



BIROn - Birkbeck Institutional Research Online

Pancs, R. and Nikandrova, Arina (2015) An optimal auction with moral hazard. Working Paper. Birkbeck College, University of London, London, UK.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/15275/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

ISSN 1745-8587



Department of Economics, Mathematics and Statistics

BWPEF 1504

An Optimal Auction with Moral Hazard

Arina Nikandrova

Birkbeck, University of London

Roman Pancs

University of Rochester

November 2015

Conjugate Information Disclosure in an Auction with Learning

Arina Nikandrova and Romans Pancs*

November 2015

Abstract

We consider a single-item, independent private value auction environment with two bidders: a leader, who knows his valuation, and a follower, who privately chooses how much to learn about his valuation. We provide sufficient conditions under which an ex-post efficient revenue-maximizing auction, which solicits bids sequentially, partially discloses the leader's bid to the follower, to influence his learning. The disclosure rule that emerges is novel in that it may reveal to the follower only a pair of bids to which the leader's actual bid belongs. We call this disclosure rule conjugate. The identified disclosure rule, relative to the first-best, induces the follower to learn less when the leader's valuation is low and more when the leader's valuation is high.

Keywords: Information Disclosure, Conjugate Disclosure, Bayesian Persuasion

JEL codes: D82, D83

1 Introduction

In early 2013, Universal Music Group was divesting Parlophone Label Group. This divestment was a part of Universal's commitment to the European Commission in exchange for clearing the Universal's merger with EMI. Among the potential buyers of Parlophone were other major music

*Nikandrova's at Birkbeck, Pancs is at ITAM.

recording companies, as well as various private equity firms. The music recording companies were well informed of their valuations of Parlophone because they eyed Parlophone a year ago, when it was on sale as a part of EMI. By contrast, most private equity firms were poorly informed and had to perform due diligence to evaluate the purchase opportunity. The sale culminated in an auction, which was preceded by protracted secretive negotiations. We will never know what (if anything) Universal revealed about the intended bids of the well informed music recording companies to the poorly informed private equity firms during these secretive negotiations. This paper is a step towards understanding what Universal should have (and so, might have) revealed. The paper's focus is thus normative.

The paper captures the essential features of the Universal's information-disclosure problem in a model in which a seller (Universal) sells an item (Parlophone) by sequentially bargaining with two bidders, one of whom knows his valuation (a major music recording company), while the other (a private equity firm) must exert a costly learning effort to form a better estimate of its valuation. Beyond the example of Universal, which is a divestment of an asset that has recently been on sale, the model applies to the sale of an asset in the presence of an industry insider,¹ to the sale of a franchise when one of the buyers is an incumbent franchise operator,² and to the procurement of a contract when one of the suppliers is the current service provider.³ In practice, in each of these cases, the sale is agreed upon either privately or through an auction. Irrespective of the ultimate mode of sale, private pre-sale negotiations, which we model, are common.

In the model, the bidders' valuations are independent and private. The seller maximizes revenue by designing, announcing, and committing to a selling mechanism. He is restricted to choosing a mechanism with an **ex-post efficient** allocation rule, which assigns the item to the bidder with the highest expected valuation (given the information available to the bidders). The mech-

¹For example, in 2012, RWE and E.ON, German utilities, were selling Horizon Nuclear Power, their British joint venture, through private negotiations. The potential buyers included nuclear operators already present in the U.K., as well as new investors from China and Japan.

²For example, in the U.K., rail passenger services are franchised for a limited period to train operating companies. An auction determines the award of the franchise to run passenger services in a certain region. The incumbent franchise operator, who knows his value for the franchise, is always allowed to bid alongside other, less informed, train operating companies.

³Government agencies routinely procure consulting services from contractors. A procurement tender can attract consultancies that have never worked with that particular agency, as well as consultancies that recently advised the agency. The consultancies that recently advised on similar projects know well the bureaucratic procedures and the private costs associated with entering the contract. By contrast, the consultancies that have not advised this government agency, in addition to learning the particulars of the tender offer, also need to study the output of the government's past external consultants, to better gauge the government's demands and expectations.

anism must also satisfy interim participation constraints that, in particular, rule out the sale of information. Admissible mechanisms are otherwise unrestricted and allow the seller to strategically disclose to either bidder the information learned from the other bidder.⁴

Because of the option value of influencing the uninformed bidder's learning decision by disclosing to him something about the informed bidder's bid, it is optimal for the seller to approach the bidders sequentially. In particular, it is optimal to approach the informed bidder first, in order to elicit his valuation. We call this informed bidder the **leader**. We call the uninformed bidder, who is approached second and is recommended by the mechanism how much effort to exert, the **follower**.

The analysis relies on the first-order approach. Using this approach, we reduce the seller's revenue-maximization problem to the problem of optimally disclosing the leader's bid to the follower. The seller's strategic disclosure distorts the follower's effort schedule relative to the effort schedule in the **first-best** mechanism, which maximizes the ex-ante expected surplus (the sum of the seller's and the bidders' payoffs).

The distortion is necessary because the policy of **full disclosure**, which always discloses the leader's valuation and is called for by the first-best mechanism, is suboptimal for the seller. So is the policy of **non-disclosure**, which always conceals the leader's valuation. Instead, under some conditions, the seller optimally partitions the leader's types (i.e., valuations, reported as bids) into pairs of **conjugate types**, and discloses to the follower the pair to which the leader's type belongs, while concealing the actual type, as is illustrated in Figure 1. An optimal partition is described by a weakly decreasing **matching function**, which pools the leader's valuations into pairs so that lower types are matched with higher types. The seller may also choose to reveal some of the leader's types without pooling them.

Once identified, the optimal disclosure rule can be equivalently restated in terms of the shape of the follower's induced effort schedule. This schedule is hump shaped and "shifted to the right" relative to the first-best schedule, also hump shaped. That is, when the leader's valuation is low, the follower acquires inefficiently little information, whereas when the leader's valuation is high,

⁴The restriction to the ex-post efficient allocation rule is motivated by applications, such as government procurement, in which the seller cannot commit to an allocation rule that is ex-post inefficient. The applications in which the seller can freely choose both the allocation rule and information disclosure are at least as important. For reasons that will be discussed, however, the techniques developed in this paper are not readily applicable to the joint determination of an optimal allocation rule and information disclosure.

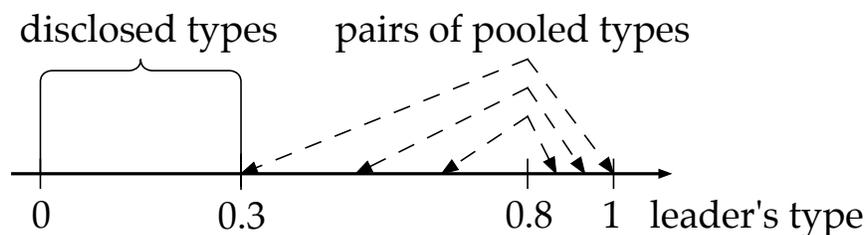


Figure 1: An example of a conjugate disclosure rule. The seller reveals the leader’s type if it is less than 0.3. The seller pools any leader’s type that exceeds 0.3 with a corresponding conjugate type to form a pair that straddles 0.8.

the follower acquires inefficiently much information. To gain intuition for the effort distortion, we consider two benchmarks: a first-best mechanism and a revenue-maximizing mechanism in which the seller can directly control the follower’s effort.

In the first-best benchmark, the follower’s information acquisition helps allocate the item in a surplus-maximizing manner. His information is most valuable when the leader’s valuation is intermediate. When the leader’s valuation is low, the follower is the likely surplus-maximizing recipient of the item. When the leader’s valuation is high, the leader is the likely surplus-maximizing recipient. In each of these two extreme cases, little would be gained from additional information about the follower’s valuation. As a result, the follower’s first-best effort schedule is hump-shaped in the leader’s valuation.

In the revenue-maximizing benchmark, in which the seller controls the follower’s effort, the seller sets the effort to zero when the leader’s valuation is below a certain threshold, and sets the effort to be maximal otherwise. The intuition for this “bang-bang” property of the effort schedule is the following. Given a leader’s valuation, denoted by θ_1 , ex-post efficiency and incentive considerations compel the seller to charge the follower amount θ_1 and charge the leader less than θ_1 . Indeed, by making the follower a take-it-or-leave-it offer at price θ_1 , the seller ensures that the follower buys if and only if the follower’s expected valuation, denoted by θ_2 , exceeds θ_1 , as ex-post efficiency requires. In the complementary case, when $\theta_2 \leq \theta_1$, the seller cannot possibly hope to sell to the leader for θ_1 . Indeed, if the leader knew that he would be charged θ_1 , he would profitably mimic a bidder with a lower valuation, thereby undermining the incentive compatibility of the seller’s mechanism. As a result, for any θ_1 , the seller would rather sell to the follower. Therefore, the seller makes the follower acquire information in a manner that increases

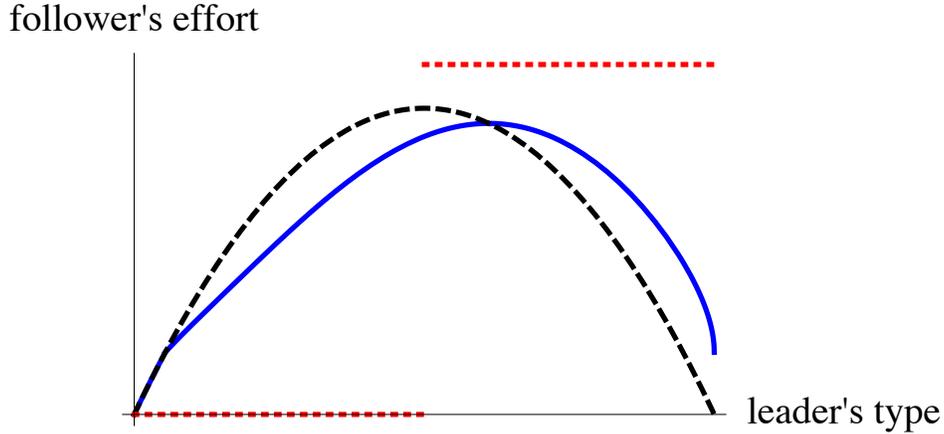


Figure 2: The follower’s optimal effort schedule (the solid hump-shaped curve) is a compromise between his first-best effort schedule (the dashed hump-shaped curve) and the bang-bang effort schedule (the two dashed horizontal line segments), which corresponds to the benchmark in which the seller controls the follower’s effort.

the probability of an ex-post efficient sale to the follower. Because information acquisition induces a mean-preserving spread in the probability distribution of θ_2 ,⁵ the event $\{\theta_2 \geq \theta_1\}$ (i.e., selling to the follower is ex-post efficient) is more likely either when θ_1 is low and the follower’s information acquisition is minimal, or when θ_1 is high and the follower’s information acquisition is maximal.

The economic forces present in both benchmarks meet in this paper’s central problem, in which the seller maximizes his revenue and cannot control the follower’s effort. In this problem, the seller can influence the follower’s effort only indirectly, by strategically revealing information about the leader’s bid. Because the follower cannot be deceived systematically, the seller’s disclosure distorts, but not beyond recognition, the first-best effort schedule. The rightward shift in the follower’s optimal effort schedule relative to the first-best schedule is a compromise between the first-best and the bang-bang schedules of the two benchmarks. This compromise is illustrated in Figure 2.

Figure 2 suggests a correspondence between the structure of an optimal disclosure rule and the shape of an optimal effort schedule. The Revelation Principle for games with private actions

⁵In the model, a higher information-acquisition effort by the follower implicitly delivers a more informative signal about a his underlying valuation. This more informative signal induces a more dispersed posterior probability distribution of the follower’s expected valuation, θ_2 . Because the follower is risk-neutral, the expected valuation is the only aspect of his posterior probability distribution that he cares about. Identifying the informativeness of a signal with the induced dispersion of the probability distribution of θ_2 is a standard modelling device (Johnson and Myatt, 2006, Ganuza and Penalva, 2010, Shi, 2012, and Roesler, 2014) and, in the presence of risk neutrality, entails no loss in generality.

(Myerson, 1982, 1986) implies that all the follower needs to be told about the leader's type is summarized in the effort that the seller would like him to exert. Then, given an optimal effort schedule, the corresponding optimal disclosure rule can be read off this schedule. Thus, given its hump shape, the optimal effort schedule prescribes the same effort for (at most) two leader's types. These types can be read off from Figure 2 by intersecting the optimal effort schedule with the horizontal line corresponding to the recommended effort (not shown). In other words, the optimal effort schedule can be implemented with the help of a conjugate-disclosure rule. Instead of deriving the optimal effort schedule first, however, it turns out to be more convenient to derive an optimal disclosure rule and only then to recover from it the implied effort schedule.

The revenue-maximizing outcome can be implemented in a sequential second-price auction with a tax (or subsidy) for the leader.⁶ The role of the tax is to encourage the leader to bid truthfully. He would do so without a tax in the standard auction model in which the follower exerts no effort. Because in our model the follower does exert an effort, the leader realizes that by biasing his bid toward intermediate values, he can encourage the follower to exert a greater effort. A greater effort translates into a more dispersed probability distribution of the follower's expected valuations, thereby benefiting the leader, whose payoff is convex in the follower's expected valuation.⁷ Hence, a corrective tax for the leader is necessary.

Our analysis is subject to two caveats. The first caveat is that the first-order approach, on which the paper relies, focuses on the seller's relaxed problem that only imposes local truth-telling constraints (captured by the Envelope Theorem) and neglects global truth-telling constraints (captured by monotonicity conditions on the probabilities of winning). Although common in mechanism design, this approach is valid if and only if, at a solution to the relaxed problem, each bidder's probability of winning is nondecreasing in his type. A limitation of our analysis is that we do not provide sufficient conditions on the model's primitives to guarantee that the monotonicity condition holds for the leader. Instead, our approach is to verify the monotonicity condition directly, once the seller's relaxed problem has been solved.⁸ The sufficient conditions prove elusive because an optimal disclosure rule, an essential determinant of the follower's effort schedule and

⁶Also the first-best outcome can be implemented in a sequential second-price auction, but with a different tax.

⁷The convexity follows because the leader benefits from the follower's lower bid by winning more often and paying less, and is insulated from the follower's higher bids by paying nothing when losing.

⁸In numerical examples, once an optimal mechanism has been identified, the monotonicity is easy to verify, which we do in the paper's leading example.

thus the leader's probability of winning, is not a simple function of the primitives. Indeed, the complexity of finding an optimal disclosure rule is comparable to the complexity of finding an optimal tax schedule in the dynamic public-finance literature (Farhi and Werning, 2013, Stantcheva, 2015, Golosov et al., 2015, and Kapicka, 2013), which also adopts the verification approach.

The second caveat is that our results rely on the assumption that—and the domain of our applications is restricted to the situations in which—the seller must choose among mechanisms with ex-post efficient allocation rules. In particular, the paper adapts and extends the results from the sender-receiver game of Rayo and Segal (2010) (henceforth referred to as RS) to an auction setting. This extension relies on rewriting the seller's objective function in a product form, which requires the uncoupling of the optimal disclosure and the optimal allocation problems. Only when the allocation rule is fixed, this uncoupling is attainable. Fixed allocation rules other than ex-post efficient can also be entertained; yet the ex-post efficient rule is the most natural in applications. The analysis of the optimal interaction of disclosure and allocation rules is beyond the scope of this paper.

The rest of the paper proceeds as follows. Section 2 discusses related literature. Section 3 describes the economic environment. Section 4 derives the first-best allocation and an auction that implements it. Section 5 establishes the suboptimality of the policies of full disclosure, non-disclosure, and interval pooling. Under additional conditions, the section partially characterizes an optimal auction by establishing that an optimal disclosure rule is conjugate, and then formulates an optimal-control problem whose solution (computed in an example) delivers the matching function underlying the optimal conjugate disclosure rule. The section also explains the departure of the follower's effort schedule from the first-best schedule. Section 6 provides additional intuition for the optimal disclosure rule by comparing it to two benchmarks. Section 7 discusses the limitations of the first-order approach and the challenges associated with departing from the ex-post efficient allocation rule. Section 8 concludes. Some proofs are relegated to the Appendix. Technical results comprise the Supplementary Appendix.

2 Related Literature

Our paper contributes to the literature on auctions in which a monopolistic seller directly or indirectly influences bidders' information structure and, more broadly, to the literature on Bayesian persuasion. The existing literature in which a seller, through his choice of a selling mechanism, affects bidders' information-acquisition effort (e.g., [Bergemann and Välimäki, 2002](#), [Persico, 2003](#), [Compte and Jehiel, 2007](#), [Crémer et al., 2009](#), and [Shi, 2012](#)) mostly focuses on auctions that are simultaneous. In these auctions, all bidders learn simultaneously, and so the issue of optimal bid-disclosure does not arise. For instance, [Persico \(2003\)](#) compares the seller's revenues in the first-price and second-price auctions. [Shi \(2012\)](#) studies a revenue-maximizing simultaneous auction. [Bergemann and Välimäki \(2002\)](#) study efficient simultaneous auctions. The optimally designed auctions of the last two papers become suboptimal (in their respective senses) if the class of admissible mechanisms is enriched to include sequential auctions, as corroborated by the present paper's results. The advantage of sequential auctions is established also by [Compte and Jehiel \(2007\)](#), who find that the revenue in the sealed-bid second-price auction with simultaneous learning is lower than the revenue in the ascending-price auction in which bidders have multiple opportunities to learn.

[Crémer et al. \(2009\)](#) admit sequential auctions and design a revenue-maximizing one. Because they assume that the seller can charge bidders for information, optimal information disclosure turns out to be trivial (e.g., full disclosure) and is not their focus. Information disclosure is the focus of the model of [Eso and Szentes \(2007\)](#), in which the seller directly designs the signals observed by the bidders, instead of motivating bidders to acquire signals themselves. Just like [Crémer et al. \(2009\)](#), [Eso and Szentes \(2007\)](#) allow the seller to charge bidders for information and find that full disclosure is optimal. The seller of [Eso and Szentes \(2007\)](#) reveals maximal information to maximize the total surplus, which he taxes away by cleverly charging for the signals that he reveals.⁹ In our paper, the critical assumption that rules out selling information and delivers the suboptimality of full disclosure is a particularly demanding interim participation constraint, according to which, upon observing his type, the follower must expect his total payoff to be nonnegative.¹⁰

⁹The seller must charge cleverly because the bidders of [Eso and Szentes \(2007\)](#), by contrast to the bidders of [Crémer et al. \(2009\)](#), already have some private information before accepting the seller's mechanism. So the charges are not simple participation fees.

¹⁰In other words, each bidder contracts with the seller only upon submitting his bid, not before.

Without this demanding participation constraint, the logic of [Eso and Szentes \(2007\)](#) would have implied the optimality of full disclosure also in our model (the various differences between the models notwithstanding).¹¹ Such stringent participation constraints are also imposed by [Ganuzza \(2004\)](#) in a second-price auction and by [Bergemann and Pesendorfer \(2007\)](#) in an optimally designed auction as they model the seller’s design of the bidders’ information structure. As a result, both [Ganuzza \(2004\)](#) and [Bergemann and Pesendorfer \(2007\)](#) arrive at the suboptimality of full disclosure; the seller seeks to avoid conferring on the bidders excessive information rents that cannot be taxed away.

Another strand of literature to which our paper contributes is the literature on Bayesian persuasion, or sender-receiver games with commitment. The two main papers in this literature are RS and [Kamenica and Gentzkow \(2011\)](#). RS assume additional structure that makes their paper especially pertinent to our analysis. We establish the relevance of RS’s results in a novel environment, an auction with costly information acquisition and, crucially, with a continuum of types.

In RS, Nature draws a “prospect” that enters the sender’s information-disclosure rule. By contrast, in our model a prospect is determined by the type the leader reports to the seller, who acts as a sender with commitment. RS’s receiver takes a binary action. Our receiver, the follower, chooses a divisible amount of information to acquire. Routine mechanism-design techniques reduce the seller’s information-disclosure problem to the seemingly unrelated sender-receiver game of RS and make their equilibrium characterization applicable, subject to one significant qualification. RS assume a discrete prospect set, whereas in our model, the prospect set is a continuum. A limit argument establishes a formal connection between RS’s model and ours, thereby paving the way for the proof of the optimality of the conjugate disclosure. This sharp characterization of optimal disclosure has no direct counterpart in the discrete model of RS.

In the context of persuasion models with the sender commitment, our conjugate disclosure is unusual in that it pools non-adjacent types. Without commitment, a sender-receiver game in which non-adjacent types are pooled is due to [Golosov et al. \(2011\)](#). They consider a dynamic

¹¹With more demanding participating constraints, the logic of [Eso and Szentes \(2007\)](#) need not apply in an otherwise unchanged model of [Eso and Szentes \(2007\)](#). Indeed, [Bergemann and Wambach \(2015\)](#) strengthen the participation constraints in the model of [Eso and Szentes \(2007\)](#) and find that, to implement the same allocation as in the auction of [Eso and Szentes \(2007\)](#), full disclosure must be sacrificed.

version of the model of Crawford and Sobel (1982).¹² Each period, an expert that is privately informed about a state sends a message to a decision maker. Golosov et al. (2011) construct an equilibrium that eventually reveals that state. A critical ingredient in their construction is a disclosure rule that initially pools faraway types into pairs. These pairs are not, however, ordered in a way that conforms with the conjugate disclosure rule that we derive. Moreover, their pairwise disclosure serves a different purpose (the eventual disclosure of the state, not profit-maximization) in a disclosure game that is substantively different from ours (their game has multiple stages of communication, no commitment, and a different economic setting).¹³

Furthermore, any auction design or agency design in which information disclosure at an early stage affects players' unenforceable (by the seller or the principal) action at a later stage features Bayesian persuasion. An example of such design is the application of the Kamenica and Gentzkow (2011) techniques to a two-player contest, by Zhang and Zhou (2015). In their contest, bids model unenforceable efforts (as opposed to enforceable payments, as in auctions). Another example of design with Bayesian persuasion is an auction followed by resale. While early resale models (Bikhchandani and Huang, 1989; Gale et al., 2000; Haile, 2003; Gupta and Lebrun, 1999) fix an auction format and study the informational linkage between the primary market and the resale market, more recent work (Calzolari and Pavan, 2006a; Zheng, 2002) adopts the mechanism-design approach. The work of Calzolari and Pavan (2006a) is especially closely related to ours.

Calzolari and Pavan (2006a) study a mechanism-design problem of a monopolistic seller who sells to a potential buyer, a leader, in the primary market and anticipates the possibility that the leader resells to another buyer, a follower, in the resale market. The seller's mechanism comprises a rule for allocating an item to the leader and a rule for disclosing information to follower. In the resale market, either the leader or the follower is randomly chosen to make a take-it-or-leave-it price offer to the other buyer. When the follower is chosen, he is the counterpart of the follower in our model; his resale offer is a private action informed by the seller's strategic disclosure of the

¹²Also Krishna and Morgan (2001, 2004), Chen (2009), and Ivanov (2011) analyze sender-receiver games without commitment in which non-adjacent types can be pooled in (so-called non-monotone) equilibria.

¹³Introducing commitment into the sender-receiver model of Crawford and Sobel (1982) does not generate anything resembling the conjugate disclosure that we find. Indeed, in the quadratic formulation of that model, if the sender can commit to a disclosure rule, full disclosure is optimal. Intuitively, the sender cannot persuade the receiver to systematically bias his action (to indulge the sender's bias), but by committing to full disclosure, the sender can make the receiver's action match the state. Such matching maximizes the sender's ex-ante expected utility because the sender is risk-averse with respect to the discrepancy between the receiver's action and the state.

leader’s reported type.

