An Optimal Auction with Moral Hazard

Arina Nikandrova and Romans Pancs*

February 2014

Abstract

We consider a single-item, independent private value auction environment with two bidders: the leader, who knows his valuation, and the follower, who exerts an effort that affects the probability distribution of his valuation, which he then learns. We provide sufficient conditions under which a revenue-maximizing auction solicits bids sequentially and partially discloses the leader's bid to the follower, thereby influencing the follower's effort. This disclosure rule, which is novel, prescribes sometimes revealing only a pair to which the leader's bid belongs and sometimes (not always) revealing the bid itself. Thus, the interval-pooling pattern of Crawford and Sobel (1982) is suboptimal. The induced effort distortion relative to the first-best is discussed.

Keywords: Information Disclosure, Conjugate Disclosure, Optimal Auction, Moral Hazard

JEL codes: D82, D83

1 Introduction

A seller seeks to maximize revenue by selling the item he holds to one of two bidders, whose valuations are independent and private. One bidder knows his valuation, whereas the other bidder must exert a costly effort to acquire information about his valuation. The seller designs, announces, and commits 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

*Nikandrova is at Birkbeck, University of London; Pancs is at the University of Rochester.

valuation (conditional on the information acquired by the initially uninformed bidder). Apart from this restriction on the allocation rule, the seller's choice of a mechanism is unconstrained.

Because of the option value of influencing the uninformed bidder's action by disclosing to him something about the informed bidder's bid, it is optimal for the seller to approach the bidders sequentially. We call the informed bidder, who is (optimally) approached first, the **leader**, and the bidder who chooses his information-acquisition effort and is (optimally) approached second, the **follower**.

This paper focuses on the following question: To maximize his revenue, how should the seller disclose the leader's bid (or report) to the follower? A related question is: How does the optimal disclosure rule distort the follower's effort schedule relative to the corresponding schedule in the **first-best** mechanism, which maximizes the ex-ante expected surplus (the sum of the seller's and the bidders' payoffs)? Restricting the seller to choosing an ex-post efficient allocation rule does not imply that the mechanism will be first-best (as it would in the canonical auction environment of Myerson, 1981), because the seller's disclosure rule may induce the follower's effort profile that fails to maximize the ex-ante expected surplus.

Answering this paper's questions is a step toward understanding a larger problem, which is intractable at its most general. This larger problem is concerned with revenue maximization by a seller who faces multiple bidders, each of whom chooses how much divisible and cumulable information to acquire. The canonical class of mechanisms has the seller approach bidders one by one to confidentially recommend how much information to acquire and to solicit a confidential report regarding the outcome of information acquisition, possibly revisiting bidders with additional disclosures about others' reports, followed by requests to acquire additional information, and finally allocating the item as a function of all past reports. This paper's model is a special case that assumes that information can be acquired by only one bidder (the follower), is divisible but not cumulable (i.e., the follower cannot acquire some information in the first instance, and then add to that information at a later stage), and that the allocation rule must be ex-post efficient.[1] Subject to these restrictions, the class of considered mechanisms is canonical. Fixing the admissible allocation rule to be ex-post efficient enables us to isolate the distortions due to the disclosure rule, which is the focus of this paper, from the distortions due to an allocation rule, which have been studied extensively in the optimal-auction literature without information acquisition.[2]

In addition to being of conceptual interest, the model captures some features of procurement auctions that are observed in practice. For example, in Great Britain, rail passenger services are franchised for a limited period to train operating companies. An auction determines the award of the franchise. This auction fits the model's assumptions. 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 (as will also be implied by 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 may 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, is consistent with the model's implications, as we will go on to show.

We show that the policies of either always concealing or always disclosing the leader's valuation are suboptimal for the seller. Moreover, under some conditions, the seller optimally partitions the leader's valuations (or types) into pairs of **conjugate types**, and discloses to the follower the pair to which the leader's valuation belongs, but conceals the actual valuation (Figure 1). An op-

[1] It is assumed that even though the seller is unable to commit to an ex-post inefficient allocation rule, he commits to a disclosure rule. Such differential commitment is common in policy design. For example, the U.K. Chief Medical Officer advises, "Women who are pregnant or trying to conceive should avoid alcohol altogether" (http://www.nhs.uk/chq/Pages/2270.aspx#close, accessed: 16 August 2013). This advice pools any number of units of alcohol under the same recommendation of "do not drink," thereby precluding a pregnant woman from making a better-informed decision of whether, how much, and when to drink. At the same time, professional ethics preclude physicians from punishing a new mother by refusing to treat babies with Fetal Alcohol Spectrum Disorder. As a result, an ex-post inefficient decision cannot be enforced, whereas an inefficient information-disclosure rule can.

[2] Although of great interest, the study of the optimal interaction of disclosure and allocation rules is outside the scope of this paper.

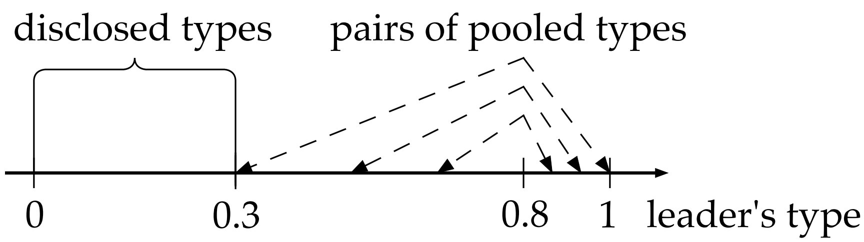

Figure 1: An example of a conjugate disclosure rule. The seller reveals the leader's type when it is less than 0.3; the seller pools any leader's type that exceeds 0.3 with a conjugate type into a pair that straddles 0.8.

timal partition is described by a weakly decreasing **matching function**, which pools the leader's valuations into pairs so that lower valuations are matched with higher valuations. The seller may also choose to reveal some of the leader's valuations without pooling them.

The paper's main finding, the optimality of the conjugate disclosure rule, can be equivalently restated in terms of the shape of the follower's effort schedule that is implied by the optimal disclosure rule. This schedule is hump shaped and "shifted to the right" relative to the first-best schedule, which is 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, the follower acquires inefficiently much information. To gain intuition for this distortion, consider two benchmarks: a first-best mechanism, in which the planner's goal is to maximize the total surplus, and a revenue-maximizing mechanism in which the seller can directly control the follower's effort.

In the first-best benchmark, the planner benefits from the follower's information because this information helps the planner efficiently allocate the item. This benefit is maximal when the leader's valuation is intermediate. When the leader's valuation is low, the follower is the likely efficient recipient of the item, even if little is known about his exact valuation. When the leader's valuation is high, he is the likely recipient. In either extreme case, the planner gains little from additional information about the follower's valuation. So the follower's first-best effort schedule is hump shaped. Moreover, this schedule can (nonuniquely) be implemented in a mechanism that reveals the leader's valuation to the follower.

