoffs and Canonical Information Structures

We now define the set of *implementable* equilibrium outcomes—that is, the payoffs that obtain in some equilibrium under some information structure—and some canonical classes of information structures.

For a game (Γ, τ), let the equilibrium payoff set be

$$\Pi(\Gamma, \tau) \equiv \{(\pi_b, \pi_s) : \exists \text{ wPBE of } (\Gamma, \tau) \text{ with payoffs } (\pi_b, \pi_s)\}.$$

Denote the class of all information structures by $\mathbf{T}$ and define

$$\mathbf{\Pi}(\Gamma) \equiv \bigcup_{\tau \in \mathbf{T}} \Pi(\Gamma, \tau).$$

That is, for environment Γ, $\mathbf{\Pi}(\Gamma)$ is the set of all equilibrium payoff pairs that obtain under some information structure.

Uninformed Seller. An information structure has *uninformed Seller* if T_s is a singleton: Seller's own signal contains no information about Buyer's value v, and hence neither about Seller's cost $c(v)$. When discussing such information structures, we write the associated distribution as just $P(t_b, v)$ and Seller's strategy as just $\sigma(p)$, omitting the argument t_s in both cases. The class of all uninformed-Seller information structure is denoted $\mathbf{T}_{us}$.

Fully-informed Buyer. An information structure has *fully-informed Buyer* if Buyer's signal fully reveals his value v. Formally, this holds if $T_b = V$ and the conditional distribution on V, $P(\cdot|t_b)$, satisfies $P(\{t_b\}|t_b) = 1$. We denote the class of fully-informed-Buyer information structures $\mathbf{T}_{fb}$.

More-informed Buyer. An information structure has *more-informed Buyer* if Buyer has more information than Seller. Formally, this holds when v and t_s are independent conditional on t_b, i.e., for any measurable $D_s \subset T_s$ and $D_v \subset V$, $P(D_s \cap D_v|t_b) = P(D_s|t_b)P(D_v|t_b)$.[10] Another way to interpret this requirement is that random variable t_b must be statistically sufficient for t_s with respect to v, i.e., t_b is more informative than t_s about v in the sense of Blackwell (1953). We denote the class of more-informed-Buyer information structures by $\mathbf{T}_{mb}$. Naturally, information structures with uninformed Seller or fully-informed Buyer are cases of more-informed Buyer: both $\mathbf{T}_{us}$ and $\mathbf{T}_{fb}$ are subclasses of $\mathbf{T}_{mb}$.

No updating from price. For more-informed-Buyer information structures, it is desirable to impose further requirements on Buyer's equilibrium belief. Since Seller's price can only depend on her own signal, and this signal contains no additional information about v given Buyer's signal, the price is statistically uninformative about v given Buyer's signal. Consequently, Buyer's posterior belief should be price independent once his signal has been conditioned upon. Formally, regardless of the price p, the equilibrium belief $\nu(\cdot|p, t_b)$ must satisfy

$$\int \nu(\mathrm{d}v|p, t_b) P(\mathrm{d}t_b, T_s, V) = \int P(\mathrm{d}t_b, T_s, \mathrm{d}v).$$

We refer to this condition as *price-independent belief*. Note that the condition is meaningful regardless of whether Buyer is more informed than Seller. In a more-informed-Buyer information structure, price-independent belief would be implied by the "no signaling what

[10] We write $\subset$ to mean "weak subset".

you don't know" requirement (Fudenberg and Tirole, 1991) frequently imposed in versions of perfect Bayesian equilibrium, and the concept of sequential equilibrium (Kreps and Wilson, 1982) in finite versions of our setting.[11]

Some notation will be helpful. Define

$$\Pi^*(\Gamma,\tau) \equiv \{(\pi_b,\pi_s) : \exists \text{ wPBE of } (\Gamma,\tau) \text{ with price-independent belief and payoffs} (\pi_b,\pi_s)\},$$

$$\mathbf{\Pi}^*(\Gamma) \equiv \bigcup_{\tau\in\mathbf{T}} \Pi^*(\Gamma,\tau),$$

$$\mathbf{\Pi}^*_i(\Gamma) \equiv \bigcup_{\tau\in\mathbf{T}_i} \Pi^*(\Gamma,\tau) \text{ for } i = us, fb, mb.$$

So Π^* and $\mathbf{\Pi}^*$ are analogous to the implementable payoff sets Π and $\mathbf{\Pi}$ defined earlier, but restricted to equilibria with price-independent beliefs. $\mathbf{\Pi}^*_{us}$, $\mathbf{\Pi}^*_{fb}$ and $\mathbf{\Pi}^*_{mb}$ are the implementable payoff sets when further restricted to uninformed-Seller, fully-informed-Buyer, and more-informed-Buyer information structures. Plainly, for any environment Γ,

$$\mathbf{\Pi}^*_{us}(\Gamma) \cup \mathbf{\Pi}^*_{fb}(\Gamma) \subset \mathbf{\Pi}^*_{mb}(\Gamma) \subset \mathbf{\Pi}^*(\Gamma) \subset \mathbf{\Pi}(\Gamma).$$

3. Main Results

Our goal is to characterize equilibrium payoff pairs across information structures in an arbitrary environment Γ. In particular, we seek to characterize the five sets $\mathbf{\Pi}(\Gamma)$, $\mathbf{\Pi}^*(\Gamma)$, $\mathbf{\Pi}^*_{mb}(\Gamma)$, $\mathbf{\Pi}^*_{us}(\Gamma)$, and $\mathbf{\Pi}^*_{fb}(\Gamma)$. Let

$$S(\Gamma) \equiv \mathbb{E}[v - c(v)]$$

be the (expected) surplus from trade in environment Γ. This quantity will play an important role.

