

Lemonade from lemons: Information design and adverse selection

NAVIN KARTIK

Department of Economics, Yale University

WEIJIE ZHONG

Graduate School of Business, Stanford University

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, noninformational 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.

KEYWORDS. Monopoly pricing, interdependent values, price discrimination, market segmentation, buyer-optimal information, Bayesian persuasion.

JEL CLASSIFICATION. D42, D82, D83, L12, L15.

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 of the work in information economics prior to the last decade took the market participants' information as given and studied properties of a particular market structure or mechanism, or tackled these properties across various mechanisms.

Our paper joins a recent wave of research—elaborated subsequently—by instead asking: 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

Navin Kartik: nkartik@gmail.com

Weijie Zhong: weijie.zhong@stanford.edu

We thank Laura Doval, David Kim, Wenhao Li, Pietro Ortoleva, Alessandro Pavan, and anonymous referees for useful comments. We also received helpful feedback from a number of audiences. Tianhao Liu and Yangfan Zhou provided excellent research assistance.

The first author was affiliated with Columbia University during most of the work on this paper.

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's 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 the 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. It is this variety of possible information that our paper focuses on.

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 $v \in [\underline{v}, \bar{v}] \subset \mathbb{R}$. 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, that is, the function $c(\cdot)$ 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 instead Seller is better informed than Buyer, signaling becomes relevant; the price can serve as a credible signal if the two parties' information is suitably correlated (e.g., [Bagwell and Riordan \(1991\)](#)). A constant $c(\cdot)$ captures an environment in which there is 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 Theorems 1/1*, 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 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

¹We 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 references to empirical evidence on both.

³As 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 a weak perfect Bayesian equilibrium; we also address refinements for our constructive arguments.




FIGURE 1. Outcome under different restrictions on information structures.

inclusion: Theorem 1’s set is the largest; Theorem 2’s is intermediate; and Theorem 3’s the smallest. Figure 1 summarizes.

In the figure’s axes, π_b and π_s represent respectively Buyer’s and Seller’s ex ante expected 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 $p > \bar{v}$, 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, that is, obtains in an equilibrium under some information structure. It is immediate that point A obtains when 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 simplicity, 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 $v = \underline{v}$. Section 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) can always be implemented by reducing Buyer's information in the sense of Blackwell (1953). We also explain in Section 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, that is, 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 $p \geq \bar{v}$ and both parties' payoffs are 0). As detailed in Remark 5 of Section 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.⁶ 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 Section 3.2. When D and F are distinct, maximizing Buyer's payoff, that is, 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.

While our results do not speak directly to the economics of privacy, recently reviewed by Acquisti, Taylor, and Wagman (2016), they do offer a notable twist. Consumer welfare

⁴We 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 do not know” requirement (Fudenberg and Tirole (1991)) that is standard in versions of perfect Bayesian equilibrium and implied by sequential equilibrium.

⁵It 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.

⁶This substantially broadens the message of Roesler and Szentes (2017) on the benefits to restricting Buyer's information.

can be higher when a monopolist has information about a consumer's valuation that the consumer does not; indeed, maximizing consumer welfare in our single-unit setting frequently necessitates that. This is an implication of our Theorems 1/1* and Theorem 2.

Relatedly, we view those results as cautioning against assuming away the possibility that sellers have information buyers do not—not only does this seem relevant in practice, as discussed earlier, but it has significant payoff consequences. An alternative perspective on Theorem 1/Theorem 1* is that they are negative results: “anything goes” without restrictions on information structures or equilibria. From that perspective, Theorem 2 reveals that what is essential to restrain payoffs to its smaller set (triangle *ADE* in Figure 1) is that prices do not provide Buyer with information. Whether this is a consequence of Buyer being better informed than Seller or a principle of equilibrium selection does not matter. Theorem 3 characterizes the additional payoff restrictions that obtain from Buyer being fully informed.

Related literature Our questions and results are most closely related to Bergemann, Brooks, and Morris (2015), Roesler and Szentes (2017), and Makris and Renou (2023, 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 $c(v) \leq \underline{v}$.⁷ We study interdependent values because of its importance in many economic environments; substantively and methodologically, we explore whether and how insights from the monopoly-pricing problem hold more generally.

Bergemann, Brooks, and Morris (2015) assume Buyer is fully informed, and hence can only vary Seller's information.⁸ 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 methodology is to construct the “isoprofit distributions,” which reduces to their “extremal markets” in monopoly pricing. An economic lesson from our analysis is that unlike in monopoly pricing, there are salient interdependent-value environments in which a fully informed Buyer can achieve all implementable payoffs, even when the Akerlof market has inefficiency; see Remark 5.

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, namely they identify the triangle *ADE* in Figure 1 as the implementable set when Seller is uninformed. Even for this result, our methodology is quite different from theirs because we do not assume a linear $c(\cdot)$ function; our methodology delivers new insights, including that noted in Remark 2, and also sheds a different light on Roesler and Szentes' payoff triangle. When we specialize to a linear $c(\cdot)$, we can obtain a sharper characterization

⁷Related to Roesler and Szentes (2017) are also Du (2018) and Libgober and Mu (2021), who consider worst-case profit guarantees for Seller in static and dynamic environments, respectively. Terstiege and Wasser (2020) qualify Roesler and Szentes (2017) by allowing Seller to supply Buyer with additional information, although Seller cannot have any private information of her own.

⁸Less directly related to our work, there are also recent papers extending the approach of Bergemann, Brooks, and Morris (2015) to monopolistic markets with multiple products (e.g., Ichihashi (2020), Haghpahan and Siegel (2023), Terstiege and Vigier (2025)), oligopolistic markets with differentiated products (e.g., Elliott, Galeotti, Koh, and Li (2022, 2023), Bergemann, Brooks, and Morris (2025)), and profit-maximizing information design by intermediaries (e.g., Yang (2022)).

of the point E , which extends [Roesler and Szentes](#)' characterization of Buyer-optimal information to an interdependent-values environment; see [Proposition 2](#).

Our [Theorems 1–3](#) establish that the “alignment” principle highlighted by [Bergemann, Brooks, and Morris \(2025\)](#)—Buyer surplus/payoff is maximized when total surplus (the sum of Buyer and Seller payoffs) is maximized—extends with a single seller beyond the settings of [Bergemann, Brooks, and Morris \(2015\)](#) and [Roesler and Szentes \(2017\)](#), both in terms of the information structures considered and to interdependent values. However, we qualify this point in [Section 4.2](#) when trade may be inefficient, as was also illustrated by example in [Roesler and Szentes \(2017, Online Appendix\)](#).

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 \(2023\)](#). As an application of their general results on “revelation principles” for information design in multi-stage games, [Makris and Renou's \(2023\) 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;⁹ we go further by establishing in [Theorem 1*](#) off-the-equilibrium-path belief consistency in the sense of sequential equilibrium ([Kreps and Wilson \(1982\)](#)). We also show in [Theorem 1/1*](#) that a single information structure implements all payoffs in the relevant triangle.

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 the other. [Levin \(2001\)](#) identifies conditions under which the volume of trade decreases when one party is kept uninformed and the other's information becomes more effective in the sense of [Lehmann \(1988\)](#); see also [Kessler \(2001\)](#). Assuming a linear payoff structure, [Bar-Isaac, Jewitt, and Leaver \(2021\)](#) consider how certain changes in Gaussian information affect the volume of trade, surplus, and a certain quantification of adverse selection.

[Dang \(2008\)](#), [Pavan and Tirole \(2025\)](#), and [Thereze \(2023\)](#) study endogenous costly information acquisition with interdependent values, using different assumptions about the nature and timing of information acquisition and the underlying economic environment. By contrast, we do not have strategic or costly information acquisition; rather, the informational environment is exogenously (and costlessly) varied.

In a model with interdependent values where they hold fixed a partially-informed buyer's information, [Deb, Pai, and Roesler \(2024\)](#) characterize the outcome—including what information the seller should have—that maximizes the buyer's payoff. They highlight that the solution typically involves the seller being partially informed. [Garcia, Teper, and Tsur \(2018\)](#) solve for socially optimal information provision in an insurance

⁹[Makris and Renou](#) use an apparatus of “sequential Bayes correlated equilibrium,” which we do not. 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.

setting with adverse selection; owing to a cross-subsidization motive, full information disclosure is typically not optimal. Pollrich and Strausz (2024) study a third-party certifier in an adverse-selection environment. Their environment corresponds to our Buyer being fully informed and facing a competitive market of sellers. Among other things, they discuss implementable payoff vectors for Buyer (conditional on type) and their certifier.

The rest of our 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. All formal proofs are in the Appendices.

2. MODEL

2.1 Primitives

There are two players, Seller and Buyer; given the assumptions that follow, Buyer can be viewed as representing a market of buyers. 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 $c(v)$. We assume $c : V \rightarrow \mathbb{R}$ is continuous, $v - c(v) \geq 0$ for all v , and $\mathbb{E}[v - c(v)] > 0$. Hence, the 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 the function $c(v)$ need not be monotonic. Section 4 extends the model to Seller's cost being stochastic even conditional on v , and considers the possibility of negative trading surplus. We call $\Gamma \equiv (c, \mu)$ 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 terminology and refer to “distribution” even though “measure” would sometimes be more precise.) Formally, there is a probability space $(\Omega, \mathcal{F}, P)$, complete and separable metric spaces T_s and T_b (equipped with their Borel sigma algebras), and an integrable function $X : \Omega \rightarrow T_s \times T_b \times V$. We hereafter suppress the probability space and define, with an abuse of notation, $P(D) = P(X^{-1}(D))$ for any measurable $D \subset T_s \times T_b \times V$.¹⁰ Each realization of random variable X is a triplet (t_b, t_s, v) , where $t_b \in T_b$ 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 T_i , with the convention $T_v \equiv V$. We require $P_v = \mu$; this is an 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 .

¹⁰We write $\subset$ for a “weak subset.”

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 .

We highlight that our notion of an information structure involves parties receiving information only at the outset. A more permissive notion would also allow Buyer to receive information after Seller posts her price, as in the literature on multistage information design (Makris and Renou (2023), Doval and Ely (2020)). Permitting that would not change some of our results, in particular Theorems 1/1* and Theorem 3, but would expand the implementable set characterized in Theorem 2. Methodologically, our interest in only ex ante information means that existing “revelation principles” do not directly apply.

Our assumption that Seller simply posts a price—rather than using more complicated mechanisms—is not restrictive for our main results. Remark 6 elaborates later.

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 .¹¹ Buyer's strategy $\alpha: \mathbb{R} \times T_b \rightarrow [0, 1]$ 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:

$$\begin{aligned}\pi_b &= \int (v - p)\alpha(t_b, p)\sigma(dp|t_s)P(dt_s, dt_b, dv), \\ \pi_s &= \int (p - c(v))\alpha(t_b, p)\sigma(dp|t_s)P(dt_s, dt_b, dv).\end{aligned}$$

Our baseline equilibrium concept is a 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 $\nu(v|p, t_b)$.

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

¹¹Here, $\sigma(\cdot|t_s)$ 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.

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:

$$\sigma \in \arg \max_{\hat{\sigma}} \int (p - c(v)) \alpha(p, t_b) \hat{\sigma}(dp|t_s) P(dt_s, dt_b, dv);$$

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

$$\int_D \nu(dv|p, t_b) \sigma(dp|t_s) P(dt_s, dt_b, V) = \int_D \sigma(dp|t_s) P(dt_s, dt_b, dv).$$

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 subsequently discuss refinements.

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 $c(v)$.¹² When discussing such information structures, we write the

¹²Among all reasonable notions of uninformed Seller (e.g., one might only require Seller to have no information about $\mathbb{E}[v]$, while permitting information about $c(v)$), we take the most restrictive one. Our results will imply that a more permissive notion would not change the relevant implementable sets—in particular, that characterized in Theorem 2. This point also applies to our notion of more informed Buyer defined shortly.

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 structures 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 by $\mathbf{T}_{fb}$. Note that a fully informed Buyer need not know Seller's signal; but that is irrelevant to Buyer, because his optimal action after any price only depends on his known value.

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 , that is, for any measurable $D_s \subset T_s$ and $D_v \subset V$, $P(D_s \times D_v|t_b) = P(D_s|t_b)P(D_v|t_b)$. Another way to interpret this requirement is that random variable t_b must be statistically sufficient for t_s with respect to v , that is, 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

$$\int_D \nu(dv|p, t_b)P(dt_b, T_s, V) = \int_D P(dt_b, T_s, dv)$$

for any measurable $D \subset T_b \times V$. We refer to this condition as *price independent beliefs*.¹³ Note that although we have motivated the condition by Buyer being more informed than Seller, the condition is meaningful even otherwise, capturing the notion of equilibria in which there is no signaling by Seller, or, more precisely, that Buyer does not learn anything about his value v from the price that he does not already learn from his own signal. In a more informed Buyer information structure, price independent beliefs would be implied by the “no signaling what you do not 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.¹⁴

¹³The condition is distinct from “passive beliefs,” which is typically used to restrict beliefs after off the equilibrium path events.

¹⁴An 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 beliefs, even though Seller's price and Buyer's belief are perfectly correlated ex ante. The point is that Buyer's belief does not depend on price conditional on his signal.

Some more notation will be helpful. Define

$$\begin{aligned}\Pi^*(\Gamma, \tau) &\equiv \{(\pi_b, \pi_s) : \exists \text{ wPBE of } (\Gamma, \tau) \text{ with price independent beliefs and payoffs } (\pi_b, \pi_s)\}, \\ \Pi^*(\Gamma) &\equiv \bigcup_{\tau \in \mathbf{T}} \Pi^*(\Gamma, \tau), \quad \text{and} \\ \Pi_i^*(\Gamma) &\equiv \bigcup_{\tau \in \mathbf{T}_i} \Pi^*(\Gamma, \tau) \quad \text{for } i = us, fb, mb.\end{aligned}$$

So, Π^* and Π^* are analogous to the implementable payoff sets Π and Π defined earlier, but restricted to equilibria with price independent beliefs. Π_{us}^* , Π_{fb}^* , and Π_{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 ,

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

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 $\Pi(\Gamma)$, $\Pi^*(\Gamma)$, $\Pi_{mb}^*(\Gamma)$, $\Pi_{us}^*(\Gamma)$, and $\Pi_{fb}^*(\Gamma)$. 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

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

To interpret this quantity, 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 $\underline{v} - \mathbb{E}[c(v)] - \varepsilon$ 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 with any information structure must be at least $\underline{\pi}_s(\Gamma)$.