In the revenue-maximizing benchmark in which the seller controls the follower's effort, he sets this effort to zero when the leader's valuation is below a certain threshold, and sets it to be maximal otherwise. The intuition for this "bang-bang" property of the effort schedule stems from the fact that, for any leader's valuation, denoted by θ_1, ex-post efficient allocation enables the seller to charge the follower amount θ_1, but the leader less than θ_1, and so the seller prefers selling to the follower. 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 is required by ex-post efficiency. When $\theta_2 < \theta_1$, the seller cannot possibly hope to sell for θ_1 to the leader; 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, the seller instructs the follower to acquire information so as to increase the probability of an ex-post efficient sale to the follower. In particular, because information acquisition induces a mean-preserving spread in the probability distribution of θ_2, the event $\theta_2 \geq \theta_1$ (corresponding to an ex-post efficient sale to the follower) is more likely either when θ_1 is low and the information-acquisition effort is minimal or when θ_1 is high and the information-acquisition effort is maximal.

The economic forces present in both benchmarks combine in this paper's central case, in which the seller maximizes his revenue and cannot directly control the follower's effort. The seller influences the follower's information-acquisition effort indirectly, by strategically revealing information about the leader's bid. Because the Bayesian follower cannot be deceived systematically, the seller's disclosure distorts, but does not disregard, the first-best effort schedule, which would have prevailed under full disclosure. Therefore, 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 the optimal effort schedule. The Revelation Principle for games with private actions (Myerson, 1982, 1986) implies that all that the follower needs to be told is the effort that the seller would like him to exert. Then, given any target effort schedule, the corresponding rule for disclosing leader's bids can be easily designed. For example, both in the first-best effort schedule and in the optimal effort schedule, the same effort is exerted for (at most) two leader's types, which can be read off Figure 2 by intersecting an effort schedule with a horizontal line corresponding to a recommended effort (not shown). This suggests the first-best effort schedule and the optimal effort schedule can both be implemented by means of some conjugate-disclosure rule. The first-best effort schedule is easy to derive, whereas deriving the optimal effort schedule is a nontrivial

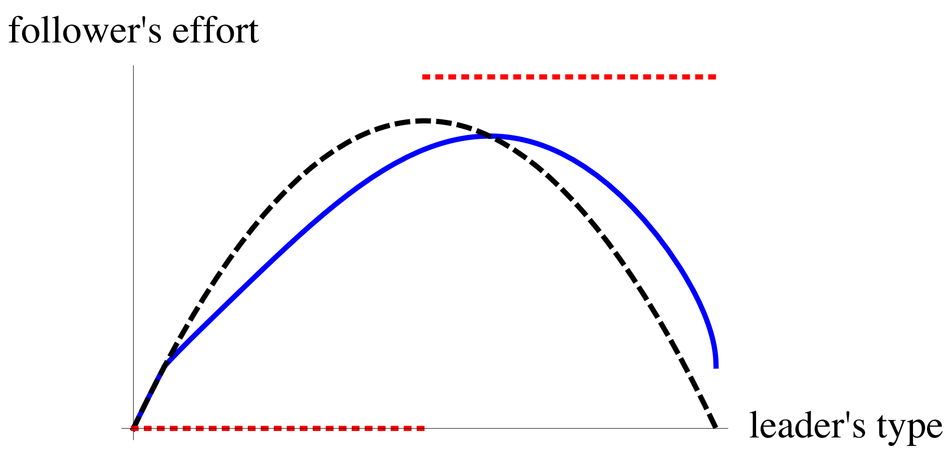

Figure 2: The follower's optimal effort schedule (the solid hump-shaped curve) is a compromise between his first-best effort schedule (the dashed hump-shaped curve) and the bang-bang effort schedule (the two dashed horizontal line segments), corresponding to the benchmark in which the seller controls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

task. Indeed, we have found that instead of deriving the optimal effort schedule directly, it is more convenient to derive the optimal disclosure rule first and then obtain the implied effort schedule.

The outcome that maximizes the seller's revenue can be implemented in a second-price auction with a tax (or subsidy) for the leader.[3] The role of the tax is to encourage the leader to bid truthfully. He would do so without a tax in a model in which the follower exerted no effort. Because 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 information-acquisition effort translates into a more dispersed probability distribution of the follower's expected valuations, which benefits the leader, whose payoff is convex in the follower's expected valuation. Hence, the corrective tax is necessary.

This paper is most closely related to two strands of literature: mechanism design with endogenous information acquisition and the models of Bayesian persuasion. Mechanism design with information acquisition tends to focus on mechanisms that are simultaneous in the sense that all bidders acquire information simultaneously, so the issues of optimal bid-disclosure do not arise. For instance, Shi (2012) studies a revenue-maximizing simultaneous auction. Bergemann and Välimäki (2002) study a Vickrey-Clark-Groves mechanism, which in the present model, corresponds to the second-price auction and is neither first-best nor profit-maximizing.

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

Methodologically, the most pertinent Bayesian-persuasion paper is by Rayo and Segal (2010), who derive a canonical class of disclosure rules in a discrete-type optimal disclosure model.[4] In deriving the seller's optimal bid-disclosure rule, our paper adapts their results to the continuous type-space auction model and exploits the special structure of our setting.

One of the paper's contributions is establishing the optimality of pooling pairs of types, never intervals of types, under the same message. To our knowledge, the only paper with a pooling pattern similar to ours is Golosov et al. (2011), who consider a dynamic version of Crawford and Sobel (1982). In that version, in every period, an expert, who is privately informed about a fixed state of the world, 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 ordered in the same way as in our paper, however. Furthermore, their pairwise disclosure serves a different purpose (i.e., the eventual separation of types) in a disclosure game that is substantively different (i.e., has multiple stages of communication and no commitment).[5]

The seller's disclosure problem is reminiscent of, but is not directly related to, the literature on optimal delegation (Krishna and Morgan, 2008; Kovac and Mylovanov, 2009; Amador and Bagwell, 2013). In this literature, an uninformed principal offers a contract designed to elicit information from an informed agent and to make him take a desired action. Optimal contracts typically involve intervals of pooled types and intervals of revealed types. By contrast, in the present model, the pooling of intervals of types is suboptimal, and pooling into conjugate pairs of types emerges instead. The difference in the results is due to a crucial difference in the settings. The optimal-delegation literature assumes that the uniformed party designs the contract, whereas in this paper, the the informed party (the seller) designs it.

Among other papers that study information acquisition in auctions are Bergemann and Välimäki (2002), Persico (2003), Compte and Jehiel, 2007, and Crémer et al. (2009). Bergemann and Välimäki

[4]Kamenica and Gentzkow (2011) address the same conceptual problem as Rayo and Segal (2010) do, but in a model with less structure, and so obtain results that are less sharp for the purposes of the present paper.