3.1. All Information Structures

Define Seller's payoff guarantee as

$$\underline{\pi}_s(\Gamma) \equiv \max\{\underline{v} - \mathbb{E}[c(v)], 0\}.$$

[11] An example clarifying our terminology may be helpful. If both Seller and Buyer are fully informed of v, then the natural equilibrium—the unique sequential equilibrium in a finite version of the game—has Seller pricing at $p = v$ and Buyer's belief being degenerate on v regardless of Seller's price. This equilibrium has price-independent belief, even though ex-ante Seller's price are Buyer's belief are perfectly correlated. The point is that Buyer's belief does not depend on price conditional on his signal.

To interpret this, observe that it is optimal for Buyer to accept the price $\underline{v}$ no matter his belief. Therefore, Seller can guarantee herself the (expected) profit $\underline{v} - \mathbb{E}[c(v)]$ no matter what the information structure is. More precisely, she can guarantee $v - \mathbb{E}[c(v)] - \varepsilon$ for any $\varepsilon > 0$, since sequential rationality requires Buyer to accept any price $\underline{v} - \varepsilon$. Similarly, Seller can also guarantee zero profit offering a price $p > \bar{v}$. Hence Seller's payoff in any equilibrium of any information structure must be at least the payoff guarantee $\underline{\pi}_s(\Gamma)$ above.

On the other hand, Buyer can guarantee himself the payoff $\pi_b = 0$ by rejecting all prices. It follows that the implementable set $\mathbf{\Pi}(\Gamma)$ must satisfy three simple constraints: 1) Seller's "individual rationality" constraint $\pi_s \geq \underline{\pi}_s(\Gamma)$; 2) Buyer's "individual rationality" constraint $\pi_b \geq 0$; and 3) the feasibility constraint $\pi_b + \pi_s \leq S(\Gamma)$.

Our first result is that these individual rationality and feasibility constraints are also sufficient for a payoff pair to be implementable.

Theorem 1. *Consider all information structures and equilibria.*

$$\mathbf{\Pi}(\Gamma) = \left\{ (\pi_b, \pi_s) : \begin{array}{c} \pi_b \geq 0 \\ \pi_s \geq \underline{\pi}_s(\Gamma) \\ \pi_b + \pi_s \leq S(\Gamma) \end{array} \right\}.$$

Theorem 1 says that the set $\mathbf{\Pi}(\Gamma)$ corresponds to the triangle AFG in Figure 1. In particular, Buyer can receive the entire surplus beyond Seller's payoff guarantee. This is surprising, as Seller has substantial bargaining power. Note that when $\underline{v} \leq \mathbb{E}[c(v)]$, an entirely reasonable condition, Seller's payoff guarantee is zero.[12]

The proof of Theorem 1 is in fact straightforward. Suppose, for expositional simplicity, $\underline{v} \geq \mathbb{E}[c(v)]$. Fix the trivial information structure in which neither player receives any information and consider the following family of strategy profiles. Seller randomizes between two prices, some $p_l \in [\underline{v}, \mathbb{E}[v]]$ and $p_h = \mathbb{E}[v]$. Buyer accepts p_l with probability one and accepts p_h with probability $\alpha(p_h)$, where $\alpha(p_h) \in [0, 1]$ is specified to make Seller indifferent between the two prices. That is, $\alpha(p_h)(p_h - \mathbb{E}[c(v)]) = p_l - \mathbb{E}[c(v)]$. The expected payoffs from this strategy profile are

$$\pi_b = \sigma(p_l)(\mathbb{E}[v] - p_l) \text{ and } \pi_s = p_l - \mathbb{E}[c(v)].$$

As p_l traverses the interval $[\underline{v}, \mathbb{E}[v]]$, Seller's payoff π_s traverses $[\underline{\pi}_s(\Gamma), S(\Gamma)]$. Given any

[12] In monopoly pricing with $c(\cdot) = \underline{v}$, Seller's payoff guarantee of zero is lower than the revenue guarantee identified in Du (2018, Section 5), which is typically positive. Du's notion is different from ours.

p_l, Buyer's payoff π_b traverses $[0, S(\Gamma) - \pi_s]$ as $\sigma(p_l)$ traverses $[0, 1]$. Therefore, the proposed strategy profiles induce all the payoff pairs stated in Theorem 1.

We are left to specify beliefs ν for Buyer. After prices p_l and p_h Buyer holds the prior belief μ. After any other (necessarily off-path) price Buyer's belief is that $v = \underline{v}$, and so Buyer rejects all prices $p \in [\underline{v}, \infty) \setminus \{p_l, p_h\}$. It is straightforward to confirm that the specified (σ, α, ν) constitute a wPBE.

To get more insight into this construction, consider its implication for the monopoly pricing with $c(\cdot) = \underline{v}$. The equilibrium with $p_l = \underline{v}$ (hence $\alpha(p_h) = 0$) and $\sigma(p_l) = 1$ corresponds to the monopolist deterministically pricing at $\underline{v}$ and Buyer purchasing. Given that both sides of the market receive no information, why doesn't the monopolist deviate to any price in $(\underline{v}, \mathbb{E}[v])$? The reason is that, in this equilibrium, the consumer will then not buy because he updates to $v = \underline{v}$. Such belief is compatible with wPBE because the equilibrium concept places no restrictions on off-path beliefs. This may seem like a game-theoretic misdirection: Buyer's belief is not consistent with "no signaling what you don't know". Put differently, since we have a (weakly) more-informed Buyer information structure, we ought to impose the price-independent belief condition described in Subsection 2.3; that would imply Buyer must purchase at any price $p < \mathbb{E}[v]$.

But the message of Theorem 1 does not rely on the permissiveness of wPBE. To illustrate, continue with the above monopoly-pricing environment, and suppose $\underline{v}$ has positive prior probability. Consider Buyer remaining uninformed but Seller learning whether $v = \underline{v}$ or $v > \underline{v}$. Now Buyer's off-path belief that $v = \underline{v}$ is consistent with "no signaling what you don't know". More generally, using richer information structures, we can prove that any payoff pair identified in Theorem 1 can be approximately implemented as a *sequential equilibrium* (Kreps and Wilson, 1982) in a suitably discretized game.