Turning to the buyer, he can guarantee himself the payoff $\pi_b = 0$ by rejecting all prices. It follows that the implementable set $\Pi(\Gamma)$ 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. *The set of implementable outcomes under all information structures and equilibria is*

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

PROOF. See Appendix A. □

Theorem 1 says that the set $\Pi(\Gamma)$ corresponds to the triangle AFG in Figure 1. In particular, Buyer can receive the entire surplus beyond Seller's payoff guarantee. This is perhaps surprising, as Seller has substantial bargaining power. Note that when $\underline{v} \leq \mathbb{E}[c(v)]$, a reasonable condition, Seller's payoff guarantee is zero; in that case, Theorem 1 implies that Buyer can obtain the entire surplus.¹⁵

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 $p_h = \mathbb{E}[v]$, with probability $\sigma(p_l) \in [0, 1]$. 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

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

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 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 the construction above, consider its implication for monopoly pricing with $c(\cdot) = \underline{v}$. The equilibrium with $p_l = \underline{v}$ (hence $\alpha(p_h) = 0$, i.e., the buyer rejects the higher price) 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 does the monopolist not 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 his belief to $v = \underline{v}$. Such updating 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 beliefs are 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 beliefs condition described in Section 2.3; that would imply Buyer must purchase at any price $p < \mathbb{E}[v]$.

¹⁵In monopoly pricing with $c(\cdot) = \underline{v}$, Seller's payoff guarantee of zero is lower than the revenue guarantee identified by Du (2018, Section 5), which is typically positive. Du's notion is different from ours.

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 do not 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 $\Delta > 0$ such that for any finite price grid with size Δ , there is a finite information structure inducing a game with a set of sequential equilibrium payoffs that is an ε -net of $\Pi(\Gamma)$, the set of implementable outcomes under all information structures and equilibria.*¹⁶

PROOF. See Appendix B. □

In fact, the proof of Theorem 1* establishes 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.¹⁷

Even when our environment is specialized to monopoly pricing, it is worth highlighting two contrasts between Theorem 1/1* and results of Bergemann, Brooks, and Morris (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, Brooks, and Morris's, where the consumer is fully informed.) We will see in Section 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 $\Pi(\Gamma)$,

¹⁶An information structure is finite if the signal spaces T_b and T_s are finite. A finite price grid of size Δ means that the set of prices is finite, with minimum price no higher than $\underline{v}$ and maximum price no lower than $\bar{v}$, and any two consecutive prices are no more than Δ apart. 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 from trading are defined directly as $(\mathbb{E}[v|t_b, t_s] - p, p - \mathbb{E}[c(v)|t_b, t_s])$.

For $Y \subset \mathbb{R}^2$ and $\varepsilon > 0$, the set $A \subset Y$ is an ε -net of Y if for each $y \in Y$ there is $a \in A$ such that $\|y - a\| < \varepsilon$, where $\|\cdot\|$ is the Euclidean distance.

¹⁷The idea behind D1 is to ask, for any off-path price, whether one type of Seller would deviate for any 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.

analogously to the construction described after Theorem 1 that used a single information structure.¹⁸ Bergemann, Brooks, and Morris (2015) and Roesler and Szentes (2017) instead 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, that is, $\tau \in \mathbf{T}_{mb}$, 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 than Seller, but that the condition is well-defined even otherwise. If Buyer is not more informed than Seller, then price independent beliefs ought to be viewed as an equilibrium restriction. Readers should bear in mind that, to reduce repetition, the qualifier “with price independent beliefs” applies for the rest of this subsection unless stated explicitly otherwise.

It is useful to define

$$\underline{\pi}_s^{us}(\Gamma) \equiv \inf\{\pi_s : \exists(\pi_b, \pi_s) \in \Pi_{us}^*(\Gamma)\} \quad (1)$$

as the infimum payoff that an *uninformed* Seller can obtain, no matter Buyer’s information (among equilibria with price independent beliefs, we stress). Plainly, $\underline{\pi}_s^{us}(\Gamma) \geq \underline{\pi}_s(\Gamma)$. In monopoly pricing with $V = [\underline{v}, \bar{v}]$ 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 uninformed and Buyer has some information, then $\underline{\pi}_s^{us} = \underline{\pi}_s = 0$. We do not have a general explicit formula for $\underline{\pi}_s^{us}$; Section 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. *The set of implementable outcomes under all information structures in equilibria with price independent beliefs is the same as the set of implementable outcomes under uninformed-Seller information structures in equilibria with price independent beliefs. Moreover:*

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

PROOF. See Appendix C. □

¹⁸In 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 $\Pi(\Gamma)$. See Proposition B.1 in the Appendix for a formal statement.

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 $\Pi^*(\Gamma) \supset \Pi_{mb}^*(\Gamma) \supset \Pi_{us}^*(\Gamma)$ is trivial. So, part 1 of the theorem amounts to establishing the reverse inclusions. The intuition for those—given part 2’s characterization of Π_{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.

Let us describe how we obtain the characterization of $\Pi_{us}^*(\Gamma)$ 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.¹⁹ 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 $\Pi_{us}^*(\Gamma)$ 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 a sufficient statistic for his decision.) For simplicity, suppose this posterior-mean distribution has a density, as depicted by the red curve in Figure 2. Fix any $(\pi_b, \pi_s) \in \Pi_{us}^*(\Gamma)$.




FIGURE 2. Construction of garbling of τ^* .

¹⁹The 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.

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^* .²⁰ 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 as the distribution depicted by the blue curve and line. There is one signal that 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.

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 than 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 $(\pi'_b, \pi'_s) \in \Pi(\Gamma)$ 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. Even when specialized to the case of monopoly pricing, this provides a different perspective on why [Roesler and Szentes \(2017\)](#) obtain a payoff triangle. More importantly, our methodology also handles the case of interdependent values—specifically, a non-linear cost function $c(\cdot)$.

REMARK 3. According to Theorems 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 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 uninformed and Buyer has some information. Conversely, $\underline{\pi}_s^{us}(\Gamma) > \underline{\pi}_s(\Gamma)$ if

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

To see why, 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

²⁰That $p^* \geq z^*$ follows from $\pi_s \geq \underline{\pi}_s^{us}(\Gamma)$, as $\underline{\pi}_s^{us}(\Gamma)$ 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^* .

expected payoff, and hence $\underline{\pi}_s^{us}(\Gamma) > 0$.²¹ But the first inequality in (2) is equivalent to $\underline{\pi}_s(\Gamma) = 0$. Hence, (2) implies $\underline{\pi}_s^{us}(\Gamma) > \underline{\pi}_s(\Gamma)$. We observe that Condition (2) is compatible with severe adverse selection resulting in very little trade when Seller is uninformed and Buyer is (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 a more informed Buyer, we maintain price independent beliefs throughout this subsection.

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

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

regardless of her information. Plainly, $\underline{\pi}_s^{fb}(\Gamma) \geq \underline{\pi}_s^{us}(\Gamma)$. In monopoly pricing with $c(\cdot) = v$, [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 uninformed, then $\underline{\pi}_s^{fb} = \underline{\pi}_s^{us} = 0$. We establish below that when Buyer is fully informed, $\underline{\pi}_s^{fb}$ is the only additional constraint on equilibrium payoffs.

THEOREM 3. *The set of implementable outcomes under fully informed Buyer information structures and equilibria with price independent beliefs is*

$$\Pi_{fb}^*(\Gamma) = \{(\pi_b, \pi_s) \in \Pi(\Gamma) : \pi_s \geq \underline{\pi}_s^{fb}(\Gamma)\}.$$

PROOF. See Appendix D. □

The payoff set characterized in Theorem 3 corresponds to the triangle ABE in Figure 1. Here is the idea behind the result. 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. Such a μ_i is analogous to an “extremal market” introduced by [Bergemann, Brooks, and Morris \(2015\)](#) in the context of monopoly pricing. To highlight the profit implication of such a distribution and because that implication is relevant across multiple information structures in our paper, we call such a μ_i an *isoprofit distribution* or IPD.

²¹More 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 $\bar{m}_v$ 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 $\bar{m}_v - \varepsilon/2$. By pricing at $\bar{m}_v - \varepsilon/2$, Seller's expected cost conditional on trade is bounded above by $\bar{m}_v - \varepsilon$, and hence Seller's profit conditional on trade is at least $\varepsilon/2 > 0$. It follows that $\underline{\pi}_s^{us} > 0$.

DEFINITION 2. ν is an *isoprofit distribution* (IPD) if

$$\int_p^{\bar{v}} (p - c(s))\nu(ds) = \text{constant} \geq 0, \quad \forall p \in \text{Supp}(\nu).$$

The [Appendix](#) provides a “greedy” algorithm to compute IPDs; the algorithm is defined for finite V , and we take limits to handle the infinite case. We can sketch how the algorithm works and construct a set of IPDs that average to the prior. Suppose $V = \{v_1, v_2, \dots, v_K\}$, with $v_i < v_{i+1}$ for $i \in \{1, \dots, K-1\}$ 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 IPD. Choose the maximum mass of type v_K for which this works. Remove that IPD—that is, take the conditional distribution after removing the masses of each type according to that IPD—and then repeat the procedure to construct the next IPD.

Crucially, whenever an IPD is removed, the price corresponding to $\underline{\pi}_s^{fb}(\Gamma)$ remains optimal in the remaining “market”; this follows from the IPD’s defining property of Seller indifference and an accounting identity. Therefore, Seller’s profit in this segmentation of IPDs 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 IPDs implements point *B* and *C* in [Figure 1](#). The entire triangle *ABC* can then be implemented by convexification: randomizing over this information structure (and the two equilibria) and full information (where Seller obtains all the surplus).

REMARK 4. In the same vein as part 3 of [Theorem 2](#), one can also establish approximately unique implementation for [Theorem 3](#)’s payoff set: for any $(\pi_b, \pi_s) \in \Pi_{fb}^*(\Gamma)$ and any $\varepsilon > 0$, there is $\tau \in \mathbf{T}_{fb}$ with $\Pi(\Gamma, \tau) \subset B_\varepsilon(\pi_b, \pi_s)$.

REMARK 5. [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 $p \geq \bar{v}$). 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; that is, triangles *ABC* and *ADE* in [Figure 1](#) are distinct if and only if triangles *ABC* and *AFG* are distinct. Or to put it another way, when (and only when) $\underline{\pi}_s^{fb}(\Gamma) > \underline{\pi}_s(\Gamma)$ there is an uninformed Seller information structure that implements some $\pi_s < \underline{\pi}_s^{fb}(\Gamma)$.²² [Theorems 2–3](#) further imply that this property also characterizes when Buyer can benefit from not

²²Pick 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

being fully informed. In the context of monopoly pricing, that can be viewed as characterizing when the buyer can benefit from strategic learning (Roesler and Szentes (2017)) rather than market segmentation (Bergemann, Brooks, and Morris (2015)).

REMARK 6. Restricting attention to posted prices is without loss for Theorems 1–3. Since we have implemented all payoffs that are feasible and individually rational for Seller, allowing Seller to use more complicated mechanisms cannot enlarge the implementable payoff sets. To see why all our payoffs can still be obtained as well, note that for Theorem 1, Seller’s deviation to any other mechanism can simply be deterred by a pessimistic belief. For Theorems 2–3, since Buyer is more informed than Seller, a posted price is optimal for Seller among all mechanisms (Myerson (1981)).

4. DISCUSSION

This section discusses some extensions and refinements of our results.

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 in $\mathbb{R}^2$. The extension of our maintained assumption of commonly known gains from trade is: for all $(c, v) \in \text{Supp}(\mu)$, $v \geq c$; and $\mathbb{E}[v - c] > 0$. An information structure is now a joint distribution $P(t_b, t_s, c, v)$ whose marginal distribution on (c, v) is μ .

The substance of Theorems 1/1*, Theorem 2 and Theorem 3 still hold.²³ 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 profit by setting either a sufficiently high price or a price (arbitrarily close to) $\underline{v}$, regardless of her signal. Abusing notation, we can define a cost function $c(v') \equiv \mathbb{E}_\mu[c|v = v'] \leq v'$. This results in an environment satisfying all the maintained assumptions of our baseline model, except that $c(\cdot)$ may not be continuous. Such continuity plays no role in proving Theorem 1 nor 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 certain convergence arguments. However, in the two-dimensional type environment, $\mathbb{E}_\nu[c]$ is still a continuous function of $\nu \in \Delta(C \times V)$. Theorem 3 uses upper semicontinuity of Seller’s profit in price; boundedness of c and $c \leq v$ is sufficient for such upper semicontinuity.

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 $\underline{\pi}_s^{fb}(\Gamma)$.

²³A caveat is that in this multidimensional setting, we do not know whether our maintained assumption that Seller only posts a price is without loss—that is, we do not rule out that certain payoff pairs are not implementable when Seller can use nonposted price mechanisms (which she might use to improve her payoff). By contrast, in our baseline one-dimensional setting, our results would not be affected if we had allowed Seller to use arbitrary mechanisms. See Che and Zhong (2025) and Deb and Roesler (2024) for work on information design in multidimensional screening problems.

4.2 Negative trading surplus

Returning to our baseline model, we next discuss what happens when 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

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

where $[\cdot]^+ \equiv \max\{\cdot, 0\}$. The function $S_\lambda(\Gamma)$ is a weighted sum of Buyer and Seller payoff assuming that trade occurs at price $p = \underline{v}$ whenever trade creates a positive weighted total payoff, and there is no trade otherwise. It is readily verified that $S_1(\Gamma) = \mathbb{E}[[v - c(v)]^+]$ and $\lim_{\lambda \rightarrow \infty} S_\lambda(\Gamma)/\lambda = \mathbb{E}[v] - \underline{v}$. Allowing negative trading surplus does not affect our definition of wpBE. So, the notation $\Pi(\Gamma)$ and $\underline{\pi}_s(\Gamma)$ still have the same meanings as before. The next proposition shows that $\Pi(\Gamma)$ 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.*

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

PROOF. See Appendix E. □

Figure 3 depicts Proposition 1. The blue triangle's frontier corresponds to total surplus under full trade.²⁴ The union of the blue and red regions is the set $\Pi(\Gamma)$. Each outer blue line has a slope $-\lambda$, with $\lambda \geq 1$, and represents a frontier $\lambda \pi_b + \pi_s = S_\lambda(\Gamma)$; the frontier of the red region is defined by their envelope.

Let us explain some of the logic underlying Proposition 1/Figure 3. Begin by observing that all the payoffs in the figure's blue triangle can be implemented analogously to our discussion of Theorem 1. It is also straightforward that some payoffs outside this set can be implemented. In particular, an information structure that publicly reveals only whether trade is efficient (i.e., whether $v \geq c(v)$ or not) can implement efficient trade with all the surplus accruing to Seller: the point $(0, S_1(\Gamma))$ in Figure 3. Why does maximizing Buyer's payoff now generally require some inefficiency (i.e., why is the red region's frontier not linear when $S_1(\Gamma) > \mathbb{E}[v - c(v)]$)? Consider, for simplicity, $\underline{v} \geq \mathbb{E}[c(v)]$, so that $\underline{\pi}_s(\Gamma) = \underline{v} - \mathbb{E}[c(v)]$. The bottom-right corner of Figure 3's blue triangle is then achieved by having trade with probability one at the price $\underline{v}$, with corresponding Buyer payoff $\mathbb{E}[v] - \underline{v}$. No higher Buyer payoff is implementable because Seller will never sell at

²⁴The figure is drawn assuming $\mathbb{E}[v - c(v)] > 0$, which ensures the blue triangle in the figure is nondegenerate. If instead $\mathbb{E}[v - c(v)] \leq 0$, then $\underline{\pi}_s(\Gamma) = \max[\underline{v} - \mathbb{E}[c(v)], 0] = 0$, and the blue triangle would be the singleton $(0, 0)$.




