

Discriminatory Information Disclosure[†]

By HAO LI AND XIANWEN SHI*

A seller designs a mechanism to sell a single object to a potential buyer whose private type is his incomplete information about his valuation. The seller can disclose additional information to the buyer about his valuation without observing its realization. In both discrete-type and continuous-type settings, we show that discriminatory disclosure—releasing different amounts of additional information to different buyer types—dominates full disclosure in terms of seller revenue. An implication is that the orthogonal decomposition technique, while an important tool in dynamic mechanism design, is generally invalid when information disclosure is part of the design. (JEL D11, D82, D83)

In many pricing problems, buyers often do not have perfect information about the valuation of the seller's product or service. The seller can have considerable control over buyers' access to additional information that they can use to refine their private estimate. For example, an auctioneer for an oil tract or a painting can choose the number and the nature of the tests that a bidder or his hired consultants can privately carry out (Eső and Szentes 2007). Similarly, a seller of a new car or a new product may offer test drives, product samples, or pre-sales technical support in an attempt to influence the amount of product information that buyers gather in order to determine how well their idiosyncratic preferences match with product characteristics (Lewis and Sappington 1994). In financial markets, the owner of a company can disseminate proprietary information about its assets (e.g., existing customer base, or internal projections on specific businesses) to potential investors who can then better evaluate the size of synergies from cross-selling business solutions (Bergemann and Pesendorfer 2007). Finally, with advances in technology, online retailers can easily

*Li: Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC, Canada, V6T 1L4 (email: hao.li@ubc.ca); Shi: Department of Economics, University of Toronto, 150 St. George Street, Toronto, ON, Canada, M5S 3G7 (email: xianwen.shi@utoronto.ca). This paper was accepted to the *AER* under the guidance of Larry Samuelson, Coeditor. We are indebted to Dirk Bergemann for his support in various stages of this project. We thank the participants of 2013 Cowles Summer Conference in Economic Theory, especially our discussant Philipp Strack, and Heski Bar-Issac, Simon Board, Maxim Ivanov, Dina Mayzlin, Marco Ottaviani, Alessandro Pavan, Mike Peters, Luis Rayo, Philipp Sadowski, Roland Strausz, Balázs Szentes, Juuso Välimäki, and Gábor Virág, as well as the three anonymous referees for many helpful comments and suggestions. Jiaqi Zou provided excellent research assistance. Both authors are grateful to the Social Sciences and Humanities Research Council of Canada for financial support. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <https://doi.org/10.1257/aer.20151743> to visit the article page for additional materials and author disclosure statement(s).

control how much product information such as online reviews or feedback to make accessible to shoppers.

This paper considers a revenue-maximizing seller who wants to sell an indivisible object to a prospective buyer. The buyer has some initial incomplete information (ex ante type) about his valuation for the object, and the seller can release an informative signal to the buyer without observing its realization. This single-buyer, two-period model of “sequential screening” (Courty and Li 2000) offers a natural and simple information environment in the framework of dynamic mechanism design to study information disclosure. The seller designs the information disclosure policy and the selling mechanism jointly in order to discriminate among different buyer types.¹ We show that discriminatory disclosure of the seller’s signal—providing different buyer types with different amounts of additional information—dominates full information disclosure in terms of the seller’s revenue. We first establish the suboptimality of full disclosure for the case where the seller’s signal is perfect regarding the buyer’s true valuation and the buyer’s ex ante type is discrete. We then extend this finding to the case where the seller’s signal is noisy and to the case where the buyer’s ex ante type is continuous.

The same disclosure problem is studied in Eső and Szentes (2007), but they take an indirect approach. They introduce an orthogonal decomposition technique to transform the signal controlled by the seller into an independent “shock” that is orthogonal to the buyer’s ex ante type. They interpret this orthogonal shock as information contained in the seller’s signal that is “new” to the buyer. This allows them to establish an “irrelevance theorem” that under some regularity conditions the maximal revenue in a “hypothetical problem” (in which the orthogonal shock is publicly observed) is attained in the original setup where the shock is private.² They argue that the irrelevance theorem implies that full disclosure is optimal because the maximal revenue in the hypothetical setting is an upper bound for the seller’s revenue in the original setting.

The orthogonal decomposition technique has since become an important tool in solving dynamic mechanism design problems with *exogenous* information. In these problems, agents (such as buyers) are often endowed with some initial private information, and after contracting receive an exogenous sequence of additional information that is correlated with their initial private information. Any sequence of such information can be transformed via the decomposition technique into a sequence of independent shocks that are orthogonal to the agents’ initial information. Although dynamic incentive compatibility and revenue maximization can be, and indeed have been, studied without this transformation (Baron and Besanko 1984; Courty and Li 2000), the decomposition technique has proven very useful because in general dynamic environments, it allows one to prove an envelope theorem (Pavan, Segal,

¹ Such combination of discriminatory pricing and information provision exists in practice. For example, online buyers clubs for women’s clothing offer members the option of paying a monthly subscription fee for personalized advice on fashion and style. Banks and wealth management companies often provide access to in-house financial advisors for clients willing to accept a different fee schedule. Start-ups or private firms seeking external financing may offer special equity shares and information access (via allocating slots on the board of directors, for example) to different types of investors.

² Eső and Szentes (2017) call it “irrelevance theorem” because, for the purpose of revenue maximization, the original dynamic problem is equivalent to the hypothetical problem in which the dynamic nature of adverse selection is irrelevant.

and Toikka 2014) and extend the irrelevance theorem (Eső and Szentes 2017).³ In these environments, one can derive revenue-maximizing mechanisms by directly working with the simpler hypothetical problem where all the subsequent orthogonal signals are publicly observed.

The optimality of full disclosure claimed in Eső and Szentes (2007) stands in contrast to our finding in the continuous type case. To reconcile this contradiction, we note that the additional information received by the buyer is chosen by the seller together with the mechanism and is thus *endogenous*. With endogenous information, selectively releasing orthogonal shocks is not the same as selectively releasing untransformed signals. We explain that the maximal revenue in the hypothetical setting is a valid upper bound in the original setting only if the seller is restricted to releasing signals garbled from the orthogonal shock (which we call *orthogonal disclosure*), but not if the seller can directly garble her untransformed signal (which we call *direct disclosure*). We also provide an explicitly solved example (Example 1) that clearly illustrates why their irrelevance theorem fails to imply the optimality of full disclosure. We show that, with direct disclosure, the seller not only avoids paying information rent for any postcontractual private information as implied by the irrelevance theorem, but also completely eliminates the information rent due to the buyer's *ex ante* type.

Our main result that full disclosure is suboptimal in general has two broad implications. First, in dynamic adverse selection environments where the seller may partially control the information flow, the dynamic nature of private information is no longer irrelevant. Both dynamic allocation efficiency and the seller's revenue can depend on the seller's information control. Second, if the seller can grant differential information access to different types of buyers, the disclosure policy becomes an additional instrument to facilitate price discrimination. There is no longer a straightforward solution to the seller's optimal information policy, as full disclosure is generally suboptimal, and how we model the seller's feasible information choices can matter in applications.

A. Related Literature

The present paper belongs in the rapidly growing literature on dynamic mechanism design. As already mentioned, the two most directly related papers are Courty and Li (2000) and Eső and Szentes (2007). Bergemann and Said (2011), Gershkov and Moldovanu (2012), Krämer and Strausz (2015a), and Pavan (forthcoming) provide excellent surveys of recent developments in this literature.

Lewis and Sappington (1994) are among the first to introduce the idea of private information disclosure to the mechanism-design literature (see also Che 1996; Anderson and Renault 2006; Johnson and Myatt 2006; Ganuza and Penalva 2010; Hoffman and Inderst 2011). The signal disclosed by our seller is private in the sense that the signal affects only the buyer's valuation and its realization is observable to

³The orthogonal decomposition technique has been applied to study managerial turnover in stochastic environments (Garrett and Pavan 2012), return policies for online auctions (Zhang 2013), implementation of handicap auctions with posterior constraints (Bergemann and Wambach 2015), and indicative bidding in optimal two-stage auctions (Lu and Ye 2016).

the buyer but not to the seller. As a result, the seller cannot contract on the realization of the released signal. This makes our paper different from the classical disclosure problem in auctions with affiliated values (Milgrom and Weber 1982; Ottaviani and Prat 2001). If there is no private information at the time of contracting, and if their participation constraints are interim so that the seller cannot charge for information, the seller faces a trade-off between disclosing more private information and thereby improving allocation efficiency on one hand, and having to elicit private information from buyers and consequently giving up more information rent on the other (Ganuza 2004; Bergemann and Pesendorfer 2007). Therefore, full disclosure is not optimal. If, as in our model, buyers have private information *ex ante* and the seller can charge fees for additional private information, the trade-off above disappears because, by the irrelevance theorem of Eső and Szentes (2007, 2017), the seller does not pay any rent for the additional private information. We show that the absence of the trade-off above does not imply the optimality of full disclosure when the seller controls access to additional private information, because the seller can use discriminatory disclosure to further reduce the information rent generated by the precontractual private information.