Theorem 1*. *Fix any $\varepsilon > 0$. There is a finite information structure and a finite price grid inducing a game with a set of sequential equilibrium payoffs that is an ε-net of $\mathbf{\Pi}(\Gamma)$.*[13]

In fact, the proof of Theorem 1* establishes even more: the sequential equilibria in the discretized games satisfy a natural version of the D1 refinement (Cho and Kreps, 1987). We relegate the logic to the Appendix, but mention here that we use imperfectly-correlated signals for Buyer and Seller.[14]

[13] An information structure is finite if the signal spaces T_b and T_s are finite. Sequential equilibrium is defined in the obvious way for the "induced" finite game where Nature directly draws (t_b, t_s), rather than first drawing v, and players' payoffs are defined directly using (t_b, t_s). For $X \subset \mathbb{R}^2$ and $\varepsilon > 0$, the set $A \subset X$ is an ε-net of X if for each $x \in X$ there is $a \in A$ such that $\|x - a\| < \varepsilon$, where $\|\cdot\|$ is the Euclidean distance.

[14] The idea behind D1 is to ask, for any off-path price, whether one type of Seller would deviate for any

Even when our environment is specialized to monopoly pricing, it is worth highlighting two contrasts between Theorem 1/1* and results of Bergemann et al. (2015) and Roesler and Szentes (2017). First, we find that by not restricting the monopolist to be uninformed, the implementable payoff set typically expands rather dramatically: trade can be efficient with the monopolist securing none of the surplus beyond her payoff guarantee ($\underline{\pi}_s$), which may be zero. (Roesler and Szentes establish, implicitly, that the implementable set with an uninformed monopolist is a superset of Bergemann et al.'s which has a fully-informed consumer.) We will see in Subsection 3.2 that what is crucial to this expansion is price-*dependent* beliefs. In particular, the proof of Theorem 1* uses an information structure in which Buyer is not better informed than Seller — if he were, then sequential equilibrium would imply price-independent beliefs. Second, Theorem 1* establishes that for a given $\varepsilon > 0$, a single information structure (and price grid) can be used to approximate the entire payoff set $\mathbf{\Pi}(\Gamma)$, analogously to the construction described after Theorem 1 that used a single information structure.[15] Bergemann et al. (2015) and Roesler and Szentes (2017), on the other hand, vary information structures to span their payoff sets.

3.2. More-informed Buyer and Price-independent Beliefs

In some economic settings it is plausible that Buyer is more informed than Seller. How does a restriction to such information structures, i.e., $\tau \in \mathbf{T}_{mb}$, affect the implementable payoff set? It turns out that what is in fact crucial is price-independent beliefs. We have explained earlier why it is desirable to impose this condition when Buyer is more informed, but that the condition is well defined even otherwise. If Buyer is not more informed than Seller, price-independent beliefs ought to be viewed as an equilibrium restriction in general. To reduce repetition, readers should be bear in mind for the rest of this subsection, the qualifier "with price-independent beliefs" applies unless stated explicitly otherwise.

It is useful to define

$$\underline{\pi}_s^{us}(\Gamma) \equiv \inf \{\pi_s : \exists(\pi_b, \pi_s) \in \mathbf{\Pi}^*_{us}(\Gamma)\} \tag{1}$$

as the lowest payoff that an *uninformed* Seller can obtain, no matter Buyer's information (among equilibria with price-independent beliefs). Plainly, $\underline{\pi}_s^{us}(\Gamma) \geq \underline{\pi}_s(\Gamma)$. In monopoly

Buyer mixed response that another type would. Our construction has multiple Buyer types that are imperfectly correlated with Seller types. So different types of Seller have different beliefs about Buyer types. This blunts the power of dominance considerations, to the point where D1 does not exclude any Seller type from the support of Buyer's off-path belief.

[15] In fact, if one lets the price grid vary with ε, then a single information structure implements exactly, rather than approximately, in sequential equilibrium all payoffs ε-away from the boundary of $\mathbf{\Pi}(\Gamma)$. See **??** in the Appendix for a formal statement.

pricing with $V = [\underline{v}, \overline{v}]$ and $c(\cdot) = \underline{v}$, Roesler and Szentes's (2017) characterization of the consumer-optimal information structure identifies $\underline{\pi}_s^{us}$, establishing that $\underline{\pi}_s^{us} > \underline{\pi}_s$. If there is no trade due to adverse selection when Seller is uniformed and Buyer has some information, then $\underline{\pi}_s^{us} = \underline{\pi}_s = 0$. We do not have a general explicit formula for $\underline{\pi}_s^{us}$; Subsection 4.3 provides it for linear $c(\cdot)$. Nonetheless, we establish next that (i) the only additional restriction on equilibrium payoffs imposed by price-independent beliefs is a lower bound of $\underline{\pi}_s^{us}$ for Seller, and (ii) uninformed-Seller information structures implement all such payoffs.

Theorem 2. *Consider equilibria with price-independent beliefs.*

1. $\mathbf{\Pi}^*(\Gamma) = \mathbf{\Pi}^*_{mb}(\Gamma) = \mathbf{\Pi}^*_{us}(\Gamma)$.
2. $\mathbf{\Pi}^*_{us}(\Gamma) = \{(\pi_b, \pi_s) \in \mathbf{\Pi}(\Gamma) : \pi_s \geq \underline{\pi}_s^{us}(\Gamma)\}$.
3. *For any* $(\pi_b, \pi_s) \in \mathbf{\Pi}^*_{us}(\Gamma)$ *with* $\pi_s > \underline{\pi}_s^{us}(\Gamma)$*, there is* $\tau \in \mathbf{T}_{us}$ *with* $\Pi(\Gamma, \tau) = \{(\pi_b, \pi_s)\}$.

Remark 1. We believe the substance of Theorem 2 would hold using discretizations and sequential equilibria, analogous to Theorem 1*. As previously noted, sequential equilibrium implies price-independent beliefs when Buyer is more informed than Seller.