[5]Introducing commitment into a canonical cheap-talk model does not generate anything resembling the conjugate disclosure that we find. Indeed, in the quadratic formulation of Crawford and Sobel (1982), if one assumes that the sender commits to a disclosure rule, full disclosure becomes 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, he can make the receiver's action match the state. Such matching maximizes the sender's ex-ante expected utility, because the seller is risk-averse with respect to the discrepancy between the receiver's action and the state.

(2002) are concerned with efficient mechanisms in environments with simultaneous information acquisition; these mechanisms are no longer efficient when information acquisition can be sequential, as Section 3 illustrates. Persico (2003) and Compte and Jehiel (2007) compare revenues in standard auctions. Compte and Jehiel, 2007 and Crémer et al. (2009) find that dynamic auctions deliver higher revenues than static ones when bidders can acquire information sequentially; this insight justifies our focus on dynamic mechanisms, which indeed are shown to dominate static ones. Crémer et al. (2009) design a revenue-maximizing auction. Because they assume that the seller can charge bidders for participation, optimal information disclosure turns out to be trivial and is not the focus of their paper. By contrast, the present paper rules out participation fees (and with them, the seller's ability to sell information about the leader's bid to the follower) by imposing interim individual rationality and focuses on the information-disclosure design.

This section concludes with an outline for the rest of the paper. Section 2 describes the economic environment. Section 3 derives the first-best allocation and an auction that implements it. Section 4 establishes the suboptimality of the policies of full disclosure, no 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 delivers the matching function underlying the conjugate disclosure rule. The section also discusses how and why the follower's effort schedule in the optimal auction departs from the corresponding schedule in the first-best auction. Section 5 discusses some immediate generalizations and extensions that do not alter the spirit (and often not even the letter) of the main results. Section 6 seeks optimal disclosure rules in two examples. Section 7 concludes. Some proofs are relegated to the Appendix. Technical results comprise the Supplementary Appendix.

2 Model

Environment

The seller must allocate an item, which he values at zero, to one of two bidders, whom he contacts sequentially. 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 on $(0, 1)$ is smooth, 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 specified shortly. The cost of effort a is $C(a) \equiv ca^2/2$, $c > 0$.

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 his payment. Then, 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 his expected underlying valuation conditional on the privately observed outcome of his information-acquisition effort. The underlying valuation is not modeled, because it affects the follower's payoff only through θ_2. For any $a \in A$, θ_2 is drawn according to a c.d.f that is linear in a:[6]

$$F(\theta_2 \mid a) \equiv aF_H(\theta_2) + (1-a) F_L(\theta_2), \qquad \theta_2 \in \Theta_2 \equiv [0,1]. \tag{1}$$

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 $(0,1)$, it 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) $\int_0^1 F_H(s)\,\mathrm{d}s = \int_0^1 F_L(s)\,\mathrm{d}s$, 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

[6] In 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).

informative signal about his underlying valuation, and with probability $1-a$, receives a less informative signal.[7] For this interpretation, part (i) requires the follower's effort not to affect his expected type.[8] In particular, part (i) rules out the situations in which the follower's effort is a value-enhancing investment. In addition, part (ii) implies that a higher effort induces a mean-preserving spread of the distribution of the follower's types.[9] 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 conditional expectations, θ_2.[10,11] When $a = 1$, the realized θ_2 can be interpreted either as an expectation of a still-uncertain underlying valuation or as the underlying valuation itself.

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

Example 1 (Information Acquisition)**.** $F(\theta_2 \mid a) = a/2 + (1-a)\,\theta_2$ for $\theta_2 < 1$, and $F(1 \mid a) = 1$.

Example 1 can be interpreted to say that with probability a, the follower observes a perfectly informative signal that reveals his underlying valuation v, which is distributed uniformly on $\{0, 1\}$, and with probability $1-a$, the follower observes a partially informative signal about v.

An example that violates Condition 1 is

Example 2 (Valuation Enhancement)**.** $F_H(\theta_2) = \theta_2^2$ and $F_L(\theta_2) = \theta_2$ for $\theta_2 \in \Theta_2$.

This example violates part (i) of Condition 1, whereas part (ii) holds with $\theta^* = 0$. Here, a

[7] A signal structure that rationalizes the probability distribution of types in Condition 1 is given in Supplementary Appendix B.1.

[8] For any c.d.f. H on $[0, 1]$, its expectation is $\int x \mathrm{d}H(x) = \int (1 - H(x))\, \mathrm{d}x$.

[9] Here is the proof that parts (i) and (ii) imply a mean-preserving spread. Let $a' > a$. Part (ii) implies $(\theta^* - s)(F(s \mid a') - F(s \mid a)) \geq 0$, which combined with part (i) gives $\int_0^x F(s \mid a')\, \mathrm{d}s \geq \int_0^x F(\theta \mid s)\, \mathrm{d}s$ for all $x \in [0, 1]$. That is, $F(\cdot \mid a)$ second-order stochastically dominates $F(\cdot \mid a')$, and the claim about the mean-preserving spread follows.

[10] For 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 $\theta_2 = v$. To summarize, anticipating a perfectly uninformative signal, one assigns probability one to $\theta_2 = 1/2$, whereas anticipating a perfectly informative signal, one views θ_2 as distributed uniformly on $[0, 1]$. In the former case, the dispersion of the probability distribution of conditional expectations, θ_2, is zero, while in the latter case, it is equal to the dispersion of the probability distribution of the underlying valuation.

[11] Characterizing the informativeness of a signal through the induced dispersion of the probability distribution of conditional expectations is common in the literature. See, for example, Ganuza and Penalva (2010), Johnson and Myatt (2006) or Shi (2012).

higher effort induces a first-order stochastic-dominance shift in the distribution of θ_2. Here, an effort is investment in enhancing the valuation, not in acquiring information.

To simplify the exposition by avoiding the corner solution $a = 1$, assume that the cost of effort is sufficiently large:[12]

$$c > \int_0^{\theta^*} \left(F_H\left(s\right) - F_L\left(s\right)\right) \mathrm{d}s. \tag{2}$$

The Seller's Problem

The seller chooses, publicly announces, and commits to a mechanism that is ex-post efficient and interim individually rational. A mechanism is **ex-post efficient** if, 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 $\mathbf{1}_{\{\}}$ is the indicator function. A mechanism is **interim individually rational** if, having observed his type θ_i, bidder expects a nonnegative payoff.