In the model of [Calzolari and Pavan \(2006a\)](#), just as in ours, optimality proscribes full disclosure because, just as we do, they assume a participation constraint that precludes the seller from expropriating the traders’ rents in the resale market.¹⁴ The assumption of [Calzolari and Pavan \(2006a\)](#) that the leader’s type is binary delivers tractability and allows them to characterize both an optimal allocation rule and an optimal disclosure rule. By contrast, our assumption of the continuum of the leader’s types enables us to study richer disclosure rules, but at the cost of fixing the allocation rule. In particular, we identify a natural economic setting in which a novel pattern of information disclosure—conjugate disclosure—optimally arises. With the binary types of [Calzolari and Pavan \(2006a\)](#), such a disclosure would be hard to detect because the message space can be restricted to just three elements. More generally, with finite types, the message space is finite without loss of generality. With a continuum of types, the optimal message space is either finite (as in the models where monotone disclosure is optimal, such as [Crawford and Sobel, 1982](#) or [Ostrovsky and Schwartz, 2010](#)) or infinite (as in our model), and thus the garbling invoked by the disclosure rule can display richer patterns.

3 Model

Environment

The seller must allocate an item, which he values at zero, to one of two bidders. Bidder 1, the **leader**, privately observes his valuation, or type, denoted by θ_1 and drawn according to a c.d.f. G with the corresponding p.d.f. g and support $\Theta_1 \equiv [0, 1]$. The c.d.f. G is smooth on $(0, 1)$, with bounded derivatives. Bidder 2, the **follower**, is unsure of his valuation, and privately exerts an effort $a \in A \equiv [0, 1]$ to acquire information, which determines his expected valuation, or type, denoted by $\theta_2 \in \Theta_2 \equiv [0, 1]$ (as will be explained shortly). The cost of effort a is the convex function $C(a) \equiv ca^2/2, c > 0$.

For any $i \in \{1, 2\}$, let $x_i \in [0, 1]$ be the probability that bidder i gets the item, and let $t_i \in \mathbb{R}$ be

¹⁴In a similar spirit, in a sequential common agency model, [Calzolari and Pavan \(2006b\)](#) identify conditions under which the upstream principal may find it strictly optimal to disclose a noisy signal about the agent’s type to the downstream principal.

his payment. The leader's payoff is

$$\theta_1 x_1 - t_1,$$

and the follower's payoff is

$$\theta_2 x_2 - t_2 - C(a).$$

Both bidders are expected-utility maximizers.

Effort Technology

The follower's type, θ_2 , is interpreted as the expectation of his underlying valuation (not modelled) conditional on the privately observed signal generated by his information-acquisition effort. The underlying valuation is not modeled, because it affects the follower's payoff only through the expectation θ_2 . The signal can be identified without the loss of generality with θ_2 , as we shall discuss.

For any $a \in A$, θ_2 is drawn according to a c.d.f that is linear in a :^{15,16}

$$F(\theta_2 | a) \equiv aF_H(\theta_2) + (1 - a)F_L(\theta_2), \quad \theta_2 \in \Theta_2 \equiv [0, 1]. \quad (1)$$

Whenever the p.d.f.s corresponding to the c.d.f.s F , F_H , and F_L exist, they are denoted by f , f_H , and f_L . The c.d.f. F may have mass points at $\{0, 1\}$, but on $(0, 1)$, F is assumed to be smooth, with bounded derivatives. Conditional on a , θ_1 and θ_2 are independent.

For a to be interpreted as an information-acquisition effort, F is assumed to satisfy

Condition 1 (Information Acquisition). (i) (equality of means) $\int_0^1 F_H(s) ds = \int_0^1 F_L(s) ds$, and (ii) (rotation) for some $\theta^* \in (0, 1)$, for all $s \in (0, \theta^*) \cup (\theta^*, 1)$, it holds that $(\theta^* - s)(F_H(s) - F_L(s)) > 0$.

According to Condition 1, the follower who exerts effort a , with probability a , receives a more informative signal about his underlying valuation, and with probability $1 - a$, receives a less in-

¹⁵In principal-agent problems, the linearity condition (1) is known as the Linear Distribution Function Condition, whose special case is the Spanning Condition of Grossman and Hart (1983, p. 25).

¹⁶The linearity condition (1) is essential in Section 5 for reducing the seller's problem to the information-disclosure problem of RS.

formative signal.¹⁷ Part (i) requires the follower’s effort not to affect his expected type.¹⁸ In particular, part (i) rules out the situations in which the follower’s effort is a value enhancing investment, which can be thought of as inducing a first-order stochastic-dominance shift in the distribution of θ_2 .¹⁹ Parts (i) and (ii) taken together imply that a higher effort induces a mean-preserving spread of the distribution of the follower’s types.²⁰

This mean-preserving spread captures a standard, but perhaps counter-intuitive, implication of Blackwell’s informativeness criterion (Blackwell, 1951, 1953). According to this criterion, a more informative signal about the follower’s (unmodelled) underlying valuation induces a greater dispersion of the probability distribution over the conditional expectation, θ_2 .^{21,22} When $a = 1$, the realized θ_2 can be equivalently interpreted either as an expectation of the underlying valuation or as the underlying valuation itself. Because the implied signal may remain informative even when $a = 0$ (if F_L nondegenerate), Condition 1 generalizes the truth-or-noise information-acquisition technology introduced by Lewis and Sappington (1994) and used by Bergemann and Välimäki (2006, Section 2.2), Johnson and Myatt (2006, Section III.B), and Shi (2012, Example 2), among others.

An example of the information-acquisition technology specified in Condition 1 is

Example 1 (Rotation Order). $F(\theta_2 | a) = a \left(\frac{1}{2} \mathbf{1}_{\{\theta_2 < 1\}} + \mathbf{1}_{\{\theta_2 = 1\}} \right) + (1 - a) \theta_2$, where $\mathbf{1}_{\{\cdot\}}$ is the indicator function.

Example 1 can be interpreted to say that, with probability a , the follower observes a perfectly

¹⁷A signal structure that delivers the probability distribution of types in Condition 1 starting from some underlying valuations is given in Supplementary Appendix B.1.

¹⁸For any c.d.f. H on $[0, 1]$, its expectation is $\int x dH(x) = \int (1 - H(x)) dx$.

¹⁹It can be shown that, under some conditions, the techniques developed in this paper apply also to environments with value enhancing investments.

²⁰To see the implication, let $a' > a$. Part (ii) implies $(\theta^* - s)(F(s | a') - F(s | a)) \geq 0$, which combined with part (i) gives the second-order stochastic dominance inequality $\int_0^x F(s | a') ds \geq \int_0^x F(s | a) ds$ for all $x \in [0, 1]$.

²¹For example, suppose that the underlying valuation, denoted by v , is distributed uniformly on $[0, 1]$. Conditional on the observation of a perfectly *uninformative* signal, the probability distribution of v is still uniform on $[0, 1]$, and so $\theta_2 = 1/2$. By contrast, conditional on the observation of a perfectly *informative* signal, the probability distribution of v is degenerate at the true underlying valuation, and so $\theta_2 = v$. Consequently, anticipating a perfectly uninformative signal, the follower assigns probability one to $\theta_2 = 1/2$, whereas anticipating a perfectly informative signal, the follower views θ_2 as distributed uniformly on $[0, 1]$. In the former case, the dispersion of the probability distribution of conditional expectations is zero, while in the latter case, it equals the dispersion of the probability distribution of the underlying valuation. In either case, because he is risk-neutral, the follower treats the realized θ_2 as if it were his underlying valuation.

²²Blackwell’s informativeness criterion implies Lehmann’s accuracy condition (Lehmann, 1988; Persico, 2003), which implies the mean-preserving-spread order on the conditional expectations (Mizuno, 2006, Proposition 1). Directly modelling a signal’s informativeness by the induced dispersion of the conditional expectation is common in economics; see, for instance, Johnson and Myatt (2006), Ganuza and Penalva (2010), Shi (2012), and Roesler (2014).

informative signal that reveals his underlying valuation, which is distributed uniformly on $\{0, 1\}$, and with probability $1 - a$, the follower observes a partially informative signal about the underlying valuation.

The linear specification (1) rules out information-acquisition technologies that let the follower choose among three or more signals, as Example 2 clarifies.

Example 2 (Nonexample). The follower chooses a tuple (a_1, a_2) in a two-dimensional probability simplex Δ^2 , and then draws θ_2 from the probability distribution with the c.d.f. $F(\theta_2 | a_1, a_2) = a_2/2 + a_1\theta_2 + (1 - a_1 - a_2)\mathbf{1}_{\{\theta_2 \geq 1/2\}}$.

In Example 2, in addition to allocating probabilities to a perfectly informative and a somewhat informative signals about the underlying valuation in $\{0, 1\}$, as in Example 1, the follower can also allocate some probability to a completely uninformative signal (with probability $1 - a_1 - a_2$). Ruling out Example 2 is economically restrictive. If the cost of information acquisition were increasing in a_1 and a_2 , one could imagine the follower prefer to set both a_1 and a_2 close to zero if he faced the price close to 0 or 1, and optimally trade off the positive a_1 and a_2 otherwise.

Condition 1 remains restrictive even conditional on the linear specification (1), as Example 3 illustrates.

Example 3 (Another Nonexample). $F(\theta_2 | a) = a\left(\frac{1}{4}\mathbf{1}_{\{\theta_2 < 1/2\}} + \frac{1}{2}\mathbf{1}_{\{1/2 \leq \theta_2 < 1\}} + \mathbf{1}_{\{\theta_2 = 1\}}\right) + (1 - a)\theta_2$.

Example 3 can be interpreted to say that, with probability a , the follower observes a signal that, with probability $1/2$, reveals his underlying valuation, which is distributed uniformly on $\{0, 1\}$, and, with probability $1/2$, reveals “nothing”; with probability $1 - a$, the follower observes the partially informative signal of Example 1. Even though $F(\cdot | 1)$ is a mean-preserving spread of $F(\cdot | 0)$, and hence is more informative in some sense (viz. Blackwell’s order on the underlying signals), $F(\cdot | 1)$ and $F(\cdot | 0)$ are not rotation ordered.

To simplify the exposition by avoiding the corner solution $a = 1$, we henceforth assume a sufficiently large cost of effort:²³

$$c > \int_{\theta^*}^1 (F_L(s) - F_H(s)) ds. \quad (2)$$

²³This condition is derived by requiring that the first-best effort $\alpha(\theta_1)$, given in (7), be less than 1 for every θ_1 .

The Seller's Problem

The seller chooses, publicly announces, and commits to a mechanism. A **mechanism** comprises an extensive game-form, a strategy to which the seller commits, and a communication device that enables players to exchange messages. The communication device is chosen by the seller. The game-form is given. Its timing is such that (i) each bidder may leave the mechanism without payment, (ii) the follower exerts his effort and observe his realized type, (iii) each bidder may once again leave the mechanism without payment, and (iv) the seller enforces a trade that is **ex-post efficient**, meaning that, for any type profile (θ_1, θ_2) , the probability that bidder $i \in \{1, 2\}$ gets the item is $x_i(\theta_i, \theta_{-i}) = \mathbf{1}_{\{\theta_i > \theta_{-i}\}}$, where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. The communication device admits arbitrary communication protocols between the game-form's stages as long as enough information is elicited to implement the ex-post efficient allocation. The described class of mechanisms is rather rich; for instance, the seller can approach the follower first, tell him something, then talk to the leader, come back to talk to the follower, wait for the follower to exert his effort, ask him about his type, revisit the leader, then revisit the follower, and only then execute a trade.

In our environment, the logic of the Revelation Principle in environments with private information and private actions (see, e.g., Myerson, 1982 and Myerson, 1986) applies. Without loss of generality, the seller can restrict attention to direct mechanisms of the form:

1. Having observed θ_1 , the leader confidentially reports $\hat{\theta}_1 \in \Theta_1$ to the seller.
2. The seller confidentially sends a message $m \in M$ to the follower according to some **disclosure rule** $\mu : \Theta_1 \rightarrow \Delta(M)$, which associates with each report of the leader a probability distribution over messages.²⁴
3. The follower exerts an effort $a^*(m) \in A$, then observes θ_2 , and confidentially reports $\hat{\theta}_2 \in \Theta_2$.
4. Bidder i with $\hat{\theta}_i > \hat{\theta}_{-i}$ gets the item. Payments $(t_1, t_2) : \Theta_1 \times \Theta_2 \rightarrow \mathbb{R}^2$ (the functions of bidders' reports) are assessed.

The logic of the Revelation Principle maintains that the seller should minimize the information revealed to bidders and maximize the information collected from them. By letting each bidder

²⁴If the seller's message is independent of the leader's report, the described mechanism is strategically equivalent to a mechanism in which the seller asks both bidders to submit their reports simultaneously.

report his type as soon as he learns it—and, in particular, by letting the leader report first—the seller maximizes the information collected from the bidders. Because the leader takes no action, the seller minimizes the information revealed to the leader by not sending any message to him. The seller minimizes the information revealed to the follower by sending a message to him only once, right before the follower exerts his (one-off) information-acquisition effort.

Further, without loss of generality, the seller can focus on mechanisms whose (perfect Bayes-Nash) equilibria that are **truthful**, meaning that $(\hat{\theta}_1, \hat{\theta}_2) = (\theta_1, \theta_2)$. Nor does any loss of generality occur if one sets $M = A$ (although a different M will be sometimes convenient).

To respect each bidder’s right to exit without payment at stage (i) of the game-form, the direct mechanism must satisfy the **ex-ante participation constraint**, meaning that each bidder must expect a nonnegative payoff from his participation in the mechanism. To respect each bidder’s right to exit without payment at stage (iii), the direct mechanism must satisfy the **interim participation constraint**, meaning that, having observed his type θ_i , each bidder i must expect a nonnegative payoff from his participation in the mechanism.

As a result, the seller’s problem consists in choosing a disclosure rule μ and a payment rule (t_1, t_2) that induce a mechanism that is ex-post efficient, truthful, and satisfies ex-ante and interim participation constraints so as to maximize the expected revenue:

$$\int_{\Theta_1} \int_M \int_{\Theta_2} (t_1(\theta_1, \theta_2) + t_2(\theta_1, \theta_2)) dF(\theta_2 | a^*(m)) d\mu(m | \theta_1) dG(\theta_1). \quad (3)$$

Two restrictions on admissible mechanisms are crucial for our results: ex-post efficiency and interim participation constraints. These restrictions limit the class of the model’s applications to situations in which the seller does not have full control over the allocation rule and cannot sell to the follower any information about the leader’s valuation. Both restrictions can be motivated by the seller’s inability to commit, respectively, to not attempting to reallocate an inefficiently allocated item and to not fabricating and selling “information” about the leader’s type in the absence of any intention of holding an auction.²⁵

The restrictions are also motivated by our desire to focus on the question that is both economically interesting and analytically tractable. When interim participation constraints are not

²⁵For example, the government may face disgruntled voters if it allocates a procurement contract to a bidder whom everyone knows not to be the best choice.

imposed, optimal information disclosure is trivial. In particular, consistent with the logic of [Eso and Szentes \(2007\)](#), it is optimal to reveal the leader’s information to the follower, as can be shown. When ex-post efficiency of the allocation rule is not imposed, the revenue-maximization problem is intractable (for us), as [Section 7](#) explains.

4 A First-Best Auction

The **first-best** outcome obtains when an omniscient and omnipotent **planner** maximizes the expected total surplus (i.e., the sum of bidders’ payoffs) while observing the leader’s type, directly controlling the follower’s effort, and observing the follower’s realized type. [Corollary 1](#) to [Theorem 1](#) shows that the first-best outcome, which the theorem reports, can be implemented in an auction even without the omniscient and omnipotent planner. Because the first-best outcome is implementable, we shall use it as a benchmark, to gauge the “distortions” introduced by the seller’s profit-maximizing motive.

The first-best outcome is characterized by the follower’s first-best effort. For every leader’s type θ_1 , the follower’s first-best effort, denoted by $\alpha(\theta_1)$, maximizes the total surplus:

$$\alpha(\theta_1) \in \arg \max_{a \in A} \left\{ \int_{\Theta_2} \max\{\theta_1, \theta_2\} dF(\theta_2 | a) - C(a) \right\}. \quad (4)$$

Integrating by parts gives:

$$\alpha(\theta_1) \in \arg \max_{a \in A} \left\{ 1 - \int_{\theta_1}^1 F(\theta_2 | a) d\theta_2 - C(a) \right\}.$$

For a given θ_1 , the planner’s marginal net benefit from an increase in the follower’s effort is the derivative of the maximand in the display above and equals

$$B(\theta_1, a) \equiv R(\theta_1, a) - C'(a), \quad (5)$$

where

$$R(\theta_1) \equiv - \int_{\theta_1}^1 \frac{\partial F(\theta_2 | a)}{\partial a} d\theta_2 \quad (6)$$

is the return to information acquisition (independent of a because F is linear in a), and $C'(a)$ is the

marginal cost of the follower's information-acquisition effort. For any a , $B(1, a) = -C'(a) < 0$ and, by part (i) of Condition 1, $B(0, a) = -C'(a) < 0$. Hence, when $\theta_1 = 1$ or $\theta_1 = 0$, the follower exerts no effort at the first-best. Indeed, in either of these extreme cases, no additional information about the follower's valuation would affect the item's efficient allocation.

Part (ii) of Condition 1 implies that $B(\theta_1, a)$ is strictly increasing in θ_1 for $\theta_1 < \theta^*$ and is strictly decreasing in θ_1 for $\theta_1 \geq \theta^*$. Hence, by the Monotone Selection Theorem (Milgrom and Shannon, 1994), the first-best effort, denoted by $\alpha(\theta_1)$, is weakly increasing in θ_1 for $\theta_1 < \theta^*$ and is weakly decreasing in θ_1 for $\theta_1 \geq \theta^*$, independently of the exact functional form of C . Moreover, the dependence is strict when $\alpha(\theta_1) \in (0, 1)$ (Edlin and Shannon, 1998), which can be shown to be the case for $\theta_1 \in (0, 1)$ as long as the convex C has $C'(0) = 0$ and $C'(1)$ sufficiently large, as indeed is implied by the quadratic C and condition (2). From the first-order condition $R(\theta_1) = C'(\alpha(\theta_1))$, the quadratic C delivers an explicit expression for the first-best effort, leading to Theorem 1.

Theorem 1. *The first-best effort*

$$\alpha(\theta_1) = \frac{1}{c} \int_{\theta_1}^1 (F_L(s) - F_H(s)) ds, \quad \theta_1 \in [0, 1] \quad (7)$$

satisfies $\alpha(0) = \alpha(1) = 0$, is strictly increasing in θ_1 when $\theta_1 < \theta^$, and is strictly decreasing in θ_1 when $\theta_1 \geq \theta^*$.*

Corollary 1 describes an auction that implements the first-best effort. In this auction, the leader who bids b pays the (possibly negative) tax

$$T(b) \equiv \int_0^b (F(s | \alpha(b)) - F(s | \alpha(s))) ds. \quad (8)$$

Corollary 1. *In a mechanism that implements the first-best outcome, the seller*

1. *Asks the leader to submit a bid, denoted by b , and charges him the tax $T(b)$.*
 2. *Discloses b to the follower and invites him to bid in the second-price auction.*
 3. *Allocates the item and assesses the payments according to the rules of the second-price auction.*
- In equilibrium, each bidder bids his type and enjoys a nonnegative expected payoff. The follower exerts the first-best effort.*

Proof. See Appendix 1. □

If not for the leader's tax, the mechanism in Corollary 1 would have been the standard second-price auction only executed sequentially, with the leader's bid made public. In this sequential second-price action, the leader would have had the incentive to manipulate the follower's information acquisition by bidding untruthfully. To see the incentive for untruthful bidding, suppose that type- θ_1 leader truthfully bids $b = \theta_1$. Then, his payoff in the second-price auction is $\mathbb{E}_{\theta_2} [\max \{0, \theta_1 - \theta_2\}]$. Because $\max \{0, \theta_1 - \theta_2\}$ is convex in θ_2 , Jensen's inequality implies that the leader gains if the distribution of the follower's valuation is more dispersed, which occurs when the follower exerts a greater effort. The leader can induce this greater effort by nudging his bid towards θ^* , thereby making the follower more uncertain about his payoff from participating in the auction and hungry for more information. By the envelope argument, this untruthful bidding would have had a second-order detrimental effect on the leader's payoff if the follower's effort had been fixed. Because the follower's effort responds to the leader's bid, however, untruthful bidding in addition has a first-order, beneficial effect on the leader's payoff.

In the mechanism of Corollary 1, the tax discourages untruthful bidding by altering the leader's marginal payoff to raising his bid. The marginal payoff is altered by the amount

$$T'(b) = \alpha'(b) \int_0^b \frac{\partial F(s | \alpha(b))}{\partial a} ds,$$

where the integral is positive for all $b \in (0, 1)$ by Condition 1. Thus, the sign of $T'(b)$ is determined by the sign of $\alpha'(b)$; it is positive when $b < \theta^*$ and negative when $b > \theta^*$. As a result, any increase in the leader's bid below θ^* and any decrease in his bid above θ^* are taxed on the margin.

5 A Seller-Optimal Auction

Without information acquisition, the seller's problem would have been trivial. By the Revenue Equivalence theorem, the ex-post efficient allocation rule, participation constraints, and the optimality of truthful reporting would have tied down the seller's expected payoff to the distribution of the bidders' types, which would have been fixed. With information acquisition, however, the revenue equivalence no longer applies because the distribution of the follower's type is no longer fixed but is affected by the seller's disclosure policy, through the follower's information-

acquisition effort. The disclosure problem thus emerges as the central part of our analysis.

Section 5.1 uses the Envelope Theorem (the tool behind the Revenue Equivalence theorem) to reduce the seller's revenue maximization problem to an information disclosure problem. Section 5.2 shows that the information disclosure problem is nontrivial; it is not solved by disclosing everything, and it is not solved by disclosing nothing. Section 5.3 analyzes the seller's disclosure problem by developing a connection between the seller's continuum-of-types disclosure problem and the finite-types disclosure problem of RS. This connection brings out the qualitative features of the optimal disclosure rule and the induced effort schedule of the follower. Section 5.4 formulates the disclosure problem as an optimal-control problem, which can be used to explore the quantitative features of optimal disclosure and verify the maintained monotonicity condition in examples.

Two remarks clarifying our approach are in order. First, the optimality analysis focuses on the seller's relaxed problem. This relaxed problem neglects the monotonicity condition that requires the leader's probability of winning to be nondecreasing in his type. This condition is necessary for truthful reporting and can be verified in applications (as we do in an example in Section 5.4).²⁶ Relaxed problems, even though restrictive, are sufficiently rich to have become focal in much of the mechanism-design literature.²⁷

Second, by the Revelation Principle (for games with private actions), without loss of generality, one can restrict attention to the disclosure rules that are (possibly stochastic) functions from the leader's types into the effort levels recommended to the follower. The thrust of our analysis lies in showing that an optimal disclosure rule is a deterministic function that maps no more than two types of the leader into the same recommended effort. As a result, one can equivalently identify the seller's message with (i) a recommended effort and (ii) one or two leader's types each of which induces the same effort. We use these two equivalent formulations interchangeably.

²⁶The analogous monotonicity condition for the follower is guaranteed to hold, as will be shown. Furthermore, both monotonicity conditions are guaranteed to hold in all the benchmarks that we consider: in the first-best mechanism of Corollary 1 and in the two benchmarks of Theorem 4, reported in Section 6.

²⁷There are exceptions. The relaxed problem is not typically unrestrictive in the weak-cartels model of McAfee and McMillan (1992) and the queueing model of Hartline and Roughgarden (2008).

5.1 Reduction of the Seller's Auction-Design Problem to an Information-Disclosure Problem

In this section, we use local truthful-reporting and participation constraints to substitute bidders' payments out of the seller's objective function. This step, a standard envelope argument, yields the seller's virtual surplus, thereby reducing the seller's problem to a disclosure problem.²⁸

The Follower's Truth-Telling, Obedience, and Participation Constraints

The incentive constraints are best introduced by rolling the game-form backwards. First, suppose that, having observed the seller's message m , the follower has already exerted some effort and observed his type θ_2 . Now the follower chooses a report that maximizes his interim expected payoff, thereby attaining the value

$$U_2(\theta_2 | m) \equiv \max_{\hat{\theta}_2 \in \Theta_2} \mathbb{E}_{\theta_1 | m} \left[\theta_2 \mathbf{1}_{\{\hat{\theta}_2 > \theta_1\}} - t_2(\theta_1, \hat{\theta}_2) \right], \quad (9)$$

where the expectation is over the leader's type θ_1 conditional on the seller's message m and, implicitly, on the disclosure rule. By inspection of (9), the follower's effort does not enter his interim expected payoff, and so he chooses his report independently of this effort. The follower's **truth-telling constraint** requires the local truth-telling constraint (implied by the Envelope Theorem applied to (9)),

$$U_2'(\theta_2 | m) \equiv \frac{dU_2(\theta_2 | m)}{d\theta_2} = \mathbb{E}_{\theta_1 | m} [\mathbf{1}_{\{\theta_2 > \theta_1\}}], \quad (10)$$

and requires the monotonicity constraint according to which the follower's probability of winning, $\mathbb{E}_{\theta_1 | m} [\mathbf{1}_{\{\theta_2 > \theta_1\}}]$, is nondecreasing in his type, θ_2 . The satisfaction of the monotonicity constraint is immediate, by inspection.