To digest Theorem 2, note that $\mathbf{\Pi}^*(\Gamma) \supset \mathbf{\Pi}^*_{mb}(\Gamma) \supset \mathbf{\Pi}^*_{us}(\Gamma)$ is trivial. So part 1 of the theorem amounts to establishing the reverse inclusions. The intuition for this—given part 2's characterization of $\mathbf{\Pi}^*_{us}$—is fairly straightforward: with price-independent beliefs, additional information cannot harm Seller, even though it could alter the set of equilibria. So Seller's lowest payoff obtains when she is uninformed.

The characterization in part 2 of payoffs with an uninformed Seller corresponds to the triangle ADE in Figure 1. Part 3 of the theorem assures "unique implementation" of all implementable payoffs satisfying $\pi_s > \underline{\pi}_s^{us}(\Gamma)$. That is, for any such payoff pair, there is an uninformed-Seller information structure such that all equilibria (with price-independent beliefs) induce exactly that payoff pair. Unique implementation is appealing for multiple reasons, one of which is that it obviates concerns about which among multiple payoff-distinct equilibria is more reasonable. For example, Ravid, Roesler, and Szentes (2019) bring to the fore the issue of equilibrium multiplicity in Roesler and Szentes's (2017) information structure.

Let us describe how we obtain the characterization of $\mathbf{\Pi}^*_{us}(\Gamma)$ and unique implementation. There are two steps. The first ensures that there is some information structure, call it $\tau^* \in \mathbf{T}_{us}$, that implements Seller's payoff $\underline{\pi}_s^{us}(\Gamma)$. That is, we ensure that the infimum in (1) is in fact a minimum.[16] While this argument is technical, knowing τ^* exists is useful in what follows. The second, and economically insightful, step is to construct information structures that implement every point in the triangle $\mathbf{\Pi}^*_{us}(\Gamma)$ by suitably garbling the information structure τ^*. The construction is illustrated in Figure 2. Consider the distribution of Buyer's posterior mean of his valuation v in information structure τ^*. (Given price-independent beliefs, Buyer's posterior mean is sufficient for his decision.) For simplicity, suppose this posterior-mean distribution has a density, as depicted by the red curve in Figure 2. Fix any $(\pi_b, \pi_s) \in \mathbf{\Pi}^*_{us}(\Gamma)$.

First, there is some number z^* such that $\pi_b + \pi_s$ is the total surplus from trading only when Buyer's posterior mean is greater than z^*. Next, there is some price $p^* \geq z^*$ such that Seller's payoff is π_s if all these trades were to occur at price p^*.[17] Note that p^* must be no larger than the expected Buyer posterior mean conditional on that being above z^*, for otherwise $\pi_b < 0$. We claim that the information structure τ^* can be garbled so that p^* is an equilibrium price and trade occurs only when Buyer's posterior mean is greater than z^*. The garbling is illustrated in Figure 2 by the blue distribution. There is one signal that

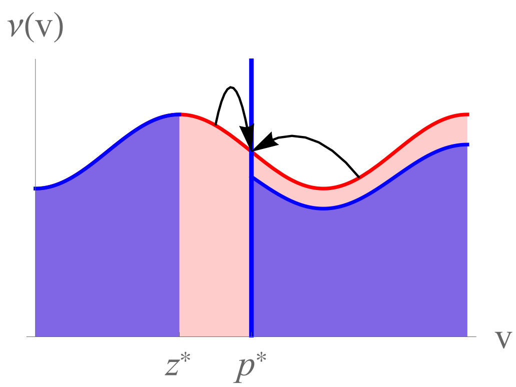

Figure 2: Construction of information structure in Theorem 2

Buyer receives when the original posterior mean is between z^* and p^*, and also receives with some probability when the original posterior mean is above p^*. The probability is chosen to make the posterior mean from this signal exactly p^*. Apart from this one new signal, Buyer receives the original signal in τ^*. Plainly, this is a garbling of τ^* and hence is feasible.

[16] The difficulty is in establishing suitable continuity. Uninformed-Seller information structures can be viewed as probability measures over Buyer's beliefs, with convergence in the sense of the weak* topology. This topology ensures continuity, with respect to probability measures, of expectations of continuous or at least Lipschitz (and bounded) functions. However, Seller's expected payoff is not the expectation of a Lipschitz function, as Seller's profit is truncated at the price she charges.

[17] That $p^* \geq z^*$ follows from $\pi_s \geq \underline{\pi}_s^{us}(\Gamma)$, as $\pi_s(\Gamma)$ itself is weakly larger than Seller's payoff from posting price z^* (and thus trading with the same set of Buyer posterior means) under information structure τ^*.

Figure 2 makes clear why the new information structure has an equilibrium with price p^* and Buyer breaking indifference in favor of trading: (i) Seller's profit from posting any price below z^* is the same as under τ^* and hence no larger than $\underline{\pi}_s^{us}(\Gamma)$; (ii) similarly, Seller's profit from posting any price above p^* is no higher than some fraction of $\underline{\pi}_s^{us}(\Gamma)$; and (iii) any price between z^* and p^* is worse that price p^*. Moreover, since Seller's profit from offering any price other than p^* is no more than $\underline{\pi}_s^{us}(\Gamma)$, it follows that when $\pi_s > \underline{\pi}_s^{us}(\Gamma)$, Buyer must break indifference as specified for Seller to have an optimal price, and the equilibrium payoffs are unique.

Remark 2. The above logic establishes that given any $\tau \in \mathbf{T}_{us}$ that implements some (π_b, π_s), τ can be garbled to uniquely implement any $(\pi_b', \pi_s') \in \mathbf{\Pi}(\Gamma)$ such that $\pi_s' > \pi_s$. That is, an uninformed Seller's payoff can always be strictly raised, and Buyer's payoff reduced (strictly, so long as it was not already zero), by garbling Buyer's information.