Our disclosure problem is also related to the persuasion problem in the sender-receiver framework, where the sender has private information and can disclose it selectively but cannot lie; that is, the disclosed information must be verifiable. The earlier literature (Grossman 1981; Milgrom 1981; Milgrom and Roberts 1986) assumes that the sender cannot commit to disclosure rules and shows that full disclosure is the unique equilibrium due to unraveling. A more recent literature (Rayo and Segal 2010; Kamenica and Gentzkow 2011; Jehiel 2015) allows the sender to commit to disclosure rules before observing private signals and shows that partial disclosure can arise in equilibrium. Our problem is different, because the seller (i.e., the sender) controls the disclosure rule but does not observe the realization, and the buyer (i.e., the receiver) has private information before contracting.⁴

I. The Model

We study a two-period sequential screening model. A seller has one object to sell to a potential buyer. The seller and the buyer are risk-neutral, and do not discount the future. The buyer's valuation for the good $\omega \in \Omega \equiv [\underline{\omega}, \bar{\omega}]$ is initially unknown to the buyer. Instead, the buyer privately observes a signal $\theta \in \Theta$ about ω , which we refer to as his *ex ante* type. We allow Θ to be either discrete or an interval $[\underline{\theta}, \bar{\theta}]$ on the real line, and introduce the notation for type distribution later. For each $\theta \in \Theta$, let $F(\cdot | \theta)$ be the conditional distribution function over Ω , which we assume has finite density $f(\cdot | \theta)$. Throughout the paper, we assume $\{F(\cdot | \theta)\}$ is ordered in first-order stochastic dominance: we say that θ is "higher" than $\bar{\theta}$ if $F(\omega | \theta) \leq F(\omega | \bar{\theta})$ for all $\omega \in [\underline{\omega}, \bar{\omega}]$, with strict inequality for a positive measure of ω . For

⁴This is also one of the main differences between our paper and recent developments on information control in sender-receiver cheap talk games: see Goltsman et al. (2009), Ivanov (2016), and references therein. Rayo and Segal (2010) also consider an application with *ex ante* private information but the privately informed party is a third party (an advertiser) rather than the receiver. Kolotilin et al. (forthcoming) introduce a privately informed receiver in the framework of Kamenica and Gentzkow (2011) and proves the equivalence between public persuasion and private (discriminatory) persuasion.

some results in the paper, we strengthen the assumption of first-order stochastic dominance. The seller's reservation value is known to be c , with $c \in (\underline{\omega}, \bar{\omega})$.

The timing of our game is as follows. In period 1, first the seller commits to a disclosure policy together with a selling mechanism, which we describe in full detail below. The buyer then decides whether to participate; if he does, the buyer reports his ex ante type to the seller. In period 2, the buyer privately receives new information about his valuation according to the seller's disclosure policy, and reports the additional information to the seller. The seller's mechanism is then implemented, which concludes the game.

A disclosure policy is a menu of signal structures, each associated with a reported type by the buyer. Formally, depending on the buyer's reported type in period 1, the disclosure policy commits the seller to releasing a signal in period 2 about the buyer's true valuation ω . Importantly, the seller does not observe the realized signal; such information policy is known as "private disclosure" in the literature. For simplicity, we also assume that all information of the buyer about ω except his ex ante type θ is under the seller's control; that is, the buyer may not acquire any additional private information about ω on his own. Further, in the main model we assume that the seller's full signal is the buyer's true valuation ω ; that is, the seller can fully disclose ω to the buyer. We relax this assumption in Corollary 1 where the seller's full signal, like the buyer's signal θ , is instead imperfect about ω . There is no disclosure cost to the seller. We focus on two classes of private disclosure policies: direct disclosure and orthogonal disclosure.

A *direct signal structure* $\langle S, \rho \rangle$ is a signal space S and a mapping $\rho : \Omega \rightarrow \Delta S$ that takes the true valuation ω to a distribution $\rho(\cdot | \omega)$ over S ; correspondingly, a *direct disclosure policy* is a menu σ that assigns a direct signal structure $\sigma(\theta)$ to each reported type θ . Direct signal structures are the same as how disclosure rules are defined in the persuasion literature (e.g., Rayo and Segal 2010), and are similar to disclosure strategies in auctions with affiliated values (Milgrom and Weber 1982), except that in our model of private disclosure, the realization of the signal is observable only to the buyer. The full signal structure can be represented by letting $S = \Omega$, and $\rho(s | \omega) = 1$ if $s = \omega$ and $\rho(s | \omega) = 0$ otherwise, while the null signal structure, with no information disclosed, can be modeled by letting S be a singleton. A simple and yet important class of direct signal structures is *binary partitions*: for any partition threshold $\kappa \in [\underline{\omega}, \bar{\omega}]$, let $S = \{s_-, s_+\}$ and let the mapping $\rho(\cdot | \omega)$ be

$$(1) \quad \rho(s | \omega) = \begin{cases} 1 & \text{if } s = s_- \text{ and } \omega < \kappa \\ 1 & \text{if } s = s_+ \text{ and } \omega \geq \kappa. \\ 0 & \text{otherwise} \end{cases}$$

We note that under the binary partition above, the probability for the buyer to receive s_+ is $1 - F(\kappa | \theta)$, which depends on his true ex ante type θ . This dependence is a general property of direct signal structures.

An alternative way of modeling signal structures is to apply the orthogonal decomposition technique introduced in Eső and Szentes (2007). Specifically, let $q = F(\omega | \theta)$ be the orthogonal transformation of the random variable ω . This is the "shock" component in the seller's full signal structure relative to the buyer's ex ante

type θ , because it has the same information content as ω to any ex ante type θ and is distributed uniformly, and thus independently of θ , over $[0, 1]$.⁵ The shock q can be interpreted as the information contained in the seller’s full signal that is “new” to the buyer. We can now define an *orthogonal signal structure* $\langle S, \vec{\rho} \rangle$ as a signal space S and a mapping $\vec{\rho} : [0, 1] \rightarrow \Delta S$ that takes the shock q to a distribution $\vec{\rho}(\cdot | q)$ over S , and correspondingly, an *orthogonal disclosure policy* as a menu σ that assigns to each reported type θ an orthogonal signal structure $\sigma(\theta)$. The full signal structure can be represented by letting $S = [0, 1]$, and $\vec{\rho}(s | q) = 1$ if $s = q$ and $\vec{\rho}(s | q) = 0$ otherwise, while the null signal structure can again be modeled by letting S be a singleton. We can also define a *binary-partition orthogonal structure* for any threshold $\kappa \in [0, 1]$:

$$\vec{\rho}(s | q) = \begin{cases} 1 & \text{if } s = s_- \text{ and } q < \kappa \\ 1 & \text{if } s = s_+ \text{ and } q \geq \kappa. \\ 0 & \text{otherwise} \end{cases}$$

Unlike binary-partition direct structure, here the buyer observes s_+ with the same probability $1 - \kappa$ regardless of his true ex ante type θ , as q is uniformly distributed.

Given any signal structure, a buyer who observes a signal s will update his belief about ω according to Bayes’ rule. Let $v(\theta, s)$ denote type- θ buyer’s posterior estimate of ω given a signal realization s . Under a direct signal structure $\langle S, \rho \rangle$,

$$v(\theta, s) = \frac{\int_{\Omega} \omega \rho(s | \omega) f(\omega | \theta) d\omega}{\int_{\Omega} \rho(s | \omega) f(\omega | \theta) d\omega}.$$

The posterior estimate v depends on θ not only directly through density $f(\omega | \theta)$, but also indirectly through disclosure rule $\rho(\cdot | \omega)$ since ω is correlated with θ . The corresponding expression under an orthogonal signal structure $\langle S, \vec{\rho} \rangle$ is

$$v(\theta, s) = \frac{\int_{[0, 1]} F^{-1}(q | \theta) \vec{\rho}(s | q) dq}{\int_{[0, 1]} \vec{\rho}(s | q) dq}.$$

Thus, under an orthogonal signal structure, v depends on θ only through $F^{-1}(q | \theta)$.

Since both the buyer and the seller are risk-neutral and disclosure is private, the buyer’s posterior estimate v of his true valuation ω (instead of the realized signal disclosed by the seller) is all that matters for any disclosure policy, regardless of his report θ . By the standard revelation principle (see for example, Myerson 1986), for a given disclosure policy σ , we can focus on direct revelation mechanisms, and on equilibria where the buyer truthfully reports θ in period 1 and v in period 2 on the

⁵To see that q is uniformly distributed, observe that $\Pr(F(\omega | \theta) \leq z | \theta) = \Pr(\omega \leq F^{-1}(z | \theta) | \theta) = F(F^{-1}(z | \theta) | \theta) = z$.

equilibrium path.⁶ A direct revelation mechanism specifies the probability of selling the good and the payment from the buyer as functions of the buyer's period 1 report θ of his ex ante type and period 2 report v of his posterior estimate realized under signal structure $\sigma(\theta)$. The seller chooses a disclosure policy σ and a direct mechanism to maximize her expected revenue.

We close this section with a few remarks on the modeling of information disclosure. Direct disclosure and orthogonal disclosure are two different ways that the seller can, depending on the buyer's report of his preexisting private information, selectively control the amount of additional private information the buyer receives to refine his estimate of the true valuation. To see how they fit in applications, consider the oil tract example mentioned at the beginning of the introduction. Imagine that the buyer's ex ante information is about his drilling cost, and the seller controls the kind of test drill the buyer may carry out to acquire information about the geology and productive capacity of the oil tract. Since the test drill is carried out by the buyer, the seller does not observe the test outcome, but she can specify a different test drill to garble the underlying full signal under her control depending on the buyer's report of his ex ante type. This is a natural model of discriminatory direct disclosure. Orthogonal disclosure is the same as direct disclosure if the drilling cost is independent of oil field geology, as the orthogonal component of the full signal is itself. However, if the drilling cost is correlated with the underlying geology, orthogonal disclosure requires the full signal to be orthogonalized first. One can imagine that, for each *reported* ex ante type, the seller specifies a test drill that performs two tasks: first orthogonalize the full signal using the buyer's *true* ex ante type as input, and then garble the orthogonalized full signal to generate a test outcome. This does not seem practical to us, and furthermore, the buyer has to trust that, when he communicates his true type for the purpose of orthogonalization, his communication is not monitored and exploited by the seller. Orthogonal disclosure is a useful theoretical construct that we have introduced to understand how our approach differs from the approach of Esó and Szentes (2007), but is not a natural model of discriminatory information disclosure.