The **interim participation constraint** holds if, even after observing θ_2 , the follower is willing to remain in the mechanism in the sense that his expected continuation payoff is nonnegative.

²⁸The envelope argument is described by [Fudenberg and Tirole \(1991, pp. 284-8\)](#). It is also summarized in the Constraint Simplification Theorem of [Milgrom \(2004\)](#), implied by the Envelope Theorem of [Milgrom and Segal \(2002\)](#).

Formally, for all $\theta_2 \in \Theta_2$ and all $m \in M$, it must be that $U_2(\theta_2 | m) \geq 0$, or equivalently,

$$U_2(0 | m) + \int_0^{\theta_2} \mathbb{E}_{\theta_1 | m} [\mathbf{1}_{\{s > \theta_1\}}] ds \geq 0, \quad (11)$$

where the equivalence holds by the Constraint Simplification Theorem of [Milgrom \(2004\)](#), which justifies the application of the fundamental theorem of calculus and rewrites U_2 in terms of its derivative from (10). Because the right-hand side of (10) is nonnegative, $U_2(\theta_2 | m)$ is nondecreasing. Thus, the follower's interim participation constraint holds if and only if

$$U_2(0 | m) \geq 0 \quad \text{for all } m \in M. \quad (12)$$

Remark 1. Supposing that $U_2(0 | m) = 0$ for all $m \in M$ (as will be the case in the optimal mechanism), the interim participation constraint rules out mechanisms that ask the follower to commit to a payment in exchange for the right to participate in the mechanism, as well as the mechanisms that offer to sell information about the leader's type before the follower decides which effort to exert.²⁹

One can now take a step back in the game-form and ask which effort is optimal for the follower who observes message m and knows that he will optimally report truthfully in future. Because the seller cannot directly control the follower's effort, the seller resigns to letting the follower exert the effort that the follower finds optimal. Equivalently, the seller is restricted to recommending efforts according to a rule that makes it optimal for the follower to obey the seller's recommendations. That is, the **obedience constraint** must hold. Formally, right after observing message m and having decided to exert effort a' , the follower expects his payoff net of the cost of effort to be

$$\int_{\Theta_2} U_2(\theta_2 | m) dF(\theta_2 | a') = U_2(0 | m) + \int_0^1 \mathbb{E}_{\theta_1 | m} [\mathbf{1}_{\{\theta_2 > \theta_1\}}] (1 - F(\theta_2 | a')) d\theta_2, \quad (13)$$

where the equality uses (10) and integration by parts.³⁰ Interchanging the order of integration and expectation (by Fubini's theorem) in the right-hand side of the above display yields the expression

²⁹In practice, shareholders may forbid managers to commit to any payments until due diligence (information acquisition) has been performed.

³⁰Integration by parts is valid even if F is discontinuous (as in Example 1), because $U_2(\cdot | m)$ has a bounded derivative everywhere and hence is continuous.

for the follower's optimal effort $a^*(m)$ as a function of the observed message m :

$$a^*(m) \in \arg \max_{a' \in A} \left(U_2(0 | m) + \mathbb{E}_{\theta_1 | m} \left[\int_{\theta_1}^1 (1 - F(\theta_2 | a')) d\theta_2 \right] - C(a') \right). \quad (14)$$

Under the maintained Condition 1, the maximization problem in (14) has the unique solution:

$$a^*(m) = \mathbb{E}_{\theta_1 | m} [\alpha(\theta_1)], \quad (15)$$

where $\alpha(\theta_1)$, defined in (7), is the first-best effort level when the leader's type is θ_1 . The obedience constraint thus requires that, if the seller's message space is the set of recommended efforts, then each recommended effort m satisfies $m = a^*(m)$.

With the knowledge that the follower will be truthful and obedient, one can take one more step back and impose the **ex-ante participation constraint**, which requires that the follower expect a nonnegative payoff from the mechanism right after he has observed the seller's message but before he has exerted any effort. Formally, for any $m \in M$, the maximand in (14) evaluated at the optimal effort $a^*(m)$ must be nonnegative:

$$U_2(0 | m) + \mathbb{E}_{\theta_1 | m} \left[\int_{\theta_1}^1 (1 - F(\theta_2 | a^*(m))) d\theta_2 \right] - C(a^*(m)) \geq 0. \quad (16)$$

Substituting the functional forms of F and C , and the expressions for a^* and α , from (15) and (7), and rearranging gives

$$U_2(0 | m) + \mathbb{E}_{\theta_1 | m} \left[\int_{\theta_1}^1 (1 - F_L(\theta_2)) d\theta_2 \right] + C(a^*(m)) \geq U_2(0 | m),$$

where the inequality follows by inspection. Moreover, the interim participation constraint (12) requires $U_2(0 | m) \geq 0$, and hence, by the display above, the ex-ante participation constraint in (16) is implied by (12). Hence, from now on, we focus on the follower's interim participation constraint and refer to it simply as his participation constraint.

The Leader's Truth-Telling and Participation Constraints

Having observed his type, the leader chooses a report that maximizes his expected payoff, thereby attaining the value

$$U_1(\theta_1) \equiv \max_{\hat{\theta}_1 \in \Theta_1} \mathbb{E}_{\theta_2 | \theta_1} \left[\theta_1 \mathbf{1}_{\{\hat{\theta}_1 > \theta_2\}} - t_1(\theta_1, \theta_2) \right]. \quad (17)$$

As in the case of the follower, by the Constraint Simplification Theorem of Milgrom (2004), the leader's **truth-telling constraint** is equivalent to the integral condition

$$U_1(\theta_1) = U_1(0) + \int_0^{\theta_1} \mathbb{E}_{m | \theta_1} [F(\theta_1 | a^*(m))] ds \quad (18)$$

and the monotonicity condition on the probability of winning:

$$\mathbb{E}_{\theta_2 | \theta_1} [\mathbf{1}_{\{\theta_1 > \theta_2\}}] = \mathbb{E}_{m | \theta_1} [F(\theta_1 | a^*(m))] \quad \text{is nondecreasing in } \theta_1. \quad (19)$$

In contrast to the follower's monotonicity condition, the leader's monotonicity condition (19) cannot be a priori argued to hold. Instead, we proceed with the analysis *assuming* that this condition holds, and then verify it in examples.

The leader's **interim participation constraint**, or simply **participation constraint**, ensures that each type of the leader is at least as well off in the mechanism as he would be if he were to refrain from participation and enjoy the payoff of zero: $U_1(\theta_1) \geq 0$ for all θ_1 . Differentiating (18) gives $U_1'(\theta_1) = \mathbb{E}_{m | \theta_1} [F(\theta_1 | a^*(m))] \geq 0$; that is, U_1 is nondecreasing, and so the participation constraint holds for all θ_1 as long as it holds for $\theta_1 = 0$:

$$U_1(0) \geq 0. \quad (20)$$

The Seller's Virtual Surplus

We can now use (18) combined with (17) and (13) both evaluated at $a^*(m)$ and combined with (9) to substitute out the bidders' transfers from the seller's objective function (3). From the seller's perspective, it is optimal to set the transfers for the lowest-type leader and the lowest-type follower so that their expected payoffs are zero. Because $U_1(0) = 0$ and, for all m , $U_2(0 | m) = 0$,

the seller's virtual surplus is^{31, 32}

$$\mathbb{E}_{m, \theta_1, \theta_2} \left[\left(\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} \right) \mathbf{1}_{\{\theta_1 > \theta_2\}} + \left(\theta_2 - \frac{1 - F(\theta_2 | a^*(m))}{f(\theta_2 | a^*(m))} \right) \mathbf{1}_{\{\theta_2 \geq \theta_1\}} \right]. \quad (21)$$

The displayed virtual surplus is the expected sum of each bidder's **virtual valuation** times the probability that he gets the item. The probability of getting the item is pinned down by the ex-post efficient allocation rule. Except for the dependence of a^* on θ_1 (through m), the virtual surplus is standard.

Integrating θ_2 out of (21) and simplifying yields a more compact expression for the virtual surplus.³³

$$\int_{\Theta_1} \mathbb{E}_{m|\theta_1} \left[\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} F(\theta_1 | a^*(m)) \right] dG(\theta_1). \quad (22)$$

To see why (22) is equivalent to (21), suppose that the leader's type is θ_1 and the seller sends a message m . Ex-post efficiency requires that the leader get the item with probability $F(\theta_1 | a^*(m))$, which is the probability of the event $\{\theta_2 < \theta_1\}$. In this case, the seller's gain is the leader's virtual valuation $\theta_1 - (1 - G(\theta_1)) / g(\theta_1)$, which is his true valuation θ_1 less the **information rent** $(1 - G(\theta_1)) / g(\theta_1)$. Analogous reasoning suggests that if the follower gets the item (which occurs with probability $1 - F(\theta_1 | a^*(m))$), the seller's gain is the follower's expected virtual valuation conditional on winning, which can be verified to be θ_1 .³⁴ Thus, the follower's information rent is implicit in (22). Expression (22) is asymmetric only because, asymmetrically, θ_2 has been designated to be integrated out. Expression (22) reflects the dynamic nature of the mechanism only in that m may depend on θ_1 .

The seller's auction-design problem has been thus reduced to the optimal-disclosure problem,

³¹If F has no density f , skip to (22), which does not rely on the existence of f .

³²The follower's information acquisition cost does not enter the virtual surplus directly, because the follower's interim participation constraint implies the ex-ante participation constraint. That is, the seller does not have to (even partially) compensate the follower for the cost of information acquisition to induce the follower to participate; the follower's expected "information rent" from trading suffices to cover his information acquisition cost.

³³The virtual surplus in (22) is non-negative because it is bounded below by $\int_{\Theta_1} [\theta_1 - (1 - G(\theta_1)) / g(\theta_1)] dG(\theta_1) = 0$. Thus, also the seller's payoff is nonnegative.

³⁴Formally, the follower's expected virtual valuation conditional on winning when the leader's type is θ_1 is

$$\int_{\theta_1}^1 \left(\theta_2 - \frac{1 - F(\theta_2 | a^*(m))}{f(\theta_2 | a^*(m))} \right) \frac{f(\theta_2 | a^*(m))}{1 - F(\theta_1 | a^*(m))} d\theta_2 = \theta_1.$$

Even though the integrand in the above display is written as if F had a positive density, f , the validity of (22) requires no such assumption.

in which the seller maximizes (22) over disclosure rules. By the Revelation Principle, the seller can restrict attention to disclosure rules with the property that, for each $m \in M$, $m = a^*(m)$; that is, the message is the recommended action. Nevertheless, because it is convenient to keep track of which types of the leader induce which recommended actions, we instead adopt the equivalent (as we show) formulation that has the property that each $m \in M$ is the set of the leader's types that induces effort $a^*(m)$.

5.2 The Suboptimality of Full Disclosure and Non-Disclosure

We show, in Theorem 2, that the seller's optimal-disclosure problem is nontrivial in that both full disclosure and non-disclosure are suboptimal. **Full disclosure** is a disclosure rule that assigns a distinct message to each type of the leader. **Non-disclosure** is a disclosure rule that pools all types of the leader under the same message.

The proof of Theorem 2 uses the seller's objective function (22) rewritten in "product form," which is also used in the proofs of subsequent results. To arrive at this form, neglect the additive term in (22) that is independent of the disclosure rule and rewrite (22) as

$$\int_{\Theta_1} \mathbb{E}_{m|\theta_1} [\pi(\theta_1) a^*(m)] dG(\theta_1), \quad (23)$$

where

$$\pi(\theta_1) \equiv \frac{1 - G(\theta_1)}{g(\theta_1)} (F_L(\theta_1) - F_H(\theta_1)) \quad (24)$$

denotes the seller's **marginal benefit** from an increase in a . This marginal benefit equals the leader's information rent times the marginal increase in the probability that the follower wins. Equation (23) is further transformed using the Law of Iterated Expectations and the expression for $a^*(m)$ in (15) to yield the **product form**

$$\mathbb{E} [\mathbb{E}_{\theta_1|m} [\pi(\theta_1)] \mathbb{E}_{\theta_1|m} [\alpha(\theta_1)]] . \quad (25)$$

The transformed objective function has the same form as the sender's objective function (equation [2]) in the model of RS.³⁵ The product structure of (25) hinges on two assumptions: the ex-post

³⁵In what follows, bracketed indices refer to equations and lemmas in RS's paper.

efficient allocation rule³⁶ and the linearity of the c.d.f. $F(\theta_2 | a)$ in the follower's effort a . Both of these assumptions are crucial for adapting RS's techniques to our setting.

To state our results, we borrow vocabulary from RS. A tuple $(\pi(\theta_1), \alpha(\theta_1))$ is called a **prospect**. The **prospect set** is the graph $\Gamma \equiv \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in \Theta_1\}$.³⁷

Showing that neither full disclosure, nor non-disclosure is optimal requires two lemmas, which are of independent interest. For Lemma 1, call functions α and π **ordered** if $(\alpha(s') - \alpha(s))(\pi(s') - \pi(s)) \geq 0$ for almost all $s, s' \in \Theta_1$.

Lemma 1. *Full disclosure is optimal if and only if α and π are ordered.*

Proof. Necessity: Suppose that α and π are not ordered. Then, an interval $I \subset \Theta_1$ exists on which α is strictly increasing and π is strictly decreasing, or the other way around. In this case,

$$\int_I \int_I (\alpha(s') - \alpha(s)) (\pi(s') - \pi(s)) ds ds' < 0.$$

In the display above, multiplying the parentheses, defining $|I| \equiv \int_I ds$, and rearranging yields³⁸

$$\int_I \alpha(s) \pi(s) ds < \frac{1}{|I|} \int_I \alpha(s) ds \int_I \pi(s) ds,$$

where the left-hand side is the seller's expected payoff from fully disclosing the types in I , and the right-hand side is the seller's expected payoff from pooling all types in I under the same message. Thus, full disclosure is suboptimal.

*Sufficiency:*³⁹ Suppose that α and π are ordered. By contradiction, suppose that, with probability-density p , prospect $(\pi(s), \alpha(s))$ occurs and induces a message m . Suppose also that, with probability-density p' , prospect $(\pi(s'), \alpha(s'))$ with $s' \neq s$ occurs and induces the same message m .⁴⁰ The seller's gain from pooling the two prospects under message m relative to revealing each of them

³⁶Section 7 shows that, as long as an allocation rule is fixed, it can depart from ex-post efficiency without compromising the product structure of the seller's objective function. If the seller were also to maximize over allocation rules, however, the product structure would be lost.

³⁷The set $\{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in \Theta_1\}$ differs from a RS prospect set only in that RS require their prospect sets to be finite, for tractability.

³⁸The obtained inequality is a continuous version of Chebyshev's sum inequality.

³⁹The sufficiency argument is Lemma [1] of RS and is included for completeness.

⁴⁰The message m may also be induced by some other prospect. Either prospect may also induce some other message.

is

$$\begin{aligned} & \frac{p\alpha(s) + p'\alpha(s')}{p + p'} \frac{p\pi(s) + p'\pi(s')}{p + p'} (p + p') - p\alpha(s)\pi(s) - p'\alpha(s')\pi(s') \\ & = -\frac{pp'(\alpha(s') - \alpha(s))(\pi(s') - \pi(s))}{p + p'} \leq 0, \end{aligned}$$

where the inequality follows because α and π are ordered. Thus, a weak improvement can be attained by revealing any two prospects that are sometimes pooled under the same message; full disclosure is optimal. \square

For Lemma 2, define a **nonincreasing line** as a straight line that is either vertical or has a nonpositive slope.

Lemma 2. *It is optimal to pool a subset S of the prospect set Γ under the same message if and only if S lies on a nonincreasing line.*

Proof. For sufficiency, suppose first that S lies on a vertical line. Then, the seller's payoff from S , $\mathbb{E}[\mathbb{E}_{|m}[\alpha] \mathbb{E}_{|m}[\pi]] = \pi \mathbb{E}[\alpha]$, is independent of the disclosure rule. Any disclosure of the elements of S is optimal, including pooling them under the same message.

If Γ is a nonincreasing line that is not vertical, then for some $k_0 \in \mathbb{R}$ and $k_1 \in \mathbb{R}_+$, every prospect $(\alpha(\theta_1), \pi(\theta_1))$ in S can be written as $\alpha(\theta_1) = k_0 - k_1\pi(\theta_1)$. The seller's payoff from S ,

$$\mathbb{E}[\mathbb{E}_{|m}[\alpha] \mathbb{E}_{|m}[\pi]] = k_0 \mathbb{E}[\pi] - k_1 \mathbb{E}[(\mathbb{E}_{|m}[\pi])^2],$$

is maximized when $\mathbb{E}[(\mathbb{E}_{|m}[\pi])^2]$ is minimized, which, by Jensen's inequality, occurs when the signal structure is least informative in Blackwell's sense, when the random variable $\mathbb{E}_{|m}[\pi]$ is least dispersed. The least dispersion is achieved by pooling all prospects in S under the same message.

To summarize, pooling all prospects on a line segment is optimal, and strictly so when $k_1 > 0$.

Necessity follows from RS's Lemma [3]. \square

Taking $S = \Gamma$ in Lemma 2 immediately yields

Corollary 2. *Non-disclosure is optimal if and only if the prospect set Γ lies on a nonincreasing line.*

One can now state and prove Theorem 2.

Theorem 2. *Under Condition 1, the policies of full disclosure and non-disclosure are suboptimal. If, in addition, c.d.f.s G , F_L , and F_H are analytic functions, it is never optimal to pool an open interval of the leader's types under the same message.⁴¹*

Proof. Recall that

$$\begin{aligned}\pi(\theta_1) &= \frac{1 - G(\theta_1)}{g(\theta_1)} (F_L(\theta_1) - F_H(\theta_1)) \\ \alpha(\theta_1) &= \frac{1}{c} \int_{\theta_1}^1 (F_L(s) - F_H(s)) ds.\end{aligned}$$

By Theorem 1, α is uniquely maximized at $\theta_1 = \theta^* \in (0, 1)$. By part (ii) of Condition 1 and by the above display, $\theta_1 < \theta^* \implies \pi(\theta_1) < 0$ and $\theta_1 > \theta^* \implies \pi(\theta_1) > 0$. Thus, α and π are not ordered, and Lemma 1 implies that full disclosure is suboptimal.

The prospect set Γ does not lie on a nonincreasing line. Indeed, Γ is not on a vertical line because the sign of $\pi(\theta_1)$ depends on θ_1 , as argued above. Nor is Γ on a decreasing line, because for each $\theta_1 < \theta^*$, there exists an $\epsilon > 0$ such that $\alpha(\theta_1) < \alpha(\theta^* + \epsilon)$ and $\pi(\theta_1) < 0 < \pi(\theta^* + \epsilon)$. Hence, Corollary 2 implies that non-disclosure is suboptimal.

The remainder of the proof establishes the suboptimality of pooling an open interval of types and relies on two standard observations about analytic functions and analytic curves.⁴²

Observation 1. Sums, products, reciprocals (if well-defined), derivatives, and integrals of analytic functions are analytic.

Observation 2. If two analytic curves coincide on any open interval, these curves are identical everywhere.

By Observation 1, π and α are analytic. Hence, the prospect set Γ is analytic.

By Lemma 2, all types in an open interval in Θ_1 can be optimally pooled under the same message only if a nonincreasing line coincides with the prospect set Γ on that interval. If so, Observation 2 implies that Γ must be a nonincreasing line, which has been shown to be false. Hence, no interval of types is optimally pooled. \square

⁴¹A function is analytic if it can be locally represented by a convergent Taylor series. Many common functions are analytic, and by the Stone–Weierstrass Theorem, any continuous function can be approximated arbitrarily well by an analytic function.

⁴²A curve is analytic if it has an analytic parametrization.

5.3 The Optimality of Conjugate Disclosure

Under an additional, convexity, assumption, this section derives the structure of an optimal information-disclosure rule, summarized in Theorem 3. Roughly speaking, this rule partitions Θ_1 into pairs and singletons and reveals to the follower only the element of the partition to which the leader's type belongs. The partition has a special structure, pooling "extreme" leader types into pairs. The two types that are pooled in a pair are called "conjugate," and so the optimal disclosure rule is called "conjugate disclosure rule." The optimal effort schedule recovered from the optimal disclosure rule is also described in Theorem 3.

The remainder of this section is concerned with stating and proving Theorem 3. Doing so calls for new definitions and intermediate results. This section is more technical than others due to the subtleties stemming from the assumption that the leader's type space is continuous.⁴³

Restriction to Convex Prospect Sets

This paper's techniques for finding an optimal disclosure rule rely on the prospect set Γ being convex in the sense of

Definition 1. A prospect set Γ is **convex** if it is a strictly convex curve.⁴⁴

Convex prospect sets are illustrated in Figure 3. Because Condition 1 does not imply (nor is implied by) the convexity of Γ , the subsequent analysis will maintain an additional assumption:

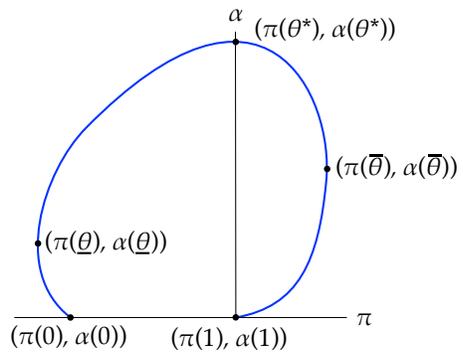
Condition 2. The prospect set Γ is convex.

One can verify that Conditions 1 and 2 imply the existence of $\underline{\theta} \in [0, \theta^*)$ and $\bar{\theta} \in (\theta^*, 1)$ (illustrated in Figure 3) such that:

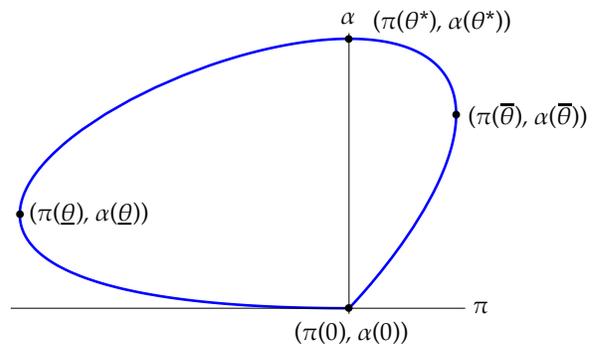
- On $(0, \underline{\theta})$, Γ is downward-sloping (α is strictly increasing; π is strictly decreasing).
- On $(\underline{\theta}, \theta^*)$, Γ is upward-sloping (both α and π are strictly increasing).

⁴³Assuming instead that Θ_1 is finite would have simplified some arguments, complicated others, and, crucially, weakened the conclusions. With a finite but large Θ_1 , each optimal disclosure rule (in case there are multiple) may fail to be conjugate, but this failure is economically insignificant. This insignificance can be ascertained by studying the limit of an increasingly finer sequence of type sets (which we do in this section). With the continuous Θ_1 , the emergent conjugate disclosure rule is easy to interpret economically and can be recovered from an optimal-control problem.

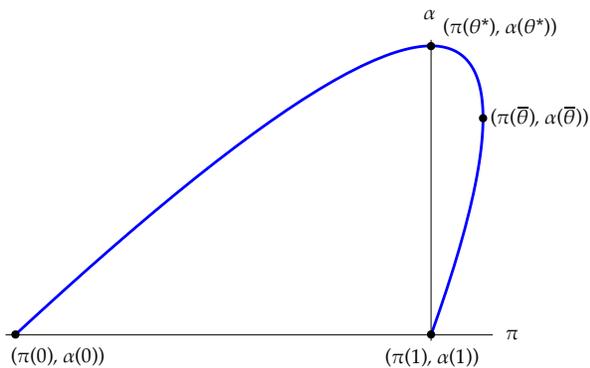
⁴⁴A strictly convex curve is a curve that intersects any line at most twice.