Remark 3. According to Theorem 1/1* and Theorem 2, uninformed-Seller information structures cannot implement all implementable payoff pairs in an environment Γ if and only if $\underline{\pi}_s^{us}(\Gamma) > \underline{\pi}_s(\Gamma)$. This inequality fails in environment Γ if $\underline{\pi}_s^{us}(\Gamma) = 0$, since that implies $\underline{\pi}_s^{us}(\Gamma) = \underline{\pi}_s(\Gamma) = 0$. An example is when there is no trade due to adverse selection when Seller is uniformed and Buyer has some information. On the other hand, it is sufficient for the inequality that

$$\underline{v} \leq \mathbb{E}[c(v)] \text{ and } \forall v \in V,\ c(v) < v. \tag{2}$$

To see why (2) implies $\underline{\pi}_s^{us}(\Gamma) > \underline{\pi}_s(\Gamma)$, notice that in any uninformed-Seller information structure, Seller can price at slightly less than Buyer's highest posterior mean valuation and guarantee trade with only (a neighborhood of) that Buyer type. If $c(v) < v$ for all v, this gives Seller a positive expected payoff, and hence $\underline{\pi}_s^{us}(\Gamma) > 0$.[18] But the first inequality in (2) is equivalent to $\underline{\pi}_s(\Gamma) = 0$. We observe that Condition (2) is compatible with severe adverse selection resulting in very little trade when Seller is uninformed and Buyer (partially or fully) informed.

3.3. Fully-Informed Buyer

We now turn to the third canonical class of information structures: Buyer is fully informed of his value v. As this is a special case of more-informed Buyer, we maintain price-independent beliefs throughout this subsection.

[18] More precisely: as V is compact, $c(v) < v$ for all v implies there exists $\varepsilon > 0$ such that $v - c(v) > \varepsilon$. Given any uninformed-Seller information structure, let $\overline{m}_v$ be the highest posterior mean valuation in the support of the posterior means induced by Buyer's signals. So there is positive probability of Buyer signals with posterior mean valuations at least $\overline{m}_v - \varepsilon/2$. By pricing at $\overline{m}_v - \varepsilon/2$, Seller's expected cost conditional on trade is bounded above by $\overline{m}_v - \varepsilon$, and hence Seller's profit conditional on trade is at least $\varepsilon/2 > 0$. It follows that $\underline{\pi}_s^{us} > 0$.

Faced with a fully informed Buyer and any sequentially rational Buyer strategy, an uninformed Seller can guarantee a profit level

$$\underline{\pi}_s^{fb}(\Gamma) \equiv \sup_p \int_p^{\bar{v}} (p - c(v))\mu(\mathrm{d}v)$$

regardless of her information. Plainly, $\underline{\pi}_s^{fb}(\Gamma) \geq \underline{\pi}_s^{us}(\Gamma)$. In monopoly pricing with $c(\cdot) = \underline{v}$, Roesler and Szentes (2017) have shown that $\underline{\pi}_s^{fb} > \underline{\pi}_s^{us}$; if there is no trade due to adverse selection when Buyer is fully informed and Seller is uniformed, then $\underline{\pi}_s^{fb} = \underline{\pi}_s^{us} = 0$. Below, we establish that when Buyer is fully informed, $\underline{\pi}_s^{fb}$ is the only additional constraint on equilibrium payoffs. Let $B_\varepsilon(\pi_b, \pi_s)$ denote an ε-ball around (π_b, π_s).

Theorem 3. *Consider a fully informed Buyer and price-independent beliefs.*

1. $\mathbf{\Pi}^*_{fb}(\Gamma) = \{(\pi_b, \pi_s) \in \mathbf{\Pi}(\Gamma) : \pi_s \geq \underline{\pi}_s^{fb}(\Gamma)\}$.
2. *For any* $(\pi_b, \pi_s) \in \mathbf{\Pi}^*_{fb}(\Gamma)$ *and any* $\varepsilon > 0$, *there is* $\tau \in \mathbf{T}_{fb}$ *with* $\Pi(\Gamma, \tau) \subset B_\varepsilon(\pi_b, \pi_s)$.

The payoff set characterized in Theorem 3 corresponds to the triangle ABE in Figure 1. Part 2 of the theorem establishes "approximately unique implementation": for any point in the triangle, there is an information structure under which all equilibria yield payoff pairs in a neighborhood of that point.

Here is the idea behind part 1 of Theorem 3. When Buyer is fully informed, an information structure can be viewed as dividing v's prior distribution, μ, into a set of μ_i that average to μ, with Seller informed of which μ_i she faces. Theorem 3 is proven by establishing that we can divide μ suitably so that against each μ_i, Seller is indifferent between pricing at all prices in the support of μ_i, including the price corresponding to $\underline{\pi}_s^{fb}$ in that environment. We call such a μ_i a (Seller-price) *incentive compatible distribution* or ICD. (Bergemann et al. (2015) call their analogous construct an extreme market.) The Appendix provides a "greedy" algorithm to compute ICDs.[19]

We can sketch how the algorithm works and construct a set of ICDs that average to the prior. Suppose $V = \{v_1, v_2, \ldots, v_K\}$, with $v_i < v_{i+1}$ for $i \in \{1, \ldots, K-1\}$ and $c(v) < v$ for all v. Given any small-enough mass of v_K, there is a unique mass of type v_{K-1} that makes Seller indifferent between charging price v_K and v_{K-1}. (If the mass is too low, Seller prefers v_K; if it is too high she prefers v_{K-1}.) Iterating down to keep Seller indifferent between all prices pins down an ICD. Choose the maximum mass of type v_K for which this works. Remove that ICD and repeat.

[19] The algorithm is defined for finite V, and we take limits to handle the infinite case.

Crucially, whenever an ICD is removed, the price corresponding to $\underline{\pi}_s^{fb}(\Gamma)$ remains optimal in the remaining "market"; this follows from the ICD's defining property of Seller indifference and an accounting identity. Therefore, Seller's profit in this segmentation of ICDs remains $\underline{\pi}_s^{fb}(\Gamma)$. Moreover, it is also optimal for Seller to always (i.e., for each μ_i) price so that there is full trade or no trade. Hence, Buyer's expected payoff can be either 0 or the entire surplus less $\underline{\pi}_s^{fb}(\Gamma)$. It follows that the fully-informed Buyer information structure defined by this set of ICDs implements point B and C in Figure 1. The entire triangle ABC can then be implemented by convexification: randomizing over this information structure (and two equilibria) and full information (where Seller obtains all the surplus).