By the Revelation Principle for multi-stage games (Myerson, 1986), no generality is lost by restricting attention to direct-revelation mechanisms of the following form:

1. The leader observes θ_1 and confidentially reports $\hat{\theta}_1 \in \Theta_1$.

2. The seller confidentially sends a message $m \in M$ to the follower according to some **disclosure rule** $\mu : \Theta_1 \to \Delta\left(M\right)$, which associates with each report of the leader a probability distribution over messages.

3. The follower exerts an effort $a^*\left(m\right) \in A$, 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 $(t_1, t_2) : \Theta_1 \times \Theta_2 \to \mathbb{R}^2$, the functions of bidders' reports, are assessed.

If the seller's message were independent of the leader's report, the described direct-revelation mechanism would be strategically equivalent to a mechanism in which the seller asked both bidders to submit their reports simultaneously.

Without loss of generality, the seller can focus on the (perfect Bayes-Nash) equilibria that are **incentive-compatible,** meaning that $\left(\hat{\theta}_1, \hat{\theta}_2\right) = (\theta_1, \theta_2)$. Nor does a loss of generality occur in

[12]This condition is derived by requiring that the first-best effort, given in (6), be less than 1.

setting $M = A$, although a different M will sometimes be convenient.

The seller chooses a disclosure rule μ and payments (t_1, t_2) that induce an ex-post efficient, interim individually rational, and incentive-compatible mechanism that maximizes the seller's expected revenue:

$$\int_{\Theta_1} \int_M \int_{\Theta_2} (t_1(\theta_1, \theta_2) + t_2(\theta_1, \theta_2)) \, \mathrm{d}F(\theta_2 \mid a^*(m)) \, \mathrm{d}\mu(m \mid \theta_1) \, \mathrm{d}G(\theta_1). \tag{3}$$

Ex-post efficiency and interim individual rationality are restrictive but essential for maintaining the focus on optimal information disclosure. When interim individual rationality is not imposed, optimal information disclosure is trivial; when ex-post efficiency is not imposed, it is conflated with the design of optimal allocation rule.[13] Moreover, ex-post efficiency and interim individual rationality are relevant for applications where the seller cannot impose fees for participation or information about the leader's type and cannot commit to an allocation rule that is inefficient given the revealed information about bidders' valuations.[14]

3 A First-Best Auction

The **first-best** outcome obtains when a fictitious planner maximizes the expected total surplus (i.e., the sum of bidders' payoffs), observes each bidder's type, and directly controls the follower's effort.

For every leader's type θ_1, the follower's first-best effort $\alpha(\theta_1)$ maximizes the total surplus:

$$\alpha(\theta_1) \in \arg\max_{a \in A} \left\{ \int_{\Theta_2} \max\{\theta_1, \theta_2\} \, \mathrm{d}F(\theta_2 \mid a) - C(a) \right\}. \tag{4}$$

Integrating by parts gives:

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

[13] Without ex-post efficiency, the full mechanism-design problem is intractable. Without interim individual rationality, the seller can extract the bidders' entire surplus by implementing a first-best auction that admits full disclosure, as in the model of Crémer et al. (2009).

[14] Fees for participation and for the information about the leader's type are ruled out by interim (but not ex-ante) participation constraints.

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 above display:

$$R(\theta_1, a) \equiv -\int_{\theta_1}^{1} \frac{\partial F(\theta_2 \mid a)}{\partial a} \mathrm{d}\theta_2 - C'(a). \tag{5}$$

Because for any $a > 0$, $R(1, a) = -C'(a) < 0$ and $R(0, a) = -C'(a) < 0$ (by part (i) of Condition 1), if the leader's valuation is the highest possible or the lowest possible, it is first-best for the follower to exert no effort. Indeed, in either of these extreme cases, no additional information about the follower's valuation affects the item's efficient allocation. So no additional information is solicited.

Part (ii) of Condition 1 implies that $R(\theta_1, a)$ is strictly increasing in θ_1 for $\theta_1 < \theta^*$ and is strictly decreasing in θ_1 for $\theta_1 \geq \theta^*$, where θ^* is defined in that condition. 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 $\theta_1 \geq \theta^*$. Moreover, the dependence is strict when $\alpha(\theta_1) \in (0, 1)$ (Edlin and Shannon, 1998), which can be shown to be the case when $\theta_1 \in (0, 1)$.[15]

The following theorem summarizes the properties of the first-best effort.

Theorem 1. *The first-best effort*

$$\alpha(\theta_1) = \frac{1}{c} \int_{\theta_1}^{1} (F_L(s) - F_H(s))\, ds, \qquad \theta_1 \in [0, 1] \tag{6}$$

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

The following corollary describes a modified second-price auction that implements the first-best effort. In this auction, the leader who bids b pays the (possibly negative) tax

$$T(b) \equiv \int_{0}^{b} (F(s \mid \alpha(b)) - F(s \mid \alpha(s)))\, \mathrm{d}s. \tag{7}$$

[15] In particular, $\theta_1 \in (0, 1)$ implies $\alpha(\theta_1) > 0$ because the quadratic cost satisfies $C'(0) = 0$; and $\alpha(\theta_1) < 1$ holds for $\theta_1 = \theta^*$, and hence for all θ_1, by (2).

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 according to the rules of the second-price auction.

In equilibrium, each bidder bids his type, and 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 a standard second-price auction executed sequentially and with the leader's bid made public. The tax offsets the leader's incentives to manipulate the follower's information acquisition by bidding untruthfully. To see the incentives for untruthful bidding, suppose that type-θ_1 leader faces no tax (i.e., $T \equiv 0$) and bids truthfully (i.e., $b = \theta_1$). Then, his payoff in the second-price auction is $\mathbb{E}_{\theta_2}\left[\max\left\{0, \theta_1 - \theta_2\right\}\right]$, the standard expression except that the leader's bid affects the probability distribution of θ_2.

Because $\max\left\{0, \theta_1 - \theta_2\right\}$ is convex in θ_2, Jensen's inequality implies that the leader gains from the follower's valuation being more spread out, which occurs when the follower exerts more effort (footnote 9). The leader can induce the follower to exert more effort by untruthfully bidding away from his valuation and toward θ^*. 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 has a first-order beneficial effect on the leader's payoff.

The role of the tax in discouraging untruthful bidding can be appreciated by considering the tax-induced reduction in the leader's return to increasing marginally his bid. This reduction is

$$T'(b) = \alpha'(b) \int_0^b \frac{\partial F(s \mid \alpha(b))}{\partial a} \mathrm{d}s,$$

where the integral is positive for all $b \in (0, 1)$ (by Condition 1). Therefore, the sign of $T'(b)$ is the same as the sign of $\alpha'(b)$, which is positive when $b < \theta^*$ and is negative when $b > \theta^*$. In other words, bids below θ^* are taxed on the margin, and bids above θ^* are subsidized on the margin.