II. Main Results

This section presents our main result that it is not optimal for the seller to release the full signal to all types of the buyer. We first consider in Section IIA the case where the buyer's ex ante type is discrete and the seller's full signal is perfect about the buyer's true valuation. In particular, we show that a direct disclosure policy with a binary partition dominates full disclosure in terms of the seller's revenue (Proposition 1). A binary partition allows the buyer to learn only whether his true valuation is above or below some partition threshold, instead of informing him of the true valuation, as under the full signal.⁷ Intuitively, binary partitions are

⁶Following Myerson (1986), we do not require the buyer to truthfully report his posterior estimate v in period 2 after lying about this ex ante type θ in period 1. This is important because the distribution of the buyer's posterior estimate under a binary partition generally does not have the same support for different ex ante types.

⁷The seller may implement this by allowing the buyer to compare her product to an industry standard in such a way to discover only which one he likes better; we thank Philipp Sadowski for this suggestion. Another practical implementation is offered in Bergemann and Wambach (2015). They assume that the seller chooses the number

effective because they provide the minimal information necessary for allocation purposes while limiting the amount of additional private information disclosed to the buyer and thus lowering the information rent. We then show that full disclosure remains suboptimal if the seller's full signal is noisy about the buyer's true valuation (Corollary 1) or if the buyer's ex ante type is continuous (Proposition 2).

A. Discrete Types

Suppose that $\Theta = \{\theta_1, \dots, \theta_n\}$, with $i > j$ implying that θ_i is higher than θ_j in first-order stochastic dominance. Let ϕ_i denote the probability of the type being θ_i , and let $\Phi_i = \sum_{j=1}^i \phi_j$.

Under the full disclosure policy, the seller allows each reported type θ_i to privately learn his true valuation ω . Suppose that the seller uses a deterministic selling mechanism.⁸ Any such mechanism can be represented as a menu of option contracts (a_i, p_i) , where a_i is the nonrefundable advance payment in period 1 and p_i is the strike price in period 2 for each reported type θ_i , $i = 1, \dots, n$. Incentive compatibility constraints in period 2 imply that, regardless of the contract (a_i, p_i) chosen by any type θ_j , the buyer buys if and only if his true valuation ω is above p_i . Given this, period 1 incentive compatibility and individual rationality under full disclosure require

$$(IR_i) \quad -a_i + \int_{p_i}^{\bar{\omega}} (\omega - p_i) dF(\omega | \theta_i) \geq 0, \quad \forall i;$$

$$(IC_{ij}) \quad -a_i + \int_{p_i}^{\bar{\omega}} (\omega - p_i) dF(\omega | \theta_i) \geq -a_j + \int_{p_j}^{\bar{\omega}} (\omega - p_j) dF(\omega | \theta_i), \quad \forall i, j.$$

Since the ex ante types are ordered by first-order stochastic dominance, the incentive compatibility constraints above require p_i to be weakly decreasing in i .⁹

We will argue that the seller can strictly increase her revenue by using a partial and discriminatory disclosure policy. This alternative disclosure policy $\hat{\sigma}$ changes the signal structure only when the buyer reports to be the lowest type θ_k served under the original mechanism: the seller allows the buyer to privately learn his true valuation ω in period 2 if he reports any type higher than θ_k in period 1, but a buyer reporting θ_k is only allowed to learn whether or not his true valuation is above p_k . Formally, the signal structure $\hat{\sigma}_k$ for the reported type θ_k is a binary partition of ω at p_k , as given by (1). To make the argument as simple as possible, we make two related assumptions about the original mechanism. First, we assume full coverage; that is, the lowest type served under the mechanism is θ_1 , with $p_1 < \bar{\omega}$. Second, we

of features of her product to reveal sequentially, and the buyer determines whether he likes each feature or not, stopping at the first one he does not like as he has lexicographical preferences. Although the buyer knows his true valuation when he stops before inspecting all offered features, from an analytical point of view, the resulting signal structure can be treated as a binary partition.

⁸In general, stochastic mechanisms may be optimal under full disclosure if there are multiple types. If ex ante type is binary, however, one can show that the optimal contract under full disclosure is indeed deterministic. See Courty and Li (2000) for more discussion on stochastic sequential screening mechanisms.

⁹To see this, add up IC_{ij} and IC_{ji} and use integration by parts to get $\int_{p_i}^{p_j} (F(\omega | \theta_j) - F(\omega | \theta_i)) d\omega \geq 0$. Then, $F(\omega | \theta_i) \leq F(\omega | \theta_j)$ implies that $p_i \leq p_j$.

assume that type θ_2 is strictly higher than type θ_1 in first-order stochastic dominance, such that

$$(2) \quad F(p_1|\theta_2) < F(p_1|\theta_1).$$

These two assumptions are sufficient for Proposition 1 but neither is necessary,¹⁰ and both are satisfied if the original mechanism (a_i, p_i) under full disclosure is revenue-maximizing.¹¹

PROPOSITION 1: *Suppose that ex ante type is discrete. For any incentive compatible and individually rational menu of option contracts (a_i, p_i) under full disclosure that satisfies $p_1 < \bar{\omega}$ and $F(p_1|\theta_2) < F(p_1|\theta_1)$, there exists an alternative menu with partial and discriminatory disclosure that yields a strictly greater revenue.*

PROOF:

Consider the alternative disclosure policy $\hat{\sigma}$ together with a modified mechanism (\hat{a}_i, \hat{p}_i) , given by

$$\begin{aligned} \hat{p}_1 &= p_1 + \delta, & \hat{a}_1 &= a_1 - \delta(1 - F(p_1|\theta_1)); \\ \hat{p}_i &= p_i, & \hat{a}_i &= a_i + \delta(F(p_1|\theta_1) - F(p_1|\theta_2)), \quad \forall i \geq 2 \end{aligned}$$

where δ satisfies

$$0 < \delta \leq \min_j \int_{p_1}^{\bar{\omega}} \frac{\omega dF(\omega|\theta_j)}{1 - F(\omega|\theta_j)} - p_1.$$

That is, δ is chosen such that all types choosing contract (\hat{a}_1, \hat{p}_1) in period 1 will buy at price \hat{p}_1 if the seller only discloses that their valuation is above p_1 in period 2. We claim that (\hat{a}_i, \hat{p}_i) is individually rational and incentive compatible under $\hat{\sigma}$.

First, consider a type- θ_1 buyer. By choice of δ , a truthful type θ_1 will buy the good at \hat{p}_1 upon learning only that his valuation is greater than p_1 in period 2. By construction, the expected payoff of type θ_1 from truthful reporting in period 1 is unchanged

¹⁰For the full coverage assumption, if θ_k is the lowest type served with $2 \leq k \leq n - 1$, then the proof goes through with θ_k playing the same role as θ_1 , as long as type θ_{k-1} strictly prefers not to participate under the original mechanism. This is because the seller can always make the increase δ in the strike price p_k sufficiently small to ensure that no type below θ_k wants to participate in θ_k 's modified contract. For the strict first-order stochastic dominance assumption, if instead $F(p_1|\theta_2) = F(p_1|\theta_1)$, slightly more involved modifications of the original mechanism would still work when $F(p_1|\theta_3) < F(p_1|\theta_1)$ and $F(p_2|\theta_3) < F(p_2|\theta_2)$. The seller can partition the true valuation ω for both reported type θ_1 and type θ_2 , with threshold p_1 and p_2 respectively, set \hat{a}_1 and \hat{a}_2 to keep their truth-telling payoffs unchanged, and uniformly raise the advance payment for all types higher than θ_2 .

¹¹Under full disclosure, if certain regularity conditions in sequential screening are satisfied, the optimal allocations are determined by point-wise maximization of the discrete-type dynamic virtual surplus function, given by

$$J_i(\omega) = \omega - c + \frac{1 - \Phi_i F(\omega|\theta_{i+1}) - F(\omega|\theta_i)}{\phi_i f(\omega|\theta_i)},$$

for each type θ_i , $i = 1, \dots, n - 1$, with $J_n(\omega) = \omega - c$ (see, e.g., Courty and Li 2000). That is, $J_i(p_i) = 0$. Then, $p_1 < \bar{\omega}$ because $J_1(\bar{\omega}) = \bar{\omega} - c > 0$; and (2) holds because otherwise we would have $J_1(p_1) = p_1 - c > 0$.

relative to the original mechanism, so the individual rationality constraint for type θ_1 remains satisfied. Moreover, type θ_1 has no incentive to mimic higher types, since the option contracts for higher types have the same strike prices as under the original menu but larger advance payments. Thus, all incentive compatibility constraints are satisfied for type θ_1 .

Next, consider a type- θ_i buyer with $i \geq 2$. Since the strike prices are unchanged and the advance payments are increased uniformly for all types higher than θ_1 , he has no incentive to mimic some other type θ_j , $j \geq 2$. The choice of δ ensures that after deviating to type θ_1 's new option contract, type θ_i will buy the good if he learns only that his true valuation ω is above p_1 , so his expected payoff from mimicking θ_1 is

$$-a_1 + \int_{p_1}^{\bar{\omega}} (\omega - p_1) dF(\omega | \theta_i) - \delta(F(p_1 | \theta_1) - F(p_1 | \theta_i)).$$

By first-order stochastic dominance and IC_{i1} under the original mechanism, the above is less than or equal to type θ_i 's expected payoff from truth-telling, given by

$$-a_i + \int_{p_i}^{\bar{\omega}} (\omega - p_i) dF(\omega | \theta_i) - \delta(F(p_1 | \theta_1) - F(p_1 | \theta_2)).$$

Finally, as under the original mechanism, IR_i is implied by IR_1 and IC_{i1} under the modified mechanism.