The above logic of incentive compatible distributions reduces to that of Bergemann et al. (2015) in the context of monopoly pricing. Even in the monopoly pricing case, Theorem 3 part 2 strengthens the conclusions of Bergemann et al. (2015) by ensuring approximately unique implementation.

Remark 4. Theorems 1–3 imply that fully-informed-Buyer information structures implement all implementable payoff pairs if and only if $\underline{\pi}_s^{fb}(\Gamma) = \underline{\pi}_s(\Gamma)$. In that case, triangles AFG and ABC coincide in Figure 1. It follows that, $\underline{\pi}_s^{fb}(\Gamma) = \underline{\pi}_s(\Gamma)$ only when a fully-informed Buyer and uninformed Seller can result in full trade ($\underline{v} \geq \mathbb{E}[c(v)]$ and Seller prices at $\underline{v}$) or no trade ($\underline{v} \leq \mathbb{E}[c(v)]$ and Seller prices at some $p \geq \overline{v}$). Interestingly, when $\underline{\pi}_s^{fb}(\Gamma) > \underline{\pi}_s(\Gamma)$, fully-informed-Buyer information structures cannot even implement all payoff pairs implementable by uninformed-Seller information structures, i.e., triangles ABC and ADE in Figure 1 are distinct. That is, when $\underline{\pi}_s^{fb}(\Gamma) > \underline{\pi}_s(\Gamma)$ there is an uninformed-Seller information structure that implements some $\pi_s < \underline{\pi}_s^{fb}(\Gamma)$.[20] Theorems 2–3 further imply in this case that Buyer can obtain a strictly higher payoff when he is not fully informed (and Seller is uninformed).

4. Discussion

4.1. Multidimensionality

Suppose Buyer and Seller's cost and valuation pair (c, v) is a two-dimensional random variable distributed according to joint distribution μ with a compact support. The extension of our maintained assumption of commonly known gains from trade is: for all

[20] Pick any $p' > \underline{v}$ such that $p' - \mathbb{E}[c(v)] < \underline{\pi}_s^{fb}(\Gamma)$. Following the construction described after Theorem 2, we can mix all valuations $v \leq p'$ with a fraction $\lambda > 0$ of valuations $v > p'$ so that the mixture has posterior mean exactly p'. The remaining fraction $1 - \lambda$ of valuations above p' are revealed to Buyer. With this uninformed-Seller information structure, consider any equilibrium in which Buyer purchases when indifferent. (Such an equilibrium with price-independent belief exists.) Seller's profit is at most $(1 - \lambda)\underline{\pi}_s^{fb}(\Gamma)$ from any price $p > p'$, and $p' - \mathbb{E}[c(v)]$ from price $p = p'$. Hence, Seller's profit is strictly less than $\underline{\pi}_s^{fb}(\Gamma)$.

$(c, v) \in \mathrm{Supp}(\mu)$, $v \geq c$; and $\mathbb{E}[v - c] > 0$. An information structure is now a joint distribution $P(t_b, t_s, v, c)$ whose marginal distribution on (v, c) is μ.

We claim that the content of our key results, Theorem 1/Theorem 1*, Theorem 2 and Theorem 3 still hold. To see why, let $\underline{v}$ be the lowest valuation in the support of μ. Seller's individual rationality constraint is now $\max\{\underline{v} - \mathbb{E}[c], 0\}$, as she can guarantee this by setting either a sufficiently high price or a price (arbitrarily close to) $\underline{v}$, regardless of her signal. We can define a cost function $c(v') \equiv \mathbb{E}_\mu[c|v = v'] \geq v'$. This results in an environment satisfying all the maintained assumptions of our baseline model, except that $c(\cdot)$ may not be continuous. Recall that such continuity is not needed for proving Theorem 1, Theorem 1*. Both Theorem 2 and Theorem 3 use continuity of $c(\cdot)$ to guarantee that $\mathbb{E}_\nu[c(v)]$ is a continuous function of $\nu \in \Delta(V)$ for a key convergence result. However, in the two-dimensional type environment, $\mathbb{E}_\nu[c]$ is still a continuous function of $\nu \in \Delta(V \times C)$. Note that Theorem 3 uses upper semi-continuity of Seller's profit in price; boundedness of c and $c \leq v$ is sufficient for such upper semi-continuity.

4.2. Negative Trading Surplus

Returning to our baseline model, we next discuss cases in which trade sometimes generates negative surplus. That is, we drop the assumption that $c(v) \leq v$; we do not require $\mathbb{E}[v - c(v)] > 0$ either. Define $S_\lambda(\Gamma)$ for $\lambda \in [1, \infty)$ as

$$S_\lambda(\Gamma) \equiv \int_{\underline{v}}^{\bar{v}} [\underline{v} - c(v) + \lambda(v - \underline{v})]^+ \, \mu(\mathrm{d}v),$$

where $[\cdot]^+ \equiv \max\{\cdot, 0\}$. It is readily verified that $S_1(\Gamma) = \mathbb{E}\left[[v - c(v)]^+\right]$ and $\lim_{\lambda \to 0} S_\lambda(\Gamma)/\lambda = \mathbb{E}[v] - \underline{v}$. Allowing negative trading surplus does not affect our definition of wPBE. So the notation $\mathbf{\Pi}(\Gamma)$ and $\underline{\pi}_s(\Gamma)$ still have the same meanings as before. The next proposition shows that $\mathbf{\Pi}(\Gamma)$ is now characterized by three constraints: as before, the two individual rationality constraints, $\pi_s \geq \underline{\pi}_s(\Gamma)$ and $\pi_b \geq 0$; and different now, a Pareto frontier defined by all $S_\lambda(\Gamma)$.