FIGURE 3. Outcome when trading surplus can be negative.

a price below $\underline{v}$, and subject to that constraint; this outcome maximizes $v - p$ for every v . In other words, the maximum implementable Buyer's payoff goes hand in hand with implementing all inefficient trade.

It remains to sketch why the frontier of $\Pi(\Gamma)$ is characterized by the lines defined by $\{\lambda\pi_b + \pi_s = S_\lambda(\Gamma)\}_{\lambda \geq 1}$. When type v trades at price $p \geq \underline{v}$, the (ex post) weighted total payoff is $p - c(v) + \lambda(v - p)$, which is weakly below $\underline{v} - c(v) + \lambda(v - \underline{v})$ because $\lambda \geq 1$ and $\underline{v} \leq p$. When trade does not happen, the weighted total payoff is 0. Therefore, the weighted total payoff is bounded above by $[\underline{v} - c(v) + \lambda(v - \underline{v})]^+$, and hence each $S_\lambda(\Gamma)$ is an upper bound for the expected weighted total payoff. The proof of Proposition 1 establishes that each of these upper bounds is tight: for each λ , there exists an information structure implementing expected weighted total payoff equal to $S_\lambda(\Gamma)$. The information structure simply publicly reveals whether the weighted total payoff at price $\underline{v}$ is negative (a “negative signal”) or not (a “positive signal”). For all v , such that $\underline{v} - c(v) + \lambda(v - \underline{v}) < 0$, it holds that $v - c(v) < (1 - \lambda)(v - \underline{v}) \leq 0$. Thus, the negative signal creates common knowledge that total surplus is negative, and hence there is no trade. After a positive signal, instead, Seller can be induced to sell at price $\underline{v}$ just as in the discussion of Theorem 1. Therefore, the equilibrium expected weighted total payoff is $\mathbb{E}[[\underline{v} - c(v) + \lambda(v - \underline{v})]^+] = S_\lambda(\Gamma)$.

We should note that in certain cases there may not be a tradeoff between maximizing Buyer's payoff and efficiency. Specifically, consider the profit level $\hat{\pi}_s$ that is the maximum of 0 and Seller's profit from efficient trade at price $\underline{v}$:

$$\hat{\pi}_s \equiv \left[\int \mathbf{1}_{v \geq c(v)} (\underline{v} - c(v)) \mu(dv) \right]^+.$$

For profit levels above $\hat{\pi}_s$, the situation is analogous to our baseline model once we use a public signal to reveal that trade is efficient, and so the implementable equilibrium payoffs with $\pi_s \geq \hat{\pi}_s$ constitute a triangle, as seen in Figure 3. Hence, if $\hat{\pi}_s = 0$, then

there is no tradeoff between efficiency and maximizing Buyer's payoff. But when $\widehat{\pi}_s > 0$, then so long as some trades generate negative surplus, $\widehat{\pi}_s > \pi_s$ and there is a tradeoff.

4.3 Affine cost function

Returning to our baseline model, recall that Seller's minimum implementable payoff $\pi_s^{us}(\Gamma)$ under price independent beliefs (Theorem 2) is not amenable to a closed-form formula in general. We now provide such a formula when the cost function $c(v)$ is affine. According to our discussion in Section 4.1, an affine $c(v)$ subsumes richer environments in which conditional expectations are affine, such as under Gaussian primitives (cf. Bar-Isaac, Jewitt, and Leaver (2021)).

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

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

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

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. We focus on a special family of IPDs (see Definition 2) whose supports are intervals $[v_*, v^*]$. Such IPDs have an analytical expression under Condition 1:

$$G(v) = \begin{cases} 0 & \text{if } v \leq v_* \\ 1 & \text{if } v \geq v^* \\ 1 - \left(\frac{(1-\lambda)v - \gamma}{(1-\lambda)v_* - \gamma} \right)^{\frac{1}{\lambda-1}} & \text{if } v \in (v_*, v^*) \text{ and } \lambda \neq 1 \\ 1 - e^{-\frac{v-v_*}{\gamma}} & \text{if } v \in (v_*, v^*) \text{ and } \lambda = 1. \end{cases} \quad (3)$$

$G(v)$ is smooth everywhere except for a mass point at v^* . Our maintained assumption that $v \geq c(v)$ implies $(1-\lambda)v \geq \gamma$. Therefore, Equation (3) defines an increasing function, that is, a well-defined CDF. Given any v_* , a higher v^* corresponds to increasing $G(v)$ in the sense of first-order stochastic dominance; hence, there is a unique v^* determined by the condition $\mathbb{E}_G[v] = \mathbb{E}_F[v]$. So, the family of IPDs is parametrized by a single parameter v_* ; accordingly, we denote such an IPD by CDF G_{v_*} with density g_{v_*} on (v_*, v^*) . The domain for v_* is $[\underline{v}, \mathbb{E}[v]]$. We separately define $G_{\mathbb{E}[v]}(v) = \mathbf{1}_{v \geq \mathbb{E}[v]}$.

It can be verified that a higher v_* lowers the corresponding v^* (i.e., the corresponding intervals $[v_*, v^*]$'s are nested) and that $G_{v_*}(v)$ is pointwise decreasing in v_* within the common support. As a result, for any two different v_* 's the corresponding G_{v_*} 's cross once. Consequently, all IPDs are ordered according to the MPC order (increasing in $v_* \in [\underline{v}, \mathbb{E}[v]]$). Not every G_{v_*} is in $D(\mu)$, but $G_{v_*} \in D(\mu)$ for all v_* larger than some threshold. The following proposition shows that this threshold pins down π_s^{us} .

PROPOSITION 2. *Assume Condition 1 and let $p_* \equiv \min\{v_* | G_{v_*} \in D(\mu)\}$. It holds that $\pi_s^{us}(\Gamma) = p_* - \mathbb{E}_\mu[c(v)]$.*

PROOF. See Appendix F. □

Proposition 2 states that an uninformed Seller's minimum payoff (with price independent beliefs) is characterized by a specific IPD of Buyer posterior means. By the Seller indifference property of IPDs and that v_* is the minimum of G_{v_*} 's support, Seller's profit when facing IPD G_{v_*} is $v_* - \mathbb{E}[c(v)]$. Proposition 2 thus implies that an uninformed Seller's minimum payoff is implemented by the most dispersed IPD 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 so that the posterior mean distribution becomes an IPD makes Seller weakly worse off. As such, it is without loss to only consider IPDs to implement $\pi_s^{us}(\Gamma)$. This makes the problem one-dimensional and tractable. To elaborate on the key step, suppose we find a distribution $G(v)$ such that: (i) $G \in D(\mu)$; (ii) G is an IPD; and (iii) there is a $p \in \text{Supp}(G)$ such that $\int_p^{\bar{v}} G(s) ds = \int_p^{\bar{v}} F(s) ds$. Such a G is the most dispersed IPD that is a MPC of the prior μ . Consider the following two identities derived using integration by parts:²⁵

$$\lambda \int_p^{\bar{v}} F(s) ds = -(1 - F(p-))(p - c(p)) + \lambda(\bar{v} - p) + \int_p^{\bar{v}} (p - c(s)) dF(s), \quad (4)$$

$$\lambda \int_p^{\bar{v}} G(s) ds = -(1 - G(p-))(p - c(p)) + \lambda(\bar{v} - p) + \int_p^{\bar{v}} (p - c(s)) dG(s). \quad (5)$$

First, by property (iii) above, the LHS of Equation (4) equals the LHS of Equation (5). Second, the MPC condition implies $\int_v^{\bar{v}} (F(s) - G(s)) ds \leq 0$ for all v and it reaches 0 when $v = p$, it holds that $F(p-) \leq G(p-)$. Therefore, the integral term on RHS of Equation (4) must be greater than that of Equation (5). Notice that the integral term on either 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 Seller's maximum profit given $F(v)$. At the same time, by the IPD property, p is an optimal price given $G(v)$. Therefore, the optimal profit under valuation distribution F must be no lower than that under G . It follows that to minimize Seller's profit, it is without loss to consider only IPDs within the set $D(\mu)$, which is a one-dimensional subspace.

The logic above generalizes that of [Roesler and Szentes \(2017\)](#). Their 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](#).²⁶

If the prior μ has binary support, then any cost function $c(v)$ is affine. This case permits an explicit solution for $\pi_s^{us}(\Gamma)$.

²⁵ $F(p-)$ is defined as the left limit of F at p , and similarly $G(p-)$. The integration by parts formula is for Lebesgue–Stieltjes integral.

²⁶In recent work, [Inostroza and Tsoy \(2025\)](#) extend Proposition 2 to a setting in which Seller does not have all the bargaining power.

COROLLARY 1. Assume μ has binary support: $V = \{v_1, v_2\}$ with $v_1 < v_2$. Let $\lambda = (c(v_2) - c(v_1))/(v_2 - v_1)$, and let p be the unique solution to

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

It holds that $\underline{\pi}^{us}(\Gamma) = \max\{p, v_1\} - \mathbb{E}[c(v)]$.

PROOF. See Appendix G. □

APPENDIX A: PROOF OF THEOREM 1

PROOF. We first show that $\{(\pi_b, \pi_s) : \pi_b \geq 0, \pi_s \geq \underline{\pi}_s(\Gamma), \pi_b + \pi_s \leq S(\Gamma)\} \subset \Pi(\Gamma)$. Consider a trivial information structure τ_0 in which both player's signal spaces are singletons. Note that τ_0 has more informed Buyer, but since we are interested in Π rather than Π^* , we do not require Buyer's belief to be price independent. For any $(\pi_b, \pi_s) \in \Pi(\Gamma)$, define strategies and beliefs as follows. Let $p_l = \pi_s + \mathbb{E}[c(v)] \in [\underline{v}, \mathbb{E}[v]]$ and $p_h = \mathbb{E}[v] > \mathbb{E}[c(v)]$.

- Buyer's strategy:

$$\alpha(p_h) = \frac{p_l - \mathbb{E}[c(v)]}{p_h - \mathbb{E}[c(v)]}, \quad \alpha(p_l) = 1, \quad \text{and} \quad \alpha(p) = \mathbf{1}_{p \leq \underline{v}} \quad \forall p \notin \{p_l, p_h\},$$

where $\mathbf{1}_{p \leq \underline{v}}$ denotes the indicator function for the set $\{p : p \leq \underline{v}\}$.

- Seller's strategy:

$$\sigma(p_l) = \frac{\pi_b}{\mathbb{E}[v] - p_l} \quad \text{and} \quad \sigma(p_h) = 1 - \frac{\pi_b}{\mathbb{E}[v] - p_l}.$$

Note that $\mathbb{E}[v] - p_l = S(\Gamma) - \pi_s \geq \pi_b$ guarantees that $\sigma(p_l), \sigma(p_h) \in [0, 1]$.

- Beliefs:

$$\nu(v|p_l) = \nu(v|p_h) = \mu(v) \quad \text{and} \quad \nu(v|p) = \delta_{\underline{v}}(v) \quad \forall p \notin \{p_l, p_h\}.$$

It is straightforward that the payoff from this strategy profile is (π_b, π_s) . So, we need only verify that (σ, α, ν) constitutes a wpBE. First, Buyer's strategy is optimal given beliefs because $\mathbb{E}[v] - p_h = 0$, $\mathbb{E}[v] - p_l \geq 0$, and for any other price, Buyer's belief is a point mass on $\underline{v}$. Second, Seller's strategy is optimal: $\alpha(p_h)$ is defined such that $\alpha(p_h)(p_h - \mathbb{E}[c(v)]) = p_l - \mathbb{E}[c(v)]$. So, Seller is indifferent between offering p_l and p_h . Any other price above $\underline{v}$ is rejected and so is no better than p_l and p_h . Seller's payoff is $p_l - \mathbb{E}[c(v)] = \pi_s \geq \underline{\pi}_s(\Gamma)$, so any price below $\underline{v}$ is also no better. Third, since Seller's strategy is type independent and $\nu = \mu$ on path, Bayes rule is satisfied on path.

It remains to prove that $\Pi(\Gamma) \subset \{(\pi_b, \pi_s) : \pi_b \geq 0, \pi_s \geq \underline{\pi}_s(\Gamma), \pi_b + \pi_s \leq S(\Gamma)\}$. Pick any signal structure and any wpBE with belief ν . Since $\text{Supp}(\nu) \subset V$, sequential rationality implies that Buyer buys with probability one after any price $p < \underline{v}$ and with probability zero when $p > \bar{v}$. Therefore, Seller must obtain payoff $\pi_s \geq \underline{\pi}_s(\Gamma) \equiv \max\{\underline{v} - \mathbb{E}[c(v)], 0\}$. It is straightforward that Buyer's payoff $\pi_b \geq 0$ and $\pi_b + \pi_s \leq S(\Gamma)$. □

APPENDIX B: PROOF OF THEOREM 1*

We first prove Proposition B.1 below, which we will use to prove Theorem 1*.

B.1 A related result

PROPOSITION B.1. *Fix any $\varepsilon > 0$. $\exists$ a finite information structure such that $\forall (\pi_b, \pi_s) \in \{(\pi_b, \pi_s) : \pi_b \geq \varepsilon, \pi_s \geq \underline{\pi}_s(\Gamma) + \varepsilon, \text{ and } \pi_b + \pi_s \leq S(\Gamma) - \varepsilon\}$, $\exists$ a finite price grid defining a game that has a sequential equilibrium with payoffs (π_s, π_b) .*

One aspect of this result is weaker than Theorem 1* because the price grid here varies with the equilibrium payoffs (π_b, π_s) . But another aspect is stronger: all payoffs ε away from the boundary of $\Pi(\Gamma)$ are obtained, rather than just an ε -net of payoffs.

PROOF. Let us initially prove the statement assuming $\mathbb{P}(\underline{v}) > 0$. First, choose any $\delta \in (0, \frac{1}{2})$ and any $\eta \in (0, \mathbb{P}(v = \underline{v}))$. For now, we keep δ and η as free parameters and we define the information structure and the corresponding equilibrium. At the end of the proof, we will verify that when δ and η go to zero, the equilibrium payoffs span the target set of payoffs in Proposition B.1.

We first define the information structure. Buyer is uninformed: $T_b = \{\emptyset\}$. Seller gets two signals: $T_s = \{l, h\}$, with distribution given by

$$P(l, \emptyset, v) = \eta \delta_{\underline{v}}(v),$$

$$P(h, \emptyset, v) = \mu(v) - \eta \delta_{\underline{v}}(v).$$

That is, $v = \underline{v}$ is revealed to Seller with probability η using signal “ l .” In the rest of the proof, we omit t_b , as Buyer is uninformed.