The allocation for all types, and hence the trade surplus created, under the modified mechanism is the same as under the original one. The expected payoff for all types other than θ_1 is lowered by $\delta(F(p_1 | \theta_1) - F(p_1 | \theta_2))$, which is strictly positive under condition (2). The seller's revenue, which is the difference between the total surplus and the buyer's expected payoff, is strictly greater. ■

Since Proposition 1 considers all deterministic, incentive compatible, and individually rational selling mechanisms under full disclosure, an immediate corollary is that full disclosure is not optimal if it leads to a deterministic mechanism. To understand why binary partitions are more effective than full disclosure in facilitating price discrimination, we first note that in our setup the buyer's final decision is either to buy or not to buy, so full disclosure of the buyer's true valuation is not necessary to implement a given allocation. For any type θ_i or set of types served under the original mechanism, if the seller discloses to reported type θ_i only whether his true valuation ω is above or below the strike price p_i , type θ_i will make the same purchase decision as if the seller discloses ω itself. Moreover, any other type contemplating a deviation to (a_i, p_i) by misreporting θ_i in period 1 will have a weakly lower deviation payoff because of the loss of information due to the partitioning of ω . With the same mechanism as under full disclosure, the partitioning of ω thus implements the same allocation and revenue for the seller.

The seller can do strictly better than full disclosure, however, by manipulating the terms of trade together with a partitioning of ω . After partitioning ω only for reported lowest type θ_1 at the original strike price p_1 , the seller can raise the strike price to $p_1 + \delta$. As long as δ is small, a truthful type- θ_1 buyer buys in period 2 if he learns only that ω is above p_1 , and thus the trading probability $1 - F(p_1 | \theta_1)$ is the same as under full disclosure. Moreover, the payoff of type θ_1 stays unchanged if the seller compensates for the increase in the strike price by lowering the period 1

advance payment a_1 by $\delta(1 - F(p_1 | \theta_1))$. Such manipulations, however, affect the deviation payoff for any higher type θ_i pretending to be θ_1 . After the deviation, type θ_i only learns that his valuation is above or below the partition threshold p_1 , so the probability of purchase is $1 - F(p_1 | \theta_i)$, which is higher than that for type θ_1 by first-order stochastic dominance. This means that the reduction in the advance payment a_1 is not enough to compensate type θ_i for the increase in the strike price p_1 , and thus incentives for type θ_i to mimic θ_1 are reduced compared to the original mechanism. The seller can then raise the advance payments uniformly for all higher types, while keeping their strike prices unchanged, without violating any of their incentive constraints. The allocation and trade surplus remain the same for all higher types, but their information rent is strictly lower, and thus the seller's revenue is strictly higher.

The power of discriminatory direct disclosure relative to full disclosure in facilitating price discrimination may be understood more broadly as follows. Under full disclosure, for each contract in the menu, the strike price simultaneously determines the allocation and defines the terms of trade, and these two roles are tied together. If the seller retains the strike price but replaces the full signal structure by a binary partition with the same strike price as the partition threshold, she can replicate the same allocation and revenue as under full disclosure. By allowing the strike price to differ from the partition threshold, the seller can now control the allocation and terms of trade separately. Due to first-order stochastic dominance, an increase in the strike price has a differential impact on different buyer types who choose the same option contract, so the seller acquires another instrument to discriminate among buyer types to improve her revenue.

We conclude the analysis of the discrete-type setting by showing that Proposition 1 remains valid if the seller's full signal $\zeta \in [\underline{\zeta}, \bar{\zeta}]$, just as the buyer's ex ante type $\theta \in \Theta$, is a noisy signal of $\omega \in [\underline{\omega}, \bar{\omega}]$. For simplicity, we restrict our analysis to two ex ante buyer types, with $\Theta = \{\theta_1, \theta_2\}$. Let $V_i(\zeta)$ be the expected valuation of a type- θ_i buyer conditional on realized ζ ; that is, $V_i(\zeta) = E[\omega | \theta_i, \zeta]$. For each $i = 1, 2$, denote by $g(\zeta | \theta_i)$ the density function of ζ conditional on the buyer's ex ante type θ_i , and let $G(\zeta | \theta_i)$ be the corresponding conditional distribution function. We assume that $V_i(\zeta)$ is strictly increasing in ζ to capture the idea that a higher realization of ζ represents a more favorable signal about ω . We continue to assume that θ_2 represents a type higher than θ_1 in the sense that $V_2(\zeta) > V_1(\zeta)$ and $G(\zeta | \theta_2) < G(\zeta | \theta_1)$ for all $\zeta \in (\underline{\zeta}, \bar{\zeta})$.¹² In this environment, we can define direct disclosure of ζ similarly as in Section I; we focus on binary partitions of ζ .

COROLLARY 1: *Suppose that $\Theta = \{\theta_1, \theta_2\}$, and that the seller's signal ζ is noisy about the buyer's true valuation ω , with $V_i(\underline{\zeta})$ strictly increasing in ζ , and $V_2(\zeta) > V_1(\zeta)$ and $G(\zeta | \theta_2) < G(\zeta | \theta_1)$ for all $\zeta \in (\underline{\zeta}, \bar{\zeta})$. If the optimal menu of contracts under full disclosure is deterministic, then there exists an alternative*

¹²The primitive of this environment is the joint density $f(\omega, \zeta, \theta)$. A sufficient condition for the ordering of θ and the strict monotonicity of $V_i(\zeta)$ in ζ , is that ω, ζ , and θ are strictly affiliated or equivalently f is strictly log-supermodular. See Milgrom (2004) for the definition of affiliation and its properties (Theorems 5.4.3 and 5.4.4). The direct disclosure model of Section I is a special case of the present formulation, with $\zeta = \omega$ and hence $V_2(\zeta) = V_1(\zeta)$. The orthogonal disclosure model can also be incorporated, by assuming that $\zeta = F(\omega | \theta)$ and is thus independent of θ , with $G(\zeta | \theta_2) = G(\zeta | \theta_1)$ for all ζ .

menu with partial and discriminatory direct disclosure of ζ that yields a strictly higher revenue.

The idea behind Corollary 1 is similar to the one behind Proposition 1. We sketch the argument below. Suppose the optimal mechanism under full disclosure is given by a menu of option contracts $(a_1, p_1; a_2, p_2)$. Since both $V_1(\zeta)$ and $V_2(\zeta)$ are increasing, a type- θ_1 buyer will buy at price p_1 if the signal realization ζ exceeds some cutoff ζ_1 satisfying $V_1(\zeta_1) = p_1$, while a deviating type- θ_2 buyer will buy if instead ζ exceeds ζ_2 implicitly defined by $V_2(\zeta_2) = p_1$. Since $V_2(\zeta) > V_1(\zeta)$ and $G(\zeta|\theta_2) < G(\zeta|\theta_1)$, optimality implies that the individual rationality constraint IR_1 for type θ_1 and the incentive compatibility constraint IC_{21} for type θ_2 bind under full disclosure, and pin down the information rent for type θ_2 . Now suppose the seller continues to fully disclose signal ζ for reported type θ_2 , but discloses only whether ζ is above or below ζ_1 for reported θ_1 . At the same time, the seller changes the strike price from p_1 to $p_1 + \delta$ with $\delta > 0$. Since $V_1(\zeta)$ is strictly increasing, for sufficiently small δ type θ_1 continues to buy the good if he learns only that the realized ζ is above ζ_1 , and does not buy otherwise. As in the case where the seller’s signal is perfect, the increase in p_1 hurts the deviating type θ_2 more than the truth-telling type θ_1 because $G(\zeta|\theta_2) < G(\zeta|\theta_1)$ for all $\zeta \in (\zeta_1, \bar{\zeta})$. As a result, when the seller reduces the advance payment for type θ_1 to bind IR_1 , the new information rent of type θ_2 becomes lower.¹³ This allows the seller to raise the advance payment for type θ_2 to bind IC_{21} . By construction, the modified mechanism with the binary-partition direct disclosure replicates the allocation and hence the total surplus under the original mechanism $(a_1, p_1; a_2, p_2)$ with full disclosure. The seller’s revenue is strictly greater, establishing that binary-partition direct disclosure of ζ dominates full disclosure.

B. Continuous Types

Suppose that the buyer’s ex ante type θ is drawn from distribution $\Phi(\cdot)$ with density $\phi(\theta) > 0$ for all $\theta \in \Theta \equiv [\underline{\theta}, \bar{\theta}]$. Types are ordered: $\theta > \tilde{\theta}$ implies θ is higher than $\tilde{\theta}$ in first-order stochastic dominance. We assume that distribution $F(\omega|\theta)$ is continuously differentiable with respect to θ and that distributions $\{F(\omega|\theta)\}$ share the same support at least “at the top”; that is, there exists $\kappa \in [\underline{\omega}, \bar{\omega})$ such that $f(\omega|\theta) > 0$ for all θ and all $\omega \in [\kappa, \bar{\omega}]$. For simplicity, we assume as in our main model that the seller’s signal is perfect about the buyer’s valuation, but the analysis again goes through with appropriate modifications as in Corollary 1 if the seller’s signal is noisy.