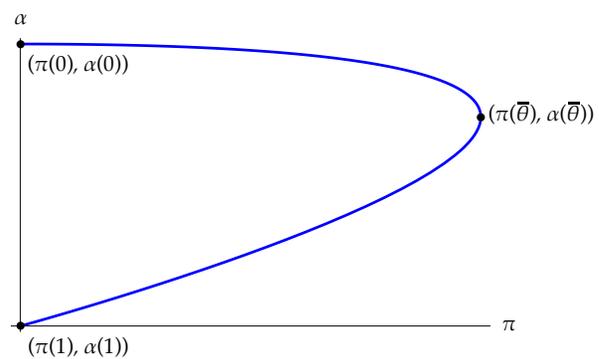
(a) A “typical” convex prospect set that satisfies Condition 1.



(b) G is uniform, and F_L and F_H are Beta-distribution c.d.f.s that satisfy Condition 1.



(c) G is uniform, and F_L and F_H are as specified in Example 1 (and hence satisfy Condition 1).



(d) A convex prospect set that violates Condition 1.

Figure 3: Convex prospect sets. Certain “critical” points have been marked on each prospect set and are referenced in the subsequent analysis. An increase in θ_1 corresponds to the clockwise movement along the prospect set.

- On $(\theta^*, \bar{\theta})$, Γ is downward-sloping (α is strictly decreasing; π is strictly increasing).
- On $(\bar{\theta}, 1)$, Γ is upward-sloping (both α and π are strictly decreasing).

The feasibility of partitioning Γ into the described segments relies on α being single-peaked (implied by Condition 1; see Theorem 1) and on π being decreasing (possibly on a degenerate interval), then increasing, and then decreasing again. This restriction on π is an additional joint restriction on G and F , embedded into Condition 2.

In applications, Condition 2 can be checked analytically. To do so, let

$$r(\theta_1) \equiv \frac{1 - G(\theta_1)}{g(\theta_1)}, \quad \theta_1 \in \Theta_1,$$

denote the inverse hazard rate of the leader's c.d.f. As is standard, $r(\theta_1)$ is interpreted as the profit that the seller forgoes—equivalently, the information rent the leader reaps—when the seller commits to sell to type- θ_1 leader.⁴⁵ In addition, recall that $R(\theta_1)$, defined in (6), denotes the planner's return to the follower's information acquisition in the first-best benchmark when the leader's type is θ_1 . This return is closely related to the follower's information-acquisition technology (in particular, $R'(\theta_1) = F_H(\theta_1) - F_L(\theta_1)$ and $R''(\theta_1) = f_H(\theta_1) - f_L(\theta_1)$) and so can be treated as a primitive.

Lemma 3. *Suppose that Condition 1 holds and $f_L(\theta^*) \neq f_H(\theta^*)$.⁴⁶ Then, a prospect set is convex if and only if*

$$r''(\theta_1) + \left(r(\theta_1) \frac{R''(\theta_1)}{R'(\theta_1)} \right)' < 0 \quad \text{for all } \theta_1 \in (0, \theta^*) \cup (\theta^*, 1). \quad (26)$$

Proof. See Appendix B. □

Inequality (26) in Lemma 3 is essentially equivalent to saying that the prospect set's curvature does not change the sign as one varies θ_1 . This curvature condition is local and hence does not quite suffice to conclude that the prospect set is convex (in the sense of Definition 1); a spiral is a counterexample. Coupled with Condition 1, however, which ensures $\alpha(0) = \alpha(1) = 0$ (and so in

⁴⁵When the seller commits to sell to type θ_1 at some price, all types higher than θ_1 may be tempted to imitate type θ_1 and buy at the same price, thereby constraining the seller in how much he can charge these higher types.

⁴⁶Condition $f_L(\theta^*) \neq f_H(\theta^*)$, which can be interpreted to hold "generically," simplifies the analytical characterization in the lemma but is not required for the convexity of the prospect set.

particular rules out the possibility that Γ is a spiral), the curvature condition in (26) is equivalent to the convexity of Γ .

The value of the calculus characterization in Lemma 3 is in providing an analytical condition to check in applications, as an alternative to plotting and eyeballing the prospect set. The characterizing inequality (26) does not have a natural economic interpretation, though. Nor does convexity itself. Instead, our two justifications for imposing the convexity assumption are that (i) it is satisfied in various “natural” examples (such as those generating Figures 3b and 3c), and (ii) it renders the seller’s optimal-disclosure problem tractable.

A Discretized Prospect Set

The analysis draws on RS’s optimal-disclosure results, which have been developed for discrete prospect sets, and which we extend to the continuous set Γ by taking an appropriate limit. For an arbitrary $n \geq 1$, we let Γ^n denote a discrete prospect set induced by an n -th finite approximation of the leader’s type space Θ_1 . In particular, for an integer $n \geq 1$, let the discretized type space be $\Theta_1^n \equiv \{y_i\}_{i=1}^{2^n}$, where $y_i = i/2^n$, $i \in \{0, 1, 2, 3, \dots, 2^n\}$. The probability of any $y_i \in \Theta_1^n$ is set equal to $G(y_i) - G(y_{i-1})$, which is the probability of interval $(y_{i-1}, y_i] \subset \Theta_1$. The approximation Θ_1^n is finer for larger values of n (i.e., $\Theta_1^n \subset \Theta_1^{n+1}$) and satisfies $\cup_{n=1}^{\infty} \Theta_1^n = \Theta_1$. The induced discrete prospect set is denoted by $\Gamma^n \equiv \{(\pi(y), \alpha(y)) : y \in \Theta_1^n\}$. Prospect $(\pi(y_i), \alpha(y_i))$ will be referred to as prospect i and denoted by (π_i, α_i) .

Optimal Disclosure with the Discrete Prospect Set

For a discrete prospect set Γ^n , the seller’s disclosure problem, denoted by \mathcal{P}^n , is a special case of the problem studied by RS, whose results we shall use to narrow down the search for an optimal disclosure rule. The results use the following jargon. A prospect is **revealed** if it induces a message that causes the follower to assign probability one to this prospect. Two prospects are **pooled** if they sometimes induce the seller to send the same message. Graphically, in the (π, α) -space of prospects, this shared message is represented by a **pooling link**, a line segment that connects two pooled prospects on the graph Γ^n . If Γ^n is derived from Γ that satisfies Condition 2, then RS’s results imply the following facts about every \mathcal{P}^n -optimal disclosure rule.

Fact 1. By RS's Lemma [1], no two prospects are pooled if (i) both lie in $\Gamma^n \cap \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in [\underline{\theta}, \theta^*]\}$ or (ii) both lie in $\Gamma^n \cap \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in [\bar{\theta}, 1]\}$.

Fact 2. By RS's Lemma [3] (also by this paper's Lemma 2), only the prospects that lie on a nonincreasing line can be pooled under the same message.

Fact 3. By RS's Lemma [3], at most two prospects can be pooled under the same message, because no more than two prospects lie on the same nonincreasing line, by Condition 2.⁴⁷

Fact 4. By RS's Lemma [4], no two pooling links intersect.

Fact 5. By RS's Proposition [1], a prospect either is always revealed or is pooled with some other prospects with probability one.

This partial characterization of a \mathcal{P}^n -optimal disclosure rule in Facts 1–5 is refined in Lemma 4, which further exploits the special structure of the seller's problem.

Lemma 4. Suppose that Condition 2 holds. Then, any discrete disclosure problem \mathcal{P}^n has an optimal disclosure rule that is partially characterized by an $s_* \in [0, \underline{\theta}]$ and an $s^* \in [\theta^*, \bar{\theta}]$ such that

(i) Any type in $[s_*, s^*] \cap \Theta_1^n$ either is always revealed or is pooled with some types in $([0, s_*] \cup [s^*, 1]) \cap \Theta_1^n$. Symmetrically, any type in $([0, s_*] \cup [s^*, 1]) \cap \Theta_1^n$ either is always revealed or is pooled with some types in $[s_*, s^*] \cap \Theta_1^n$. The types are pooled so that, in the prospect space, the pooling links never intersect.

(ii) The optimal effort is single-peaked and, if $s^* \in \Theta_1^n$, is maximal at type s^* . (If $s^* \notin \Theta_1^n$, the optimal effort is maximal "close" to s^* , either at type $\max\{[0, s^*] \cap \Theta_1^n\}$ or at type $\min\{[s^*, 1] \cap \Theta_1^n\}$.)

Proof. See Appendix A.3. □

Part (i) of Lemma 4 defines types s_* and s^* , both in Θ_1 (not necessarily in Θ^n), such that each pooling link intersects the line that passes through prospects $(\pi(s_*), \alpha(s_*))$ and $(\pi(s^*), \alpha(s^*))$, as shown in Figure 4. Part (ii) of the lemma shows that the optimal effort schedule is single-peaked in θ_1 , with the peak, at s^* , to the right of the peak of the first-best effort schedule, at θ^* .

⁴⁷Fact 3 does not rule out the situation in which a prospect probabilistically invokes one of multiple messages.

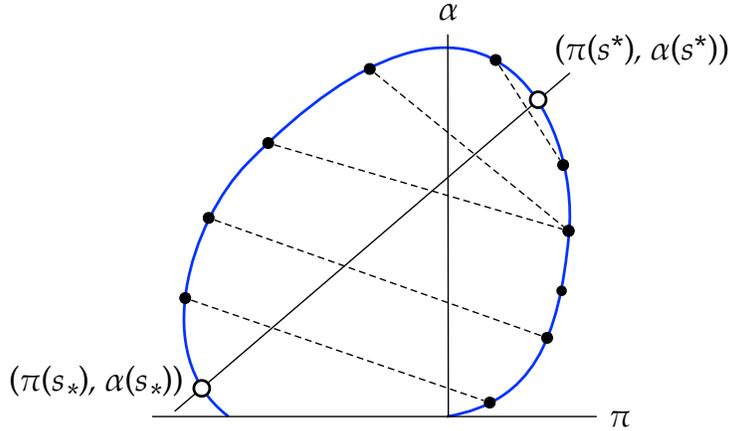


Figure 4: A discretized prospect set Γ^n is the union of solid dots. Without loss of generality, the seller can restrict attention to disclosure rules such as the one illustrated here. Each dashed link denotes a message that pools two prospects. These links never intersect and are oriented so that one can draw an upward-sloping line (passing through points $(\pi(s_*), \alpha(s_*))$ and $(\pi(s^*), \alpha(s^*))$, both marked by empty dots) that intersects each of the pooling links. The isolated prospect is revealed.

Optimal Disclosure with the Continuous Prospect Set: The Main Result

The \mathcal{P}^n -optimal disclosure rule of Lemma 4 either reveals a prospect or pools it under the same message with another prospect. Lemma 4 does not rule out situations in which a prospect probabilistically invokes multiple messages (i.e., several pooling links could emanate from a single prospect). Theorem 3 shows, however, that, in the continuous problem, denoted by \mathcal{P} , there is no loss of generality in focusing on disclosure rules that deterministically associate each prospect with a unique message.

The formal argument proceeds in two steps. Lemma 5 shows that, starting from a \mathcal{P}^n -optimal disclosure rule, one can construct a disclosure rule for \mathcal{P} that pools prospects deterministically and delivers a payoff close to the optimal payoff in \mathcal{P}^n . Roughly, when optimality in \mathcal{P}^n calls for probabilistically pooling a prospect under multiple messages, the disclosure rule in \mathcal{P} splits the corresponding “prospect” into nearby prospects each of which is pooled deterministically. This splitting exploits the continuity of the type space. Thus constructed disclosure rule is then verified to be optimal in \mathcal{P} by using a limit argument of Lemma 6. The results of Lemma 5 and Lemma 6 combine in Theorem 3, which describes \mathcal{P} -optimal disclosure and the associated effort schedule.

To state and prove Theorem 3, we need two more definitions (Definitions 2 and 3). Definition 2

describes a function that partitions the prospect set into revealed singletons and pooled pairs.

Definition 2. A **matching function** τ takes one of two forms: (i) $\tau : [0, s^*] \rightarrow [s^*, 1]$ is weakly decreasing, with $\tau(0) = 1$ and $\tau(s^*) = s^*$, $0 < s^* < 1$, or (ii) $\tau : [s_*, s^*] \rightarrow (0, s_*] \cup [s^*, 1]$ is weakly decreasing on $[s_*, s')$ and on $[s', s^*]$, with $\tau(s_*) = s_*$, $\lim_{s \uparrow s'} \tau(s) = 0$, $\tau(s') = 1$, and $\tau(s^*) = s^*$, $0 < s_* \leq s' < s^* < 1$.

In Definition 2, case (i) is equivalent to case (ii) with $s' = s^*$. Case (i) prevails in examples in which the distribution of the follower's underlying valuations is binary, and the information-acquisition technology grants probabilistic access to a perfectly informative signal, as in Example 1. In this case, the follower's c.d.f. F has mass points at 0 and 1. Example 1, coupled with the assumption of the monotone increasing hazard rate for the leader's c.d.f. G , yields the prospect set in Figure 3c. This prospect set's critical feature is that it slopes upwards near $\theta_1 = 0$ (that is, both α and π are increasing in θ_1 near 0) and so $\underline{\theta} = 0$. That α is increasing near 0 follows from Theorem 1. That π is increasing near 0 follows by taking an arbitrarily small $\varepsilon > 0$ and evaluating

$$\pi'(\varepsilon) = -r'(\varepsilon)(F_H(\varepsilon) - F_L(\varepsilon)) + r(\varepsilon)(f_L(\varepsilon) - f_H(\varepsilon)) > 0,$$

where the inequality follows because $r'(\varepsilon) < 0$ (the hazard-rate condition on G), $r(\varepsilon) > 0$, $F_H(\varepsilon) = 1/2 > F_L(\varepsilon) = \varepsilon$ (the mass point that corresponds to probabilistically learning that the underlying valuation is 0), and $f_L(\varepsilon) = 1 > f_H(\varepsilon) = 0$ (made possible by F_H 's mass point at 0).

Figure 5 illustrates how a downward-sloping segment for Γ near $\theta_1 = 0$ is necessary for case (ii) not to collapse into case (i). The figure also illustrates the role played by s' . In particular, the matching function takes values in $(0, s_*]$ when $s \in [s_*, s')$ and in $[s^*, 1]$ when $s \in [s', s^*]$. Point s' is the point of discontinuity where the matching function jumps upward from 0 to 1, thereby switching from taking values in one to taking values in the other interval of its codomain.⁴⁸ An example of case (ii) is illustrated in Figure 3b; c.d.f.s F_H and F_L are Beta distributions chosen to satisfy Condition 1. Then, $F_H(0) = F_L(0) = 0$. Furthermore, one can (merely to simplify the argument) choose F_H and F_L so that $f_H(0) > f_L(0)$. As a result, $\pi'(0) = r(0)(f_L(0) - f_H(0)) < 0$; π is decreasing near 0. Because α is increasing near 0, the prospect set is downward-sloping

⁴⁸Mathematically, case (ii) can be viewed as isomorphic to case (i) if the matching function's domain is "offset" by s_* , so that s_* is the "new zero," and every point in $(0, s^*)$ is "greater than" every point in $(s^*, 1)$.

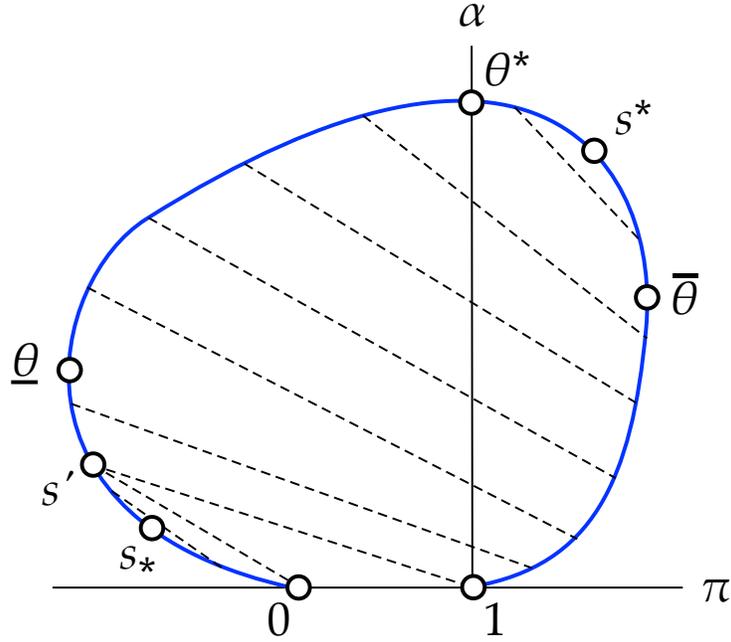


Figure 5: The convex curve is the prospect set Γ . The circles mark the prospects induced by the leader's types in $0, s_*, s', \underline{\theta}, \theta^*, s^*, \bar{\theta}$, and 1 . The dashed links comprise a subset of the links that pool prospects into messages. Type s' demarcates the leader's types that are pooled with types in $[0, s_*)$ and those that are pooled with types in $(s^*, 1]$.

near 0 .

A disclosure rule that reveals an element of the partition described by a matching function is called conjugate:

Definition 3. Under the **conjugate disclosure rule** induced by a matching function τ , the seller who receives leader's report θ_1 sends message s to the follower if $\theta_1 \in \{s, \tau(s)\}$ for some s ; otherwise, the seller sends message $s = \theta_1$.

According to Definition 3, a conjugate disclosure rule either fully discloses the leader's type or pools it with one other type. When τ is differentiable at s , the seller's announcement $\{s, \tau(s)\}$ induces the follower to assign probability $g(s) / (g(s) + |\tau'(s)| g(\tau(s)))$ to the event $\theta_1 = s$ and the complementary probability to the event $\theta_1 = \tau(s)$, by Bayes's rule. When $\tau'(s) = 0$, the seller's announcement $\{s, \tau(s)\}$ induces the follower to assign probability one to $\theta_1 = s$.⁴⁹ Finally, when the range of τ omits some type, the seller reveals this type.

⁴⁹Informally, when $\tau'(s) = 0$, the seller effectively pools a "small positive-measure interval" of types near s with the infinitesimal-measure type $\tau(s)$. The infinitesimal measure of $\tau(s)$ is further spread thinly over the positive measure of types near s , thereby endowing each message $\{s, \tau(s)\}$ with the infinitesimal odds of $\tau(s)$ relative to s .

Lemma 5 shows that, for \mathcal{P} , one can construct a conjugate disclosure rule that delivers to the seller a payoff approximately equal to the seller's optimal payoff in \mathcal{P}^n when n is large. The lemma's statement uses the big-O notation, in which O stands for a function that satisfies $\limsup_{n \rightarrow \infty} |O(2^{-n}) / 2^{-n}| < \infty$.

Lemma 5. *Suppose that V^n is the seller's optimal payoff in the discrete disclosure problem \mathcal{P}^n . A conjugate disclosure rule exists that delivers to the seller payoff $V^n + O(2^{-n})$ in the continuous disclosure problem \mathcal{P} .*

Proof. See Appendix A.4. □

Lemma 5 does not rule out a discontinuity: the possibility that, in \mathcal{P} , the seller can improve upon the conjugate disclosure rule that is the limit of \mathcal{P}^n -optimal disclosure rules as n goes to infinity. Lemma 6 rules out this discontinuity by showing that the value in \mathcal{P} is no greater than the limit of the values in \mathcal{P}^n as n increases.

Lemma 6. *The continuous disclosure problem \mathcal{P} has a solution, which induces the value denoted by V^* . The discrete disclosure problems in the sequence $\{\mathcal{P}^n : n \geq 1\}$ have solutions, which induce the corresponding sequence of values denoted by $\{V^n\}$. Furthermore, $V^* \leq \liminf_{n \rightarrow \infty} V^n$.*

Proof. See Appendix A.5. □

Theorem 3 collects Lemmas 5 and 6, as well as earlier observations, to deliver our main result: a characterization of an optimal information-disclosure rule.

Theorem 3. *Under Conditions 1 and 2, the seller's disclosure problem \mathcal{P} has a solution that*

(i) *is a conjugate disclosure rule;*

(ii) *induces the follower's effort schedule a^* that is maximized at an s^* with $s^* \geq \theta^*$ and $a^*(s^*) \leq \alpha(\theta^*)$ (where α is the first-best effort schedule, and θ^* is its maximizer) and whose expectation is the same as that of the first-best effort: $\mathbb{E}[a^*(m)] = \mathbb{E}[\alpha(\theta_1)]$.*

Proof. Lemmas 5 and 6 imply the optimality of the conjugate disclosure rule and, with it, part (i) of the theorem.

For part (ii) of the theorem, note that Lemma 4's part (ii) and Lemma 6 imply that the induced effort schedule a^* is maximized at s^* , which satisfies $s^* \geq \theta^*$.

Because, $a^*(m) = \mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]$ for any message m (by (15)), $a^*(s^*) \leq \max_{\theta_1} \alpha(\theta_1) = \alpha(\theta^*)$, where the equality is by Theorem 1. Hence, $a^*(s^*) \leq \alpha(\theta^*)$.

Furthermore, $a^*(m) = \mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]$ implies $\mathbb{E}[a^*(m)] = \mathbb{E}[\mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]] = \mathbb{E}[\alpha(\theta_1)]$, where the last equality is by the Law of Iterated Expectations. \square

According to part (i) of Theorem 3, the optimality of the conjugate disclosure rule defies the pooling pattern common in many models of strategic disclosure, in which all types in a certain interval are pooled (e.g., Crawford and Sobel, 1982). In their footnote [11], RS conjecture that, with a continuum of prospects, one would be unable to dismiss interval pooling as nongeneric. By contrast, our model dismisses interval pooling for the prospect set that, while “nongeneric,” emerges in an economically interesting setting.⁵⁰

According to part (ii) of Theorem 3, the seller’s strategic information disclosure distorts the follower’s effort schedule by shifting its peak to the right of the peak of the first-best effort schedule. The overall, expected, payoff is the same as in the first best.

Corollary 3 shows that the allocation induced by the optimal disclosure rule can be implemented in a second-price auction with a tax. To define the tax schedule, let $m(\theta_1)$ denote the message optimally induced by type θ_1 . This message motivates the follower to take his uniquely optimal action, denoted by $a^*(m(\theta_1))$. The leader who bids b is taxed in the amount

$$T^*(b) \equiv \int_0^b (F(s | a^*(m(b))) - F(s | a^*(m(s)))) ds. \quad (27)$$

Corollary 3. *In an optimal mechanism, the seller*

1. *Asks the leader to submit a bid, denoted by b , and charges him tax $T^*(b)$, defined in (27).*
2. *Discloses to the follower the message prescribed by the optimal conjugate disclosure rule of Theorem 3 and asks him to submit a bid.*
3. *Allocates the item and assesses the payments (in addition to the tax $T^*(b)$) according to the rules of the second-price auction.*

In equilibrium, each bidder bids his type, the follower exerts the optimal effort, and the seller collects the optimal revenue.

⁵⁰Our prospect set is nongeneric if only because it is a one-dimensional curve in a two-dimensional space. On top of that, this curve is also convex, by Condition 2.

Proof. See Appendix A.2. □

Remark 2. To modify the auction so that the leader pays the tax if and only if he wins, replace $T^*(b)$ with $T^*(b) / F(b | a^*(m(b)))$.

By contrast to the first-best mechanism of Corollary 1, the optimal mechanism of Corollary 3 fails to fully disclose the leader's bid and specifies a different tax schedule. As in the first-best case, the leader's tax is increasing in his bid if the follower's effort is increasing in the leader's bid, and is decreasing otherwise. Thus, as in the first-best case, the leader's tax countervails the leader's motive to manipulate his bid so as to induce the follower to acquire more information.

5.4 The Seller's Optimal-Control Problem

By Theorem 3, the seller can restrict attention to disclosure rules induced by matching functions. As a result, the seller's problem can be represented as an optimal-control problem. This representation is useful for the numerical analysis of optimal disclosure.