4 A Seller-Optimal Auction

Without information acquisition, the seller's problem would have been trivial. The ex-post efficient allocation rule, interim individually rational participation, and incentive compatibility would have uniquely determined the seller's expected payoff, by the revenue equivalence theorem. With information acquisition, however, the allocation is not uniquely determined by the ex-post efficiency requirement, because the follower's effort influences the allocation. So the seller's problem is non-trivial and will be analyzed in three steps, comprising three subsections.

In Section 4.1, incentive compatibility is used to substitute payments out of the seller's objective function. This step relies on the standard envelope argument for local incentive compatibility and yields a virtual surplus.[16] The seller's problem then reduces to an information-disclosure problem in which the policies of full disclosure and no disclosure are suboptimal.

In Section 4.2, under additional assumptions, it is demonstrated that the search for a seller-optimal disclosure rule can be restricted to so-called conjugate disclosure rules, described by a matching function that pools pairs of the leader's valuations. The justification for focusing on conjugate disclosure relies on results of Rayo and Segal (2010), which are adapted for the continuous type space in the seller's disclosure problem. The shape of the optimal effort schedule is then derived from the optimal disclosure rule.

In Section 4.3, the seller's disclosure problem is formulated as an optimal-control problem.

Two remarks clarifying our approach are in order. First, the optimality analysis focuses on the seller's relaxed problem that neglects the monotonicity condition that requires the leader's probability of winning the item to be nondecreasing in his type. This monotonicity is required by incentive compatibility and must be verified in applications (e.g., as in Section 6).[17] The relaxed problem, even though restrictive, is sufficiently rich to have become focal in the mechanism-design literature.

Second, by the Revelation Principle, the canonical disclosure rule takes the form of the effort schedule that recommends an effort to the follower as a (possibly stochastic) function of the

[16] For local incentive compatibility, see Fudenberg and Tirole (1991, pp. 284-8) or the Constraint Simplification Theorem of Milgrom (2004), implied by the Envelope Theorem of Milgrom and Segal (2002).

[17] The analogous monotonicity condition for the follower is guaranteed to hold. The monotonicity condition is guaranteed to hold also for the leader in the mechanism that implements the first-best outcome and in the various benchmarks considered in this paper.

leader's type. In order to find this schedule, it is analytically convenient to begin with a non-canonical disclosure rule, which admits the possibility that distinct messages induce the same effort. The derived optimal (non-canonical) disclosure rule identifies the optimal effort schedule, which, in turn, implies a disclosure rule in the canonical form, where pairs of leader's types are pooled under the same recommended effort.

4.1 The Seller's Information-Disclosure Problem

The Follower's Optimal Effort

When devising a disclosure rule, the seller anticipates its effect on the follower's effort schedule, which maps the leader's type into the follower's information-acquisition effort. For a given disclosure rule, this schedule is determined as follows. Reasoning backwards, suppose that the follower has observed his type θ_2 and the seller's message m. The follower then chooses the report that maximizes his expected payoff:

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

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 the Envelope Theorem, incentive compatibility requires

$$U_2'\left(\theta_2 \mid m\right) \quad \equiv \quad \frac{\mathrm{d}U_2\left(\theta_2 \mid m\right)}{\mathrm{d}\theta_2} = \mathbb{E}_{\theta_1 \mid m}\left[\mathbf{1}_{\{\theta_2 > \theta_1\}}\right]. \tag{8}$$

Then, the follower's expected payoff from exerting effort a' is

$$\int_{\Theta_2} U_2\left(\theta_2 \mid m\right) \mathrm{d}F_2\left(\theta_2 \mid a'\right) \quad = \quad U_2\left(0 \mid m\right) + \int_0^1 \mathbb{E}_{\theta_1 \mid m}\left[\mathbf{1}_{\{\theta_2 > \theta_1\}}\right]\left(1 - F\left(\theta_2 \mid a'\right)\right) \mathrm{d}\theta_2,$$

where the equality uses (8) and follows by integration by parts.[18] Interchanging the order of integration and expectation (by Fibula's theorem) in the right-hand side of the above display yields

[18]Integration by parts is valid even if even F is discontinuous (e.g., as in Example 1), because $U_2\left(\cdot \mid m\right)$ has a bounded derivative everywhere and hence is continuous.

the follower's choice-of-effort problem:

$$a^*(m) \in \arg\max_{a' \in A} \left(\mathbb{E}_{\theta_1 \mid m} \left[\int_{\theta_1}^{1} \left(1 - F\left(\theta_2 \mid a'\right)\right) \mathrm{d}\theta_2 \right] - C\left(a'\right) \right). \tag{9}$$

Under the maintained Condition 1, (9) has a unique solution:

$$a^*(m) = \mathbb{E}_{\theta_1 \mid m}\left[\alpha\left(\theta_1\right)\right], \tag{10}$$

where $\alpha(\theta_1)$, defined in (6), is the first-best effort level when the leader's type is θ_1.

The Seller's Virtual Surplus

The seller's objective function (3) is transformed into the **virtual surplus** by using a standard mechanism-design technique: bidders' transfers are substituted out using bidders' local incentive-compatibility constraints, such as the one derived for the follower in (8):[19,20]

$$\mathbb{E}_{m,\theta_1,\theta_2}\left[\left(\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} \right) \mathbf{1}_{\{\theta_1 > \theta_2\}} + \left(\theta_2 - \frac{1 - F\left(\theta_2 \mid a^*(m)\right)}{f\left(\theta_2 \mid a^*(m)\right)} \right) \mathbf{1}_{\{\theta_2 \geq \theta_1\}} \right]. \tag{11}$$

The displayed virtual surplus is standard; it is the expected sum of each bidder's virtual valuation times the probability that he gets the item. Here, the probability of getting the item is pinned down by the ex-post efficient allocation rule. Integrating θ_2 out and simplifying yields:

$$\int_{\Theta_1} \mathbb{E}_{m \mid \theta_1}\left[\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} F\left(\theta_1 \mid a^*(m)\right) \right] \mathrm{d}G(\theta_1), \tag{12}$$

where $a^*(m)$ is given by (9).[21] If not for the possibility that a^* depended on θ_1 (through m), the virtual surplus (12) would have been standard.

To understand (12), 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 $F(\theta_1 \mid a^*(m))$, which is the probability that $\theta_2 < \theta_1$. 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** $(1 - G(\theta_1))/g(\theta_1)$.

[19] The analogue of (8) for the leader is $U_1'(\theta_1) = \mathbb{E}_{m \mid \theta_1}\left[F\left(\theta_1 \mid a^*(m)\right)\right]$.

[20] If F has no density f, skip to (12), which does not rely on the existence of f.

[21] The virtual surplus in (12) is non-negative because it is bounded below by $\int_{\Theta_1}\left[\theta_1 - (1 - G(\theta_1))/g(\theta_1)\right] \mathrm{d}G(\theta_1) = 0$.