¹³There are three cases in comparing the new information rent \hat{U}_2 with the old rent U_2 for type θ_2 . If after misreporting his type as θ_1 , type θ_2 buys at $p_1 + \delta$ only when he learns that his signal ζ is above ζ_1 , then the rent difference $U_2 - \hat{U}_2$ is

$$\delta(G(\zeta_1|\theta_1) - G(\zeta_1|\theta_2)) + \int_{\zeta_2}^{\zeta_1} (V_2(\zeta) - p_1) g(\zeta|\theta_2) d\zeta,$$

where the first term is positive because $G(\zeta_1|\theta_1) > G(\zeta_1|\theta_2)$, and the second term is positive because $V_2(\zeta) > V_1(\zeta)$ implies $\zeta_1 > \zeta_2$. If type θ_2 always buys, then the rent difference $U_2 - \hat{U}_2$ becomes

$$\delta(G(\zeta_1|\theta_1) - G(\zeta_1|\theta_2)) + \int_{\zeta}^{\zeta_2} (p_1 - V_2(\zeta)) g(\zeta|\theta_2) d\zeta,$$

which is again positive because $G(\zeta_1|\theta_1) > G(\zeta_1|\theta_2)$. Finally, if type θ_2 never buys, $\hat{U}_2 = 0$ and so $U_2 - \hat{U}_2 > 0$.

Under full disclosure, the seller reveals the true valuation ω to each reported type θ . Consider any menu of incentive compatible and individually rational option contracts $(a(\theta), p(\theta))_{\theta \in \Theta}$, as in the sequential screening setting of Courty and Li (2000). Incentive compatibility implies that

$$-a(\theta) + \int_{p(\theta)}^{\bar{\omega}} (\omega - p(\theta)) f(\omega | \theta) d\omega \geq -a(\tilde{\theta}) + \int_{p(\tilde{\theta})}^{\bar{\omega}} (\omega - p(\tilde{\theta})) f(\omega | \theta) d\omega$$

for all $\theta, \tilde{\theta} \in \Theta$. As in the discrete-type setting, under the assumption of first-order stochastic dominance, the incentive compatibility constraints above imply that $p(\theta)$ is weakly decreasing in θ and thus differentiable almost everywhere.

We show below that there is a discriminatory direct disclosure policy consisting of a binary-partition signal structure with threshold $p(\theta)$ for each θ that, together with a modified menu of option contracts, strictly improves the seller’s revenue. The binary-partition signal structure for each reported type θ reveals to the buyer whether his true valuation ω is above or below the original strike price $p(\theta)$. In contrast to the discrete-type setting where only downward incentive constraints bind, with continuous ex ante type, downward and upward incentive constraints coincide. Moreover, unlike in Proposition 1 where it is sufficient to change the signal structure just for a single buyer type because each type has a strictly positive measure, here the new disclosure policy has to replace full disclosure by a binary partition for all types. With continuous ex ante type, it is harder to ensure that the modified menu of option contracts is incentive compatible. Our proof requires a strengthening of first-order stochastic dominance to hazard rate dominance. Formally, we assume that ex ante types are ordered in hazard rates under full disclosure: for $\theta \geq \tilde{\theta}$,

$$\frac{f(\omega | \theta)}{1 - F(\omega | \theta)} \leq \frac{f(\omega | \tilde{\theta})}{1 - F(\omega | \tilde{\theta})}$$

for all $\omega \in [\underline{\omega}, \bar{\omega}]$ with strict inequality for a positive measure of ω .¹⁴

PROPOSITION 2: *Suppose that ex ante types are ordered in hazard rate dominance. If a menu of option contracts $(a(\theta), p(\theta))_{\theta \in \Theta}$ with differentiable $p(\theta)$ is incentive compatible and individually rational under full disclosure, the set $\{\theta : \partial F(p(\theta) | \theta) / \partial \theta < 0\}$ has a positive measure and $p(\theta) < \bar{\omega}$ for all θ , then there exists a binary-partition direct disclosure policy that strictly increases the seller’s revenue.*

PROOF:

Consider the direct disclosure policy of assigning to each reported type $\theta \in \Theta$ a binary partition of ω at threshold $p(\theta)$, together with a modified menu $(\hat{a}(\theta), \hat{p}(\theta))_{\theta \in \Theta}$. The new strike price $\hat{p}(\theta)$ is given by

$$(3) \quad \hat{p}(\theta) = p(\theta) + \delta,$$

¹⁴Hazard rate dominance, however, is not necessary for our result, as one can see from Example 1 where ex ante types are ordered only in first-order stochastic dominance but full disclosure is not optimal.

where $\delta > 0$ is sufficiently small so that $\hat{p}(\theta) < \bar{\omega}$ for all θ ; this is feasible because $p(\theta) < \bar{\omega}$ for all θ by assumption. The new advance payment $\hat{a}(\theta)$ is given by

$$(4) \quad \hat{a}(\theta) = \int_{p(\theta)}^{\bar{\omega}} (1 - F(\omega|\theta)) d\omega - (1 - F(p(\theta)|\theta)) \delta - U(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} \left(\int_{p(t)}^{\bar{\omega}} \left(-\frac{\partial F(\omega|t)}{\partial t} \right) d\omega + \frac{\partial F(p(t)|t)}{\partial t} \delta \right) dt,$$

where $U(\underline{\theta})$ is the expected payoff of $\underline{\theta}$ in the original menu.

First, consider the purchase decision of type- θ buyer after he chooses any contract $(\hat{a}(\tilde{\theta}), \hat{p}(\tilde{\theta}))$ from the modified menu. If he learns that his valuation is below $p(\tilde{\theta})$, it is optimal not to buy because $\hat{p}(\tilde{\theta}) > p(\tilde{\theta})$. If his valuation is revealed to be above $p(\tilde{\theta})$, then for sufficiently small but positive δ , the buyer will buy at $\hat{p}(\tilde{\theta})$, as for any $\theta, \tilde{\theta}$,

$$p(\tilde{\theta}) < \int_{p(\tilde{\theta})}^{\bar{\omega}} \frac{\omega f(\omega|\theta) d\omega}{1 - F(p(\tilde{\theta})|\theta)}.$$

Second, we argue that for positive and sufficiently small δ such that

$$1 - \frac{f(p(\tilde{\theta})|\theta)}{1 - F(p(\tilde{\theta})|\theta)} \delta > 0$$

for all $\tilde{\theta}, \theta$, the modified menu is incentive compatible in period 1. Since $p(\theta) < \bar{\omega}$ for all θ , such δ exists. The expected payoff to a type- θ buyer from reporting $\tilde{\theta}$ is

$$\hat{U}(\theta, \tilde{\theta}) = -\hat{a}(\tilde{\theta}) + \int_{p(\tilde{\theta})}^{\bar{\omega}} (1 - F(\omega|\theta)) d\omega - (1 - F(p(\tilde{\theta})|\theta)) \delta.$$

Suppose that $\tilde{\theta} < \theta$. Since $p(\tilde{\theta})$ is decreasing, we have

$$\begin{aligned} \frac{\partial \hat{U}(\theta, \tilde{\theta})}{\partial \tilde{\theta}} &= -\frac{d\hat{a}(\tilde{\theta})}{d\tilde{\theta}} - \left(1 - \frac{f(p(\tilde{\theta})|\theta)}{1 - F(p(\tilde{\theta})|\theta)} \delta \right) (1 - F(p(\tilde{\theta})|\theta)) \frac{dp(\tilde{\theta})}{d\tilde{\theta}} \\ &\geq -\frac{d\hat{a}(\tilde{\theta})}{d\tilde{\theta}} - \left(1 - \frac{f(p(\tilde{\theta})|\tilde{\theta})}{1 - F(p(\tilde{\theta})|\tilde{\theta})} \delta \right) (1 - F(p(\tilde{\theta})|\tilde{\theta})) \frac{dp(\tilde{\theta})}{d\tilde{\theta}}, \end{aligned}$$

where the inequality follows from first-order stochastic dominance and hazard rate dominance. One can easily verify from the construction of the advance payments (4) that the last line above is zero. By integration we have $\hat{U}(\theta) \equiv \hat{U}(\theta, \theta) \geq \hat{U}(\theta, \tilde{\theta})$. The case of $\theta < \tilde{\theta}$ can be argued analogously.

Third, we claim that for sufficiently small δ , the new menu $(\hat{a}(\theta), \hat{p}(\theta))_{\theta \in \Theta}$ is individually rational. By the envelope theorem, from the expression of $\hat{U}(\theta, \theta)$ we have

$$(5) \quad \frac{d\hat{U}(\theta)}{d\theta} = \int_{p(\theta)}^{\bar{\omega}} \left(-\frac{\partial F(\omega|\theta)}{\partial \theta} \right) d\omega + \frac{\partial F(p(\theta)|\theta)}{\partial \theta} \delta.$$

By construction, $\hat{U}(\underline{\theta}) = U(\underline{\theta}) \geq 0$. For sufficiently small δ , we have $d\hat{U}(\theta)/d\theta \geq 0$, and thus $\hat{U}(\theta) \geq 0$ for all θ .