The Optimal-Control Formulation

For the parsimony of exposition, we illustrate the optimal-control representation when the optimal matching function takes the form in case (i) of Definition 2. In this case, $s_* = 0$, and the seller's objective function (23) can be shown to be

$$\int_0^{s^*} \pi(s) a^*(s) g(s) ds + \int_0^{s^*} \pi(\tau(s)) a^*(s) g(\tau(s)) \beta(s) ds + \int_{[s^*, 1] \setminus \text{range}(\tau)} \alpha(s) \pi(s) ds, \quad (28)$$

where $\beta \equiv -\tau'$ denotes the derivative of a matching function τ , and $a^*(s)$ denotes the follower's optimal action when the leader's type is $s \in [0, s^*]$.⁵¹ The first integral in (28) is the seller's payoff from the prospects induced by the leader's types in $[0, s^*]$. The second integral is the seller's payoff from the types in $[s^*, 1]$ that the matching function pools with the types in $[0, s^*]$. The third integral is the seller's payoff from the types in $[s^*, 1]$ that are fully revealed.

⁵¹Equation (28) implicitly normalizes the set of messages to $[0, s^*] \cup ([s^*, 1] \setminus \text{range}(\tau))$, with any type $\theta_1 \in [0, s^*]$ generating message $m(\theta_1) = \theta_1$, any fully revealed type $\theta_1 \in (s^*, 1]$ (i.e., in $(s^*, 1]$ but not in the range of τ) generating message $m(\theta_1) = \theta_1$, and any pooled type $\theta_1 \in [s^*, 1]$ generating message $m(\theta_1) \in \tau^{-1}(\theta_1) \subset [0, s^*]$.

The first and last integrals in (28) copy the corresponding terms from the seller’s objective function (23). The second integral is derived using the change-of-variables formula, according to which, for any type $s \in (0, s^*)$ and a “small” $ds > 0$, the interval $(s, s + ds)$, whose probability is approximately $g(s) ds$, is pooled with interval $(\tau(s + ds), \tau(s))$, whose probability is approximately $g(\tau(s)) \beta(s) ds$.⁵² The follower’s equilibrium effort that enters (28) is computed from (15) by appealing to Bayes’s rule and the change-of-variables argument:

$$a^*(s) = \frac{g(s) \alpha(s) + g(\tau(s)) \beta(s) \alpha(\tau(s))}{g(s) + g(\tau(s)) \beta(s)}, \quad s \in [0, s^*]. \quad (29)$$

One can now formulate the seller’s optimal-control problem.

Definition 4. When $s_* = 0$, the **seller’s optimal-control problem** consists of maximizing (28) over s^* , a piecewise-continuous function $\beta : [0, s^*] \rightarrow \mathbb{R}_+$, and the implied piecewise-differentiable function $\tau : [0, s^*] \rightarrow [s^*, 1]$, which together induce a^* from (29), subject to $\tau(0) = 1$, $\tau(s^*) = s^*$, and, for almost all $s \in (0, s^*)$, $\tau'(s) = -\beta(s)$.

A Numerical Example

Here, we report an outcome of a search for a numerical solution to the seller’s problem in Example 1. To arrive at a solution, we apply Hamiltonian techniques to the optimal-control problem of Definition 4. Because the Hamiltonian analysis imposes the additional assumption of piecewise *continuous* differentiability of the matching function, and because we have been unable to show that the problem in Definition 4 is convex, the numerical “solution” we report is an informed guess,⁵³ which has been verified to improve upon full disclosure and non-disclosure.⁵⁴ The solution also satisfies the leader’s monotonicity condition; that is, $F(\theta_1 | a^*(\theta_1))$ is nondecreasing in θ_1 .

Figure 6 plots selected pooling links in an optimal mechanism. These links are derived from the optimal matching function that solves the problem in Definition 4. Each pooling link, for some $\theta_1 < s^*$, connects the prospects induced by θ_1 and by $\tau(\theta_1)$, and induces the follower’s effort

⁵²Indeed, the Taylor expansion implies $G(\tau(s) + ds) \approx G(\tau(s)) + G'(\tau(s)) \tau'(s) ds$.

⁵³An analogous caveat applies to this section’s references to “optimality.” In the remainder, the quotation marks are suppressed.

⁵⁴So we have not identified a minimum. Indeed, we conjecture we have identified the maximum.

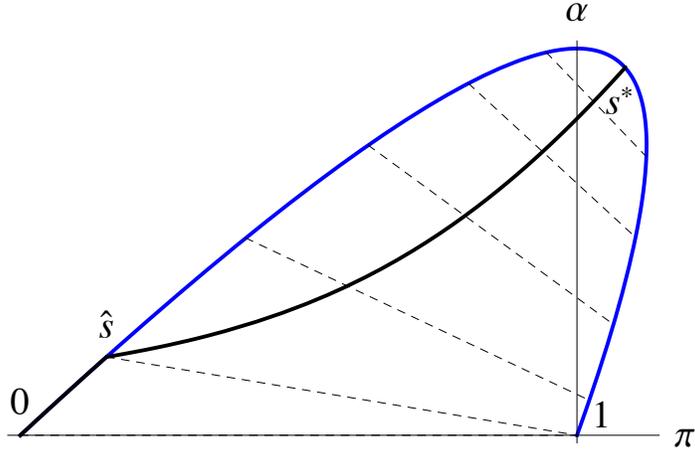


Figure 6: An optimal disclosure rule in Example 1. The solid blue arc is the prospect set, with selected prospects labelled by the leader’s types that induce them (0 , $\hat{s} = 0.054$, $s^* = 0.61$, and 1). Each point on the solid black curve (the optimal-prospect path) is a tuple containing the follower’s optimal action and the seller’s associated expected marginal benefit from that action. Each tuple is induced by pooling the prospects at the endpoints of the dashed link that passes through that tuple. The leader’s types that are revealed are in $[0, \hat{s}]$ and lie at the intersection of the solid black curve and the prospect set.

$a^*(\theta_1)$, which is the ordinate of the intersection point of the solid black curve and the pooling link. The corresponding abscissa is the seller’s expected marginal benefit from the follower’s action when the message corresponding to that pooling link has been sent. The collection of the induced efforts and the corresponding expected benefits (i.e., the solid black curve) is an **optimal-prospect path**.

Any optimal-prospect path is nondecreasing, which is a necessary condition for optimality. If the path had a strictly decreasing segment, then by Lemma [1] in RS, it would be optimal to pool under a single message all messages that induced that segment. Whenever a prospect is revealed, this prospect belongs both to the optimal-prospect path and the prospect set.

Figure 7 plots the follower’s effort schedule. Constrained by the Bayes plausibility, the seller designs the disclosure rule so as to better align the follower’s effort, a^* , with the marginal benefit from effort, π . Consistent with Theorem 3, doing so involves inducing a “rightward shift” in the follower’s effort schedule relative to the first-best effort schedule.

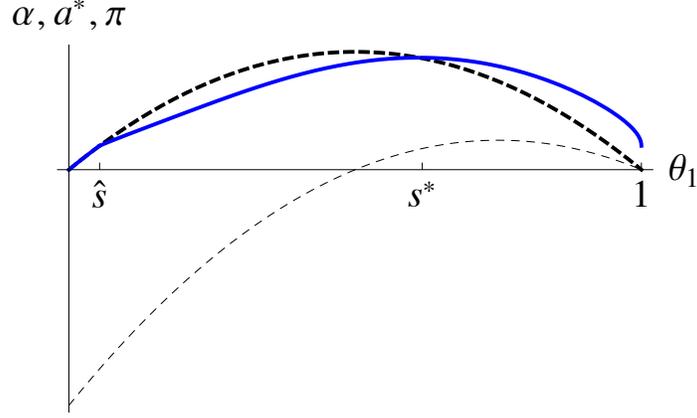


Figure 7: The solid curve is the optimal effort, a^* , in Example 1. The thick dashed curve is the first-best effort, α . The thin dashed curve is the seller's marginal benefit from the follower's effort, π . The seller uses strategic disclosure to better (i.e., assortatively) align a^* with π . The areas under the solid thick and the dashed thick curves coincide; that is, $\mathbb{E}[a^*(m)] = \mathbb{E}[\alpha(\theta_1)]$.

6 Two More Benchmarks for the Seller's Disclosure Problem

To gain additional intuition for the seller's choice of the optimal disclosure rule (derived in Theorem 3), it is instructive to consider two more benchmarks, in addition to the first-best benchmark. These two benchmarks are summarized in Theorem 4.

Theorem 4. (i) *If the seller observed the leader's valuation, any disclosure rule would be optimal. In particular, fully disclosing the leader's valuation would be optimal, in which case the first-best outcome would be achieved.*

(ii) *If the seller could choose for the follower any effort from an interval $[0, \bar{a}]$ with some $\bar{a} < 1$,⁵⁵ any disclosure rule would be optimal, and the seller would choose the bang-bang effort schedule $\bar{a}\mathbf{1}_{\{\theta_1 \geq \theta^*\}}$, for all $\theta_1 \in \Theta_1$.*

Proof. For part (i), if the seller observes the leader's valuations, he can extract the leader's entire information rent by charging him θ_1 if the follower refuses to buy at θ_1 . The leader's virtual valuation is replaced by his valuation, θ_1 , and the implied virtual surplus (22),

$$\int_{\Theta_1} \mathbb{E}_{m|\theta_1}[\theta_1] dG(\theta_1),$$

is independent of the follower's effort.

⁵⁵The requirement $\bar{a} < 1$ ensures that the distribution of the follower's type has the full support $[0, 1]$, thereby justifying the envelope argument on which the analysis relies.

For part (ii), if the seller can directly and costlessly choose the follower’s effort $a \in [0, \bar{a}]$, he will choose it so as to minimize pointwise the probability that the leader is the ex-post efficient recipient of the item, by inspection of (22). Formally, for $\theta_1 < \theta^*$, the seller sets $a = 0$, and for $\theta_1 \geq \theta^*$, the seller sets $a = \bar{a}$.⁵⁶ (The argument relies on the allocation rule being ex-post efficient, just as (22) does.) When the follower has no control over his action, any disclosure rule is optimal because truthful reporting can be made a dominant strategy for the follower (at no additional cost to the seller) by making a take-it-or-leave-it offer to the follower at price θ_1 . \square

Part (i) of Theorem 4 indicates that minimizing the information rent left to the leader is the sole rationale for the seller’s strategic disclosure. Part (ii) shows that the follower’s information-acquisition effort is an effective instrument for doing so—at least as long as this effort can be controlled directly. Figure 2 juxtaposes the seller-controlled effort schedule of part (ii) of Theorem 4 with the first-best effort schedule of Theorem 1.

Part (ii) of Theorem 4 exposes the force that, in a subdued form, influences the design of an optimal disclosure rule. By inspection of (22), the seller seeks to minimize the expected weighted probability with which the leader buys (where the weights are the leader’s information rents). To minimize the leader’s expected weighted probability of winning, the seller encourages the follower to become “stronger,” thereby intensifying the competition that the leader faces. This encouragement is accomplished directly in part (ii) of Theorem 4 and indirectly, through the strategic choice of a disclosure rule, in an optimal mechanism. Whether a better or a worse informed follower is “stronger” depends on the leader’s valuation. When the leader’s valuation is high, a better informed follower has a more dispersed distribution of his types and thus stands a better chance of outbidding the leader. The opposite is true when the leader’s valuation is low.

7 Discussion of Two Critical Assumptions

Two assumptions are at the heart of our analysis. First, we assume that the identified solution to the relaxed problem, which ignores the leader’s monotonicity constraint, also solves the full problem, which retains this constraint. Second, we fix the allocation rule to be ex-post efficient,

⁵⁶The leader’s monotonicity constraint holds; $F(\theta_1 | \mathbf{1}_{\{\theta_1 \geq \theta^*\}})$ is weakly increasing in θ_1 . Indeed, $F(\theta_1 | 0)$ and $F(\theta_1 | 1)$ are both weakly increasing in θ_1 , and $\lim_{\theta_1 \uparrow \theta^*} F(\theta_1 | 0) \leq \lim_{\theta_1 \uparrow \theta^*} F(\theta_1 | 1) \leq F(\theta^* | 1)$.

instead of maximizing also with respect to the allocation rule.

The Monotonicity Assumption

The analysis has been performed under the hypothesis that the first-order approach characterizes the solution to the seller's problem. The second-order conditions that verify this hypothesis are two monotonicity conditions, each of which asserts that the corresponding bidder's interim probability of winning is nondecreasing in his type. We do not provide sufficient conditions on the model's primitives to guarantee that the leader's monotonicity constraint holds.⁵⁷ The source of our difficulties is that the optimal disclosure rule is not a simple function of the model's primitives. The disclosure rule, in turn, affects the follower's effort schedule a^* and thus affects the leader's probability of winning.

A type- θ_1 leader wins if his type exceeds the follower's type, which occurs with probability

$$F(\theta_1 | a^*(\theta_1)) = a^*(\theta_1) F_H(\theta_1) + (1 - a^*(\theta_1)) F_L(\theta_1),$$

which depends on the leader's type θ_1 both directly and indirectly, through the follower's action. The leader's monotonicity condition requires $F(\theta_1 | a^*(\theta_1))$ to be weakly increasing in θ_1 . If a^* and thus the probability distribution of θ_2 had been fixed, the leader's monotonicity condition would have been satisfied automatically because F_H and F_L are c.d.f.s and hence weakly increasing. The distribution of θ_2 is not fixed, however. Instead, it depends on θ_1 through a^* . This dependence threatens monotonicity (albeit need not overturn it, as one can ascertain in examples).

To see the threat to monotonicity, suppose that $\theta_1 > \theta^*$, so that $F_H(\theta_1) < F_L(\theta_1)$. In addition, suppose that $\theta_1 < s^*$, so that, in the neighborhood of θ_1 , a^* is increasing in the leader's type. In this case, a small increase in θ_1 , while (weakly) increasing the values of both F_H and F_L , shifts the weight in F towards F_H , the smaller of the two constituent c.d.f.s. So the overall direction of change in $F(\theta_1 | a^*(\theta_1))$ is ambiguous unless a^* is known and $F(\theta_1 | a^*(\theta_1))$ can be evaluated precisely.⁵⁸

Intuitively, when $\theta_1 \in (\theta^*, s^*)$, the seller attempts to take advantage of the leader's relatively high valuation by increasing competition, which is accomplished by asking the follower to ac-

⁵⁷The analogous monotonicity condition for the follower is guaranteed to hold.

⁵⁸In principle, it would suffice to bound the steepness of a^* , but doing so seems as hard as solving for a^* .

quire more information. Even though a more informed follower is not stronger in the first-order stochastic-dominance sense, he is stronger in the sense of having a higher chance of an extremely high type realization, which is what is needed to outbid a high-type leader. Thus, a stronger leader with a type in (θ^*, s^*) has a stronger follower and so *a priori* need not be more likely to win.

In absence of sufficient conditions for monotonicity, our approach is to recommend ex-post verification of monotonicity in examples.⁵⁹ Ex-post verification of monotonicity is a common practice in the dynamic public finance literature, in which optimal contracts are notoriously hard to characterize explicitly.⁶⁰

Towards More General Allocation Rules

The paper focuses on the problem in which the allocation rule is ex-post efficient. This focus is justified by the features of natural economic applications (e.g., government divestment), by the interest in isolating the distortions due to the disclosure rule, and by tractability. Our approach to characterizing an optimal disclosure rule can be extended to any allocation rule that is fixed. However, the joint determination of an optimal allocation rule and an optimal disclosure rule is a major challenge. Below, we describe how to generalize our approach to the fixed allocation rules that are not ex-post efficient and then explain why this generalization does not deliver a recipe for solving the more general problem in which both the disclosure rule and the allocation rule are optimally chosen by the seller.

Suppose the seller would like to find an optimal disclosure rule when the allocation rule is fixed at

$$x_i(\theta_1, \theta_2) = \mathbf{1}_{\{\theta_i > \phi_i(\theta_{-i})\}}, \quad i = 1, 2, \quad (30)$$

for some weakly increasing $\phi_i : \Theta_{-i} \rightarrow [0, 1]$, $i = 1, 2$, such that $x_1(\theta_1, \theta_2) + x_2(\theta_1, \theta_2) \leq 1$. The rule in (30) reduces to the ex-post efficient allocation rule when ϕ_1 and ϕ_2 are identity functions. For another example, the rule in (30) is ex-post efficient save for reserve prices if $\phi_i(\theta_{-i}) = \max\{r_i, \theta_{-i}\}$, $i = 1, 2$, where r_1 and r_2 are some positive reserve types.

⁵⁹In the numerical examples that we have explored, monotonicity holds.

⁶⁰Among the papers advocating ex-post verification of monotonicity are [Farhi and Werning \(2013\)](#), [Stantcheva \(2015\)](#), [Golosov et al. \(2015\)](#), and [Kapicka \(2013\)](#).

Because ϕ_1 is increasing only weakly, it need not be invertible in the usual sense. Define the generalized inverse of ϕ_1 to be

$$\phi_1^{-1}(z) \equiv \sup \{ \theta_2 \in \Theta_2 \mid \phi_1(\theta_2) = z \}, \quad z \in [0, 1].$$

The seller's virtual surplus is

$$\mathbb{E}_{m, \theta_1, \theta_2} \left[\left(\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} \right) \mathbf{1}_{\{\theta_1 > \phi_1(\theta_2)\}} + \left(\theta_2 - \frac{1 - F(\theta_2 \mid a^*(m \mid \phi_2))}{f(\theta_2 \mid a^*(m \mid \phi_2))} \right) \mathbf{1}_{\{\theta_2 \geq \phi_2(\theta_1)\}} \right], \quad (31)$$

where $a^*(m \mid \phi_2) = \mathbb{E}_{\theta_1 \mid m} [\alpha(\phi_2(\theta_1))]$ generalizes the follower's optimal action in (15).⁶¹ The virtual surplus in (31) differs from the virtual surplus in (21) only in the indicator functions, which reflect the more general allocation rule.

Integrating θ_2 out of the display above gives

$$\mathbb{E}_{m, \theta_1} \left[\left(\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} \right) F(\phi_1^{-1}(\theta_1) \mid a^*(m)) + \phi_2(\theta_1) (1 - F(\phi_2(\theta_1) \mid a^*(m))) \right],$$

a counterpart of (22). Substituting the functional form for F , and ignoring the terms that are independent of the seller's choice of the disclosure rule gives the seller's objective function

$$\mathbb{E}_{m, \theta_1} [\tilde{\pi}(\theta_1 \mid \phi_1, \phi_2) a^*(m \mid \phi_2)],$$

where

$$\begin{aligned} \tilde{\pi}(\theta_1 \mid \phi_1, \phi_2) \equiv & \left(\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} \right) \left(F_H(\phi_1^{-1}(\theta_1)) - F_L(\phi_1^{-1}(\theta_1)) \right) \\ & - \phi_2(\theta_1) \left(F_H(\phi_2(\theta_1)) - F_L(\phi_2(\theta_1)) \right) \end{aligned} \quad (32)$$

generalizes the seller's marginal benefit in (24). The seller's objective function is further transformed using the Law of Iterated Expectations and the expression for $a^*(m \mid \phi_2)$ to yield the product form

$$\mathbb{E}_{\theta_1} [\mathbb{E}_{\theta_1 \mid m} [\tilde{\pi}(\theta_1 \mid \phi_1, \phi_2)] \mathbb{E}_{\theta_1 \mid m} [\alpha(\phi_2(\theta_1))]], \quad (33)$$

⁶¹Fixing the allocation rule, when the leader's type is θ_1 , the follower's first-best effort is $\alpha(\phi_2(\theta_1))$.

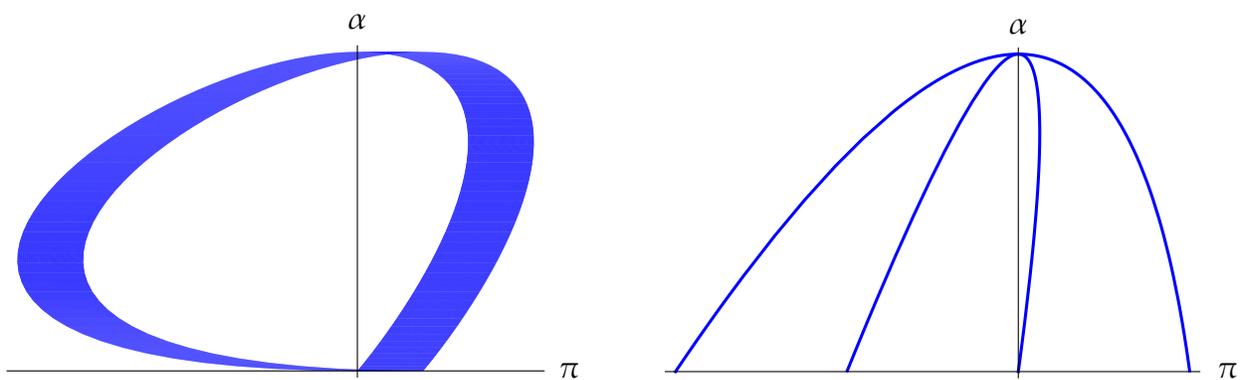
which generalizes (25). One can now seek an optimal disclosure rule following the techniques developed in this paper.⁶²

As soon as one attempts to optimize over both the allocation rule (as parametrized by ϕ_1 and ϕ_2) and the disclosure rule, the analytical convenience of the product structure of (33) vanishes; the techniques developed in this paper become inapplicable. Naively, one might have been tempted to address this more general problem in two steps: (i) for each disclosure rule, pointwise maximize the virtual surplus (31) over ϕ_1 and ϕ_2 so as to allocate the item to the bidder with the highest nonnegative virtual valuation, (ii) pick that disclosure rule which gives the highest value of the virtual surplus in step (i). The problem with this naive approach is that the virtual valuations in step (i) depend on the allocation rule itself, through ϕ_1 and ϕ_2 . The allocation rule that insists on selecting the highest nonnegative virtual valuation may be an allocation rule that induces particularly low virtual valuations, thereby casting doubt on the optimality of the naive approach. In this case, the allocation rule cannot be pointwise maximized out, as is common in mechanism design, and the seller's problem cannot be reduced to the problem studied in this paper. Instead, for a given disclosure policy, the optimal ϕ_1 and ϕ_2 would solve an appropriately formulated optimal control problem. This problem's analysis is beyond the scope of this paper.

8 Concluding Remarks

This paper's primary purpose is to reinterpret and extend the techniques of the Bayesian persuasion model of RS to study, in a simple auction model with information acquisition, the distortions that the seller's strategic bid-disclosure introduces into an otherwise efficient auction. The mapping of the seller's optimal-auction problem into the optimal-disclosure problem of RS relies on three assumptions: (i) the seller must choose an ex-post efficient allocation rule, (ii) the follower's information-acquisition effort is the probability with which he gains access to a more precise signal about his underlying valuation, and (iii) the probability with which the leader wins is weakly increasing in his type even if the seller does not explicitly heed this monotonicity constraint when designing the auction. With these assumptions, the agents truth-telling, obedience, and participation constraints enable the seller to rewrite his objective function in a product form, as the

⁶²These techniques apply provided the prospect set $\{(\bar{\pi}(\theta_1 | \phi_1, \phi_2), \alpha(\phi_2(\theta_1))) | \theta_1 \in \Theta_1\}$ is convex and the requisite monotonicity condition can be verified to hold.



(a) The prospect set is “thick.” It has been obtained by augmenting the prospect set in Figure 3b by assuming that the seller’s private benefit from awarding the item to the leader varies within some range independently of the leader’s type.

(b) The prospect set is not a convex curve. It has been obtained by assuming the uniform G , and F_L and F_H as specified in Example 3 (in violation of Condition 1).

Figure 8: More than two prospects may lie on the same line. The results of RS cannot be applied.

expectation of the follower’s expected effort conditional on the seller’s message times the seller’s conditional expected marginal benefit from the follower’s effort. The three listed assumptions ensure that the seller’s marginal benefit is independent of—and thus his expected revenue is linear in—the follower’s effort. The seller’s marginal benefit in our model corresponds to the search engine’s benefit (in revenues from advertisers) from a consumer’s click on an online advertisement link in the model of RS. The follower’s information-acquisition effort corresponds to the probability with which the consumer clicks.