Analogous reasoning suggests that if the follower gets the item (which occurs with probability $1 - F(\theta_1 \mid a^*(m))$), the seller's gain is the follower's expected virtual valuation conditional on winning, which can be verified to be θ_1.[22] That is, the follower's information rent is implicit in (12). The expression for the virtual surplus is asymmetric because θ_2 has been integrated out; the dynamic nature of the mechanism is reflected only in the fact that m may depend on θ_1.

By (12), the seller's goal is to design a disclosure rule so as to minimize the expected weighted probability with which the leader buys (where the weights are the leader's information rents). The asymmetric treatment of bidders in the seller's objective function stems solely from the fact that the seller can affect the distribution of the follower's valuations, but not the leader's. By aiming to minimize the leader's expected weighted probability of winning, the seller aims to encourage the follower to become "stronger," thereby intensifying bidder competition.[23] Whether a better or a worse informed follower is "stronger" depends on the leader's valuation. When the leader's valuation is high, the follower stands a better chance of outbidding the leader if better informed, with the distribution of his types more dispersed. The opposite is true when the leader's valuation is low.

Two Benchmarks

The expression for the virtual surplus, the maximand in the seller's disclosure problem, suggests two benchmarks that are useful for building an intuition for an optimal disclosure rule:

Theorem 2. *(i) If the seller observed the leader's valuation, any disclosure rule would be optimal. In particular, fully disclosing to the follower 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 $\bar{a} < 1$, any disclosure rule would be optimal and the seller would choose the effort $\bar{a}\mathbf{1}_{\{\theta_1 \geq \theta^*\}}$, $\theta_1 \in \Theta_1$.

Proof. For part (i), if the seller observes the leader's valuations, he can extract the leader's entire

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

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

Even though the integrand is written as if F had a positive density, f, the validity of (12) requires no such assumption.

[23] The bidder asymmetry does not suggest that if he could, the seller would wish to sabotage the leader's valuation, so that the leader loses more often.

information rent by charging him θ_1 when the follower refuses to buy at θ_1. In this case, the leader's virtual valuation in (11) is replaced simply by his valuation, θ_1, and hence the implied (12),

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

is independent of the follower's effort.

For part (ii), if the seller can directly and costlessly choose the follower's effort $a \in [0, \bar{a}]$,[24] he does so to minimize pointwise the probability that the leader is the ex-post efficient recipient of the item, by (12). Formally, for $\theta_1 < \theta^*$, the seller sets $a = 0$, and for $\theta_1 \geq \theta^*$, the seller sets $a = \bar{a}$.[25] When the follower has no control over his action, any disclosure rule can be optimal, as truthful reporting can be made a dominant strategy to the follower (at no additional cost to the seller), for instance, by making a take-it-or-leave-it offer to the follower at the price θ_1. □

Part (i) of Theorem 2 indicates that managing the leader's information rent is the sole rationale for the seller's strategic disclosure. Part (ii) of the theorem shows that the follower's information-acquisition effort is an effective instrument for doing so—at least if this effort can be controlled directly. Figure 2 juxtaposes the ideal controlled effort schedule and the first-best effort schedule.

The Suboptimality of Full Disclosure and Non-disclosure

The following theorem shows that the restriction that the seller must persuade, but cannot coerce, the follower to exert an intended effort does not render the seller's disclosure problem trivial. In the theorem, **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.

Theorem 3. *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*

[24] The 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.

[25] The leader's monotonicity (or global incentive-compatibility) constraint holds; that is, the leader's expected allocation $F\left(\theta_1 \mid \mathbf{1}_{\{\theta_1 \geq \theta^*\}}\right)$ is weakly increasing in his type θ_1. Indeed, $F(\theta_1 \mid 0)$ and $F(\theta_1 \mid 1)$ are both weakly increasing in θ_1, and $\lim_{\theta_1 \uparrow \theta^*} F(\theta_1 \mid 0) \leq \lim_{\theta_1 \uparrow \theta^*} F(\theta_1 \mid 1) \leq F(\theta^* \mid 1)$. The follower's monotonicity constraint also holds, trivially.

leader's types under the same message.[26]

The proof of Theorem 3 uses the seller's objective function (12) rewritten in a canonical form, which is also used in the proofs of subsequent results. To arrive at the canonical form, neglect the additive term in (12) that is independent of the disclosure rule. Then, the objective function becomes

$$\int_{\Theta_1} \mathbb{E}_{m|\theta_1} \left[\pi(\theta_1) a^*(m)\right] \mathrm{d}G(\theta_1), \tag{13}$$

where

$$\pi(\theta_1) \equiv \frac{1 - G(\theta_1)}{g(\theta_1)} \left(F_L(\theta_1) - F_H(\theta_1)\right) \tag{14}$$

is the seller's **marginal benefit** from an increase in a. The marginal benefit π equals the leader's information rent times the marginal increase in the probability that the follower gets the item.

The seller's objective function (13) is further transformed by using the Law of Iterated Expectations and the relationship $a^*(m) = \mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]$, derived in (10), yielding the **canonical form**

$$\mathbb{E}\left[\mathbb{E}_{\theta_1|m}[\pi(\theta_1)]\, \mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]\right]. \tag{15}$$

The transformed objective function is canonical in the sense that it coincides with the sender's objective function (equation [2]) in the optimal information-disclosure model of Rayo and Segal (2010), henceforth abbreviated as RS.[27] For the rest of the paper, a tuple $(\pi(\theta_1), \alpha(\theta_1))$ is called a **prospect**.[28] The **prospect set** is the graph $\Gamma \equiv \{(\alpha(\theta_1), \pi(\theta_1)) : \theta_1 \in \Theta_1\}$.

The proof of Theorem 3 relies on two lemmas, which are of independent interest. For the first lemma, call functions α and π **co-monotone** 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 co-monotone.*

Proof. Necessity: Suppose that α and π are not co-monotone. Then, an interval $I \subset \Theta_1$ exists on

[26] A 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 that is not analytic can be approximated arbitrarily well by an analytic function.

[27] In what follows, bracketed indices will refer to equations and lemmas in RS's paper.