Finally, we argue that the seller’s revenue is strictly higher under the modified menu $(\hat{a}(\theta), \hat{p}(\theta))_{\theta \in \Theta}$, because it generates the same trade surplus but leads to lower information rent for each type than the original menu. Under the original menu, the expected payoff to type θ from the contract $(a(\tilde{\theta}), p(\tilde{\theta}))$ is

$$U(\theta, \tilde{\theta}) = -a(\tilde{\theta}) + \int_{p(\tilde{\theta})}^{\bar{\omega}} (1 - F(\omega | \theta)) d\omega.$$

Defining $U(\theta) \equiv U(\theta, \theta)$, from the envelope theorem we have

$$(6) \quad \frac{dU(\theta)}{d\theta} = \int_{p(\theta)}^{\bar{\omega}} \left(-\frac{\partial F(\omega | \theta)}{\partial \theta} \right) d\omega.$$

Since $U(\underline{\theta}) = \hat{U}(\underline{\theta})$ and since the set $\{\theta : \partial F(p(\theta) | \theta) / \partial \theta < 0\}$ has a positive measure, it follows from (5) and (6) that $U(\theta) \geq \hat{U}(\theta)$ for all $\theta > \underline{\theta}$, and $U(\theta) > \hat{U}(\theta)$ for a positive measure of types. ■

The intuition behind Proposition 2 is similar to the one behind Proposition 1. When we replace the full signal for each reported θ by a binary partition with partition threshold $p(\theta)$, the same trade surplus for a truthful type- θ buyer is generated under the original menu. With the strike price uniformly raised by δ to a slightly higher $\hat{p}(\theta)$, the trade surplus remains unchanged because the truthful type θ buys after learning only that his valuation is above $p(\theta)$. For sufficiently small δ , the modified menu remains incentive compatible. The assumption of full coverage in the original menu (i.e., $p(\theta) < \bar{\omega}$ for all θ) and the assumption of differentiability of $p(\theta)$ are used to facilitate the verification of incentive compatibility.¹⁵

The rate of increase $d\hat{U}(\theta)/d\theta$ in buyer’s information rent under the modified menu is related to $dU(\theta)/d\theta$ under the original menu as follows:

$$\frac{d\hat{U}(\theta)}{d\theta} = \frac{dU(\theta)}{d\theta} + \frac{\partial F(p(\theta) | \theta)}{\partial \theta} \delta.$$

Due to higher strike prices ($\delta > 0$) and first-order stochastic dominance, the buyer’s information rent increases more slowly under the modified menu. Since under the modified menu the lowest type $\underline{\theta}$ earns the same rent, the overall information rent is lower and thus the seller’s revenue is strictly higher, if the set $\{\theta : \partial F(p(\theta) | \theta) / \partial \theta < 0\}$ has a positive measure.

¹⁵These two assumptions are satisfied if $(a(\theta), p(\theta))$ is revenue-maximizing under full disclosure. Under regularity conditions, the continuous-type dynamic virtual surplus function $J(\omega, \theta)$ under full disclosure, given by

$$\omega - c + \frac{1 - \Phi(\theta)}{\phi(\theta)} \frac{\partial F(\omega | \theta) / \partial \theta}{f(\omega | \theta)},$$

is monotone in both ω and θ (see Courty and Li 2000 and Eső and Szentes 2007). Then, optimal $p(\theta)$ is determined by $J(p(\theta), \theta) = 0$. Full coverage follows because $J(\underline{\theta}, \bar{\omega}) = \bar{\omega} - c > 0$ implies that $p(\underline{\theta}) < \bar{\omega}$, and because $p(\theta)$ is weakly decreasing. The differentiability of $p(\theta)$ follows from the implicit function theorem if $J(\omega, \theta)$ is continuously differentiable and $\partial J(p(\theta), \theta) / \partial \omega \neq 0$.

III. Discussion

Our main results, Proposition 2 in particular, stand in contrast to Eső and Szentes (2007) who argue that full disclosure is optimal in a sequential screening framework similar to ours, with the buyer's ex ante type continuous and ordered by first-order stochastic dominance. Their indirect approach is based on an irrelevance theorem that the seller's hypothetical revenue when she can observe any released signal is attainable in the original problem under full disclosure, and their claim that the hypothetical revenue is an upper bound in the original problem. In Section IIIA, after reviewing their arguments, we explain that their result is due to the implicit restriction that the seller can only use orthogonal disclosure as defined in Section I. We use an explicit example to show that orthogonal disclosure cannot replicate direct disclosure. In Section IIIB, we further illustrate the difference between our approach and that of Eső and Szentes, using a discrete-type setting where the irrelevance theorem of Eső and Szentes (2007) has been shown to fail (Krähmer and Strausz 2015b). With two ex ante types, discriminatory direct disclosure yields revenue strictly greater than the hypothetical revenue when the two types become sufficiently "close" to each other. This result shows that the suboptimality of full disclosure is independent of whether or not the irrelevance theorem holds. It also explains why, even though the revenue gap between full disclosure and the hypothetical setting disappears when the type distribution becomes continuous and the irrelevance theorem holds, the revenue gap between full disclosure and discriminatory direct disclosure persists.

A. Relation to Eső and Szentes (2007)

In a sequential screening framework as in Section IIB, Eső and Szentes (2007) argue that full disclosure is optimal. Their argument takes the following four steps. First, using the orthogonal decomposition approach, they define shock $q = F(\omega | \theta)$ as the new information available to the buyer. The shock q is uniformly distributed over $[0, 1]$ for all ex ante type θ , so it is orthogonal to θ , and has the same information content as the true valuation ω for any type θ . Second, they define a *hypothetical* setting where the seller can release, *and observe*, the realized orthogonal shock q to the buyer. The seller cannot infer anything about the buyer's ex ante type θ by observing q as q is orthogonal to θ , while the buyer has the same private ex ante information as in the original setting but any additional information he receives in period 2 is public. The seller's hypothetical problem is to find a menu of contracts, $(x(\theta, q), y(\theta, q))_{\theta \in \Theta, q \in [0, 1]}$, to maximize her revenue, where $x(\theta, q)$ and $y(\theta, q)$ are the allocation and transfer for each reported type θ conditional on the realized q , respectively. Third, they establish an irrelevance theorem that, under certain regularity conditions, the optimal mechanism in the hypothetical setting is implementable in the original setting,¹⁶ and thus the seller's maximal revenue in the hypothetical

¹⁶Formally, the regularity conditions require the type distribution $\Phi(\theta)$ to have a monotone hazard rate and the informativeness measure $(\partial F(\omega | \theta) / \partial \theta) / f(\omega | \theta)$ to be increasing in both ω and θ . As a result, the virtual surplus function under full disclosure is increasing in both ω and θ (see footnote 15). Essentially, these regularity conditions ensure that the solution to the seller's hypothetical problem is implementable in the original setting under full disclosure.

setting is attained in the original setting under full disclosure. Fourth, they claim that the irrelevance theorem indirectly establishes the optimality of full disclosure, because the seller’s revenue from full disclosure in the hypothetical setting is an upper bound on what the seller can achieve in the original setting.

Since full disclosure of the buyer’s shock q leads to the same maximal revenue to the seller as full disclosure of the true valuation ω , the claim of Eső and Szentes (2007) that full disclosure is optimal directly contrasts with our Proposition 2. With an indirect approach, Eső and Szentes (2007) do not explicitly model partial or discriminatory disclosure policies. However, if they allow discriminatory disclosure by having the seller garble the shock q according to the reported type in the way we define in Section I (i.e., orthogonal disclosure), then the hypothetical revenue is indeed an upper bound on the seller’s revenue in the original setting. This follows from two observations. First, for any orthogonal disclosure policy in the original setting where the seller garbles the shock q depending on the type report, there is a corresponding hypothetical setting in which the seller observes and garbles the realized shock q . Since the seller can always commit to not using the information about the shock, she cannot do worse in the corresponding hypothetical setting. Second, in the hypothetical setting, there is no loss in assuming full disclosure of the shock.¹⁷ Thus, in our language, what Eső and Szentes (2007) show is that their irrelevance theorem implies the optimality of full disclosure among all orthogonal disclosure policies with continuous ex ante buyer type.

The irrelevance theorem of Eső and Szentes (2007) does *not* imply, however, that full disclosure is optimal when the seller is not restricted to orthogonal disclosure. This is already established by our Proposition 2, but our approach is direct. In terms of their indirect argument for the optimality of full disclosure reviewed at the beginning of this subsection, the issue lies in the last step if the seller is not restricted to orthogonal disclosure. The maximal revenue in the hypothetical setting is generally not an upper bound on the seller’s revenue if direct disclosure policies are also available, as demonstrated in the following example.

Example 1: Suppose the seller’s reservation value $c = 1/2$. The buyer’s ex ante type θ is uniformly distributed, with distribution function $\Phi(\theta) = 2\theta - 1$ and support $[1/2, 1]$. Suppose the distribution of a type- θ buyer’s true valuation ω is also uniform with $F(\omega|\theta) = 1 - (1 - \omega)\theta$ and support $[1 - 1/\theta, 1]$. Clearly, distributions $\{F(\cdot|\theta)\}$ are ordered by first-order stochastic dominance. Moreover, it is easy to verify that the regularity conditions in Eső and Szentes (2007) are satisfied.

In the hypothetical setting where the seller observes the realization of $q = F(\omega|\theta)$ and offers a menu of contracts conditioned on q and reported θ , one can show that the optimal mechanism sets allocation $x(\theta, q) = 1$ if $q \geq 1 - \theta^2/2$ and 0 otherwise. Note that $q \geq 1 - \theta^2/2$ is equivalent to $\omega \geq 1 - \theta/2$, so there is allocation

¹⁷ More precisely, because shock q is independent of the buyer’s ex ante type, any incentive-compatible contract based on some garbled shock q remains incentive compatible if the seller replaces it with one based on the original q . To see this, suppose that the seller garbles q through a mapping $\bar{p} : [0, 1] \rightarrow \Delta S$ for reported type θ , with corresponding contract $(x(\theta, s), y(\theta, s))$ that conditions the probability of selling the good and the payment to the seller on $s \in S$. Then, the expected payoff from choosing the contract $(x(\theta, s), y(\theta, s))$ for any type is unaffected if the seller changes the contract to $(\hat{x}(\theta, q), \hat{y}(\theta, q))$, where $\hat{x}(\theta, q) = \int_S x(\theta, s) \bar{p}(s|q) ds$ and $\hat{y}(\theta, q) = \int_S y(\theta, s) \bar{p}(s|q) ds$. We are indebted to Roland Strausz for providing this argument in a private communication.

distortion everywhere except at the top and the resulting hypothetical revenue is strictly below the expected full surplus associated with efficient allocation. The same distorted allocation and hypothetical revenue can be implemented by a menu of option contracts $(a(\theta), p(\theta))_{\theta \in \Theta}$ with strike price $p(\theta) = 1 - \theta/2$ in the original setting under full disclosure. Thus, the irrelevance theorem holds.