We next specify certain prices and a property of the finite price grid. Choose any $\sigma_h \in [\delta, 1 - \delta]$. Define

$$\bar{p}_h = \frac{\mathbb{E}[v] - \eta \underline{v}}{1 - \eta} \quad \text{and} \quad \underline{p}_h = \frac{\eta \underline{v} + (1 - \eta) \sigma_h \bar{p}_h}{\eta + (1 - \eta) \delta}.$$

That is, $\bar{p}_h$ is $\mathbb{E}[v|t_s = h]$ and $\underline{p}_h$ is the expectation of v conditional on the event that pools σ_h proportion of $t_s = h$ with all $t_s = l$. Pick any $p_l \in [\max\{\frac{\mathbb{E}[c(v)] - \eta c(\underline{v})}{1 - \eta}, \underline{v}\}, \frac{\eta \underline{v} + (1 - \eta) \delta \bar{p}_h}{\eta + (1 - \eta) \delta}]$.²⁷ It holds that $\underline{v} \leq p_l < \underline{p}_h < \mathbb{E}[v] < \bar{p}_h$. Consider any finite grid of prices that contains $\{p_l, \underline{p}_h, \bar{p}_h\}$, and is otherwise arbitrary.

Now we specify the strategy profile and beliefs, and verify equilibrium.

- *Case 1:* $c(\underline{v}) > \mathbb{E}[c(v)]$. Seller's and Buyer's strategies, σ and α , and Buyer's beliefs ν are, respectively,

$$\sigma(\underline{p}_h|h) = \sigma_h, \quad \sigma(p_l|h) = 1 - \sigma_h, \quad \text{and} \quad \sigma(\underline{p}_h|l) = 1;$$

²⁷Fixing $\delta > 0$, the interval is nonempty when η is sufficiently close to 0.

$$\alpha(\underline{p}_h) = \frac{p_l - \frac{\mathbb{E}[c(v)] - \eta c(\underline{v})}{1 - \eta}}{\underline{p}_h - \frac{\mathbb{E}[c(v)] - \eta c(\underline{v})}{1 - \eta}}, \quad \alpha(p_l) = 1, \quad \text{and} \quad \forall p \notin \{p_l, \underline{p}_h\}, \alpha(p) = \mathbf{1}_{p \leq \underline{v}};$$

$$\nu(v|\underline{p}_h) = \frac{\eta(1 - \sigma_h)\delta_{\underline{v}}(v) + \sigma_h\mu(v)}{\eta(1 - \sigma_h) + \sigma_h},$$

$$\nu(v|p_l) = \frac{\mu(v) - \eta\delta_{\underline{v}}(v)}{1 - \eta}, \quad \text{and} \quad \forall p \notin \{p_l, \underline{p}_h\}, \quad \nu(v|p) = \delta_{\underline{v}}(v).$$

That is, Seller with signal h randomizes between prices $\underline{p}_h$ and p_l , while after signal l she chooses $\underline{p}_h$. Buyer randomizes after price $\underline{p}_h$, accepts p_l , and off-path accepts prices below $\underline{v}$ and rejects otherwise.

Let us verify that (σ, α, ν) is a sequential equilibrium. Buyer's sequential rationality is straightforward, as $\mathbb{E}_\nu[v|p_l] = \bar{p}_h > p_l$, $\mathbb{E}_\nu[v|\underline{p}_h] = \underline{p}_h$ and $\mathbb{E}_\nu[v|p] = \underline{v}$ for any other p . For Seller, note that by definition of $\alpha(\underline{p}_h)$, Seller with signal h is indifferent between offering p_l and $\underline{p}_h$. Since $c(\underline{v}) > \mathbb{E}[c(v)]$ by hypothesis, Seller with signal l finds it strictly better offering $\underline{p}_h$ than p_l . Any other price p is worse than offering p_l for both Seller types. Finally, for consistency of Buyer's belief: Bayes' rule is straightforward on the equilibrium path. The off-path belief can be derived from the limit of Seller's fully mixed strategy $\tilde{\sigma}_n(p|h) = \frac{n^2-1}{n^2} \sigma(p|h) + \frac{1}{n^2 \times k}$ and $\tilde{\sigma}_n(p|l) = \frac{n-1}{n} \sigma(p|l) + \frac{1}{n \times k}$, where k is the number of prices in the grid.

- *Case 2:* $c(\underline{v}) \leq \mathbb{E}[c(v)]$. Now consider

$$\sigma(\bar{p}_h|h) = \sigma_h, \quad \sigma(p_l|h) = 1 - \sigma_h, \quad \text{and} \quad \sigma(p_l|l) = 1;$$

$$\alpha(\bar{p}_h) = \frac{p_l - \frac{\mathbb{E}[c(v)] - \eta c(\underline{v})}{1 - \eta}}{\bar{p}_h - \frac{\mathbb{E}[c(v)] - \eta c(\underline{v})}{1 - \eta}}, \quad \alpha(p_l) = 1, \quad \text{and} \quad \forall p \notin \{p_l, \bar{p}_h\}, \alpha(p) = \mathbf{1}_{p \leq \underline{v}};$$

$$\nu(v|\bar{p}_h) = \frac{\mu(v) - \eta\delta_{\underline{v}}(v)}{1 - \eta},$$

$$\nu(v|p_l) = \frac{\eta\sigma_h\delta_{\underline{v}}(v) + (1 - \sigma_h)\mu(v)}{\eta\sigma_h + (1 - \sigma_h)}, \quad \text{and} \quad \forall p \notin \{p_l, \bar{p}_h\}, \quad \nu(v|p) = \delta_{\underline{v}}(v).$$

That is, Seller with signal h randomizes between prices $\bar{p}_h$ and p_l , while after signal l she chooses p_l . Buyer randomizes after price $\bar{p}_h$, accepts p_l , and off-path accepts prices below $\underline{v}$ and rejects otherwise.

Let us verify that (σ, α, ν) is a sequential equilibrium. Buyer's sequential rationality is straightforward, as $\mathbb{E}_\nu[v|p_l] \geq \frac{\eta\underline{v} + (1-\eta)\delta_{\underline{p}_h}}{\eta + (1-\eta)\delta} > p_l$, $\mathbb{E}_\nu[v|\bar{p}_h] = \bar{p}_h$, and $\mathbb{E}_\nu[v|p] = \underline{v}$ for any other p . For Seller, note that by definition of $\alpha(\bar{p}_h)$, Seller with signal h is indifferent between offering p_l and $\bar{p}_h$. Since $c(\underline{v}) \leq \mathbb{E}[c(v)]$ by hypothesis, Seller with signal l finds it weakly better offering p_l than $\bar{p}_h$. Any other price is worse than offering p_l for both Seller types. Finally, for consistency of Buyer's belief: Bayes'

rule is straightforward on the equilibrium path. The off-path belief can be derived from the limit of Seller's fully mixed strategy $\tilde{\sigma}_n(p|h) = \frac{n^2-1}{n^2} \sigma(p|h) + \frac{1}{n^2 \times k}$ and $\tilde{\sigma}_n(p|l) = \frac{n-1}{n} \sigma(p|l) + \frac{1}{n \times k}$, where k is the number of prices in the grid.

Now we calculate the players' payoffs in the above equilibria.

- *Case 1:* $c(\underline{v}) > \mathbb{E}[c(v)]$. In this case, it is optimal for Seller to offer p_l after signal h and $\underline{p}_h$ after signal l . Therefore, in equilibrium

$$\pi_s = (1 - \eta) p_l + \eta \underline{p}_h - \mathbb{E}[c(v)] - \eta(1 - \alpha(\underline{p}_h))(\underline{p}_h - c(\underline{v})).$$

Note that $\underline{p}_h$ depends on σ_h but p_l does not. For any $\sigma_h \in [\delta, 1 - \delta]$, when $p_l = \max\{\frac{\mathbb{E}[c(v)] - \eta c(\underline{v})}{1 - \eta}, \underline{v}\}$,

$$\begin{aligned} \pi_s &\leq p_l - \mathbb{E}[c(v)] + \eta(\underline{p}_h - p_l) \\ &\leq \max\left\{\frac{\eta}{1 - \eta}(\mathbb{E}[c(v)] - c(\underline{v})), \underline{v} - \mathbb{E}[c(v)]\right\} + \eta(\mathbb{E}[v] - \underline{v}). \end{aligned}$$

When $p_l = \frac{\eta \underline{v} + (1 - \eta) \delta \bar{p}_h}{\eta + (1 - \eta) \delta}$,

$$\pi_s \geq p_l - \mathbb{E}[c(v)] - \eta(\mathbb{E}[v] - \mathbb{E}[c(v)]) = \frac{\eta \underline{v} + (1 - \eta) \delta \bar{p}_h}{\eta + (1 - \eta) \delta} - \mathbb{E}[c(v)] - \eta S(\Gamma).$$

Therefore, when p_l traverses its domain, π_s traverses a set containing the interval

$$\begin{aligned} I_s &= \left[\max\left\{\frac{\eta}{1 - \eta}(\mathbb{E}[c(v) - c(\underline{v})]), \underline{v} - \mathbb{E}[c(v)]\right\} \right. \\ &\quad \left. + \eta(\mathbb{E}[v] - \underline{v}), \frac{\eta \underline{v} + (1 - \eta) \delta \bar{p}_h}{\eta + (1 - \eta) \delta} - \mathbb{E}[c(v)] - \eta S(\Gamma) \right). \end{aligned}$$

In other words, $\forall \pi_s \in I_s$ and $\forall \sigma_h \in [\delta, 1 - \delta]$, there exists $p_l(\sigma_h)$ such that Seller's payoff is π_s . Now consider Buyer's payoff holding π_s fixed. When σ_h traverses $[\delta, 1 - \delta]$, π_b changes continuously. If $\sigma_h = 1 - \delta$, then $\pi_b = (1 - \eta) \delta (\bar{p}_h - p_l) \leq \delta(\bar{v} - \underline{v})$. If $\sigma_h = \delta$, then with at most $\eta + \delta - \eta \delta$ probability the offer is rejected, and hence $\pi_s + \pi_b \geq S(\Gamma) - (\eta + \delta)(\bar{v} - \inf c(v))$.

- *Case 2:* $c(\underline{v}) \leq \mathbb{E}[c(v)]$. In this case, it is optimal for both Seller to offer p_l no matter her signal, which induces Buyer to accept with probability 1. Therefore, in equilibrium $\pi_s = p_l - \mathbb{E}[c(v)]$. Buyer is indifferent between accepting the offer or not at $\bar{p}_h$. So, Buyer gets positive payoff only when the price offered is p_l , and hence $\pi_b = (\eta \sigma_h \underline{v} + (1 - \sigma_h) \mathbb{E}[v]) - (\eta \sigma_h + (1 - \sigma_h)) p_l$. Similar to Case 1, we can calculate that as p_l traverses its domain, π_s traverses the interval

$$\left[\max\left\{\frac{\eta}{1 - \eta}(\mathbb{E}[c(v) - c(\underline{v})]), \underline{v} - \mathbb{E}[c(v)]\right\}, \frac{\eta \underline{v} + (1 - \eta) \delta \bar{p}_h}{\eta + (1 - \eta) \delta} - \mathbb{E}[c(v)] \right).$$

Holding any p_l fixed, as σ_h traverses the interval $[\delta, 1 - \delta]$, π_b traverses

$$[\eta(1 - \delta) \underline{v} + \delta \mathbb{E}[v] - (\eta(1 - \delta) + \delta) p_l, \eta \delta \underline{v} + (1 - \delta) \mathbb{E}[v] - (\eta \delta + (1 - \delta)) p_l].$$

It follows that in either case, as η and δ converge to zero (with the order η first and δ second), Seller's payoff that obtain across the family of equilibria we have constructed converges to $(\max\{0, \underline{\pi}_s(\Gamma)\}, S(\Gamma))$. For any such π_s , Buyer's payoff that obtain converges (uniformly) to $(0, S(\Gamma) - \pi_s)$. This completes the proof of Proposition B.1 when $\mathbb{P}(\underline{v}) > 0$.

When $\mathbb{P}(\underline{v}) = 0$, we first modify the original environment by pooling a small mass of valuations near $v = \underline{v}$ (which is feasible since $\underline{v}$ is the lowest value in the support of V). Call this modified environment $\tilde{\Gamma}$. Plainly, $S(\tilde{\Gamma}) = S(\Gamma)$ and $\underline{\pi}_s(\tilde{\Gamma}) \approx \underline{\pi}_s(\Gamma)$. Therefore, $\forall \varepsilon > 0$, there exists such $\tilde{\Gamma}$ such that $\Pi(\tilde{\Gamma})$ covers all payoffs in $\Pi(\Gamma)$ that are more than $\frac{1}{2}\varepsilon$ away from the boundary of $\Pi(\Gamma)$. We can now apply the previous argument with a positive probability of the lowest valuation and find an information structure $\tilde{\tau}$ that implements all payoffs in $\Pi(\tilde{\Gamma})$ that are more than $\frac{1}{2}\varepsilon$ away from the boundary of $\Pi(\tilde{\Gamma})$. The proof is completed by converting $\tilde{\tau}$ to an information structure for the original environment Γ . $\square$

B.2 Proof of Theorem 1*

PROOF. We utilize the construction in the proof of Proposition B.1. First, $\forall \varepsilon > 0$, choose δ and η as the corresponding parameters derived in Proposition B.1 such that the implementable payoffs cover all points $\varepsilon/2$ away from the boundary of $\Pi(\Gamma)$. Then $\bar{p}_h = \frac{\mathbb{E}[v] - \eta v}{1 - \eta}$. Choose grid size $\Delta \in (0, \frac{1}{2}|\bar{p}_h - \mathbb{E}[v]|)$. Construct an arbitrary grid of $[v, \bar{v}]$ with grid size Δ . By the definition of grid size, there exists an on-grid price $\bar{p}'_h \in [\bar{p}_h - \Delta, \bar{p}_h]$. Now choose $\eta' \leq \eta$ such that $\bar{p}'_h = \frac{\mathbb{E}[v] - \eta' v}{1 - \eta'}$. Note that $\bar{p}'_h \geq \bar{p}_h - \Delta$ implies $\frac{1 - \eta'}{\eta'(\mathbb{E}[v] - v)} \leq \frac{1}{\frac{\eta}{1 - \eta}(\mathbb{E}[v] - v) - \Delta}$. From now on, we fix Δ , η' , δ , $\bar{p}'_h$ and the grid.