RS’s results can be mapped back into our setting because, in spite of our assumption of the continuum of prospects (which are the marginal revenue and effort pairs), we preserve RS’s critical feature: no three prospects lie on the same line. This feature is preserved because (i) the seller’s private information, the leader’s elicited type, is “one-dimensional,” and (ii) the follower’s information-acquisition technology satisfies the rotation order (Condition 1). To see the role of (i), consider a minimal departure from the one-dimensionality: the seller is privately informed about his private benefit from allocating the item to the leader.⁶³ This private benefit affects the seller’s marginal benefit without affecting the follower’s effort (we still require that the highest-type bidder win). As a result, the prospect set is “thick,” as in Figure 3b; uncountably many prospects

⁶³In the case of government procurement, this private benefit could reflect the administrative cost of switching the contractor or a kickback from the incumbent, none of which is allowed to be reflected in the choice of the allocation rule, which is still required to allocate the contract to the highest-type bidder.

may lie on the same line. To see the role of (ii), consider Example 3, which violates the rotation order. The implied prospect set, in Figure 8b, fails to be a convex curve because the follower's first best-effort is not single-peaked in the leader's type.

The paper finds that the profit-maximizing seller induces the follower to acquire inefficiently little information when the leader's valuation is low, and to acquire inefficiently much information otherwise. Carefully pooling the pairs of the leader's extreme bids under the same messages accomplishes this distortion. In practice, information disclosure typically occurs in private pre-sale negotiations and hence is largely unobservable by outsiders. This lack of observability limits the scope for testing the model's predictions and makes the paper's focus normative.

Nevertheless, one can tentatively ask the positive question of whether anything resembling the derived disclosure pattern is ever observed in practice. The auctions of rail passenger service franchises in the U.K. fit the model's assumptions and may be interpreted to feature the disclosure similar to the optimal disclosure derived in the paper. In particular, in the U.K., the rail passenger services are franchised for a limited time to train operating companies. An auction determines the award of the franchise. An incumbent and a potential entrant bid for the right to run passenger services in a certain region. The incumbent (the leader) is likely to know his valuation for running the services, whereas the entrant (the follower) can choose how much information to acquire about his valuation. In practice (and also in the model), the incumbent's bid is distinct from his payment because the incumbent may request an operating subsidy from the government. Any details about this request that leak to the entrant are a noisy signal about the incumbent's valuation; this signal guides the entrant's information acquisition. For instance, the entrant can interpret the incumbent's request for a large subsidy in two ways: (i) the incumbent is weak and needs help to continue operating, or (ii) the incumbent is strong and plans to invest in new trains and services. An incumbent who does not request a subsidy can be interpreted as mediocre. This interpretation, which pools the extremes, resembles the conjugate disclosure rule prescribed by the model.

A Appendix: Omitted Proofs

A.1 Proof of Corollary 1

For the follower, it is a weakly dominant strategy to bid his type in the second-price auction. If he bids his type, he also finds it optimal to exert the first-best effort; his expected payoff in

the mechanism described in the corollary has been constructed to coincide with the planner's maximand in the surplus-maximization problem (4). The follower's payoff is nonnegative because he can obtain a nonnegative payoff by bidding in the second-price auction without having exerted any effort.

The leader chooses his bid, b , to maximize

$$\int_{\Theta_2} \mathbf{1}_{\{b > \theta_2\}} (\theta_1 - \theta_2) dF(\theta_2 | \alpha(b)) - T(b).$$

Integration by parts and the substitution of T from (8) transforms the display above into

$$\int_0^{\theta_1} F(s | \alpha(s)) ds + \int_{\theta_1}^b (F(s | \alpha(s)) - F(b | \alpha(b))) ds,$$

which is maximized at $b = \theta_1$ because $F(s | \alpha(s))$ is increasing in s , as we now show.

That $F(s | \alpha(s))$ is increasing in s can be seen by letting $s' > s$ and writing

$$F(s' | \alpha(s')) - F(s | \alpha(s)) = [F(s' | \alpha(s')) - F(s | \alpha(s'))] + \int_{\alpha(s)}^{\alpha(s')} \frac{\partial F(s | a)}{\partial a} da.$$

In the display above, the bracketed term is nonnegative because F is a c.d.f. and $s' > s$. To see that the integral in the display above is positive, we consider three cases: (i) if $s < s' \leq \theta^*$, then $\alpha(s) < \alpha(s')$ (Theorem 1) and $\partial F(s | a) / \partial a > 0$ (Condition 1), and so the integral is positive, (ii) if $\theta^* \leq s < s'$, then $\alpha(s) > \alpha(s')$ (Theorem 1) and $\partial F(s | a) / \partial a < 0$ (Condition 1), and so the integral is positive, and (iii) if $s < \theta^* < s'$, then considering the change in the leader's type from s to θ^* and applying case (i) and then considering the change in the leader's type from θ^* to s' and applying case (ii) delivers the positivity of the integral.

When $b = \theta_1$, the leader's expected payoff is $\int_0^{\theta_1} F(s | \alpha(s)) ds$ and hence nonnegative.

A.2 Proof of Corollary 3

For the follower, it is a weakly dominant strategy to bid his valuation in the second-price auction. The follower participates because he can obtain a nonnegative payoff by exerting no effort and then bidding in the second-price auction.

The leader chooses his bid by solving

$$\max_b \left[\int_{\Theta_2} \mathbf{1}_{\{b > \theta_2\}} (\theta_1 - \theta_2) dF(\theta_2 | a^*(m(b))) - T^*(b) \right],$$

where $a^*(m(b))$ is the follower's optimal action conditional on the message $m(b)$, which the seller sends when the leader bids b . Integration by parts and the substitution of T^* from (27) transforms

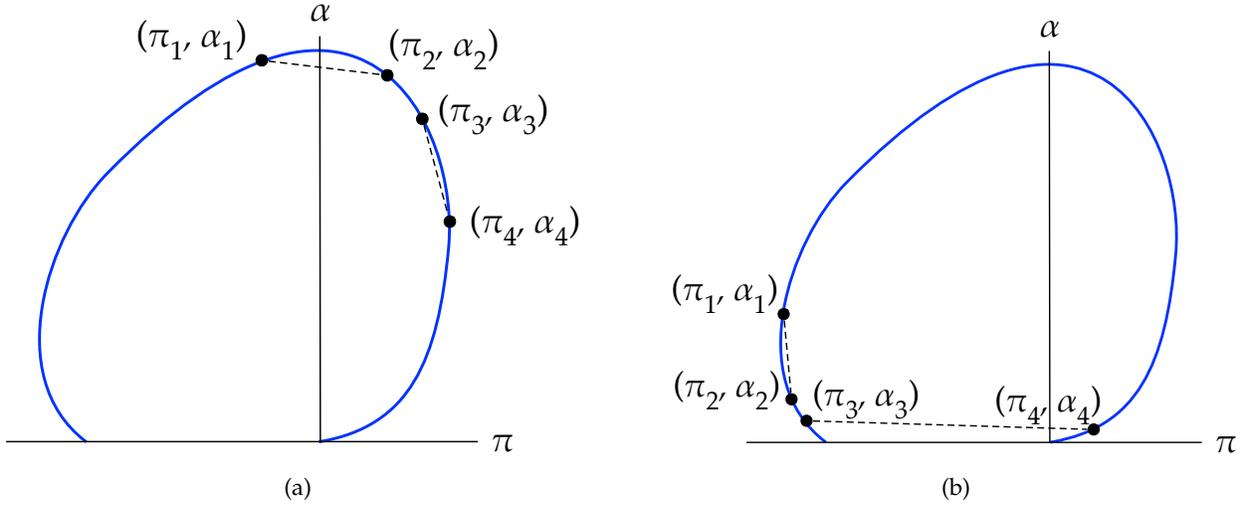


Figure 9: Contradiction hypotheses for Step 1 in the proof of Lemma 4. Each dashed link denotes a pair of prospects that are pooled under the same message. In neither panel can the two links be crossed by an upward-sloping line.

the maximand above into

$$\int_0^{\theta_1} F(s | a^*(m(s))) ds + \int_{\theta_1}^b (F(s | a^*(m(s))) - F(b | a^*(m(b)))) ds,$$

which is maximized at $b = \theta_1$ because $F(s | a^*(m(s)))$ is weakly increasing in s , which is the monotonicity condition required by the incentive compatibility of truthful reporting and implied by the hypothesis that the mechanism is optimal.

A.3 Proof of Lemma 4

The proof is constructive and proceeds in three steps. Step 1 rules out the pooling patterns depicted in both panels of Figure 9. Step 2 combines Facts 1–5 with Step 1 to construct an $s_* \in [0, \underline{\theta}]$ and an $s^* \in [\underline{\theta}^*, \bar{\theta}]$ satisfying part (i) of the lemma. Step 3 establishes part (ii).

Step 1: Take any pooling link that has a northwest prospect, denoted by (π_3, α_3) .⁶⁴ Then, optimality rules out the existence of a link that pools two prospects, say (π_1, α_1) and (π_2, α_2) , such that each of these prospects lies to the northwest of (π_3, α_3) . (See both panels of Figure 9.)

Prospects are optimally pooled so that the pooling links are nonincreasing (Fact 2) and never intersect (Fact 4). Hence, to prove the claim in Step 1, it suffices to show that the pooling patterns depicted in Figure 9 are never optimal. A single argument rules out both patterns.

By contradiction, suppose that one can pick four prospects $\{(\pi_i, \alpha_i)\}_{i=1,2,3,4}$ such that prospects (π_1, α_1) and (π_2, α_2) are optimally pooled under some message, say, m ; prospects (π_3, α_3) and

⁶⁴We define the cardinal directions in the (π, α) -space in the obvious manner. For instance, a point (π_3, α_3) is **northwest** of point (π_4, α_4) if $\pi_3 < \pi_4$ and $\alpha_3 > \alpha_4$.

(π_4, α_4) are optimally pooled under another message, say m' ; and either (a) $\pi_4 \geq \pi_3 \geq \pi_2 > \pi_1$ and $\alpha_1 \geq \alpha_2 \geq \alpha_3 > \alpha_4$ or (b) $\pi_4 > \pi_3 \geq \pi_2 \geq \pi_1$ and $\alpha_1 > \alpha_2 \geq \alpha_3 \geq \alpha_4$ holds.⁶⁵ For each $i = 1, 2, 3, 4$, let p_i denote the joint probability that (i) type x_i —which is assumed to invoke prospect (π_i, α_i) —is realized, and (ii) type x_i induces either message m or m' (whichever is appropriate).

The seller's expected gain from using the distinct messages m and m' (as in Figure 9) relative to pooling all four prospects under a single message is

$$\Delta \equiv \frac{(p_1\pi_1 + p_2\pi_2)(p_1\alpha_1 + p_2\alpha_2)}{p_1 + p_2} + \frac{(p_3\pi_3 + p_4\pi_4)(p_3\alpha_3 + p_4\alpha_4)}{p_3 + p_4} - \frac{(p_1\pi_1 + p_2\pi_2 + p_3\pi_3 + p_4\pi_4)(p_1\alpha_1 + p_2\alpha_2 + p_3\alpha_3 + p_4\alpha_4)}{p_1 + p_2 + p_3 + p_4},$$

which can be rearranged to give

$$\Delta = \frac{(p_1 + p_2)(p_3 + p_4)}{p_1 + p_2 + p_3 + p_4} \left(\frac{p_3\pi_3 + p_4\pi_4}{p_3 + p_4} - \frac{p_1\pi_1 + p_2\pi_2}{p_1 + p_2} \right) \left(\frac{p_3\alpha_3 + p_4\alpha_4}{p_3 + p_4} - \frac{p_1\alpha_1 + p_2\alpha_2}{p_1 + p_2} \right) < 0, \quad (\text{A.1})$$

where the inequality follows either from $\pi_4 \geq \pi_3 \geq \pi_2 > \pi_1$ and $\alpha_1 \geq \alpha_2 \geq \alpha_3 > \alpha_4$ or from $\pi_4 > \pi_3 \geq \pi_2 \geq \pi_1$ and $\alpha_1 > \alpha_2 \geq \alpha_3 \geq \alpha_4$. Because $\Delta < 0$, the pooling patterns in the both panels of Figure 9 are suboptimal, and the claim in Step 1 follows.

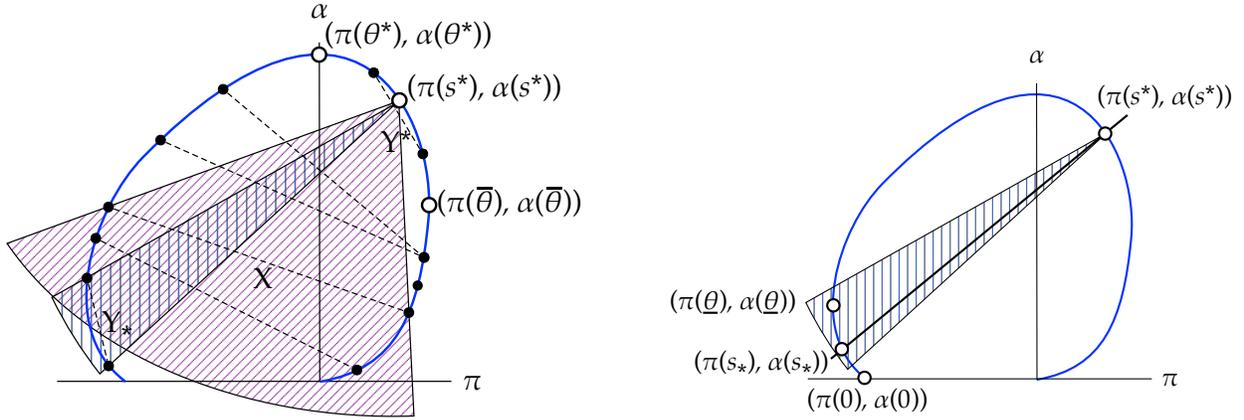
Step 2: *There exist an $s_* \in [0, \underline{\theta}]$ and an $s^* \in [\theta^*, \bar{\theta}]$ that satisfy part (i) of the lemma.*

Here we describe a procedure for constructing the sought s_* and s^* . This procedure uses a strict, complete, and transitive “smaller-than” order on optimal pooling links. We denote this order by \prec . To define \prec , take an arbitrary pooling link and call it X . The unique line passing through X is X 's hyperplane that splits the (π, α) -space in two half-spaces. The upper half-space of X is the closed half-space comprising the points each of which is weakly greater in the (Cartesian) product order than some point on the X 's hyperplane. X is said to be smaller than some other pooling link Y (and Y is greater than X)—denoted by $X \prec Y$ —if Y lies in the upper half-space of X . Thus defined \prec is complete because, by Facts 1–5, optimal pooling links never intersect. The order is strict because two distinct links cannot share a hyperplane. Finally, it is immediate that the order is transitive.

Because Γ^n is finite, the number of pooling links is finite, and so there exists the unique \prec -maximal pooling link, which we denote by Y^* . Because the pooling links are nonincreasing, the arc of Γ that lies in the upper half-space of Y^* contains at least one prospect induced by a type in $[\theta^*, \bar{\theta}]$. Define s^* to be an arbitrary such type.

We now turn to the construction of s_* . Draw a sequence of pairs of secants of Γ . Each secant in the pair passes through the prospect $(\pi(s^*), \alpha(s^*))$ and either endpoint of a pooling link. Each secant pair delimits a cone in the (π, α) -space as depicted in Figure 10a. By Step 1, and because Γ is a convex curve, \prec induces an inclusion order on the cones generated by the pooling links.

⁶⁵Cases (a) and (b) differ in the placement of the strict and weak inequalities. Case (a) prevails if the leader's types x_1, x_2, x_3 , and x_4 in Θ_1^n satisfy $0 < x_1 < x_2 \leq x_3 < x_4 < 1$, $x_2 > \theta^*$, and $x_3 < \bar{\theta}$. Case (b) prevails if the leader's types satisfy either $\theta^* > x_1 > x_2 \geq x_3 > x_4 > 0$ or $x_4 > \theta^* > x_1 > x_2 \geq x_3 > 0$ and $x_2 < \bar{\theta}$.



(a) Type $s^* \in [\theta^*, \bar{\theta}]$ is chosen to induce a prospect, $(\pi(s^*), \alpha(s^*)) \in \Gamma$, in the upper half-space of the \prec -maximal pooling link, Y^* . The outer shaded cone originates at prospect $(\pi(s^*), \alpha(s^*))$ and straddles some pooling link X . The inner shaded cone originates at prospect $(\pi(s^*), \alpha(s^*))$ and straddles the \prec -minimal pooling link, Y_* .

(b) Type $s_* \in [0, \underline{\theta}]$ is chosen to induce a prospect, $(\pi(s_*), \alpha(s_*))$, in the lower half-space of the \prec -minimal pooling link (not shown). The shaded cone straddles the \prec -minimal pooling link and is also the (nonempty) intersection of all cones that straddle pooling links. This intersection property ensures that the secant that passes through prospects $(\pi(s_*), \pi(s_*))$ and $(\pi(s^*), \pi(s^*))$ traverses every pooling link (not shown).

Figure 10: The construction of s_* and s^* in Step 2 of the proof of Lemma 4. The solid dots mark prospects in the discretized prospect set Γ^n ; the empty dots mark “critical” prospects in the continuous prospect set Γ .

In particular, the cones associated with \prec -smaller pooling links are smaller in the inclusion sense. As a result, the intersection of all the cones (depicted in Figure 10b) is nonempty and is the cone induced by the \prec -minimal pooling link, denoted by Y_* .⁶⁶ Consequently, the intersection of all cones contains the arc of Γ that lies in the lower half-space of Y_* . Because Y_* is nonincreasing, this arc contains at least one prospect induced by some type in $[0, \underline{\theta}]$. Define s_* to be an arbitrary such type.

The line that passes through the prospects induced by s^* and s_* constructed above (Figure 10b) is nondecreasing and traverses all pooling links. This line partitions the prospects in those induced by types in $[s_*, s^*] \cap \Theta_1^n$ and those induced by types in $([0, s_*] \cup [s^*, 1]) \cap \Theta_1^n$. By construction of s^* and s_* , every prospect in Γ^n is either pooled with a prospect in the other element of this partition or is fully revealed. Hence, part (i) of the lemma follows.

Step 3: *The optimal effort is single-peaked and, if $s^* \in \Theta_1^n$, is maximal at type s^* . (If $s^* \notin \Theta_1^n$, the optimal effort is maximal “close” to s^* , either at type $\max \{[0, s^*] \cap \Theta_1^n\}$ or at type $\min \{[s^*, 1] \cap \Theta_1^n\}$).*

Take any four leader’s types x_1, x_2, x_3 , and x_4 in Θ_1^n such that $\{x_1, x_2\}$ are pooled under some message, $\{x_3, x_4\}$ are pooled under some other message, and either $x_1 \neq x_2$ or $x_3 \neq x_4$, or both. Relabel the pairs of types so that the link connecting the prospects induced by points $\{x_1, x_2\}$ is

⁶⁶Just as Y^*, Y_* exists because Γ^n , and thus the number of pooling links, is finite.

\prec -smaller than the link connecting the prospects induced by $\{x_3, x_4\}$.⁶⁷ For the described pooling to be optimal, it must be, in particular, that the seller does not gain from pooling $\{x_1, x_2, x_3, x_4\}$ all under the same message. That is, the inequality in (A.1) must be reversed to yield $\Delta \geq 0$. The only way to satisfy $\Delta \geq 0$ is to have⁶⁸

$$\frac{p_3\alpha_3 + p_4\alpha_4}{p_3 + p_4} \geq \frac{p_1\alpha_1 + p_2\alpha_2}{p_1 + p_2} \quad \text{and} \quad \frac{p_3\pi_3 + p_4\pi_4}{p_3 + p_4} \geq \frac{p_1\pi_1 + p_2\pi_2}{p_1 + p_2}.$$

The first inequality in the above display implies in particular that the message corresponding to a \prec -larger pooling link (or a fully revealed prospect) induces a weakly larger equilibrium effort. Equivalently, because of the orientation of pooling links relative to s^* and s_* reported in part (i), the induced effort increases as the leader's type increases away from s_* and towards s^* —as long as either $x_1 \neq x_2$ or $x_3 \neq x_4$ —corroborating part (ii) of the lemma.

It remains to show that, also when two fully revealed prospects are compared (instead of two links or a link and a fully revealed prospect), the larger of the two leader's types in $[s_*, s^*]$ induces the higher effort. Once again appealing to $\Delta \geq 0$, now with $x_1 = x_2$ and $x_3 = x_4$, conclude that, for any two arbitrary fully revealed prospects (π_1, α_1) and (π_3, α_3) ,

$$(\pi_3 - \pi_1)(\alpha_3 - \alpha_1) \geq 0.$$

That is, because fully revealed, the two prospects lie on the upward-sloping segments of the arc $\{(\pi(s), \alpha(s)) \in \Gamma \mid s \in [s_*, s^*]\}$. On this segment, it is indeed the case that a higher type of the leader induces a higher effort of the follower. The proof of part (ii) of the lemma is thus complete.

A.4 Proof of Lemma 5

Let \mathcal{P} denote the seller's disclosure problem when the type space is Θ_1 . Let \mathcal{P}^n denote the seller's disclosure problem when the type space is the discrete Θ_1^n . By Lemma 4 and Facts 1–5, a \mathcal{P}^n -optimal disclosure rule can be represented by a matrix $\mathbf{p}^n \equiv [p_{ij}]_{i,j \in \{1, \dots, 2^n\}}$, whose typical element p_{ij} is the joint probability that prospect i arises and that it induces the message that pools prospects i and j . The probability that prospect i arises and is fully revealed is denoted by p_{ii} . The probability that prospect i arises is denoted by p_i and equals $\sum_{j \in \{1, \dots, 2^n\}} p_{ij}$, which is the joint probability that prospect i arises and either is pooled with any other prospect or is fully revealed. Because, by Fact 5, a prospect cannot be fully revealed sometimes and pooled at other times,

⁶⁷If $x_3 = x_4$, extend \prec (to compare a link and a fully revealed prospect) so that the $\{x_1, x_2\}$ -pooling link is \prec -smaller than the $\{x_3, x_4\}$ -prospect if x_3 is in the upper half-space of the $\{x_1, x_2\}$ -pooling link. Similarly, if $x_1 = x_2$, extend \prec so that the $\{x_1, x_2\}$ -prospect is \prec -smaller than the $\{x_3, x_4\}$ -pooling link if x_1 is in the lower half-space of the $\{x_3, x_4\}$ -pooling link.

⁶⁸Indeed, $\left(\frac{p_i\pi_i + p_{i+1}\pi_{i+1}}{p_i + p_{i+1}}, \frac{p_i\alpha_i + p_{i+1}\alpha_{i+1}}{p_i + p_{i+1}}\right)$ is an average point on the link connecting prospects (π_i, α_i) and $(\pi_{i+1}, \alpha_{i+1})$, $i = 1, 3$. By $\Delta \geq 0$, the two averages must be product-ordered. Because, by normalization, the link connecting (π_1, α_1) and (π_2, α_2) is \prec -smaller than the link connecting (π_3, α_3) and (π_4, α_4) , the only way the averages can be ordered (and a picture makes this clear) is if the average for $i = 1$ is weakly smaller than the average for $i = 3$, which is what the displayed pair of inequalities advocates.

$p_{ii} > 0$ implies $p_{ii} = p_i$ (i.e., $p_{ij} = 0$ for every $j \neq i$). A prospect can be pooled with more than one other prospect (depending on the realized message); that is, $p_{ij} > 0$ does not imply $p_{ij} = p_i$.