Now consider the following partial disclosure policy and selling mechanism. The seller discloses to all buyer types whether ω is above or below $1/2$, and charges price $3/4$ in period 2. This disclosure policy, together with the posted price, implements the efficient allocation and extracts the full surplus.

In the example above, the regularity conditions in Eső and Szentes (2007) are satisfied, but full disclosure is not optimal in the original setting. Their proof of the irrelevance theorem requires that distributions $\{F(\omega|\theta)\}$ share a common support, which is not satisfied in this example. It is not important for our argument, however, because we are able to explicitly construct a menu of contracts and directly verify their irrelevance theorem: the maximal revenue in the hypothetical setting is indeed attained through full disclosure in the original setting. But the seller can do better in the original setting than in the hypothetical setting if she is allowed to use direct disclosure policies that are partitions of the buyer's true valuation ω .

The reason that the hypothetical revenue is generally not a revenue upper bound for a seller with access to direct disclosure policies, and binary partitions of the buyer's true valuation ω in particular, can be understood as follows. As we explained in Section I, the full orthogonal signal $q = F(\omega|\theta)$ is uniformly distributed and thus independent of ex ante true type θ , so even though orthogonal disclosure allows the seller to garble q depending on reported type $\tilde{\theta}$, any signal structure under orthogonal disclosure is independent of θ . Due to this independence, orthogonal disclosure cannot generally replicate direct disclosure where signal structures can depend on θ indirectly through the true valuation ω . In Example 1, the direct disclosure policy assigns the same binary partition for every reported buyer type, revealing whether their true valuation ω is above or below $1/2$. It is equivalent to a signal structure which discloses to a type- θ buyer whether the (type-independent) shock q is above or below $1 - \theta/2$. Therefore, in order for a signal structure generated by garbling q to replicate the binary-partition direct structure, it would have to garble q according to the true type θ . This is impossible by the definition of orthogonal disclosure.

We can also verify that orthogonal disclosure generally cannot replicate the distributions of posterior estimates under direct disclosure, which is what matters to mechanism design. In Example 1, under the binary-partition direct structure with threshold $1/2$, the posterior estimate v has a two-point distribution: $3/4$ with probability $\theta/2$, and $(3\theta - 2)/(4\theta)$ with probability $1 - \theta/2$. Suppose that under some orthogonal disclosure rule $\langle S, \tilde{p} \rangle$, the posterior estimate v has the same two-point distribution. Without loss of generality, we can assume a binary signal space $S = \{s_+, s_-\}$, where s_+ is the more favorable signal that leads to the posterior estimate $3/4$. Then the probability of having posterior estimate $3/4$ under $\langle S, \tilde{p} \rangle$ is $\int_0^1 \tilde{p}(s_+|q) dq$, which must be equal to $\theta/2$. However, this is impossible because $\int_0^1 \tilde{p}(s_+|q) dq$ is independent of θ .

In principle, one could always transform any type-dependent signal structure resulting from our direct disclosure into a type-independent one through orthogonal decomposition. If the resulting signal structure can be replicated by garbling the full

orthogonal signal structure, then the irrelevance theorem of Eső and Szentes (2007) implies that full disclosure is optimal. As we have argued above using binary partitions, however, replication is generally impossible for direct disclosure policies. A signal structure resulting from first garbling true valuation ω and then orthogonally decomposing the garbled signal may not be replicated by first orthogonally decomposing ω and then garbling the full orthogonalized signal $q = F(\omega|\theta)$. In other words, the order of garbling and orthogonal transformation generally matters. In this case, even though one can analogously define the hypothetical setting under direct disclosure through orthogonal transformation, the resulting signal structure may not be replicated by garbling the full orthogonal signal.

B. Further Comments

Our result that direct disclosure can generate revenue greater than the hypothetical revenue does not depend on the assumption that the buyer's ex ante type is continuously distributed. To illustrate this, suppose the ex ante type is binary, $\Theta = \{\theta_1, \theta_2\}$, with $F(\omega|\theta_1) > F(\omega|\theta_2)$ for $\omega \in (\underline{\omega}, \bar{\omega})$. To save notation, we denote the inverse of the conditional distribution function as $Q_i(q) \equiv F^{-1}(q|\theta_i)$ for each $i = 1, 2$; we have $Q_1(q) < Q_2(q)$ for all $q \in (0, 1)$. In Figure 1, we plot $Q_1(q)$ and $Q_2(q)$ for two cases, when they are "far apart" and when they are "close."

As shown in Krämer and Strausz (2015b), the irrelevance theorem of Eső and Szentes (2007) does not hold here: the hypothetical revenue cannot be attained in the original sequential-screening setting under full disclosure.¹⁸ In the latter setting, given any inefficient allocation of type θ_1 determined by a strike price $p_1 \in (c, \bar{\omega})$, it is optimal for the seller to choose θ_1 's advance payment a_1 to bind his individual rationality constraint IR_1 and θ_2 's advance payment a_2 to bind θ_2 's incentive compatibility constraint IC_{21} . Combining IR_1 with IC_{21} and using integration by parts, we have the following information rent U_2 for type θ_2 :

$$(7) \quad \begin{aligned} U_2 &= \int_{p_1}^{\bar{\omega}} (F(\omega|\theta_1) - F(\omega|\theta_2)) d\omega \\ &= \int_{F(p_1|\theta_2)}^{q_1} (Q_2(q) - p_1) dq + \int_{q_1}^1 (Q_2(q) - Q_1(q)) dq, \end{aligned}$$

where $q_1 = F(p_1|\theta_1)$. Figure 1 illustrates the decomposition of U_2 above, due to Krämer and Strausz (2015b), where the first part is marked with horizontal lines and the second part is marked with vertical lines. In the hypothetical problem, the seller observes the realized q and can implement "trade if and only if $q \geq q_1$," both for the truth-telling type θ_1 and for the deviating type θ_2 . As a result, the seller only has to pay type θ_2 the second part of U_2 in (7), which we denote as \bar{U}_2 . The irrelevance theorem of Eső and Szentes (2007) fails because $\bar{U}_2 < U_2$ for any allocation of type θ_1 determined by p_1 .

¹⁸In the working paper version, we offer an independent proof of this result, as well as a proof of the convergence of the full-disclosure revenue to the hypothetical revenue in the limit when the type distribution becomes continuous.

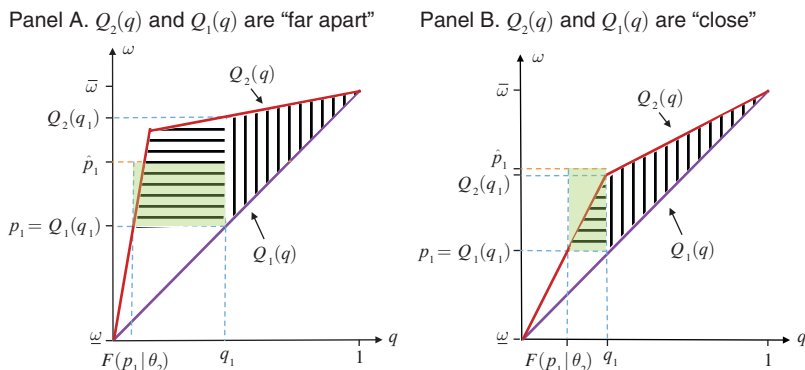


FIGURE 1. HYPOTHETICAL REVENUE VERSUS DIRECT DISCLOSURE

Now, consider direct disclosure. In particular, suppose that for reported type θ_1 the seller only discloses whether ω is above or below p_1 , raises the strike price slightly from p_1 to some \hat{p}_1 and simultaneously reduces the advance payment to keep IR_1 binding. Then the deviating type θ_2 would still trade with probability $1 - F(p_1 | \theta_2)$ but at worse terms, because an increase in strike price hurts him more than type θ_1 as his probability of trading is greater due to first-order stochastic dominance. The difference in trading probabilities is $q_1 - F(p_1 | \theta_2)$, so the rent reduction $U_2 - \hat{U}_2$ is

$$U_2 - \hat{U}_2 = (q_1 - F(p_1 | \theta_2))(\hat{p}_1 - p_1),$$

which is marked as the shaded area in both panels of Figure 1. The rent reduction is increasing in \hat{p}_1 , and is maximized at $\hat{p}_1 = E[Q_1(q) | q \geq q_1]$, which binds IR_1 with the new advance payment of $\hat{a}_1 = 0$.¹⁹

Taking the difference between the rent reductions $U_2 - \bar{U}_2$ and $U_2 - \hat{U}_2$ (the difference between the area shaded with horizontal lines and the solid shaded area in the two panels of Figure 1), we have $\hat{U}_2 < \bar{U}_2$, and thus partitioning ω can generate a revenue strictly higher than the hypothetical revenue if

$$(8) \quad \int_{F(p_1 | \theta_2)}^{q_1} (Q_2(q) - \hat{p}_1) dq < 0.$$

A sufficient condition for (8) is $\hat{p}_1 > Q_2(q_1)$, which is satisfied if the two ex ante types are not too different in the sense that $F(p_1 | \theta_1)$ is close to $F(p_1 | \theta_2)$ so that $Q_2(q_1)$ is not much higher than p_1 (as shown in panel B of Figure 1). More generally, compare the rent reductions $U_2 - \bar{U}_2$ and $U_2 - \hat{U}_2$ for fixed distribution $Q_1(q)$ and partitioning threshold q_1 (and hence p_1 and \hat{p}_1). As $Q_2(q)$ moves closer to $Q_1(q)$, the rent reduction $U_2 - \bar{U}_2$ shrinks at quadratic speed, as it is bounded from above by the

¹⁹The derivation implicitly assumes that the conditional mean valuation above p_1 is greater for type θ_2 than for type θ_1 , so that type θ_2 buys the good at the maximum new strike price if he learns that his true valuation is above p_1 . This assumption is satisfied regardless of p_1 if we strengthen first-order stochastic dominance to hazard rate dominance, which holds in Example 2. If this assumption is not satisfied, type θ_2 's rent is reduced to zero with the partitioning of ω , and direct disclosure clearly dominates orthogonal disclosure.

product of $q_1 - F(p_1|\theta_2)$ and $Q_2(q_1) - p_1$. In contrast, $U_2 - \hat{U}_2$ shrinks at linear speed, as $\hat{p}_1 - p_1$ is fixed. Thus, if the two ex ante types are sufficiently “close,” partitioning ω can generate revenue stricter than the hypothetical revenue. Since the latter remains a revenue upper bound for the seller in the original problem under full disclosure, binary-partition direct disclosure dominates full disclosure.