Pick any $(\pi_b, \pi_s) \in \Pi(\Gamma)$ that is $\varepsilon/2$ away from the boundary of $\Pi(\Gamma)$. Note that reducing η to η' expands the set of implementable payoffs in Proposition B.1. Therefore, given δ , η' , the construction in Proposition B.1 defines an information structure such that (π_b, π_s) is an equilibrium payoff pair. Let (p_l, σ_h) define the constructed equilibrium.²⁸ Now we modify p_l and σ_h to “snap” the on-path prices onto the grid. Choose p'_l to be the on-grid price no greater than and closest to p_l . So, $p_l - p'_l < \Delta$. Let $\underline{p}'_h$ be the on-grid price closest to $\underline{p}_h$ such that $\sigma'_h = \frac{\eta'}{1 - \eta'} \frac{\underline{p}'_h - v}{\bar{p}'_h - \underline{p}'_h} \in [\delta, 1 - \delta]$ (note that since σ'_h is increasing in $\underline{p}'_h$, this is achieved by one of the two grid points to the left and right of $\underline{p}_h$). It can be easily verified that $p'_l < \underline{p}'_h < \mathbb{E}[v] < \bar{p}'_h$. Observe that

$$\begin{aligned} \left| \frac{d\sigma_h}{d\underline{p}'_h} \right| &= \frac{\eta'}{(1 - \eta')^2} \frac{\mathbb{E}[v] - v}{(\bar{p}'_h - \underline{p}'_h)^2} \leq \frac{\eta'}{(1 - \eta')^2} \frac{\mathbb{E}[v] - v}{(\bar{p}'_h - \mathbb{E}[v])^2} \\ &= \frac{1}{\eta'(\mathbb{E}[v] - v)} \leq \frac{1}{1 - \eta'} \frac{1}{\frac{\eta}{1 - \eta}(\mathbb{E}[v] - v) - \Delta}, \end{aligned}$$

where the last inequality is from $\frac{1 - \eta'}{\eta'(\mathbb{E}[v] - v)} \leq \frac{1}{\frac{\eta}{1 - \eta}(\mathbb{E}[v] - v) - \Delta}$. Therefore, $|p_h - p'_h| \leq \Delta$ implies $|\sigma_h - \sigma'_h| \leq \frac{1}{1 - \eta'} \frac{\Delta}{\frac{\eta}{1 - \eta}(\mathbb{E}[v] - v) - \Delta}$.

²⁸All other parameters defining the equilibrium are calculated from η' , p_l , σ_h . Recall that the on-path prices are p_l , $\bar{p}'_h$, $\underline{p}'_h$.

Take the information structure and equilibrium from the proof of Proposition B.1 corresponding to parameters η' , p'_l , and σ'_h . Now we calculate the equilibrium payoffs and compare that to (π_b, π_s) . We discuss the two cases separately:

In case 1 ($c(\underline{v}) > \mathbb{E}[c(v)]$), we first bound $|\alpha'(\underline{p}'_h) - \alpha(\underline{p}_h)|$:

$$\begin{aligned} & |\alpha'(\underline{p}'_h) - \alpha(\underline{p}_h)| \\ & \leq \frac{|p'_l - p_l|}{\underline{p}_h - \mathbb{E}[c(v)|v > \underline{v}]} + \left| \frac{p'_l - \mathbb{E}[c(v)|v > \underline{v}]}{\underline{p}_h - \mathbb{E}[c(v)|v > \underline{v}]} - \frac{p'_l - \mathbb{E}[c(v)|v > \underline{v}]}{\underline{p}_h - \mathbb{E}[c(v)|v > \underline{v}]} \right| \\ & \leq \frac{2\Delta}{\underline{p}_h - \mathbb{E}[c(v)|v > \underline{v}]} \end{aligned}$$

The second inequality follows from $\underline{p}'_h > p'_l$ and $|p'_l - p_l|, |\underline{p}'_h - \underline{p}_h| < \Delta$. In this case, Seller's payoff is $\alpha'(\underline{p}'_h)(\underline{p}'_h - \mathbb{E}[c(v)])$ (note that Seller always finds $\underline{p}'_h$ optimal, which is accepted with probability α'). Therefore,

$$|\pi_s - \pi'_s| \leq |\alpha'(\underline{p}'_h) - \alpha(\underline{p}_h)|(\underline{p}_h - \mathbb{E}[c(v)]) + \alpha'(\underline{p}'_h)|\underline{p}'_h - \underline{p}_h| \leq 3\Delta,$$

where the last inequality uses $\mathbb{E}[c(v)] > \mathbb{E}[c(v)|v > \underline{v}]$. Buyer's payoff is $\pi'_b = \mathbb{E}[v] - \underline{p}'_h + (1 - \eta')(1 - \sigma'_h)(\underline{p}'_h - p'_l)$ (note that Buyer always finds accepting the on-path prices optimal). Therefore,

$$\begin{aligned} |\pi_b - \pi'_b| & \leq |\underline{p}_h - \underline{p}'_h| + (1 - \eta')(|\sigma'_h - \sigma_h|(\underline{p}_h - p_l) + (1 - \sigma'_h)(|\underline{p}_h - \underline{p}'_h| + |p_l - p'_l|)) \\ & \leq 3\Delta + (\bar{v} - \underline{v}) \frac{1}{1 - \eta'} \frac{\Delta}{\frac{\eta}{1 - \eta}(\mathbb{E}[v] - \underline{v}) - \Delta}. \end{aligned}$$

In case 2 ($c(\underline{v}) \leq \mathbb{E}[c(v)]$), Seller finds it optimal to always offer p_l ; hence, $\pi_s = p_l - \mathbb{E}[c(v)]$. Therefore,

$$|\pi_s - \pi'_s| \leq |p_l - p'_l| \leq \Delta.$$

Buyer's payoff is $\pi'_b = \mathbb{E}[v] - p'_l - (1 - \eta')\sigma'_h(\bar{p}'_h - p'_l)$ (note that Buyer always finds accepting the on-path prices optimal). Therefore,

$$\begin{aligned} |\pi_b - \pi'_b| & \leq |p_l - p'_l| + (1 - \eta')(|\sigma'_h - \sigma_h|(\underline{p}_h - p_l) + \sigma'_h(|\underline{p}_h - \underline{p}'_h| + |p_l - p'_l|)) \\ & \leq 3\Delta + (\bar{v} - \underline{v}) \frac{1}{1 - \eta'} \frac{\Delta}{\frac{\eta}{1 - \eta}(\mathbb{E}[v] - \underline{v}) - \Delta}. \end{aligned}$$

In either case,

$$\|(\pi_b, \pi_s) - (\pi'_b, \pi'_s)\| \leq 6\Delta + \frac{2\Delta(\bar{v} - \underline{v})}{\eta(\mathbb{E}[v] - \underline{v}) - (1 - \eta)\Delta}.$$

By choosing Δ sufficiently small, we bound $\|(\pi_b, \pi_s) - (\pi'_b, \pi'_s)\|$ above by ε .

To summarize, $\forall \varepsilon > 0$, there exist parameters δ, η, η' , and Δ such that for an arbitrary grid of $[\underline{v}, \bar{v}]$ with grid size Δ , for any $(\pi_b, \pi_s) \in \Pi(\Gamma)$, we construct an information structure with sequential equilibrium payoff within the ε -neighborhood of (π_b, π_s) . That is, the set of payoffs from sequential equilibria corresponding to some information structure is an ε -net of $\Pi(\Gamma)$. $\square$

APPENDIX C: PROOF OF THEOREM 2

We prove the theorem via three lemmas.

LEMMA C.1. $\forall (\pi_b^*, \pi_s^*) \in \Pi_{us}^*(\Gamma)$, $\forall \pi_s \in [\pi_s^*, S(\Gamma)]$, and $\pi_b \in [0, S(\Gamma) - \pi_s]$, there exists $\tilde{\tau} \in \mathbf{T}_{us}$ such that $(\pi_b, \pi_s) \in \Pi^*(\Gamma, \tilde{\tau})$.

In words, this lemma says that the set $\Pi_{us}^*(\Gamma)$ consists of all payoff pairs in $\Pi(\Gamma)$ such that Seller's payoff is above some floor. By definition, the floor is $\underline{\pi}_{us}^*(\Gamma)$ defined in (1). Hart and Reny (2019, Theorem 12, part (2)) implies that the floor is achieved. Hence, the lemma implies part 2 of Theorem 2.

PROOF. Let $\gamma = S(\Gamma) - (\pi_b + \pi_s)$ be the loss of total surplus for payoff pair (π_b, π_s) . We construct an information structure $\tilde{\tau}$ such that the efficiency loss is γ , Seller's payoff is π_s and Buyer's payoff is π_b . Let $P(t_b, v)$ be the joint distribution specified by an Seller-uninformed information structure τ for which $(\pi_b^*, \pi_s^*) \in \Pi^*(\Gamma, \tau)$.

First, we determine the types that are not traded. For this, we find a threshold value z^* such that trading all expected valuations strictly below z^* and some fraction of expected valuation z^* generates surplus γ . Consider the function

$$y(z) = \int_{\mathbb{E}[v|t_b] < z} (v - c(v))P(dt_b, dv),$$

which is well-defined because the domain of integration is measurable. The set $\{t_b | \mathbb{E}[v|t_b] < z\}$ expands when z increases. So, $y(z)$ is increasing in z . Moreover, $y(\infty) = S(\Gamma)$ and $y(-\infty) = 0$. So, there exists z^* such that $y(z) \leq (\geq) \gamma$ for $z < (>) z^*$. By definition, $\{t_b | \mathbb{E}[v|t_b] < z\} = \bigcup_{\varepsilon > 0} \{t_b | \mathbb{E}[v|t_b] < z - \varepsilon\}$, so $y(z)$ is a left-continuous function. Hence, $y(z^*) \leq \gamma$. Define β by

$$\gamma = y(z^*) + \beta \int_{\mathbb{E}[v|t_b] = z^*} (v - c(v))P(dt_b, dv)$$

The RHS is $y(z^*) \leq \gamma$ when $\beta = 0$ and $\lim_{z \rightarrow z^{*+}} y(z) \geq \gamma$ when $\beta = 1$. So, $\beta \in [0, 1]$. In words, excluding all t_b that induces $\mathbb{E}[v|t_b] < z^*$ and β portion of t_b inducing $\mathbb{E}[v|t_b] = z^*$ leads to efficiency loss γ .

Next, we construct $\tilde{\tau} \in \mathbf{T}_{us}$ such that all remaining surplus is realized and Seller gets payoff π_s . If Seller sells at price p and trades with all remaining types, Seller's payoff is

$$\int_{\mathbb{E}[v|t_b] > z^*} (p - c(v))P(dt_b, dv) + (1 - \beta) \int_{\mathbb{E}[v|t_b] = z^*} (p - c(v))P(dt_b, dv).$$

Therefore, Seller's payoff is π_s when trading with all remaining types at the price²⁹

$$p^* = \frac{\pi_s + \int_{\mathbb{E}[v|t_b] > z^*} c(v)P(dt_b, dv) + (1 - \beta) \int_{\mathbb{E}[v|t_b] = z^*} c(v)P(dt_b, dv)}{\int_{\mathbb{E}[v|t_b] > z^*} P(dt_b, dv) + (1 - \beta) \int_{\mathbb{E}[v|t_b] = z^*} P(dt_b, dv)}.$$

To ensure that all nonexcluded Buyer types accept price p^* , we construct $\tilde{\tau}$ by pooling all nonexcluded types such that $\mathbb{E}[v|t_b] < p^*$ and a λ fraction of those with signal t_b such that $\mathbb{E}[v|t_b] \geq p^*$. The fraction λ is determined as follows:

$$\begin{aligned} & \lambda \int_{\mathbb{E}[v|t_b] \geq p^*} (v - p^*)P(dt_b, dv) + \int_{z^* < \mathbb{E}[v|t_b] < p^*} (v - p^*)P(dt_b, dv) \\ & + (1 - \beta) \int_{\mathbb{E}[v|t_b] = z^*} (v - p^*)P(dt_b, dv) = 0 \\ \implies \lambda = & \frac{\int_{z^* < \mathbb{E}[v|t_b] < p^*} (p^* - v)P(dt_b, dv) + (1 - \beta) \int_{\mathbb{E}[v|t_b] = z^*} (p^* - z^*)P(dt_b, dv)}{\int_{\mathbb{E}[v|t_b] > p^*} (v - p^*)P(dt_b, dv)}, \end{aligned}$$

where $\lambda \in [0, 1]$ follows from the fact that the LHS of the first equality traverses from negative to positive when λ traverses $[0, 1]$.

Let $\tilde{T}_b = T_b \cup \{t_\emptyset\}$, where $t_\emptyset$ is topologically disjoint from T_b . The information structure $\tilde{\tau} \in \mathbf{T}_{us}$ is given by the following distribution:

$$\tilde{P}(t_b, v) = \begin{cases} (1 - \lambda)P(t_b, v) & \forall t_b \text{ such that } \mathbb{E}[v|t_b] \geq p^* \\ P(t_b, v) & \forall t_b \text{ such that } \mathbb{E}[v|t_b] < z^* \\ \beta P(t_b, v) & \forall t_b \text{ such that } \mathbb{E}[v|t_b] = z^*; \end{cases}$$

$$\tilde{P}(t = t_\emptyset, v) = \lambda \int_{\mathbb{E}[v|t_b] \geq p^*} P(dt_b, v) + \int_{z^* < \mathbb{E}[v|t_b] < p^*} P(dt_b, v) + (1 - \beta) \int_{\mathbb{E}[v|t_b] = z^*} P(dt_b, v).$$

(It can be verified that $\tilde{P}$ defines a valid information structure.)

Now we define Buyer's strategy $\tilde{\alpha}$. Let α be Buyer's strategy corresponding to the wPBE of game (Γ, τ) . Define $\tilde{\alpha}(p, t_b) = \alpha(p, t_b)$ when $t_b \neq t_\emptyset$, and $\tilde{\alpha}(p, t_\emptyset) = \mathbf{1}_{p \leq p^*}$. Sequential rationality of $\tilde{\alpha}$ is straightforward.

It remains only to verify that pricing at p^* is optimal for Seller. There is no profitable deviation to any higher price because

$$\sup_{p > p^*} \int \tilde{\alpha}(p, t_b)(p - c(v))\tilde{P}(dt_b, dv) = (1 - \lambda) \int \alpha(p, t_b)(p - c(v))P(dt_b, dv) \leq (1 - \lambda)\pi_s^*.$$

There is no profitable deviation to any price lower than z^* because

$$\sup_{p \leq z^*} \int \tilde{\alpha}(p, t_b)(p - c(v))\tilde{P}(dt_b, dv) \leq \sup_{p \leq z^*} \int \alpha(p, t_b)(p - c(v))P(dt_b, dv) \leq \pi_s^*.$$

²⁹Note that $\pi_s \leq S(\Gamma) - \gamma$ guarantees that $p^* \geq z^*$.

By construction, there is no t_b that induces a belief with $\mathbb{E}_{\tilde{\nu}}[v|t_b] \in (z^*, p^*)$. Therefore, it is suboptimal for Seller to post any price in (z^*, p^*) . It follows that it is optimal (strictly optimal when $\pi_s > \pi_s^*$) for Seller to offer p^* and get payoff π_s . Buyer's payoff is all the remaining surplus: $S(\Gamma) - \gamma - \pi_s = \pi_b$.

REMARK 7. If $\pi_b = S(\Gamma) - \pi_s$, the market is efficient, $z^* = -\infty$, and $p^* = \inf_{t_b} \mathbb{E}[v|t_b]$. $\square$

LEMMA C.2. $\Pi_{us}^*(\Gamma) = \Pi^*(\Gamma)$.

In words, this lemma says that uninformed Seller information structures implement all payoff pairs implementable with price independent beliefs under any information structure. As it is trivial that $\Pi_{us}^*(\Gamma) \subset \Pi_{mb}^*(\Gamma) \subset \Pi^*(\Gamma)$, this establishes part 1 of Theorem 2.