The value of problem \mathcal{P}^n is denoted by V^n . To define this value, first define n_* and n^* , the prospect indices that correspond to the threshold types s_* and s^* defined in Lemma 4:

$$n_* \equiv \min \{i : y_i \geq s_*\} \quad (\text{A.2})$$

and

$$n^* \equiv \max \{i : y_i \leq s^*\}. \quad (\text{A.3})$$

With this notation,

$$V^n \equiv \max_{\mathbf{p}^n} \left\{ \sum_{i=1}^{2^n} p_{ii} \tau(y_i) \alpha(y_i) + \sum_{i=n_*}^{n^*} \sum_{j \in \{1, \dots, n_*-1\} \cup \{n^*+1, \dots, 2^n\}} (p_{ij} + p_{ji}) \frac{p_{ij} \tau(y_i) + p_{ji} \tau(y_j)}{p_{ij} + p_{ji}} \frac{p_{ij} \alpha(y_i) + p_{ji} \alpha(y_j)}{p_{ij} + p_{ji}} \right\}, \quad (\text{A.4})$$

where the first term is the payoff from revealed prospects, and the second term is the payoff from pooled prospects. The maximization is over disclosure rules.

Henceforth, let \mathbf{p}^n be the \mathcal{P}^n -optimal disclosure rule, which attains the maximum in (A.4). This rule will be used to construct an approximately \mathcal{P} -optimal disclosure rule. Roughly, if in \mathcal{P}^n , \mathbf{p}^n pools prospects i and j only with each other (i.e., $\sum_k p_{ik} = p_{ij}$ and $\sum_k p_{jk} = p_{ji}$), then in \mathcal{P} , the intervals $(y_{i-1}, y_i]$ and $(y_{j-1}, y_j]$ will be “linked” pointwise, by pooling every element in $(y_{i-1}, y_i]$ with a corresponding element in $(y_{j-1}, y_j]$ according to some matching function.⁶⁹ If in \mathcal{P}^n , \mathbf{p}^n only sometimes pools prospects i and j with each other (i.e., $\sum_k (p_{ik} + p_{jk}) > p_{ij} + p_{ji}$), then in \mathcal{P} , the intervals $(y_{i-1}, y_i]$ and $(y_{j-1}, y_j]$ are divided into subintervals and only one subinterval in $(y_{i-1}, y_i]$ is linked pointwise with a subinterval in $(y_{j-1}, y_j]$.

To make the linking procedure precise, define $P_i \equiv \{j : p_{ij} > 0\}$ to be the set of prospects that \mathbf{p}^n pools with prospect i , $i \in \{1, \dots, 2^n\}$. If $P_i = \{i\}$, prospect i is revealed in \mathcal{P}^n . Partition interval $(y_{i-1}, y_i]$ into a collection of $|P_i|$ subintervals⁷⁰

$$\mathcal{C}_i \equiv \left\{ \left(\underline{b}_{ij}, \bar{b}_{ij} \right) : j \in P_i \right\} \quad (\text{A.5})$$

so that $G(\bar{b}_{ij}) - G(\underline{b}_{ij}) = p_{ij}$, and so that whenever \mathbf{p}^n pools prospects i and j , one can draw a link between element $(\underline{b}_{ij}, \bar{b}_{ij}]$ in \mathcal{C}_i and element $(\underline{b}_{ji}, \bar{b}_{ji}]$ in \mathcal{C}_j in such a manner that no two links intersect. If $|P_i| = 1$, the only subinterval is the interval $(y_{i-1}, y_i]$ itself, which either is linked to some (sub)interval or remains unlinked. The construction of the links between (sub)intervals is illustrated in Figure 11.

The rule that pools prospects in \mathcal{P} is described by a matching function τ . This matching function is constructed according to the following algorithm, which is initialized by setting $i = n_*$

⁶⁹Recall that type space Θ_1^n that induces disclosure problem \mathcal{P}^n partitions Θ_1 into 2^n subintervals $\{(y_{i-1}, y_i]\}_{i=1}^{2^n}$ so that prospect i in \mathcal{P}^n “corresponds” to the interval of types $(y_{i-1}, y_i]$ in Θ_1 .

⁷⁰Here, $|P_i|$ denotes the number of elements in the set P_i .

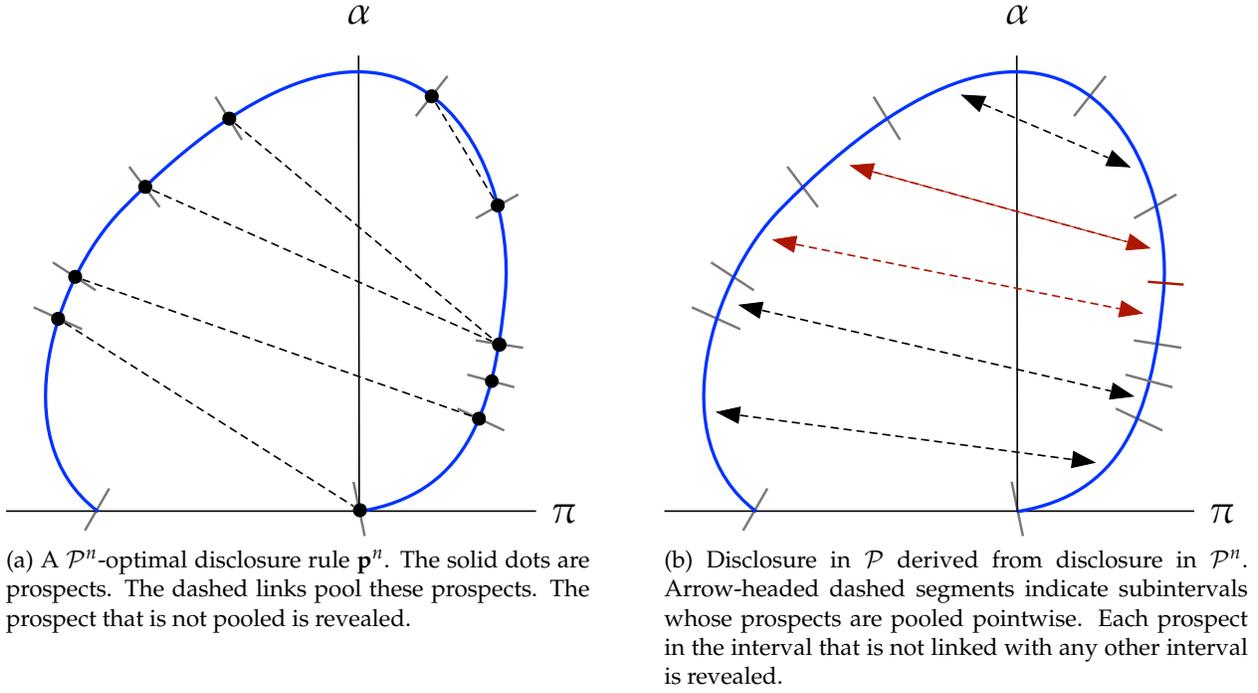


Figure 11: An optimal disclosure rule in the discrete problem \mathcal{P}^n is used to construct a disclosure rule in the continuous problem \mathcal{P} .

(where n_* is defined in (A.2)):

1. If no subinterval in \mathcal{C}_i is linked to any other subinterval, set $\tau(\theta_1) = \tau(y_{i-1})$ for all $\theta_1 \in (y_{i-1}, y_i]$, with the convention that $\tau(y_{n_*-1}) = 1$ if $s_* = 0$ and $\tau(y_{n_*-1}) = y_{n_*-1}$ if $s_* > 0$.
2. If a subinterval $(\underline{b}_{ij}, \bar{b}_{ij}]$ in \mathcal{C}_i is linked to some subinterval $(\underline{b}_{ji}, \bar{b}_{ji}]$ in \mathcal{C}_j , define a strictly decreasing τ and the corresponding derivative $\beta \equiv -\tau'$ so that for all $s \in [\underline{b}_{ij}, \bar{b}_{ij}]$, $g(s) / (\beta(s) g(\tau(s))) = p_{ij} / p_{ji}$, and $\tau(\underline{b}_{ij}) = \bar{b}_{ji}$ and $\tau(\bar{b}_{ij}) = \underline{b}_{ji}$.⁷¹ Otherwise, go to Step 3.
3. If $i < n^*$ (where n^* is defined in (A.3)), increment i by 1 and go to Step 1; otherwise, termi-

⁷¹When the c.d.f. G is uniform, the sought τ is linear: $\tau(s) = \bar{b}_{ji} - (s - \underline{b}_{ij}) p_{ji} / p_{ij}$, where $p_{ij} = \bar{b}_{ij} - \underline{b}_{ij}$ and $p_{ji} = \bar{b}_{ji} - \underline{b}_{ji}$. For a general G , set up the initial-value problem $\tau' = -p_{ji}g(s) / (p_{ij}g(\tau))$ on $[\underline{b}_{ij}, \bar{b}_{ij}]$ subject to $\tau(\underline{b}_{ij}) = \bar{b}_{ji}$. Because the right-hand side of the problem's ordinary differential equation (ODE) is continuous in (s, τ) , the Peano existence theorem implies the existence of a solution. The solution is strictly decreasing because the right-hand side of the ODE is negative. To see that the solution satisfies $\tau(\bar{b}_{ij}) = \underline{b}_{ji}$, rewrite the ODE as $-g(\tau(s)) \tau'(s) = g(s) p_{ji} / p_{ij}$ and integrate to obtain

$$-\int_{\underline{b}_{ij}}^{\bar{b}_{ij}} g(\tau(s)) \tau'(s) ds = p_{ji}.$$

The displayed integral can be rewritten equivalently by changing the variable of integration from s to $z \equiv \tau(s)$:

$$\int_{\tau(\bar{b}_{ij})}^{\tau(\underline{b}_{ij})} g(z) dz = p_{ji}.$$

Integrating gives $G(\tau(\bar{b}_{ij})) - G(\tau(\underline{b}_{ij})) = p_{ji}$, which implies $\tau(\bar{b}_{ij}) = \underline{b}_{ji}$, as desired.

nate.⁷²

Any interval $(y_{i-1}, y_i]$ whose elements are revealed contributes to the seller's payoff (25) amount

$$\int_{y_{i-1}}^{y_i} \pi(s) \alpha(s) g(s) ds \equiv p_{ii} \pi(z_i) \alpha(z_i), \quad (\text{A.6})$$

where the identity uses $p_{ii} = G(y_i) - G(y_{i-1})$ and implicitly (and not necessarily uniquely) defines $z_i \in (y_{i-1}, y_i)$ by appealing to the First Mean Value Theorem for Integration.⁷³

Any pair of linked intervals $(\underline{b}_{ij}, \bar{b}_{ij}]$ and $(\underline{b}_{ji}, \bar{b}_{ji}]$ contributes to the seller's payoff amount

$$\begin{aligned} & \int_{\underline{b}_{ij}}^{\bar{b}_{ij}} \frac{g(s) \pi(s) + \beta(s) g(\tau(s)) \pi(\tau(s))}{g(s) + \beta(s) g(\tau(s))} \frac{g(s) \alpha(s) + \beta(s) \alpha(\tau(s))}{g(s) + \beta(s) g(\tau(s))} (g(s) + \beta(s) g(\tau(s))) ds \\ & \equiv (p_{ij} + p_{ji}) \frac{g(z_{ij}) \pi(z_{ij}) + \beta(z_{ij}) g(z_{ji}) \pi(z_{ji})}{g(z_{ij}) + \beta(z_{ij}) g(z_{ji})} \frac{g(z_{ij}) \alpha(z_{ij}) + \beta(z_{ij}) \alpha(z_{ji})}{g(z_{ij}) + \beta(z_{ij}) g(z_{ji})} \\ & = (p_{ij} + p_{ji}) \frac{p_{ij} \pi(z_{ij}) + p_{ji} \pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij} \alpha(z_{ij}) + p_{ji} \alpha(z_{ji})}{p_{ij} + p_{ji}}, \quad (\text{A.7}) \end{aligned}$$

where the identity uses $p_{ij} = G(\bar{b}_{ij}) - G(\underline{b}_{ij})$ and $p_{ji} = G(\bar{b}_{ji}) - G(\underline{b}_{ji})$, and implicitly (and not necessarily uniquely) defines $z_{ij} \in (\underline{b}_{ij}, \bar{b}_{ij})$ by appealing to the First Mean Value Theorem for Integration; furthermore, $z_{ji} \equiv \tau(z_{ij})$. The equality in the last line of the above display follows by construction of τ . Because τ is strictly decreasing, $z_{ji} \in (\underline{b}_{ji}, \bar{b}_{ji})$.

Assembling the contributions (A.6) and (A.7) gives the value of the seller's objective function (25), for problem \mathcal{P} , under the disclosure rule induced by τ , constructed form \mathbf{p}^n :

$$\hat{V}^n \equiv \sum_{i=1}^{2^n} p_{ii} \pi(z_i) \alpha(z_i) + \sum_{i=n_*}^{n^*} \sum_{j \in \{1, \dots, n_*-1\} \cup \{n^*+1, \dots, 2^n\}} (p_{ij} + p_{ji}) \frac{p_{ij} \pi(z_{ij}) + p_{ji} \pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij} \alpha(z_{ij}) + p_{ji} \alpha(z_{ji})}{p_{ij} + p_{ji}}. \quad (\text{A.8})$$

By construction of $\{z_i\}$ and $\{z_{ij}\}$, $|z_i - y_i| \leq y_i - y_{i-1}$ and $|z_{ij} - y_i| \leq y_i - y_{i-1}$. Because α and π are twice continuously differentiable on $(0, 1)$ with bounded derivatives (which occurs because g , F_L , and F_H are twice continuously differentiable with bounded derivatives), the Taylor theorem implies the following for $z \in (0, 1)$:

$$\begin{aligned} \pi(z) &= \pi(y) + \pi'(y)(z - y) + O((z - y)^2) \\ \alpha(z) &= \alpha(y) + \alpha'(y)(z - y) + O((z - y)^2). \end{aligned}$$

By construction, $y_i - y_{i-1} = 1/2^n$. Hence, $y_i - y_{i-1} = O(2^{-n})$, and so $z_i - y_i = O(2^{-n})$ and

⁷²The non-linked intervals indexed by $i > n^*$ or $i < n_*$ do not affect the matching function; they automatically translate into discontinuities. The linked intervals indexed by $i > n^*$ or $i < n_*$ are accounted for when the intervals indexed by $i \in \{n_*, n_* + 1, \dots, n^*\}$ are considered.

⁷³See http://en.wikipedia.org/wiki/Mean_value_theorem#First_mean_value_theorem_for_integration

$z_{ij} - y_i = O(2^{-n})$. Using the standard properties of O , one can write:

$$\begin{aligned} \pi(z_i) \alpha(z_i) &= \pi(y_i) \alpha(y_i) + O(2^{-n}) \\ \frac{p_{ij}\pi(z_{ij}) + p_{ji}\pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij}\alpha(z_{ij}) + p_{ji}\alpha(z_{ji})}{p_{ij} + p_{ji}} &= \frac{p_{ij}\pi(y_i) + p_{ji}\pi(y_j)}{p_{ij} + p_{ji}} \frac{p_{ij}\alpha(y_i) + p_{ji}\alpha(y_j)}{p_{ij} + p_{ji}} + O(2^{-n}), \end{aligned}$$

which are substituted into (A.8) to obtain

$$\hat{V}^n = V^n + O(2^{-n})$$

as desired.

A.5 Proof of Lemma 6

Preliminary Definitions

Normalize the set of the seller's messages by setting it equal to the set of the follower's posterior probability distributions: $M \equiv \Delta\Theta_1$, where $\Delta\Theta_1$ denotes the set of Borel probabilities on the space of leader's types $\Theta_1 = [0, 1]$. The space $\Delta\Theta_1$ is a compact metric space when endowed with the topology of weak convergence.⁷⁴ Let $\Delta(\Delta\Theta_1)$ denote the space of probability measures on the subsets of $\Delta\Theta_1$. Like $\Delta\Theta_1$, the space $\Delta(\Delta\Theta_1)$ is also a compact metric space when endowed with the topology of weak convergence.

In the disclosure problem with a continuum of prospects, the seller can induce any probability distribution $\hat{\nu} \in \Delta(\Delta\Theta_1)$ over posterior probability distributions as long as $\hat{\nu}$ is Bayes plausible, that is, as long as the expected posterior probability distribution equals the prior probability distribution:

$$\int_{\Delta\Theta_1} P d\hat{\nu} = P^0,$$

where P^0 is the prior probability measure over the leader's types. The prior P^0 is derived from the c.d.f. $G: P^0\{\theta_1 : \theta_1 \leq s\} = G(s), s \in \Theta_1$. The necessity of Bayes plausibility follows from Bayes's rule, and the sufficiency has been shown by [Kamenica and Gentzkow \(2011\)](#).

Formally, when the prospect set is Γ , the seller's disclosure problem is:

$$V^* \equiv \max_{\hat{\nu} \in \Delta(\Delta\Theta_1)} \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\hat{\nu} \quad \text{s.t.} \quad \int_{\Delta\Theta_1} P d\hat{\nu} = P^0, \quad (\text{A.9})$$

where

$$\bar{\pi}(P) = \int_{\Theta_1} \pi(\theta_1) dP \quad \text{and} \quad a^*(P) = \int_{\Theta_1} \alpha(\theta_1) dP.$$

Note that both $\bar{\pi}(P)$ and $a^*(P)$ are continuous in P . To see this, take an arbitrary sequence $\{P_k\}$ of probability measures on Θ_1 that converge weakly to P . Because Lebesgue measurable functions π and α can have at most a countable number of discontinuity points, the set of disconti-

⁷⁴The topology of weak convergence is metrizable under the Prohorov metric.

nities is of measure zero, and thus by the Mapping Theorem (Billingsley, 1968, Theorem 5.1, p. 30), $\lim_k \bar{\pi}(P_k) = \bar{\pi}(P)$ and $\lim_k a^*(P_k) = a^*(P)$. The continuity of the integrand together with Proposition 3 on p. 10 of the Online Appendix to Kamenica and Gentzkow (2011) imply that the solution to problem (A.9) exists. Let ν^* denote this solution.

Similarly, when the prospect set is Γ^n , the seller solves

$$V^n \equiv \max_{\hat{\nu}_n \in \Delta(\Delta\Theta_1^n)} \int_{\Delta\Theta_1^n} \bar{\pi}(P_n) a^*(P_n) d\hat{\nu}_n \quad \text{s.t.} \quad \int_{\Delta\Theta_1^n} P_n d\hat{\nu}_n = P_n^0, \quad (\text{A.10})$$

where P_n is a probability measure on Θ_1^n , $\hat{\nu}_n$ is a probability measure on $\Delta\Theta_1^n$, and P_n^0 is the prior probability measure over the leader's types given the discretization Θ_1^n :

$$P_n^0 \equiv P(B_1^n) \delta_{y_1} + P(B_2^n) \delta_{y_2} + \dots + P(B_{2^n}^n) \delta_1,$$

where $B_i^n \equiv (y_{i-1}, y_i]$ and δ_{y_i} denotes the Dirac measure at $y_i \in [0, 1]$ (i.e., $\delta_{y_i}(B) = \mathbf{1}_{\{y_i \in B\}}$, $B \subset \Theta_1$). Let ν_n denote a solution to the discrete problem (A.10). The solution exists by Proposition 1 and Corollary 1 of Kamenica and Gentzkow (2011).

Consider a sequence of solutions $\{\nu_n\}$ and note that $\{\nu_n\}$ is a sequence of measures over because for each n , $P_n \in \Delta\Theta_1^n \subseteq \Delta\Theta_1$ and $\nu_n \in \Delta(\Delta\Theta_1^n) \subseteq \Delta(\Delta\Theta_1)$.

The Proof of the Lemma

Because the space of $\Delta(\Delta\Theta_1)$ is a compact metric space, it is sequentially compact under the topology of weak convergence, and thus the sequence of $\{\nu_n\}$ has a subsequence $\{\nu_{n'}\}$ such that as $n' \rightarrow \infty$, $\nu_{n'}$ converges weakly to some limit ν . Because $\bar{\pi}(P)$ and $a^*(P)$ are continuous, bounded, and real-valued functions defined on $\Delta\Theta_1$, by definition of weak convergence,

$$\int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\nu_{n'} \rightarrow \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\nu. \quad (\text{A.11})$$

By contradiction, suppose that ν , the limit of $\{\nu_{n'}\}$, does not solve the continuous problem (A.9) and that:

$$\epsilon \equiv \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\nu^* - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\nu > 0. \quad (\text{A.12})$$

Let $\mathcal{N} \equiv \{n'_1, n'_2, \dots\}$ be the set indexing the convergent sequence $\{\nu_{n'}\}$. Because of convergence (A.11), one can choose an $N \in \mathcal{N}$ such that for all $n' \geq N$, $n' \in \mathcal{N}$:

$$\left| \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\nu_{n'} - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\nu \right| < \frac{\epsilon}{2}. \quad (\text{A.13})$$

Because space $\Delta(\Delta\Theta_1)$ is separable and because ν^* solves the seller's maximization problem (A.9) with the leader's type space Θ_1 , one can choose an $\hat{N} \in \mathcal{N}$ such that for any $n' \geq \hat{N}$, $n' \in \mathcal{N}$,

there exists an approximation $v_{n'}^*$ to v^* with a support on $\Theta_1^{n'}$:

$$0 \leq \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv^* - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{n'}^* < \frac{\epsilon}{2}. \quad (\text{A.14})$$

Proof that such an approximation exists is in the Supplementary Appendix B.

Take $\bar{N} = \max \{N, \hat{N}\}$. Then,

$$\begin{aligned} & \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{\bar{N}} - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{\bar{N}}^* \\ = & \left(\int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{\bar{N}} - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv \right) - \left(\int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv^* - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv \right) \\ & + \left(\int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv^* - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{\bar{N}}^* \right) < \frac{\epsilon}{2} - \epsilon + \frac{\epsilon}{2} = 0, \end{aligned}$$

where the term in the first parenthesis is less than $\epsilon/2$ by (A.13), the term in the second parenthesis equals ϵ by the contradiction hypothesis (A.12), and the term in the third parenthesis is less than $\epsilon/2$ by (A.14). The inequality is a contradiction, however, because $v_{\bar{N}}$ solves the seller's discrete maximization problem with the leader's type space $\Theta_1^{\bar{N}}$. Hence,

$$\int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv \geq \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv^* \equiv V^*.$$

The argument above shows that every convergent subsequence $\{v_{n'}\}$ converges weakly to a limit that delivers a payoff at least as high as V^* . Consequently, it must be the case that $\liminf_{n \rightarrow \infty} V^n \geq V^*$. If not, there exists an $\epsilon > 0$ such that infinitely many v_n deliver payoff $V^n < V^* - \epsilon$. Then it is possible to pick a subsequence $\{v_{n_k}\}$ for which no term delivers payoff $V_{n_k} \geq V^* - \epsilon$. By sequential compactness, $\{v_{n_k}\}$ has a convergent subsequence, which is a subsequence of the original $\{v_n\}$, and by construction, the limit of this subsequence delivers a payoff strictly below V^* . This payoff contradicts the earlier established fact that every convergent subsequence of $\{v_n\}$ must converge weakly to a limit that delivers at least V^* .

References

- Bergemann, Dirk and Achim Wambach**, "Sequential information disclosure in auctions," *Journal of Economic Theory*, 2015, 159, 1074–1095.
- **and Juuso Välimäki**, "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 2002, 70 (3), 1007–1033.
- **and –**, "Information in Mechanism Design," in Richard Blundell, Whitney K. Newey, and Torsten Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Vol. I, Cambridge University Press, 2006, chapter 5, pp. 186–221.