[28] The set $\{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in \Theta_1\}$ differs from RS's set of prospects only in that RS require it to be finite, for tractability.

which α is strictly increasing and π is strictly decreasing, or the other way around. In this case,

$$\int_I \int_I \left(\alpha\left(s'\right) - \alpha\left(s\right)\right) \left(\pi\left(s'\right) - \pi\left(s\right)\right) \mathrm{d}s\mathrm{d}s' < 0.$$

In the above display, multiplying the parentheses, defining $|I| \equiv \int_I \mathrm{d}s$, and rearranging yields[29]

$$\int_I \alpha\left(s\right) \pi\left(s\right) \mathrm{d}s < \frac{1}{|I|} \int_I \alpha\left(s\right) \mathrm{d}s \int_I \pi\left(s\right) \mathrm{d}s,$$

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:[30] Suppose that α and π are co-monotone. By contradiction, suppose that, with probability density p, prospect $\left(\alpha\left(s\right), \pi\left(s\right)\right)$ occurs and induces a message m. Suppose also that, with probability density p', prospect $\left(\alpha\left(s'\right), \pi\left(s'\right)\right)$ with $s' \neq s$ occurs and induces the same message m.[31] The seller's gain from pooling the two prospects under message m relative to revealing each of them is

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

where the inequality is by co-monotonicity. Thus, a weak improvement can be attained by revealing any two prospects that are sometimes pooled under the same message; full disclosure is optimal. □

For the second lemma, 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.*

[29] The obtained inequality is a continuous version of Chebyshev's sum inequality.

[30] The argument is as in Lemma [1] of Rayo and Segal (2010) and is included here for completeness.

[31] The message m may be also induced by some other prospect. Either prospect may also induce some other message.

Proof. For sufficiency, suppose first that S lies on a vertical line. Then, the seller's payoff from S, $\mathbb{E}\left[\mathbb{E}_{|m}[\alpha]\,\mathbb{E}_{|m}[\pi]\right] = \pi\mathbb{E}[\alpha]$, 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 $(\pi(\theta_1), \alpha(\theta_1))$ in S can be written as $\alpha(\theta_1) = k_0 - k_1\pi(\theta_1)$. The seller's payoff from S,

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

is maximized when $\mathbb{E}\left[\left(\mathbb{E}_{|m}[\pi]\right)^2\right]$ 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. 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 Rayo and Segal (2010, Lemma 3, p. 960). $\square$

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 prove Theorem 3.

Proof. (of Theorem 3) Recall that

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

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 co-monotone, 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 well-known observations about analytic functions and analytic curves.[32]

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

Observation 2. If two analytic curves coincide on an open interval, the two curves are identical.

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. $\square$

4.2 The Optimality of Conjugate Disclosure

Under an additional regularity assumption, this section derives the structure of an optimal information-disclosure rule. This section is more technical than the rest of the paper due to the subtleties stemming from the assumption that the leader's type space is continuous.[33] Hence, the section begins by stating the main result, whose proof is deferred until later in the section.

Main Result

Theorem 4 refers to a "conjugate disclosure rule," which, roughly speaking, 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. This partition has a special structure, roughly pooling "extreme" leader types into pairs. The two types that are pooled in a pair are called "conjugate," hence the name of the disclosure rule. The theorem's conclusion holds under a regularity condition, Condition 2, described shortly.

Theorem 4. *Under Conditions 1 and 2, the seller's disclosure problem has a solution such that*

(i) a conjugate disclosure rule is optimal;

[32] A curve is analytic if it has an analytic parametrization.

[33] Assuming instead that Θ_1 is finite would have simplified some arguments, complicated others, and led to weaker conclusions. In particular, with a finite but large Θ_1, the optimal disclosure rule fails to be conjugate, but this failure is economically insignificant. This insignificance can be ascertained by considering an increasing sequence of type sets, which the remainder of this section does. The conjugate disclosure rule is easy to interpret economically and leads to a one-dimensional optimal-control problem for finding an optimal disclosure rule.

(ii) the follower's induced effort schedule a is maximized at an s^* that satisfies $s^* \geq \theta^*$ and $a^*(s^*) \leq \alpha(\theta^*)$, where θ^* is the maximizer of the the first-best effort schedule, α. Moreover, the follower's unconditional expected optimal and expected efficient efforts are the same, that is, $\mathbb{E}[a^*(m)] = \mathbb{E}[\alpha(\theta_1)]$.

According to part (i) of Theorem 4, 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 (Crawford and Sobel, 1982). According to part (ii) of the theorem, the seller's strategic information disclosure distorts the follower's effort schedule by shifting its peak to the right relative to the first-best effort schedule. This finding is consistent with the intuition that the profit-maximizing effort schedule is a compromise between the first-best effort schedule and the unconstrained-optimal effort schedule of part (ii) of Theorem 2.

Corollary 3 shows that the allocation induced by the optimal disclosure rule can be implemented in a second-price auction with taxes. In particular, by part (i) of Theorem 4, any leader's type θ_1 deterministically induces a message, denoted by $m(\theta_1)$. This message induces the follower to take a uniquely optimal action, denoted by $a^*(m(\theta_1))$. Let the bidder who bids b be taxed in the amount

$$T^*(b) \equiv \int_0^b \left(F(s \mid a^*(m(b))) - F(s \mid a^*(m(s)))\right) \mathrm{d}s. \tag{16}$$

Then,

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

1. asks the leader to submit a bid, denoted by b, and charges him the tax $T^*(b)$;

2. discloses to the follower the message prescribed by the optimal conjugate disclosure rule and invites him to bid in the second-price auction;

3. allocates the item according to the rules of the second-price auction.

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

Proof. See Appendix A.2. □

Remark 1. One can modify the auction so that the leader pays the tax if and only if he wins the item. For this modification, replace $T^*(b)$ with $T^*(b)/F(b \mid a^*(m(b)))$.

By contrast to the efficient mechanism of Corollary 1, the optimal mechanism of Corollary 3 fails to disclose the leader's bid and specifies a different tax schedule. Analogous to the efficient case, the sign of $\mathrm{d}T^*(b)/\mathrm{d}b$ is the same as the sign of $\mathrm{d}a^*(m(b))/\mathrm{d}b$; the tax counteracts the leader's tendency to manipulate his bid to induce the follower to acquire more information. Because $T^*(b) = T^*(b')$ does not imply $a^*(m(b)) = a^*(m(b'))$, the seller generally cannot disclose to the follower the leader's tax $T^*(b)$ in lieu of the of the message required by an optimal disclosure policy.

The remainder of this section is concerned with proving Theorem 4.

Restriction to "Regular" Prospect Sets

Recall that the set of prospects is $\Gamma \equiv \{(\alpha(\theta_1), \pi(\theta_1)) : \theta_1 \in \Theta_1\}$. This paper's techniques for finding an optimal disclosure rule rely on Γ being regular in the sense of

Definition 1. A prospect set is **regular** if no three prospects lie on the same weakly decreasing line. Equivalently, a prospect set Γ is regular if no segment of Γ is vertical, and if every downward-sloping segment of Γ is nonlinear and lies on the boundary of the convex hull of Γ.[34]

Regular prospect sets are illustrated in Figure 3.[35] Because Condition 1 does not imply (nor is implied by) the regularity of the prospect set, the subsequent analysis will maintain an additional assumption:

Condition 2. The prospect set is regular.