PROOF. $\Pi_{us}^*(\Gamma) \subset \Pi^*(\Gamma)$ is trivial, so we need only prove the opposite direction. Suppose that under some information structure τ there is a wPBE (σ, α, ν) with price independent beliefs and payoffs (π_b, π_s) . Consider an information structure $\tau' \in \mathbf{T}_{us}$ defined by $Q(t_b, v) = P(T_s, t_b, v)$. ν is a consistent belief system given information structure τ and strategy σ . Now we verify that $\forall \sigma'$, ν is a consistent belief system given τ' and σ' . For every measurable rectangle $D \subset \mathbb{R} \times T_b \times V$,

$$\begin{aligned} & \int_D \nu(dv|p, t_b) \sigma'(dp) Q(dt_b, V) \\ &= \int_{D_p} \sigma'(dp) \cdot \int_{D_{t_b, v}} \nu(dv|p, t_b) P(T_s, dt_b, V) \\ &= \int_{D_p} \sigma'(dp) \cdot \int_{D_{t_b, v}} P(T_s, dt_b, dv) \\ &= \int_D \sigma'(dp) Q(dt_b, dv), \end{aligned}$$

where D_p and $D_{t_b, v}$ are the projection of D on dimension p and t_b, v , respectively. The first and third equalities use the definition of measure Q . The second equality is the definition of price independent belief. Since the product sigma algebra is uniquely defined by the product of sigma algebras, verifying on all rectangular D guarantees that ν is a consistent belief system. Therefore, α remains a best response for Buyer. Moreover,

$$\begin{aligned} & \sup_{\sigma'} \int (p - c(v)) \alpha(p, t_b) \sigma'(dp) Q(dt_b, dv) \\ &= \sup_{\sigma'} \int (p - c(v)) \alpha(p, t_b) \sigma'(dp) P(dt_s, dt_b, dv) \\ &\leq \int (p - c(v)) \alpha(p, t_b) \sigma(dp|t_s) P(dt_s, dt_b, dv) = \pi_s. \end{aligned}$$

The first line is achievable by a Seller's strategy when α is modified to break ties in favor of Seller. Therefore, $\pi_s \geq \underline{\pi}_{us}^*(\Gamma)$, and hence $(\pi_b, \pi_s) \in \Pi_{us}^*(\Gamma)$. $\square$

LEMMA C.3. *For any $(\pi_b, \pi_s) \in \Pi_{us}^*(\Gamma)$ with $\pi_s > \underline{\pi}_s^{us}(\Gamma)$, there is $\tau \in \mathbf{T}_{us}$ with $\Pi(\Gamma, \tau) = \{(\pi_b, \pi_s)\}$.*

This “unique implementation” lemma corresponds to part 3 of Theorem 2.

PROOF. We have established that $\underline{\pi}_s^{us}(\Gamma)$ is achieved in an equilibrium. Use $\underline{\pi}_s^{us}(\Gamma)$ as the π_s^* in the proof of Lemma C.1 and construct the corresponding information structure. Note that given the information structure, Seller’s payoff from any deviation to a price other than p^* is bounded above by $\pi_s^* < \pi_s$. As a result, p^* is the uniquely optimal price given Buyer’s best response $\tilde{\alpha}$.

Now we show that for any other α' that is sequentially rational, Seller’s payoff is still bounded above by π_s^* . Suppose not, to contradiction. Then there is p such that $\int \alpha'(p, t_b)(p - c(v))P(dt_b, dv) > \pi_s^*$. Let T_p be the subset of all Buyer’s signals t_b for which $\mathbb{E}[v|t_b] = p$ —signals making Buyer indifferent between buying or not. Note that any two sequentially rational Buyer strategies differ only on T_p . We have

$$\begin{aligned} & \lim_{p' \rightarrow p^-} \int \tilde{\alpha}(p', t_b)(p' - c(v))P(dt_b, dv) \\ &= \int_{\mathbb{E}[v|t_b] \geq p} (p - c(v))P(dt_b, dv) \\ &= \int \alpha'(p, t_b)(p - c(v))P(dt_b, dv) + \int_{t_b \in T_p} (1 - \alpha'(p, t_b))(p - c(v))P(dt_b, dv) \\ &\geq \int \alpha'(p, t_b)(p - c(v))P(dt_b, dv) > \pi_s^*. \end{aligned}$$

The first two equalities are from the fact that $\tilde{\alpha}$ and α' differ from $\mathbf{1}_{\mathbb{E}[v|t_b] \geq p}$ only on T_p . The inequality is from $\mathbb{E}[v|t_b] = p$ on T_p , $v \geq c(v)$ and $\alpha' \leq 1$. This implies that there exists $p' < p$ giving Seller payoff strictly above π_s^* , which is a contradiction.

Therefore, when $\pi_s > \pi_s^*$, the information structure constructed in Lemma C.1 implements the unique equilibrium payoff pair (π_b, π_s) . The result follows from choosing $\pi_s^* = \underline{\pi}_s^{us}(\Gamma)$. $\square$

APPENDIX D: PROOF OF THEOREM 3

PROOF. We first introduce some notation. Given the continuous cost function $c(v)$, any Buyer belief $\nu \in \Delta(V)$, and any price $p \in V$, define

$$\begin{aligned} \pi(c, \nu, p) &= \int_{v \geq p} (p - c(v))\nu(dv), \\ \pi^*(c, \nu) &= \max_{p \in V} \pi(c, \nu, p), \\ \sigma^*(c, \nu) &= \arg \max_{p \in V} \pi(c, \nu, p). \end{aligned}$$

In words, π is Seller’s payoff from offering an arbitrary price p , π^* is Seller’s payoff from an optimal price, and σ^* is the set of optimal prices. Our assumption that $v - c(v) \geq 0$

implies that $\pi(c, \nu, p)$ is a left-continuous function of p that only jumps up, and hence it is upper semicontinuous. Therefore, π^* is well-defined and $\sigma^*(c, \nu)$ is nonempty and compact.

We prove Theorem 3 in 4 steps. In step 1, we define a discretized environment for a grid size d . In step 2, we construct a distribution of IPDs for the discretized environment. In step 3, we show that as $d \rightarrow 0$, the distributions in step 2 converge to a distribution of IPDs whose expectation is the prior μ . In step 4, we construct an information structure and equilibrium for the original environment utilizing the distribution derived in step 3.

Step 1. We discretize the problem. Pick any $d > 0$. Discretize the support V to a grid $V' = \{v_1, \dots, v_n\}$ such that $v_{i+1} - v_i < d$, $v_1 \leq \underline{v}$ and $v_n > \bar{v}$. Let p^* be an element of $\sigma^*(c, \mu)$ and include p^* in V' . Define

$$\begin{aligned}\mu'_i &= \int \mathbf{1}_{v_i \leq v < v_{i+1}} \mu(dv), \\ c'(v_i) &= \int \mathbf{1}_{v_i \leq v < v_{i+1}} c(v) \mu(dv).\end{aligned}$$

Now consider a new environment $\Gamma' = (c', \mu')$ with the discrete support V' . Γ' augments Γ by grouping all Buyer types in interval $[v_i, v_{i+1})$ and assuming Buyer behaves as if the valuation is v_i . A key property of the environment Γ' is that $\forall v_i \in V'$, $\pi(\Gamma', v_i) = \pi(\Gamma, v_i)$, that is, Seller's payoff from offering on-grid prices is invariant under the environment change. Since $p^* \in V'$, $\pi^*(\Gamma') = \pi^*(\Gamma)$.

Step 2. The following lemma—whose proof is provided after the current proof is completed—implies that there exist IPDs $\{\nu_j\}_{j=1}^J$ and $\{p_j\} \in \Delta(J)$ such that $\sum p_j \nu_j = \mu'$ and $\sigma^*(c', \mu') \subset \cap \sigma^*(c', \nu_j)$.

LEMMA D.1. *When $\text{Supp}(\mu)$ is finite, there exists IPDs $\{\nu_j\}_{j=1}^J$ and $\{q_j\} \in \Delta(J)$ such that $\sum q_j \nu_j = \mu$ and $\sigma^*(c, \mu) \subset \sigma^*(c, \nu_j)$.*

Step 3. For each $d_n = \frac{1}{2n}$, go through steps 1–2 and construct a collection $\{p^j, \nu^j\}$. This collection resembles a probability measure $P_n \in \Delta^2(V)$. By construction, any $\nu \in \text{Supp}(P_n)$ is an IPD satisfying $p^* \in \sigma^*(c, \nu)$. We use the following lemma—whose proof is provided after the current proof is completed—to construct a measure P^* whose support contains only those IPDs such that p^* is an optimal price (P^* is a limit point of P_n under the weak topology).

LEMMA D.2. *Suppose the sequence (P_n) in $\Delta^2(V)$ satisfies $\int \nu P_n(d\nu) \xrightarrow{w} \mu$ and $\forall \nu \in \text{Supp}(P_n)$, ν is an IPD satisfying $p^* \in \sigma^*(c, \nu)$. Then $\exists P^* \in \Delta^2(V)$ such that $\int \nu P^*(d\nu) = \mu$ and $\forall \nu \in \text{Supp}(P^*)$, ν is an IPD satisfying $p^* \in \sigma^*(c, \nu)$.*

Then

$$\int \pi^*(c, \nu) P^*(d\nu) = \int \pi(c, \nu, p^*) P^*(d\nu) = \pi^*(c, \mu). \quad (\text{D.1})$$

Step 4. Now we define a fully informed Buyer information structure that implements any $(\pi_b, \underline{\pi}_s^{fb}(\Gamma)) \in \Pi(\Gamma)$. Let $\beta \in [0, 1]$ satisfy $\pi_b = \beta(S(\Gamma) - \underline{\pi}_s^{fb}(\Gamma))$. Take the signal

space $T_s = \Delta(V)$ and define the signal distribution by $\int_D P(dt_s, dv) = \int_D t_s(dv)P^*(dt_s)$. That is, the information structure $\tau = (T_s, P) \in \mathbf{T}_{fb}$ induces Seller's belief ν according to distribution $P^*(\nu)$.

Buyer's strategy is $\alpha(v, p) = \mathbf{1}_{v \geq p}$, which is obviously optimal. Seller's strategy is $\sigma(p|t_s = \nu) = \beta \delta_{p=\min \text{Supp}(\nu)} + (1 - \beta) \delta_{p=\max \text{Supp}(\nu)}$. Then $\forall \sigma'$:

$$\begin{aligned} \int (p - c(v)) \mathbf{1}_{v \geq p} \sigma'(dp|t_s) P(dt_s, dv) &= \int (p - c(v)) \mathbf{1}_{v \geq p} \sigma'(dp|\nu) \nu(dv) P^*(d\nu) \\ &= \int \pi(c, \nu, p) \sigma'(dp|\nu) P^*(d\nu) \\ &\leq \int \pi^*(c, \nu) P^*(d\nu) \quad \text{because } \pi(c, \nu, p) \leq \pi^*(c, \nu), \end{aligned}$$

where the equalities are accounting identities. Meanwhile, Seller's payoff using strategy σ is

$$\begin{aligned} &\int (p - c(v)) \mathbf{1}_{v \geq p} \sigma(dp|t_s) P(dt_s, dv) \\ &= \int \beta \pi(c, \nu, \min \text{Supp}(\nu)) + (1 - \beta) \pi(c, \nu, \max \text{Supp}(\nu)) P^*(d\nu) \\ &= \int \pi^*(c, \nu) P^*(d\nu) = \underline{\pi}_s^{fb}(\Gamma), \end{aligned}$$

where the second equality is from ν being an IPD and the third equality is from Equation (D.1). Therefore, σ is optimal for Seller and Seller's equilibrium payoff is $\underline{\pi}_s^{fb}$. Buyer's payoff is

$$\begin{aligned} &\int (v - p) \mathbf{1}_{v \geq p} \sigma(dp|t_s) P(dt_s, dv) \\ &= \int \beta (v - \min \text{Supp}(\nu)) \mathbf{1}_{v \geq \min \text{Supp}(\nu)} \\ &\quad + (1 - \beta) (v - \max \text{Supp}(\nu)) \mathbf{1}_{v = \max \text{Supp}(\nu)} \nu(dv) P^*(d\nu) \\ &= \int \beta (v - c(v) - (\min \text{Supp}(\nu) - c(v))) \nu(dv) P^*(d\nu) \\ &= \beta \left(S(\Gamma) - \int \pi^*(c, \nu) P^*(d\nu) \right) = \pi_b, \end{aligned}$$

where second equality is from $v - \max \text{Supp}(\nu) \mathbf{1}_{v = \max \text{Supp}(\nu)} = 0$, the third equality is from $\nu \in \text{Supp}(P^*)$ being an IPD, and the last equality is from Equation (D.1).

To sum up, we construct $\tau \in \mathbf{T}_{fb}$ such that $(\pi_b, \underline{\pi}_s^{fb}(\Gamma)) \in \Pi^*(\Gamma, \tau)$. Since $(0, S(\Gamma))$ can be implemented by perfect revealing v to Seller, any other (π_b, π_s) in $\Pi_{fb}^*(\Gamma)$ can be implemented by public randomization, which means Buyer is better informed in the strong sense that his information is a refinement of Seller's. $\square$

PROOF OF LEMMA D.1. We prove the result by induction. When $|\text{Supp}(\mu)| = 1$, the statement is trivially true. Now we assume by induction that the statement is true for

$|\text{Supp}(\mu)| \leq n$ and prove it for $|\text{Supp}(\mu)| = n + 1$. Let $V = \text{Supp}(\mu) = \{v_1, \dots, v_{n+1}\}$. We discuss two cases separately:

- *Case 1:* $v_i > c_i$ for all $i \leq n$. Define $\hat{v}_{n+1} = 1$ and recursively define

$$\hat{v}_i = \frac{\sum_{j=i+1}^{n+1} \hat{v}_j \cdot (v_{i+1} - v_i)}{v_i - c_i}$$

for $i = n \dots 1$. Normalize $\{\hat{v}_i\}$ to a probability vector $\nu = \frac{1}{\sum_i \hat{v}_i} \hat{\nu}$. Then it is easy to verify that $\nu \in \Delta V$ and ν is an IPD:

$$\pi(c, \nu, v_{i+1}) - \pi(c, \nu, v_i) = \sum_{j=i+1}^{n+1} \nu_j \cdot (v_{i+1} - v_i) - v_i \cdot (v_i - c_i) = 0, \quad \forall i.$$

Therefore, $\sigma^*(c, \nu) \supset V$.

- *Case 2:* $v_i = c_i$ for some $i \leq n$. Let i_0 be the smallest i such that this is true. Define $\hat{v}_{i_0} = 1$ and recursively define $\hat{v}_i = \frac{\sum_{j=i+1}^{n+1} \hat{v}_j \cdot (v_{i+1} - v_i)}{v_i - c_i}$ for $i = 1, \dots, i_0$. Normalize $\{\hat{v}_i\}$ to $\nu = \frac{1}{\sum_i \hat{v}_i} \hat{\nu}$. Then the exactly same argument as in Case 1 implies that $\nu \in \Delta V$ and ν is an IPD. Moreover, since $v_{i_0} = c_{i_0}$, $\pi^*(c, \nu) = 0$. Therefore, $\sigma^*(c, \nu) \supset \text{Supp}(\nu) \cup [v_{i_0}, +\infty) \supset V$.