Proposition 1. *Consider all information structures and equilibria when trade can generate negative surplus.*

$$\mathbf{\Pi}(\Gamma) = \left\{ (\pi_b, \pi_s) : \begin{array}{l} \pi_b \geq 0 \\ \pi_s \geq \underline{\pi}_s(\Gamma) \\ \lambda \pi_b + \pi_s \leq S_\lambda(\Gamma), \ \forall \lambda \geq 1 \end{array} \right\}.$$

Proposition 1 shows that when the trade can be inefficient, it may be impossible to lower

Seller's profit without losing total surplus (i.e., $\pi_b + \pi_s$). This result is intuitive at the extremes. Suppose for simplicity $\underline{v} > \mathbb{E}[c(v)]$. If $\pi_s = \underline{\pi}_s(\Gamma)$, then since $\underline{\pi}_s(\Gamma) = \underline{v} - \mathbb{E}[c(v)]$, Seller must find it optimal to price at $\underline{v}$ and trade with probability one, including when trade generates negative surplus. So Buyer's payoff is $\mathbb{E}[v] - \underline{v} = \lim_{\lambda \to \infty} S_\lambda(\Gamma)/\lambda$, and total surplus is $\mathbb{E}[v - c(v)]$, which includes the maximum amount of inefficient trade. On the other hand, an information structure that publicly reveals (only) whether trade is efficient (i.e., whether $v \geq c(v)$) can implement efficient trade with all the surplus accruing to Seller. In this case Seller obtains $S_1(\Gamma)$ while Buyer obtains zero. The proposition establishes that in between $S_1(\Gamma)$ and $\underline{\pi}_s(\Gamma)$, information structures and equilibria that lower Seller's profit must reduce total surplus, as the Pareto frontier $S_\lambda(\Gamma)$ is a concave function with maximal slope -1.

The function capturing the Pareto frontier $S_\lambda(\Gamma)$ is the maximal weighted sum of Buyer and Seller payoff assuming that Seller will trade at the low price $p = \underline{v}$ whenever it creates positive weighted total payoff. Since we place higher weight on buyer ($\lambda \geq 1$), $S_1(\Gamma)$ is the upper bound for the weighted total payoff. In the proof of Proposition 1, we show that these bounds are tight — any payoff pair on the frontier can be implemented by some information structure. The key observation is that whenever $\underline{v} - c(v) + \lambda(v - \underline{v}) < 0$, it follows that $v - c(v) < (1 - \lambda)(v - \underline{v}) \leq 0$. That is, it is incentive compatible for Seller to not sell when it is common knowledge that v creates negative weighted total payoff. Therefore, we can simply release a public signal indicating the sign of weighted total payoff, and if positive, induce Seller to sell at $\underline{v}$, which implies all bounds $S_\lambda(\Gamma)$.

4.3. Linear Cost

When Seller's cost function $c(v)$ is linear, we can explicitly characterize an information structure that implements Seller's minimum implementable payoff $\underline{\pi}_s^{us}(\Gamma)$ when Seller is uninformed (or, per Theorem 2, when Buyer is better informed), subject to price-independent beliefs. We note that given the discussion in Subsection 4.1, a linear $c(v)$ subsumes richer environments in which conditional expectations are linear, including Gaussian environments (cf. Bar-Isaac et al., 2018).

Condition 1. $c(v) = \lambda v + \gamma$, *for some* $\lambda \in \mathbb{R}$.

Let $F(v)$ be the cumulative distribution function (CDF) corresponding to the prior measure μ. Let $D(\mu)$ be the set of all distributions whose CDF G satisfies

$$\int_V v \mathrm{d}G(v) = \int_V v \mathrm{d}F(v) \quad \text{and} \quad \int_{\underline{v}}^{v} G(s)\mathrm{d}s \leq \int_{\underline{v}}^{v} F(s)\mathrm{d}s,\ \forall v \in V.$$

That is, $D(\mu)$ contains all distributions that are mean-preserving-contractions (MPC) of μ. It is well known that $D(\mu)$ characterizes the set of distributions of Buyer posterior means that can be generated by any (uninformed-Seller) information structure.

Definition 2. G is an *incentive compatible distribution* (ICD) if

$$\int_{s\geq p}(p-c(s))G(\mathrm{d}s)=\text{constant}\geq 0,\ \forall p\in \mathrm{Supp}(G).$$

That is, when Buyer's (mean) valuation is distributed according to an ICD G, Seller is indifferent among all prices in the support of G, assuming Buyer buys when indifferent. We focus on a special family of ICDs whose supports are intervals $[v_*, v^*]$. Given Condition 1, such ICDs have analytical an expression:

$$G(v)=\begin{cases}1-\left(\dfrac{(1-\lambda)v-\gamma}{(1-\lambda)v_*-\gamma}\right)^{\frac{1}{\lambda-1}} & \text{if } \lambda\neq 1\\ 1-e^{\frac{v-v_*}{\gamma}} & \text{if } \lambda=1.\end{cases} \tag{3}$$

Given any v_* and the above $G(v)$, an upper bound v^* is pinned down by the condition $\mathbb{E}_G[v]=\mathbb{E}_F[v]$. So a family of ICDs is parametrized by a single parameter v_*, and hence an ICD is denoted by CDF $\widetilde{G}_{v^*}(v)$ with density $\widetilde{g}_{v_*}(v)$. It is not difficult to verify that $\widetilde{G}_{v_*}(v)$ is monotonic in v_* for v in the support. Consequently, all ICDs are ordered according to the MPC order. For $v_*\in[\underline{v},\mathbb{E}[v]]$, when v_* increases, the ICD increases in the MPC order.

Proposition 2. *Suppose Condition 1 is satisfied and let* $p_*\equiv\min\left\{v_*\middle|\widetilde{G}_{v_*}\in D(\mu)\right\}$. *It holds that* $\underline{\pi}_s^{us}(\Gamma)=p_*-\mathbb{E}_\mu[c(v)]$.