If we insert additional ex ante types between θ_2 and θ_1 so that adjacent types become closer, the difference between the hypothetical revenue and the full disclosure revenue in the original setting would diminish. Kräbmer and Strausz (2015b) show that this difference disappears in the continuous type limit. In contrast, as implied by Proposition 2, the revenue gain of discriminatory direct disclosure over full disclosure does not diminish as the set of ex ante types becomes an interval, so full disclosure remains dominated by binary-partition direct disclosure.

We conclude the discussion by presenting a parameterized example where the revenue generated by optimal binary-partition direct disclosure is strictly greater than the hypothetical revenue when $F(\omega|\theta_2)$ and $F(\omega|\theta_1)$ are “close.” This is the example that we use to generate Figure 1.

Example 2: Suppose $\phi_1 = \phi_2 = \frac{1}{2}$, $c = 0$, $F(\omega|\theta_1) = \omega$ for $\omega \in [0, 1]$, and

$$F(\omega|\theta_2) = \begin{cases} \omega\varepsilon/(1 - \varepsilon) & \text{if } \omega \in [0, 1 - \varepsilon] \\ 1 - (1 - \omega)(1 - \varepsilon)/\varepsilon & \text{if } \omega \in (1 - \varepsilon, 1] \end{cases}$$

with $\varepsilon \in (0, 1/2)$. As $\varepsilon \rightarrow 0$, $F(\omega|\theta_2)$ converges to a mass point at 1; as $\varepsilon \rightarrow 1/2$, $F(\omega|\theta_2)$ converges to $F(\omega|\theta_1)$. With some algebra, it can be verified that the optimal binary-partition direct structure for θ_1 has threshold $p_1 = (1 - 2\varepsilon)/(2 - 2\varepsilon)$, and the corresponding seller’s revenue is $((1 - 2\varepsilon)/(4 - 4\varepsilon))^2 + 1/2$. The hypothetical revenue is $1/2$ if $\varepsilon > 1/3$ and is $(1 - \varepsilon)/(8 - 12\varepsilon) + (1 - \varepsilon)/2$ if $\varepsilon \leq 1/3$. Therefore, there exists a cutoff value $\varepsilon^* \in (0, 1/3)$ such that, for all $\varepsilon \in (\varepsilon^*, 1/2)$, binary-partition direct disclosure can generate revenue strictly greater than the hypothetical revenue.

IV. Concluding Remarks

In this paper we have incorporated information disclosure in the simplest dynamic mechanism design framework of sequential screening, where a buyer is initially privately informed of the distribution function of his valuation for a seller’s good. The seller controls the selling mechanism and the buyer’s access to any private post-contractual information. We find that full information disclosure is generally suboptimal if the seller can directly garble her signal and grant different buyer types differential access to information.

Although binary-partition direct disclosure is very effective in facilitating discrimination by providing minimal information for allocation decisions, it is not optimal in general because binary partition sometimes can be too informative for high buyer types. In future research, we plan to characterize the optimal direct disclosure policy and natural restrictions on disclosure policies for full disclosure to be optimal.

REFERENCES

- Anderson, Simon P., and Regis Renault.** 2006. "Advertising Content." *American Economic Review* 96 (1): 93–113.
- Baron, David P., and David Besanko.** 1984. "Regulation and Information in a Continuing Relationship." *Information Economics and Policy* 1 (3): 267–302.
- Bergemann, Dirk, and Martin Pesendorfer.** 2007. "Information Structures in Optimal Auctions." *Journal of Economic Theory* 137 (1): 580–609.
- Bergemann, Dirk, and Maher Said.** 2011. "Dynamic Auctions: A Survey." In *Wiley Encyclopedia of Operations Research and Management Science*, edited by James J. Cochran. New York: Wiley.
- Bergemann, Dirk, and Achim Wambach.** 2015. "Sequential Information Disclosure in Auctions." *Journal of Economic Theory* 159: 1074–95.
- Che, Yeon-Koo.** 1996. "Customer Return Policies for Experience Goods." *Journal of Industrial Economics* 44 (1): 17–24.
- Courty, Pascal, and Hao Li.** 2000. "Sequential Screening." *Review of Economic Studies* 67 (4): 697–717.
- Esó, Peter, and Balázs Szentes.** 2007. "Optimal Information Disclosure in Auctions and the Handicap Auction." *Review of Economic Studies* 74 (3): 705–31.
- Esó, Peter, and Balázs Szentes.** 2017. "Dynamic Contracting: An Irrelevance Theorem." *Theoretical Economics* 12 (1): 109–39.
- Ganuza, Juan-José.** 2004. "Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations." *RAND Journal of Economics* 35 (3): 583–98.
- Ganuza, Juan-José, and Jose S. Penálva.** 2010. "Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions." *Econometrica* 78 (3): 1007–30.
- Garrett, Daniel F., and Alessandro Pavan.** 2012. "Managerial Turnover in a Changing World." *Journal of Political Economy* 120 (5): 879–925.
- Gershkov, Alex, and Benny Moldovanu.** 2012. "Dynamic Allocation and Pricing: A Mechanism Design Approach." *International Journal of Industrial Organization* 30 (3): 283–86.
- Goltsman, Maria, Johannes Horner, Gregory Pavlov, and Francesco Squintani.** 2009. "Mediation, Arbitration and Negotiation." *Journal of Economic Theory* 144 (4): 1397–1420.
- Grossman, Sanford J.** 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics* 24 (3): 461–83.
- Hoffmann, Florian, and Roman Inderst.** 2011. "Pre-Sale Information." *Journal of Economic Theory* 146 (6): 2333–55.
- Ivanov, Maxim.** 2016. "Dynamic Learning and Strategic Communication." *International Journal of Game Theory* 45 (3): 627–53.
- Jehiel, Philippe.** 2015. "On Transparency in Organizations." *Review of Economic Studies* 82 (2): 736–61.
- Johnson, Justin P., and David P. Myatt.** 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review* 96 (3): 756–84.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590–2615.
- Kolotilin, Anton, Ming Li, Tymofiy Mylovanov, and Andriy Zapechelnyuk.** Forthcoming. "Persuasion of a Privately Informed Receiver." *Econometrica*.
- Krähmer, Daniel, and Roland Strausz.** 2015a. "Dynamic Mechanism Design." In *An Introduction to the Theory of Mechanism Design*, edited by Tilman Börgers, 204–34. Oxford: Oxford University Press.
- Krähmer, Daniel, and Roland Strausz.** 2015b. "Ex Post Information Rents in Sequential Screening." *Games and Economic Behavior* 90: 257–73.
- Lewis, Tracy R., and David E. M. Sappington.** 1994. "Supplying Information to Facilitate Price Discrimination." *International Economic Review* 35 (2): 309–27.
- Lu, Jingfeng, and Lixin Ye.** 2016. "Optimal Two-Stage Auctions with Costly Information Acquisition." <http://www.econ.ohio-state.edu/lixinye/Research/alpha2.pdf> (accessed September 12, 2017).
- Milgrom, Paul.** 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12 (2): 380–91.
- Milgrom, Paul.** 2004. *Putting Auction Theory to Work*. Cambridge, UK: Cambridge University Press.
- Milgrom, Paul, and John Roberts.** 1986. "Relying on Information of Interested Parties." *RAND Journal of Economics* 17 (1): 18–32.
- Milgrom, Paul R., and Robert J. Weber.** 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–1122.
- Myerson, Roger.** 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.
- Ottaviani, Marco, and Andrea Prat.** 2001. "The Value of Public Information in Monopoly." *Econometrica* 69 (6): 1673–83.

- Pavan, Alessandro.** Forthcoming. "Dynamic Mechanism Design: Robustness and Endogenous Types." In *Advances in Economics and Econometrics: Eleventh World Congress*, edited by Bo Honoré et al. Cambridge, UK: Cambridge University Press.
- Pavan, Alessandro, Ilya Segal, and Juuso Toikka.** 2014. "Dynamic Mechanism Design: A Myersonian Approach." *Econometrica* 82 (2): 601–53.
- Rayo, Luis, and Ilya Segal.** 2010. "Optimal Information Disclosure." *Journal of Political Economy* 118 (5): 949–87.
- Zhang, Jun.** 2013. "Revenue Maximizing with Return Policy When Buyers Have Uncertain Valuations." *International Journal of Industrial Organization* 31 (5): 452–61.