Next, we “remove ν from μ ” to reduce its support size. Let $q = \min\{\frac{\mu_i}{\nu_i}\}$ and $\hat{\mu} = \mu - q \cdot \nu$. By definition, $|\text{Supp}(\hat{\mu})| \leq n$. Normalize $\hat{\mu}$ to $\mu' = \frac{1}{\sum_i \hat{\mu}_i} \hat{\mu}$. Then $\forall i, i' \in (1, \dots, n+1)$,

$$\begin{aligned} & \left(\sum_{\ell} \hat{\mu}_{\ell} \right) (\pi(c, \mu', v_i) - \pi(c, \mu', v_{i'})) \\ &= \pi(c, \mu, v_i) - \pi(c, \mu, v_{i'}) - q \cdot (\pi(c, \nu, v_i) - \pi(c, \nu, v_{i'})) \\ &= \pi(c, \mu, v_i) - \pi(c, \mu, v_{i'}) \\ &\implies \sigma^*(c, \mu') = \sigma^*(c, \mu). \end{aligned}$$

The first equality is from the linearity of π and the second equality is from $\sigma^*(c, \nu) \supset V$. Then, by induction, there exists IPDs ν_j and q_j such that $\sum q_j \nu_j = \mu'$ and $\sigma^*(c, \mu) = \sigma^*(c, \mu') \subset \sigma^*(c, \nu_j)$. Therefore, the statement is proved by appending ν to (ν_j) , and normalizing the probability to $(q_j \cdot \sum \hat{\mu}_i, q)$. $\square$

PROOF OF LEMMA D.2. By Prokhorov’s theorem, there exists a convergent subsequence of P_n ; without loss we suppose $P_n \xrightarrow{w} P^*$ (i.e., weak convergence, which is implied by

convergence in the Prokhorov metric). Let $\mu_n = \int \nu P_n(d\nu)$. By assumption, $\mu_n \xrightarrow{w} \mu$. It follows that $\int \nu P^*(d\nu) = \mu$.³⁰

Now we show that $\forall \nu \in \text{Supp}(P^*)$, ν is an IPD. First, Lemma D.3 shows that there exists a subsequence $n_k \rightarrow \infty$ and $\nu_{n_k} \in \text{Supp}(P_{n_k})$ such that $\nu_{n_k} \xrightarrow{w} \nu$. Then $\forall p \in \text{Supp}(\nu)$, there exists a subsequence $n_{k_s} \rightarrow \infty$ and $p_{n_{k_s}} \in \text{Supp}(\nu_{n_{k_s}})$ such that $p_{n_{k_s}} \rightarrow p$. Lemma D.4 proves that $\pi(c, \nu, p) \geq \overline{\lim} \pi(c, \nu_{n_{k_s}}, p_{n_{k_s}}) = \overline{\lim} \pi^*(c, \nu_{n_{k_s}})$. The equality is from $\nu_{n_{k_s}}$ being IPD and $p_{n_{k_s}}$ being in its support. Hart and Reny (2019, Theorem 12, part 1) proves that $\pi^*(c, \nu) \leq \underline{\lim} \pi^*(c, \nu_{n_{k_s}})$. Therefore, $\pi(c, \nu, p) = \pi^*(c, \nu)$, and hence ν is an IPD.

In the previous analysis, if we pick $p = p^*$, then since $p^* \in \text{Supp}(\nu_{n_k})$, it follows that $p_{n_k} = p^*$, and hence trivially $p_{n_k} \rightarrow p^*$. Therefore, $\pi(c, \nu, p^*) = \pi^*(c, \nu)$. $\square$

LEMMA D.3. *Let (S, ρ) be a separable metric space, $\{P_n\} \subset \Delta(S)$ and $P_n \xrightarrow{w} P$. Then $\forall s \in \text{Supp}(P)$, $\exists$ sequence $s_{n_k} \in \text{Supp}(P_{n_k})$ such that $n_k \rightarrow \infty$ and $s_{n_k} \xrightarrow{\rho} s$.*

PROOF. For any $s \in \text{Supp}(P)$, suppose towards contradiction that the statement is not true. Then we claim that $\exists \varepsilon > 0$, $N \in \mathbb{N}$ such that $\forall n \geq N$ $\text{Supp}(P_n) \cap B_\varepsilon(s) = \emptyset$. Otherwise, $\forall \varepsilon > 0$, $N \in \mathbb{N}$ exists $n \geq N$ such that $\text{Supp}(P_n) \cap B_\varepsilon(s) \neq \emptyset \implies$ pick any $N = k$ and $\varepsilon = \frac{1}{k}$, there exists $n_k \geq k$ and $s_{n_k} \in \text{Supp}(P_{n_k})$ such that $\rho(s, s_{n_k}) < \frac{1}{k}$, and hence the assumption is not true.

Since $\exists \varepsilon > 0$ and N such that $\forall n \geq N$ $\text{Supp}(P_n) \cap B_\varepsilon(s) = \emptyset$, this implies $\underline{\lim} P_n(B_\varepsilon(s)) = 0 \geq P(B_\varepsilon(s))$ (by the Portmanteau theorem). This contradicts the assumption that $s \in \text{Supp}(P)$. $\square$

LEMMA D.4. *Let $c \in C(V)$, $\{\nu_n\} \subset \Delta(V)$, $\{p_n\} \subset V$. If $\nu_n \xrightarrow{w} \nu$ and $p_n \rightarrow p$, then*

$$\pi(c, \nu, p) \geq \overline{\lim} \pi(c, \nu_n, p_n).$$

PROOF. Define

$$h_{\delta, p}(v) = \frac{v - p + \delta}{\delta} \wedge [0, 1], \quad (\text{D.2})$$

where $\cdot \wedge [0, 1]$ is the truncation functional on $[0, 1]$. Then $h_{\delta, p}$ is a continuous and bounded function and $\mathbf{1}_{v \geq p} \leq h_{\delta, p}(v) \leq \mathbf{1}_{v \geq p - \delta}$. Then $\forall \eta > \delta > 0$:

$$\begin{aligned} \int_{p-\eta}^{\infty} (p - c(v)) \nu(dv) &\geq \int h_{\delta, p-\eta+\delta}(v) (p - c(v)) \nu(dv) \\ &= \lim_{n \rightarrow \infty} \int h_{\delta, p-\eta+\delta}(v) (p - c(v)) \nu_n(dv) \end{aligned}$$

³⁰For any continuous $h(v)$, $\nu \mapsto \int h(v) \nu(dv)$ is a bounded and continuous function on $\Delta(V)$ under the Prokhorov metric. Therefore, since $\mu_n \xrightarrow{w} \mu$ and $P_n \xrightarrow{w} P^*$,

$$\int h(v) \mu_n(dv) \rightarrow \int h(v) \mu(dv) \quad \text{and} \quad \iint h(v) \nu(dv) P_n(d\nu) \rightarrow \iint h(v) \nu(dv) P^*(d\nu).$$

Since $\int h(v) \mu_n(dv) = \iint h(v) \nu(dv) P_n(d\nu)$, it follows that $\int \nu P^*(d\nu) = \mu$.

$$\begin{aligned}
&\geq \overline{\lim}_{n \rightarrow \infty} \int_{p-\eta+\delta}^{\infty} (p - c(v)) \nu_n(dv) \\
&\geq \overline{\lim}_{n \rightarrow \infty} \int_{p-\eta+\delta}^{\infty} (p + \eta - c(v)) \nu_n(dv) - \eta \\
&\geq \overline{\lim}_{n \rightarrow \infty} \int_{p_n}^{\infty} (p + \eta - c(v)) \nu_n(dv) - \eta \\
&\geq \overline{\lim}_{n \rightarrow \infty} \int_{p_n}^{\infty} (p_n - c(v)) \nu_n(dv) - \eta.
\end{aligned}$$

The first inequality above is because $h_{\delta, p-\eta+\delta}(v) \leq \mathbf{1}_{v \geq p-\eta}$ and $\forall v \in [p - \eta, p - \eta + \delta]$, $p > v \geq c(v)$. The first equality is from $\nu^n \xrightarrow{w} \nu$ and the integrand being continuous and bounded. The second inequality is from $h_{\delta, p-\eta+\delta}(v) \geq \mathbf{1}_{v \geq p-\eta+\delta}$ and $\forall v \in [p - \eta, p - \eta + \delta]$ $p > v \geq c(v)$. The third inequality is straightforward. The fourth inequality is from $\lim p_n > p - \eta + \delta$. The last inequality is from $\lim p_n < p + \eta$.

Letting $\eta \rightarrow 0$, we obtain $\pi(c, \nu, p) \geq \overline{\lim} \pi(c, \nu_n, p_n)$. $\square$

APPENDIX E: PROOF OF PROPOSITION 1

PROOF. We first show that $\Pi(\Gamma)$ is included in the set defined in Proposition 1. $\forall (\pi_b, \pi_s) \in \Pi(\Gamma)$, it is clear that $\pi_b \geq 0$ and $\pi_s \geq \underline{\pi}_s(\Gamma)$. Now we prove that the inequality $\lambda \pi_b + \pi_s \leq S_\lambda(\Gamma)$ is satisfied for any $\lambda \in [1, \infty)$. Let τ be the information structure and (σ, α, ν) be an equilibrium with payoff (π_b, π_s) . We define the following Borel measure β : for any Borel set $V' \subset \mathbb{R}$,

$$\beta(V') = \int_{v \in V'} \alpha(p, t_b) \sigma(dp|t_s) P(dt_b, dt_s, dv).$$

In words, β calculates the trading probability for a given set of types V' . By definition,

$$\begin{aligned}
\pi_b &= \int (v - \mathbb{E}[p|v]) \beta(dv), \\
\pi_s &= \int (\mathbb{E}[p|v] - c(v)) \beta(dv),
\end{aligned}$$

and hence

$$\begin{aligned}
\lambda \pi_b + \pi_s &= \int (\mathbb{E}[p|v] - c(v) + \lambda(v - \mathbb{E}[p|v])) \beta(dv) \\
&= \int (\lambda v - c(v) - (\lambda - 1)\mathbb{E}[p|v]) \beta(dv) \\
&\leq \int (\lambda v - c(v) - (\lambda - 1)\underline{v}) \beta(dv) \\
&\leq S_\lambda(\Gamma),
\end{aligned}$$

where the first inequality uses any on-path price being no lower than $\underline{v}$ and $\lambda \geq 1$.

Now we show that all payoff pairs (π_b, π_s) satisfying the inequality constraints can be implemented by equilibrium payoffs in $\Pi(\Gamma)$. $\forall \lambda \in [1, \infty)$ and $\alpha \in [0, 1]$, define $\beta_\alpha(v) = \mathbf{1}_{\lambda v + \underline{v} > c(v) + \lambda \underline{v}} + \alpha \mathbf{1}_{\lambda v + \underline{v} = c(v) + \lambda \underline{v}}$. Let us ignore the individual rationality constraint $\pi_s \geq 0$ for now.

Construct the following information structure: a public signal is sent to both players indicating whether $\beta_\alpha(v) = 0$. Following the positive signal, construct an information structure as in Theorem 1 that induces Seller selling with probability one at $p = \underline{v}$ almost surely in the subgame.³¹ When $\beta_\alpha(v) = 0$, trading surplus is nonpositive, as $c(v) \geq v + (\lambda - 1)(v - \underline{v})$ and so an equilibrium with no trade exists. The realized weighted total surplus is exactly $S_\lambda(\Gamma)$, with Buyer's payoff $\int \beta_\alpha(v)(v - \underline{v})\mu(dv)$ and Seller's the remaining $\int \beta_\alpha(v)(\underline{v} - c(v))\mu(dv)$. Note that Buyer's payoff is continuously increasing in α . Therefore, $\forall \lambda \in [1, \infty)$, we can implement an interval (possibly degenerate) on the frontier $S_\lambda(\Gamma)$ defined by $\{(\pi_b^\alpha, \pi_s^\alpha) = (\int \beta_\alpha(v)(v - \underline{v})\mu(dv), \int \beta_\alpha(v)(\underline{v} - c(v))\mu(dv))\}_{\alpha \in [0, 1]}$. Since we construct the equilibria explicitly, this interval satisfies all other constraints.

Next, we show two key properties of the interval $\{(\pi_b^\alpha, \pi_s^\alpha)\}_{\alpha \in [0, 1]}$.

- $\forall \lambda \geq 1, \forall \delta > 0, (\pi_b', \pi_s') = (\pi_b^1 + \delta, \pi_s^1 - \lambda \delta)$ violates some frontier $S_{\lambda'}(\Gamma)$ with $\lambda' > \lambda$. It is straightforward to calculate

$$\lambda' \pi_b' + \pi_s' = (\lambda' - \lambda)(\pi_b^1 + \delta) + S_\lambda(\Gamma),$$

and hence

$$\frac{\lambda' \pi_b' + \pi_s' - S_\lambda(\Gamma)}{\lambda' - \lambda} = (\pi_b^1 + \delta).$$

Now we calculate $S_{\lambda'}(\Gamma)$:

$$\begin{aligned} \frac{S_{\lambda'}(\Gamma) - S_\lambda(\Gamma)}{\lambda' - \lambda} &= \frac{\int_{\lambda(v - \underline{v}) + \underline{v} - c(v) \geq 0} (\lambda' - \lambda)(v - \underline{v})\mu(dv)}{\lambda' - \lambda} \\ &\quad + \frac{\int_{\lambda(v - \underline{v}) + \underline{v} - c(v) \in [(\lambda' - \lambda)(\underline{v} - v), 0]} (\lambda'(v - \underline{v}) + \underline{v} - c(v))\mu(dv)}{\lambda' - \lambda} \\ &\rightarrow \int \beta_1(v)(v - \underline{v})\mu(dv) = \pi_b^1 \quad \text{when } \lambda' \rightarrow \lambda. \end{aligned}$$

The limit is derived by canceling out $(\lambda' - \lambda)$ in the first line and observing the integrand is bounded by $(\lambda' - \lambda)(v - \underline{v})$ in the second line. Since $\delta > 0$, we have that when $\lambda' - \lambda$ is sufficiently small, $\lambda' \pi_b' + \pi_s' > S_{\lambda'}(\Gamma)$, violating the frontier $S_{\lambda'}(\Gamma)$.

- $\forall \lambda > 1, \forall \delta > 0, (\pi_b', \pi_s') = (\pi_b^0 - \delta, \pi_s^0 + \lambda \delta)$ violates some frontier $S_{\lambda'}(\Gamma)$ with $\lambda' < \lambda$. The argument is symmetric.

³¹In the subgame following $\beta_\alpha(v) > 0$, the support of v might not contain $\underline{v}$. This does not affect the consistency of off-path beliefs as we use wPBE as the equilibrium notion (without imposing subgame perfection).