Proposition 2 states that an uninformed Seller's minimum payoff (with price-independent beliefs) is characterized by the ICDs that are MPCs of the prior distribution μ. By definition of ICDs, Seller's profit when facing ICD $\widetilde{G}_{v_*}$ is $v_*-\mathbb{E}[v]$. Seller's minimum payoff is implemented by the most dispersed ICD that is a MPC of the prior distribution.

In proving Proposition 2, the key step is to show that given the prior μ, garbling Buyer's information such that the posterior mean distribution is an ICD makes Seller weakly worse off. It is then without loss to only consider ICDs to implement $\underline{\pi}_s^{us}(\Gamma)$. We can thus work with a tractable one-dimensional problem. Consider the set $D(\mu)$. Suppose we find $G(v)$ such that: (i) $G\in D(\mu)$; (ii) G is an ICD; and (iii) $\exists p$ such that $\int_{\underline{v}}^{p}G(s)\mathrm{d}s=\int_{\underline{v}}^{p}F(s)\mathrm{d}s$. Such a G is the most-dispersed ICD that is a mean preserving contraction of F. Consider the following two manipulations:

$$\lambda \int_{p}^{\bar{v}} F(s)\mathrm{d}s = \int_{p}^{\bar{v}} (p - c(s))\mathrm{d}F(s) - (1 - F(v))\,(p - c(v)) + \lambda(\bar{v} - p), \tag{4}$$

$$\lambda \int_{p}^{\bar{v}} G(s)\mathrm{d}s = \int_{p}^{\bar{v}} (p - c(s))\mathrm{d}G(s) - (1 - G(v))\,(p - c(v)) + \lambda(\bar{v} - p). \tag{5}$$

First, by property (iii) above, the LHS of Equation 4 equals the LHS of Equation 5. Second, since $\int_{p}^{\bar{v}}(F(s) - G(s))\mathrm{d}s$ is minimized at $\overline{v}$, this implies $F(p) \leq G(p)$. Therefore, the first term on RHS of Equation 4 must be larger than that of Equation 5. Notice that the first term on RHS is Seller's profit when offering price p. This implies that the profit from offering p given Buyer mean-valuation distribution $G(v)$ is lower than that given $F(v)$, which is lower than Seller's maximum profit given $F(v)$. On the other hand, by the ICD property, p is already Seller's optimal price given $G(v)$. The optimal profit given valuation distribution F must be no lower than that of G. Therefore, we have proved that to minimize Seller's profit, it is without loss to consider only ICDs within the set $D(\mu)$, which is a one dimensional subspace. The detailed proof that such ICDs exist is relegated to the appendix.

The monopoly-pricing environment with $c(\cdot) = \underline{v}$ is covered by Condition 1 with $\lambda = 0$ and $\gamma = \underline{v}$. The distribution G in (3) then reduces to that identified by Roesler and Szentes (2017).

If the prior μ has binary support, then the cost function $c(v)$ is linear. In this special case, we can solve for $\underline{\pi}_s^{us}(\Gamma)$ explicitly.

Corollary 1. *Suppose μ has binary support: $V = \{v_1, v_2\}$. Let $\lambda = \frac{c(v_2)-c(v_1)}{v_2-v_1}$, and let p be the unique solution to*

$$(p - c(p))^{\frac{1}{\lambda-1}}(p - \mathbb{E}[c(v)]) = (v_2 - c(v_2))^{\frac{\lambda}{\lambda-1}}. \tag{6}$$

Then $\underline{\pi}_s^{us}(\Gamma) = \max\{p, v_2\} - \mathbb{E}[c(v)]$.

References

AKERLOF, G. A. (1970): "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84, 488–500.

BAGWELL, K. AND M. H. RIORDAN (1991): "High and Declining Prices Signal Product Quality," *American Economic Review*, 81, 224–239.

BAR-ISAAC, H., I. JEWITT, AND C. LEAVER (2018): "Adverse Selection, Efficiency and the Structure of Information," Unpublished.

BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): "The Limits of Price Discrimination," *American Economic Review*, 105, 921–57.

BERGEMANN, D. AND S. MORRIS (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

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

CHO, I.-K. AND D. KREPS (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179–221.

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

DURRETT, R. (1995): *Probability: Theory and Examples*, Duxbury Press, 2nd ed.

EINAV, L. AND A. FINKELSTEIN (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 25, 115–38.

FUDENBERG, D. AND J. TIROLE (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236–260.

GARCIA, D., R. TEPER, AND M. TSUR (2018): "Information Design in Insurance Markets: Selling Peaches in a Market for Lemons," Unpublished.

GREENWALD, B. C. (1986): "Adverse Selection in the Labour Market," *Review of Economic Studies*, 53, 325–347.

JOVANOVIC, B. (1982): "Favorable Selection with Asymmetric Information," *Quarterly Journal of Economics*, 97, 535–539.

KESSLER, A. S. (2001): "Revisiting the Lemons Market," *International Economic Review*, 42, 25–41.

KREPS, D. AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.

LEHMANN, E. L. (1988): "Comparing Location Experiments," *The Annals of Statistics*, 16, 521–533.

LEVIN, J. (2001): "Information and the Market for Lemons," *RAND Journal of Economics*, 32, 657–666.

LIBGOBER, J. AND X. MU (2019): "Informational Robustness in Intertemporal Pricing," Working paper.

MAKRIS, M. AND L. RENOU (2018): "Information Design in Multi-Stage Games," Working paper.

MILGROM, P. R. AND R. J. WEBER (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*, 10, 619–632.

MUSSA, M. AND S. ROSEN (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301–317.

RAVID, D., A.-K. ROESLER, AND B. SZENTES (2019): "Learning Before Trading: On the Inefficiency of Ignoring Free Information," Working paper.

ROESLER, A.-K. AND B. SZENTES (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107, 2072–80.

TERSTIEGE, S. AND C. WASSER (2019): "Buyer-optimal Extensionproof Information," Working paper.

Appendices omitted from this version, but are available on request.