- **and Martin Pesendorfer**, “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 2007, 137 (1), 580–609.
- Bikhchandani, Sushil and Chi-Fu Huang**, “Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets,” *The Review of Financial Studies*, 1989, 2 (3), 311–339.
- Billingsley, Patrick**, *Convergence of probability measures*, Wiley, 1968.
- Blackwell, David**, “Comparisons of Experiments,” *Proceedings of the Second Berkeley Symposium in Mathematical Statistics*, 1951, pp. 93–102.
- , “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 1953, 24 (2), 265–272.
- Calzolari, Giacomo and Alessandro Pavan**, “Monopoly with Resale,” *RAND Journal of Economics*, 2006, 32 (2), 362–375.
- **and —**, “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, 2006, 130 (1), 168–204.
- Chen, Ying**, “Communication with Two-sided Asymmetric Information,” 2009.
- Compte, Olivier and Philippe Jehiel**, “Auctions and information acquisition: sealed bid or dynamic formats?,” *The RAND Journal of Economics*, 2007, 38 (2), 355–372.
- Crawford, Vincent P. and Joel Sobel**, “Strategic Information Transmission,” *Econometrica*, 1982, 50 (6), 1431–1451.
- Crémer, Jacques, Yossi Spiegel, and Charles Z. Zheng**, “Auctions with Costly Information Acquisition,” *Economic Theory*, 2009, 38 (1), 41–72.
- Edlin, Aaron S. and Chris Shannon**, “Strict Monotonicity in Comparative Statics,” *Journal of Economic Theory*, 1998, 81 (1), 201–219.
- Eso, Peter and Balázs Szentes**, “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 2007, 74 (3), 705–31.
- Farhi, Emmanuel and Iván Werning**, “Insurance and Taxation over the Life Cycle,” *Review of Economic Studies*, 2013, 80, 596–635.
- Fudenberg, Drew and Jean Tirole**, *Game Theory*, The MIT Press, 1991.
- Gale, Ian L., Donald B. Hausch, and Mark Stegeman**, “Sequential Procurement with Subcontracting,” *International Economic Review*, 2000, 41 (4), 989–1020.
- Ganuzza, Juan-José**, “Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations,” *RAND Journal of Economics*, 2004, 35 (3), 583–598.

- **and José S. Penalva**, “Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions,” *Econometrica*, 2010, 78 (3), 1007–1030.
- Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski**, “Redistribution and Social Insurance,” 2015.
- , **Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson**, “Dynamic Strategic Information Transmission,” 2011.
- Grossman, Sanford J. and Oliver D. Hart**, “An analysis of the principal-agent problem,” *Econometrica*, 1983, 51 (1), 7–45.
- Gupta, Madhurima and Bernard Lebrun**, “First Price Auctions with Resale,” *Economics Letters*, 1999, 64 (2), 181–185.
- Haile, Philip A.**, “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory*, 2003, 108 (1), 72–110.
- Hartline, Jason D. and Tim Roughgarden**, “Optimal Mechanism Design and Money Burning,” *STOC*, 2008.
- Ivanov, Maxim**, “Dynamic Informational Control,” 2011.
- Johnson, Justin P. and David P. Myatt**, “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 2006, 96 (3), 756–784.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, 2011, 101 (6), 2590–2615.
- Kapicka, Marek**, “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach,” *Review of Economic Studies*, 2013, 01, 1–31.
- Krishna, Vijay and John Morgan**, “A Model of Expertise,” *The Quarterly Journal of Economics*, 2001, 116 (2), 747–775.
- **and –**, “The art of conversation: Eliciting information from experts through multi-stage communication,” *Journal of Economic Theory*, 2004, 117 (2), 147–179.
- Lehmann, E. L.**, “Comparing Location Experiments,” *The Annals of Statistics*, 1988, 16 (2), 521–533.
- Lewis, Tracy R. and David E. M. Sappington**, “Supplying information to facilitate price discrimination,” *International Economic Review*, 1994, 35 (2), 309–327.
- McAfee, R. Preston and John McMillan**, “Bidding Rings,” *The American Economic Review*, 1992, 82 (3), 579–599.
- Milgrom, Paul**, *Putting Auction Theory to Work*, Cambridge University Press, 2004.

- **and Chris Shannon**, “Monotone Comparative Statics,” *Econometrica*, 1994, 62 (1), 157–180.
- Milgrom, Paul R. and Ilya Segal**, “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 2002, 70 (2), 583–601.
- Mizuno, Toshihide**, “A relation between positive dependence of signal and the variability of conditional expectation given signal,” *Journal of applied probability*, 2006, 43 (4), 1181–1185.
- Myerson, Roger B.**, “Optimal Coordination Mechanisms in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 1982, 10 (1), 67–81.
- Myerson, Roger B.**, “Multistage Games with Communication,” *Econometrica*, 1986, 54 (2), 323–358.
- Ostrovsky, Michael and Michael Schwartz**, “Information Disclosure and Unraveling in Matching Markets,” *American Economic Journal: Microeconomics*, 2010, 2 (2), 34–63.
- Persico, Nicola**, “Information Acquisition in Auctions,” *Econometrica*, 2003, 68 (1), 135–148.
- Rayo, Luis and Ilya Segal**, “Optimal Information Disclosure,” *Journal of Political Economy*, 2010, 118 (5), 949–987.
- Roesler, Anne-Katrin**, “Information Disclosure in Market Design: Auctions, Contests and Matching Markets,” 2014.
- Shi, Xianwen**, “Optimal Auctions with Information Acquisition,” *Games and Economic Behavior*, 2012, 74 (2), 666–686.
- Stantcheva, Stefanie**, “Optimal Taxation and Human Capital Policies over the Life Cycle,” 2015.
- Zhang, Jun and Junjie Zhou**, “Information Disclosure in Contests: A Bayesian Persuasion Approach,” *The Economic Journal*, 2015, (forthcoming).
- Zheng, Charles Zhoucheng**, “Optimal Auctions with Resale,” *Econometrica*, 2002, 70 (6), 2197–2224.

B Supplementary Appendix: Omitted Technical Details

B.1 A Signal Structure Rationalizing the Information-Acquisition Technology of Section 3

The model's description, in Section 3, *could* have been specified as follows. The follower exerts an effort a . This effort affects the precision of a signal z . This signal's realization induces a conditional probability distribution μ_z of the underlying valuation v . This conditional probability distribution implies the expected conditional valuation $\theta_2 \equiv \mathbb{E}^{\mu_z} [v]$. Before the realization of z has been observed, μ_z and θ_2 are random variables.

The alternative (but equivalent) approach taken in Section 3 makes direct assumptions on how a affects the probability distribution of θ_2 . It would have been a mere normalization to identify the set of signal realizations with the set of conditional (on this signal) probability distributions by setting $z = \mu_z$ (Kamenica and Gentzkow, 2011). Because each player is an expected-utility maximizer, however, each cares only about θ_2 , and so it is appropriate to identify the set of signal realizations with the set of conditional expectations by setting $z = \theta_2$. The underlying signal structure that induces the probability distribution of θ_2 has been left implicit in the paper's main body, but can be recovered.

For concreteness, this appendix shows how the dependence of θ_2 on a assumed in Condition 1 can be (non-uniquely) rationalized with an appropriate joint probability distribution for v and z . Assume that each c.d.f. F_j in Condition 1 has a p.d.f. f_j , $j = L, H$. Let the follower's underlying valuation be $v \in \{0, 1\}$ with $\Pr \{v = 1\} = p$, where $p \equiv \int_0^1 s dF_H(s) = \int_0^1 s dF_L(s)$. Then, by construction, $\Pr \{v = 1\} = \mathbb{E} [\theta_2 | a]$ for all $a \in A$, meaning that the probability that the follower assigns to $v = 1$ before observing z equals his expectation of the conditional (on z) probability that $v = 1$, which is also his conditional expectation of v , denoted by θ_2 . This Bayesian consistency condition is necessary and sufficient for θ_2 to represent the follower's conditional expectation of his underlying valuation (Kamenica and Gentzkow, 2011).

Assume that the signal z can be either more precise, with probability a , or less precise, with probability $1 - a$. The realizations of the more and the less precise signals are governed by the conditional p.d.f.s $\sigma_H(z | v)$ and $\sigma_L(z | v)$, where

$$\sigma_j(z | v) \equiv \frac{z^v (1 - z)^{1-v}}{p^v (1 - p)^{1-v}} f_j(z), \quad j \in \{H, L\}, v \in \{0, 1\}, z \in [0, 1]. \quad (\text{B.1})$$

The Law of Total Probability applied to (B.1) implies that, conditional on signal technology j , z is distributed according to the c.d.f. F_j ; that is, the probability that the signal realization does not exceed z is

$$\int_{s \leq z} [p \sigma_j(s | 1) + (1 - p) \sigma_j(s | 0)] ds = F_j(z),$$

which immediately implies that unconditionally, for some effort a , z is distributed with the c.d.f. $F(\cdot | a)$.

Bayes's rule implies that z is also the expectation of v conditional on z and on signal technology j :

$$\mathbb{E}[v | z, j] = \Pr\{v = 1 | z, j\} = \frac{\sigma_j(z | 1) p}{\sigma_j(z | 1) p + \sigma_j(z | 0) (1 - p)} = z,$$

which immediately implies the expectation that is conditional only on z :

$$\theta_2 \equiv \mathbb{E}[v | z] = a \mathbb{E}[v | z, j = H] + (1 - a) \mathbb{E}[v | z, j = L] = z.$$

Hence, because z is distributed according to the c.d.f. $F(\cdot | a)$, so is θ_2 , as desired.

B.2 Justifying Equation (A.14) in the Proof of Lemma 6

To justify equation (A.14), Lemma 7 demonstrates that one can approximate any $\nu \in \Delta(\Delta\Theta_1)$ by a probability measure that puts some mass only on discrete measures in a countable set. The proof proceeds in two steps. First, it shows that by choosing n sufficiently large, any probability measure in $\Delta\Theta_1$ can be approximated by a probability measure that puts some mass on a countable set $\{\frac{1}{2^n}, \frac{2}{2^n}, \dots, 1\}$ in Θ_1 . Then, a similar argument is repeated to show that if one chooses n sufficiently large, any measure in $\Delta(\Delta\Theta_1)$ can be approximated by a measure that puts positive mass only on discrete measures with support $\{\frac{1}{2^n}, \frac{2}{2^n}, \dots, 1\}$. This second half is slightly trickier because it requires finding a countable set of non-overlapping neighborhoods in $\Delta\Theta_1$ which almost cover space $\Delta\Theta_1$.

Lemma 7. *Fix an arbitrary measure $\nu \in \Delta(\Delta\Theta_1)$. For every $\epsilon > 0$, there exists N such that for $n \geq N$,*

$$\left| \int_{\Delta\Theta_1} f(P) d\nu - \int_{\Delta\Theta_1} f(P) d\nu_n \right| < \epsilon,$$

where $f(P)$ is an arbitrary real-valued uniformly continuous, bounded function, and ν_n is a probability measure that puts some mass only on discrete measures in the countable set

$$\mathcal{D}_n \equiv \left\{ \alpha_1 \delta_{1/2^n} + \alpha_2 \delta_{2/2^n} + \dots + \alpha_{2^n} \delta_1 : \alpha_1, \dots, \alpha_{2^n} \in \mathbb{Q} \cap [0, 1], \sum_{j=1}^{2^n} \alpha_j = 1 \right\} \subset \Delta\Theta_1,$$

where \mathbb{Q} denotes the set of rational numbers and $\delta_{k/2^n}$ denotes the Dirac measure at $k/2^n \in [0, 1]$ (i.e., $\delta_{k/2^n}(B) = \mathbf{1}_{\{k/2^n \in B\}}$, $B \subset \Theta_1$). Set \mathcal{D}_n contains probability measures that put some (rational) mass on a countable set $\{\frac{1}{2^n}, \frac{2}{2^n}, \dots, 1\}$ in Θ_1 .

Proof. The proof proceeds in two steps.

Step 1: It is possible to approximate any measure μ in $\Delta\Theta_1$ with a measure in \mathcal{D}_n by choosing n sufficiently high.

Let $B_j^n \equiv \left[\frac{j-1}{2^n}, \frac{j}{2^n} \right)$ for $j = 1, 2, \dots, 2^n$, so that the family of disjoint sets $\{B_1^n, \dots, B_{2^n}^n\}$ completely

covers Θ_1 . Note that it is possible to approximate a discrete measure

$$\mu(B_1^n) \delta_{1/2^n} + \dots + \mu(B_{2^n}^n) \delta_1$$

by

$$\mu_n \equiv \alpha_1^n \delta_{1/2^n} + \dots + \alpha_{2^n}^n \delta_1,$$

where $\alpha_j^n \in [0, 1] \cap \mathbb{Q}$ such that $\sum_{j=1}^{2^n} \alpha_j^n = 1$ and

$$\sum_{j=1}^{2^n} \left| \mu(B_j^n) - \alpha_j^n \right| < \frac{1}{2^n}.$$

Such choice of $\{\alpha_j^n\}$ is possible because rationals are dense in reals. Then for each n , $\mu_n \in \mathcal{D}_n$. Moreover, as $n \rightarrow \infty$, $\mu_n \Rightarrow \mu$, where " \Rightarrow " denotes weak convergence and $\mu \in \Delta\Theta_1$.

To show that $\mu_n \Rightarrow \mu$, take a uniformly continuous bounded function g on $\Theta_1 = [0, 1]$.¹ Let $\|g\|_\infty \equiv \sup_{x \in \Theta_1} g(x)$ denote the supremum norm. Then

$$\begin{aligned} \left| \int g d\mu_n - \int g d\mu \right| &= \left| \sum_{j=1}^{2^n} \alpha_j^n g\left(\frac{j}{2^n}\right) - \int g d\mu \right| \\ &< \left| \sum_{j=1}^{2^n} \mu(B_j^n) g\left(\frac{j}{2^n}\right) - \int g d\mu \right| + \frac{1}{2^n} \sup_j \left| g\left(\frac{j}{2^n}\right) \right| \\ &\leq \left| \int \sum_{j=1}^{2^n} g\left(\frac{j}{2^n}\right) \mathbf{1}_{\{B_j^n\}} d\mu - \int g d\mu \right| + \frac{1}{2^n} \|g\|_\infty \\ &\leq \left| \sum_{j=1}^{2^n} \int \left(g\left(\frac{j}{2^n}\right) - g \right) \mathbf{1}_{\{B_j^n\}} d\mu \right| + \frac{1}{2^n} \|g\|_\infty \\ &\leq \left| \sum_{j=1}^{2^n} \sup_{x \in B_j^n} \left| g\left(\frac{j}{2^n}\right) - g(x) \right| \mu(B_j^n) \right| + \frac{1}{2^n} \|g\|_\infty. \end{aligned}$$

Note that $\left| \frac{j}{2^n} - x \right| < \frac{1}{2^n}$ for each $x \in B_j^n$. Because g is uniformly continuous, for every $\epsilon > 0$, there exists a $\delta > 0$ such that whenever $|x - y| < \delta$, $|g(x) - g(y)| < \epsilon$. Take some $\epsilon > 0$, then for n such that $\frac{1}{2^n} \leq \delta$, $\left| g\left(\frac{j}{2^n}\right) - g(x) \right| < \epsilon$ for all $x \in B_j^n$ and all j . Then, from previous calculations, it follows that

$$\left| \int g d\mu_n - \int g d\mu \right| \leq \epsilon + \frac{1}{2^n} \|g\|_\infty.$$

Because g is bounded, the second term on the right-hand side can be made arbitrarily small by

¹Statements that $\mu_n \Rightarrow \mu$ and that $\lim \int g d\mu_n = \int g d\mu$ for all uniformly continuous, bounded functions are equivalent because

1. The set of bounded Lipschitz functions on a metric space is dense in the set of continuous bounded functions on that space (Dudley, R. M., *Real Analysis and Probability* 2002, Theorem 11.2.4), which implies that instead of a wider class of bounded continuous function in the definition of the weak convergence, one may actually consider a smaller class of bounded Lipschitz functions (this fact is sometimes stated as a part of Portmanteau theorem).
2. Every Lipschitz function between two metric spaces is uniformly continuous.

choosing n sufficiently large, whereas ϵ is arbitrary. Hence, $\int g d\mu_n \rightarrow \int g d\mu$ as $n \rightarrow \infty$, which implies that $\mu_n \Rightarrow \mu$.

Step 2: It is possible to approximate any measure ν in $\Delta(\Delta\Theta_1)$ with a measure that puts some mass only on measures in \mathcal{D}_n .

Let

$$\mathcal{V} \equiv \left\{ \sum_{n=0}^k \sum_{j=1}^{\infty} \beta_{nj} \mu_n^j : \mu_n^j \in \mathcal{D}_n, \beta_{0j}, \dots, \beta_{kj} \in \mathbb{Q} \cap [0, 1], \sum_{n=0}^k \sum_{j=1}^{\infty} \beta_{nj} = 1, k = 0, 1, 2, \dots \right\}$$

be a countable subset of probability measures in $\Delta(\Delta\Theta_1)$. It contains measures that put positive mass only on measures in a countable set $\mathcal{D} \equiv \cup_{n=0}^{\infty} \mathcal{D}_n$. It will be demonstrated that \mathcal{V} is dense in $\Delta(\Delta\Theta_1)$, and thus an arbitrary measure in $\Delta(\Delta\Theta_1)$ can be approximated by some measure in \mathcal{V} .

Let $\nu \in \Delta(\Delta\Theta_1)$ and

$$B(\mu_n^j, 1/m) \equiv \left\{ \mu \in \Delta\Theta_1 : d_p(\mu_n^j, \mu) < 1/m \right\}$$

be an open ball in $\Delta(\Delta\Theta_1)$ with radius $1/m$ centered around measure $\mu_n^j \in \mathcal{D}_n$, where $d_p(\mu_n^j, \mu)$ denotes Prohorov distance between measures μ_n^j and μ . For each $m \geq 1$,

$$\cup_{j=1}^{\infty} B\left(\mu_0^j, \frac{1}{m}\right) \subset \cup_{j=1}^{\infty} B\left(\mu_1^j, \frac{1}{m}\right) \subset \dots \text{ and } \lim_{n \rightarrow \infty} \cup_{j=1}^{\infty} B\left(\mu_n^j, \frac{1}{m}\right) = \Delta(\Delta\Theta_1).$$

Take N and J such that

$$\nu\left(\cup_{j=1}^J B\left(\mu_N^j, \frac{1}{m}\right)\right) \geq 1 - 1/m.$$

Modify the balls $B\left(\mu_N^j, \frac{1}{m}\right)$ into disjoint sets by taking

$$B_1^m \equiv B\left(\mu_N^1, \frac{1}{m}\right), B_k^m \equiv B\left(\mu_N^k, \frac{1}{m}\right) \setminus \left[\cup_{j=1}^{k-1} B\left(\mu_N^j, \frac{1}{m}\right) \right], k = 2, \dots, J.$$

Then B_1^m, \dots, B_J^m are disjoint and $\cup_{k=1}^j B_k^m = \cup_{k=1}^j B\left(\mu_N^k, \frac{1}{m}\right)$ for all j . Consequently,

$$\nu\left(\cup_{k=1}^J B_k^m\right) = \nu\left(\cup_{k=1}^J B\left(\mu_N^k, \frac{1}{m}\right)\right) \geq 1 - 1/m. \quad (\text{B.2})$$

It is possible to approximate

$$\nu(B_1^m) \delta_{\mu_N^1} + \dots + \nu(B_J^m) \delta_{\mu_N^J}$$

by

$$\nu_m \equiv \beta_1^m \delta_{\mu_N^1} + \dots + \beta_J^m \delta_{\mu_N^J}$$

where $\beta_j^m \in [0, 1] \cap \mathbb{Q}$ is such that $\sum_{j=1}^J \beta_j^m = 1$ and

$$\sum_{j=1}^J \left| \nu(B_j^m) - \beta_j^m \right| < \frac{2}{m}.$$

Since rationals are dense in reals, such choice of $\{\beta_j^m\}$ is always possible through an appropriate rescaling. Then for each $m, \nu_m \in \mathcal{D}$.

To show that $\nu_m \Rightarrow \nu$, take a uniformly continuous bounded function f on $\Delta\Theta_1$. Then,

$$\begin{aligned} \left| \int f d\nu_m - \int f d\nu \right| &= \left| \sum_{j=1}^J \beta_j^m f(\mu_N^j) - \int f d\nu \right| \\ &\leq \left| \sum_{j=1}^J \nu(B_j^m) f(\mu_N^j) - \int f d\nu \right| + \frac{2}{m} \sup_j |f(\mu_N^j)| \\ &\leq \left| \int \sum_{j=1}^J f(\mu_N^j) \mathbf{1}_{\{B_j^m\}} d\nu - \int f d\nu \right| + \frac{2}{m} \|f\|_\infty \\ &\leq \left| \sum_{j=1}^J \int (f(\mu_N^j) - f) \mathbf{1}_{\{B_j^m\}} d\nu + \int f \mathbf{1}_{\{(\cup_{j=1}^J B_j^m)^c\}} d\nu \right| + \frac{2}{m} \|f\|_\infty \\ &\leq \left| \sum_{j=1}^J \sup_{\mu \in B_j^m} |f(\mu_N^j) - f(\mu)| \nu(B_j^m) + \|f\|_\infty \nu\left(\left(\cup_{j=1}^J B_j^m\right)^c\right) \right| + \frac{2}{m} \|f\|_\infty. \end{aligned}$$

Each B_j^m is contained in a ball with radius $1/m$ around μ_N^j , and thus $d_p(\mu_N^j, \mu) < \frac{1}{m}$ for each $\mu \in B_j^m$. Because f is uniformly continuous, for every $\epsilon > 0$, there is a $\delta > 0$ such that whenever $d_p(\mu, \nu) < \delta$, $|f(\mu) - f(\nu)| < \epsilon$. Take some $\epsilon > 0$; then, for $m \geq 1/\delta$, $|f(\mu_N^j) - f(\mu)| < \epsilon$ for all $\mu \in B_j^m$ and all j . Then, from previous calculations

$$\left| \int f d\nu_m - \int f d\nu \right| \leq \epsilon + \frac{1}{m} \|f\|_\infty + \frac{2}{m} \|f\|_\infty.$$

Because f is bounded, the last two terms on the right-hand side can be made arbitrarily small by choosing m sufficiently large, whereas ϵ is arbitrary. Hence, $\int f d\nu_m \rightarrow \int f d\nu$ as $m \rightarrow \infty$, which implies that $\nu_m \Rightarrow \nu$. \square

B.3 Proof of Lemma 3

A prospect set is a parametrically given plane curve $\Gamma \equiv \{(\pi(\theta_1), \alpha(\theta_1)) \mid \theta_1 \in \Theta_1\}$. The signed curvature of Γ at θ_1 is given by²

$$\kappa(\theta_1) \equiv \frac{\alpha''(\theta_1) \pi'(\theta_1) - \alpha'(\theta_1) \pi''(\theta_1)}{\left((\pi'(\theta_1))^2 + (\alpha'(\theta_1))^2 \right)^{3/2}}, \quad (\text{B.3})$$

²The curvature of Γ at a point is the reciprocal of the radius of the circle osculating Γ at that point; see the Wikipedia entry on curvature: <https://en.wikipedia.org/wiki/Curvature>.

where primes refer to derivatives with respect to θ_1 . Because Γ is simple³ and regular,⁴ curve Γ is strictly convex if and only if κ is either always positive or always negative. Because the denominator in (B.3) is always positive, requiring that κ does not change the sign is equivalent to requiring that the numerator in (B.3) not change the sign.

When $\theta_1 = \theta^*$, the numerator in (B.3) is negative, or $-r(\theta^*) (f_H(\theta^*) - f_L(\theta^*))^2 / c < 0$, because $f_H(\theta^*) \neq f_L(\theta^*)$ by the lemma's hypothesis and $F_L(\theta^*) = F_H(\theta^*)$ by part (ii) of Condition 1. Thus, the strict convexity of Γ is equivalent to the numerator in (B.3) being always negative.⁵

$$\alpha''(\theta_1) \pi'(\theta_1) - \alpha'(\theta_1) \pi''(\theta_1) < 0.$$

Substituting the definitions of α and π into the above display, dividing by $(R'(\theta_1))^2$, which is positive when $\theta_1 \neq \theta^*$, and rearranging gives the sought inequality (26) of Lemma 3.

This curvature condition captured by (26) is local and alone does not suffice to conclude that the prospect set is convex (in the sense of Definition 1); a spiral is a counterexample. Condition 1, however, which ensures $\alpha(0) = \alpha(1) = 0$, thereby ruling out a spiral and ensuring that the curvature condition in (26) is equivalent to the convexity of Γ .

³A curve is simple if it does not intersect itself.

⁴A curve Γ is regular if its derivative $(\alpha', \pi') \neq (0, 0)$ for all $\theta_1 \in \Theta_1$, which holds in our model.

⁵In general, the sign of the curvature κ indicates the direction in which the unit tangent vector rotates as a function of the parameter along the curve. If the unit tangent rotates counterclockwise, then $\kappa > 0$. If it rotates clockwise, then $\kappa < 0$. In our model, as θ_1 increases, the unit tangent vector of Γ rotates clockwise and thus κ must be negative everywhere.