Any point on the frontier $S_1(\Gamma)$ between $(0, S_1(\Gamma))$ and (π_b^1, π_s^1) can be implemented by public randomization. Therefore, any payoff pair on the envelope of all frontiers is implementable (while ignoring Seller's individual rationality constraint). Then we can just truncate below by the extra constraint $\pi_s \geq 0$. The implementation of $(0, \underline{\pi}_s(\Gamma))$ is trivial. Then public randomization implements all other points in the set. $\square$

APPENDIX F: PROOF OF PROPOSITION 2

PROOF. As discussed in the main text, it is sufficient to show the existence of cdf $G(v)$ such that (1) $G \in D(\mu)$, (2) G is an IPD, (3) $\exists p \in \text{Supp}(G)$ such that $\int_{\underline{v}}^p G(s) ds = \int_{\underline{v}}^p F(s) ds$.

First, we show that $\forall v_* \in [\underline{v}, \mathbb{E}_\mu[v]]$ such that $v_* > c(v_*)$ and $v_* \geq c(\mathbb{E}[v])$, IPD G_{v_*} exists.

Case 1: $\lambda \neq 1$. The indifference condition of IPD is equivalent to

$$\begin{aligned} \frac{d}{dv} \int_v^{v^*} (v - c(s)) dG_{v_*}(s) &= 0 \\ \iff -(v - c(v))g_{v_*}(v) + (1 - G_{v_*}(v)) &= 0 \\ \iff (c(v) - v) d \log(1 - G_{v_*}(v)) &= 1 \\ \iff G_{v_*}(v) = 1 - C(c(v) - v)^{\frac{1}{\lambda-1}}. \end{aligned}$$

Using condition $G_{v_*}(v_*) = 0$, we can pin down C :

$$G_{v_*}(v) = 1 - \left(\frac{v - c(v)}{v_* - c(v_*)} \right)^{\frac{1}{\lambda-1}}.$$

Lastly, v^* can be pinned down using the following condition:³²

$$\begin{aligned} (1 - G_{v_*}(v^*))(v^* - c(v^*)) &= v_* - c(\mathbb{E}_\mu[v]) \\ \implies (1 - \lambda)v^* - \gamma &= (v_* - c(\mathbb{E}[v]))^{\frac{\lambda-1}{\lambda}} ((1 - \lambda)v_* - \gamma)^{\frac{1}{\lambda}}. \end{aligned}$$

Note that if $v_* \rightarrow \mathbb{E}[v]$, then $v^* \rightarrow \mathbb{E}[v]$. One can also verify that v^* increases when v_* decreases:

$$\frac{dv^*}{dv_*} = (v_* - \mathbb{E}[v])(v_* - c(\mathbb{E}[v]))^{-\frac{1}{\lambda}} ((1 - \lambda)v_* - \gamma)^{\frac{1-\lambda}{\lambda}} \leq 0.$$

Case 2: $\lambda = 1$ (hence $\gamma < 0$). The indifference condition of IPD is equivalent to

$$\begin{aligned} \frac{d}{dv} \log(1 - G_{v_*}(v)) &= \frac{1}{\gamma} \\ \iff G_{v_*}(v) &= 1 - C \cdot e^{\frac{v}{\gamma}}. \end{aligned}$$

³²It is easy to verify that the condition is equivalent to $\int v dG_{v_*}(v) = \mathbb{E}_\mu[v]$.

We pin down C and v^* using the CDF at v_* and the mean preserving spread condition:

$$\begin{cases} C = e^{-\frac{v_*}{\gamma}} \\ v^* = v_* + \gamma \log\left(\frac{v_* - c(\mathbb{E}_\mu[v])}{-\gamma}\right). \end{cases}$$

Second, we show that there exists v_* such that the corresponding IPD $G_{v_*}(v)$ satisfies condition (1) and (3). It is trivial that if $v_* = \mathbb{E}_\mu[v]$, then G_{v_*} has unit mass at $\mathbb{E}_\mu[v]$, which is included in $D(\mu)$. By the linearity of c , $v = c(v)$ can only happen on the boundaries of V . Therefore, $v_* = \inf\{v' \in (\max\{\underline{v}, c(\mathbb{E}[v])\}, \mathbb{E}_\mu[v]) \mid G_{v'} \in D(\mu)\}$ is well-defined. We now consider three cases separately:

- *Case 1:* $v_* > \max\{c(v_*), c(\mathbb{E}[v])\}$. In this case v_{v_*} is well-defined. By the formula of $G_{v'}(v)$, it is continuous in v' for each v . Obviously, CDFs are uniformly bounded. So, by dominated convergence theorem, $\forall q$:

$$\int_{\underline{v}}^v G_{v_*}(s) ds = \lim_{v' \rightarrow v_*^+} \int_{\underline{v}}^v G_{v'}(s) ds.$$

This implies $G_{v_*} \in D(\mu)$. Now we claim that there exists $p \in [v_*, v^*]$ such that $\int_{\underline{v}}^p F(s) ds = \int_{\underline{v}}^p G_{v_*}(s) ds$. If not, this implies $\int_{\underline{v}}^p F(s) ds < \int_{\underline{v}}^p G_{v_*}(s) ds \forall p \in [v_*, v^*]$. Then choosing v_* slightly smaller, G_{v_*} is still in $D(\mu)$, contradiction.

- *Case 2:* $v_* = c(v_*) \geq c(\mathbb{E}[v])$. We show that this case is never possible. Since c is linear, this can happen only when $v_* = \underline{v}$ and $\lambda \leq 0$. Consider $v'_* = \underline{v} + \varepsilon$ where $\varepsilon > 0$. Then $G_{v'_*} \in D(\mu)$ when ε is very small. However, v^* is pinned down by

$$(1 - \lambda)v^* = \gamma + (v'_* - c(\mathbb{E}[v]))^{\frac{\lambda-1}{\lambda}} ((1 - \lambda)\varepsilon)^{\frac{1}{\lambda}}.$$

When $\varepsilon \rightarrow 0$, $v^* \rightarrow \infty$, so $G_{v'_*} \notin D(\mu)$ for sufficiently small ε , contradiction.

- *Case 3:* $v_* = c(\mathbb{E}[v]) > \underline{v}$. In this case, G_{v_*} gives Seller zero profit and $G_{v_*} \in D(\mu)$. Then $\underline{\pi}_s^{us}(\Gamma) = 0$. So, the proof of Proposition 2 is already done. $\square$

APPENDIX G: PROOF OF COROLLARY 1

PROOF. It straightforward to verify that the solution to Equation (6) is unique; denote it by p . When V is binary, $c(v)$ is trivially affine. Let $\lambda = \frac{c(v_2) - c(v_1)}{v_2 - v_1}$ and $\gamma = c(v_1) - \lambda v_1$. Then Condition 1 is satisfied and Proposition 2 applies. Let v_* be the corresponding parameter defining profit minimizing IPD. Since V is binary, $\int_{\underline{v}}^v F(s) ds$ is a piecewise linear function with two kinks at v_1, v_2 . Meanwhile, $\int_{\underline{v}}^v G_{v_*}(s) ds$ is strictly convex on its support (v_*, v^*) . Therefore, $\int_{\underline{v}}^v F(s) ds$ cannot intersect $\int_{\underline{v}}^v G_{v_*}(s) ds$ at any $v \in (v_*, v^*)$. So, either $(v_* = v_1 \text{ and } v^* \leq v_2)$ or $(v^* = v_2 \text{ and } v_* \geq v_1)$.

We begin with the conjecture that $v^* = v_2$. This implies

$$\begin{aligned} (1 - G_{v_*}(v_2))(v_2 - c(v_2)) &= v_* - \mathbb{E}_\mu[c(v)] \\ \implies (v_* - c(v_*))^{\frac{1}{\lambda-1}} (v_* - \mathbb{E}_\mu[c(v)]) &= (v_2 - c(v_2))^{\frac{1}{\lambda-1}}, \end{aligned}$$

that is, v_* solves Equation (6) and $v_* = p$. Therefore, only when $p \geq v_1$ the conjecture is valid, in which case p is an optimal price and $\underline{\pi}_s^{us}(\Gamma) = p - \mathbb{E}[c(v)]$.

Otherwise, if $p < v_1$, the conjecture $v^* = v_2$ is not valid, so v_1 is an optimal price and $\underline{\pi}_s^{us}(\Gamma) = v_1 - \mathbb{E}[c(v)]$. To sum up, $\underline{\pi}_s^{us}(\Gamma) = \max\{p, v_1\} - \mathbb{E}[c(v)]$. $\square$

REFERENCES

- Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016), “The economics of privacy.” *Journal of Economic Literature*, 54, 442–492. [0284]
- Akerlof, George A. (1970), “The market for “Lemons”: Quality uncertainty and the market mechanism.” *Quarterly Journal of Economics*, 84, 488–500. [0281]
- Bagwell, Kyle and Michael H. Riordan (1991), “High and declining prices signal product quality.” *American Economic Review*, 81, 224–239. [0282]
- Bar-Isaac, Heski, Ian Jewitt, and Clare Leaver (2021), “Adverse selection, efficiency and the structure of information.” *Economic Theory*, 72, 579–614. [0286, 0302]
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015), “The limits of price discrimination.” *American Economic Review*, 105, 921–957. [0285, 0286, 0293, 0294, 0297, 0299]
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2025), “On the alignment of consumer surplus and total surplus under competitive price discrimination.” *American Economic Journal: Microeconomics*, 17, 234–259. [0285, 0286]
- Blackwell, David (1953), “Equivalent comparisons of experiments.” *Annals of Mathematical Statistics*, 24, 265–272. [0282, 0284, 0290]
- Che, Yeon-Koo and Weijie Zhong (2025), “Robustly-optimal mechanism for selling multiple goods.” *Review of Economic Studies*, 92, 2923–2951. [0299]
- Cho, In-Koo and David Kreps (1987), “Signaling games and stable equilibria.” *Quarterly Journal of Economics*, 102, 179–221. [0293]
- Dang, Tri Vi (2008), “Bargaining with endogenous information.” *Journal of Economic Theory*, 140, 339–354. [0286]
- Deb, Rahul, Mallesh Pai, and Anne-Katrin Roesler (2024), “How informed do you want your principal to be?” unpublished. [0286]
- Deb, Rahul and Anne-Katrin Roesler (2024), “Multi-dimensional screening: Buyer-optimal learning and informational robustness.” *Review of Economic Studies*, 91, 2744–2770. [0299]
- Doval, Laura and Jeffrey C. Ely (2020), “Sequential information design.” *Econometrica*, 88, 2575–2608. [0288]
- Du, Songzi (2018), “Robust mechanisms under common valuation.” *Econometrica*, 86, 1569–1588. [0285, 0292]

- Durrett, Richard (1995), *Probability: Theory and Examples*, second edition. Duxbury Press. [0288]
- Einav, Liran and Amy Finkelstein (2011), “Selection in insurance markets: Theory and empirics in pictures.” *Journal of Economic Perspectives*, 25, 115–138. [0282]
- Elliott, Matthew, Andrea Galeotti, Andrew Koh, and Wenhao Li (2022). Market segmentation through Information. Unpublished. [0285]
- Elliott, Matthew, Andrea Galeotti, Andrew Koh, and Wenhao Li (2023), “Matching and information design in marketplaces.” Unpublished. [0285]
- Fudenberg, Drew and Jean Tirole (1991), “Perfect Bayesian equilibrium and sequential equilibrium.” *Journal of Economic Theory*, 53, 236–260. [0284, 0290]
- Garcia, Daniel, Roe Teper, and Matan Tsur (2018), “Information design in insurance markets: Selling peaches in a market for lemons.” unpublished. [0286]
- Greenwald, Bruce C. (1986), “Adverse selection in the labour market.” *Review of Economic Studies*, 53, 325–347. [0281]
- Haghpanah, Nima and Ron Siegel (2023), “Pareto improving segmentation of multi-product markets.” *Journal of Political Economy*, 131, 1546–1575. [0285]
- Hart, Sergiu and Philip J. Reny (2019), “The better half of selling separately.” *ACM Transactions on Economics and Computation*, 7, 1–18. [0310, 0317]
- Ichihashi, Shota (2020), “Online privacy and information disclosure by consumers.” *American Economic Review*, 110, 569–595. [0285]
- Inostroza, Nicolas and Anton Tsoy (2025) “Optimal Information and Security Design.” Unpublished. [0303]
- Jovanovic, Boyan (1982), “Favorable selection with asymmetric information.” *Quarterly Journal of Economics*, 97, 535–539. [0282]
- Kessler, Anke S. (2001), “Revisiting the lemons market.” *International Economic Review*, 42, 25–41. [0286]
- Kreps, David and Robert Wilson (1982), “Sequential equilibria.” *Econometrica*, 50, 863–894. [0283, 0286, 0290, 0293]
- Lehmann, Eric L. (1988), “Comparing location experiments.” *The Annals of Statistics*, 16, 521–533. [0286]
- Levin, Jonathan (2001), “Information and the market for lemons.” *RAND Journal of Economics*, 32, 657–666. [0286]
- Libgober, Jonathan and Xiaosheng Mu (2021), “Informational robustness in intertemporal pricing.” *Review of Economic Studies*, 88, 1224–1252. [0285]
- Makris, Miltiadis and Ludovic Renou (2023), “Information design in multi-stage games.” *Theoretical Economics*, 18, 1475–1509. [0285, 0286, 0288]

- Milgrom, Paul R. and Robert J. Weber (1985), “Distributional strategies for games with incomplete information.” *Mathematics of Operations Research*, 10, 619–632. [0288]
- Mussa, Michael and Sherwin Rosen (1978), “Monopoly and product quality.” *Journal of Economic Theory*, 18, 301–317. [0281]
- Myerson, Roger B. (1981), “Optimal auction design.” *Mathematics of Operations Research*, 6, 58–73. [0299]
- Pavan, Alessandro and Jean Tirole (2025), “Knowing your lemon before you dump it.” Unpublished. [0286]
- Pollrich, Martin and Roland Strausz (2024), “The irrelevance of fee structures for certification.” *American Economic Review: Insights*, 6, 55–72. [0287]
- Roesler, Anne-Katrin and Balázs Szentes (2017), “Buyer-optimal learning and monopoly pricing.” *American Economic Review*, 107, 2072–2080. [0284, 0285, 0286, 0293, 0294, 0296, 0297, 0299, 0303]
- Terstiege, Stefan and Adrien Vigier (2025), “Market segmentation and product steering.” *American Economic Journal: Microeconomics*, 17, 126–146. [0285]
- Terstiege, Stefan and Cédric Wasser (2020), “Buyer-optimal extensionproof information.” *Journal of Economic Theory*, 188, 105070. [0285]
- Thereze, Joao (2023), “Adverse selection and endogenous information.” Unpublished. [0286]
- Yang, Kai Hao (2022), “Selling consumer data for profit: Optimal market-segmentation design and its consequences.” *American Economic Review*, 112, 1364–1393. [0285]

Co-editor Juuso Toikka handled this manuscript.

Manuscript received 6 February, 2025; final version accepted 8 September, 2025; available online 10 September, 2025.

All authors assume responsibility for all aspects of the paper.