Conditions 2 and 1 imply the existence of $\underline{\theta} \in [0, \theta^*)$ and $\bar{\theta} \in (\theta^*, 1)$ (marked in Figure 3) such that:[36]

- On $(0, \underline{\theta})$, Γ is downward sloping (α is strictly increasing; π is strictly decreasing) and lies on the boundary of the convex hull of Γ.
- On $(\underline{\theta}, \theta^*)$, Γ is upward sloping (both α and π are strictly increasing).
- On $(\theta^*, \bar{\theta})$, Γ is downward sloping (α is strictly decreasing; π is strictly increasing) and lies on the boundary of the convex hull of Γ.

[34] Analytically, requiring that no three prospects lie on the same weakly decreasing line is equivalent to requiring that for every $s \in \Theta_1$ such that $\alpha'(s)\,\pi'(s) < 0$, $(\pi'(s)/\alpha'(s))' > 0$, where primes denote derivatives, assuming that they exist.

[35] The functions α and π underlying the prospect sets in Figures 3c and 3d are plotted in Figure 6.

[36] When Condition 1 is not imposed, θ^* is defined as the unique maximizer of α, the first-best effort.

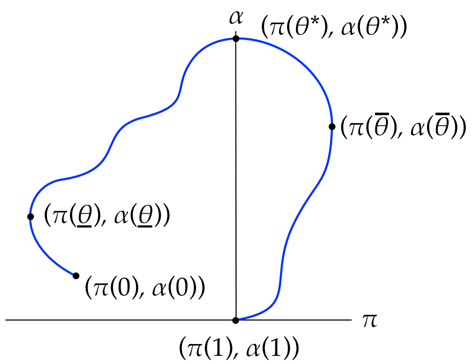

(a) A "generic" regular prospect set. Condition 1 fails because $\alpha(0) > \alpha(1)$.

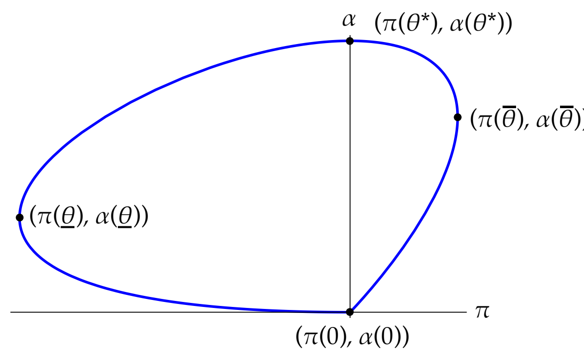

(b) A regular prospect set that satisfies Condition 1.

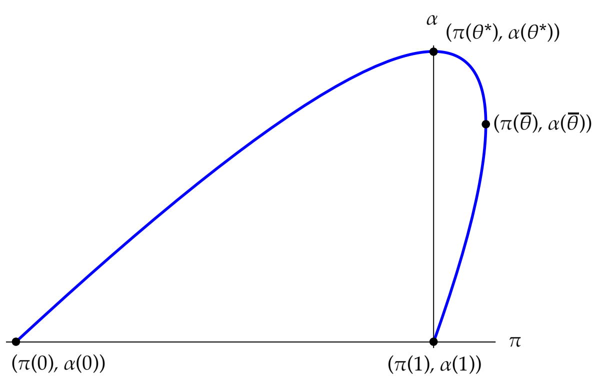

(c) Example 1 with uniform G.

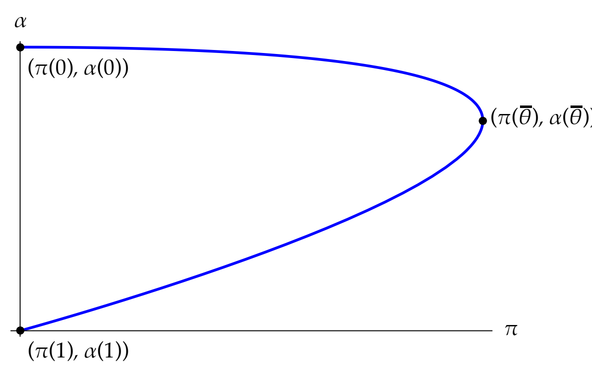

(d) Example 2 with uniform G.

Figure 3: Regular prospect sets. Certain "critical" points have been marked, to be referenced in subsequent analysis. An increase in θ_1 corresponds to the clockwise movement along the prospect set.

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

The partition of Γ into described segments relies on α being single-peaked (which is implied by Condition 1, as stated in 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, as is the requirement for any downward-sloping segment of Γ to be nonlinear and to lie on the boundary of the convex hull of Γ. The justification for Condition 2 is that it makes the analysis tractable while being satisfied for many common distributions G and F.

The restrictions on G and F that are sufficient for the regularity of a prospect set are listed below.

Lemma 3. *A prospect set induced by α and π is regular if*
(i) the inverse hazard rate $(1-G)/g$ is weakly decreasing and weakly concave;
(ii) Condition 1 holds;.
(iii) the (signed) Arrow-Pratt curvature of the follower's first-best effort, $-\alpha''/\alpha'$, is strictly increasing.

Proof. See Supplementary Appendix B. □

Part (i) of Lemma 3 requires the leader to be "strong," likely to have high valuations. The condition holds, for example, for uniform and truncated exponential distributions. The condition combines the standard monotonicity of the hazard rate with an additional concavity requirement.

Part (ii) of Lemma 3 ensures that the first-best effort schedule peaks at θ^* (Theorem 1). Then, part (iii) of the lemma can be interpreted as requiring that the first-best effort schedule become "more concave" as θ_1 increases and approaches θ^*, and "more convex" once θ^* has been passed. A "bell-shaped" effort schedule satisfies this condition. Because the first-best effort coincides in the model with the social planner's marginal benefit (gross of the effort cost) from an increase in the follower's effort, this condition can also be interpreted as describing the dependence of the planner's return to the follower's effort on the leader's type.[37]

A Discretized Prospect Set

The analysis will draw on the existing optimal-disclosure results developed for a discrete prospect set and then extend these results to the continuous set Γ by taking an appropriate limit. The

[37]In terms of restrictions on the c.d.f. F, part (iii) of Lemma 3 is equivalent to $((f_L - f_H)/(F_H - F_L))' > 0$.

discrete prospect set, denoted by Γ^n, is induced by the n-th finite approximation of the leader's type space. 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 $y_i = i/2^n$, $i \in \{1, 2, 3, .., 2^n\}$. 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 $(y_{i-1} y_i] \subset \Theta_1$ under the c.d.f. G. 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 $\Gamma^n \equiv \{(\pi(y), \alpha(y)) : y \in \Theta_1^n\}$. A 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 is denoted by $\mathcal{P}_n$ and is a special case of the problem studied by RS (Rayo and Segal, 2010), whose results we shall use. Call a prospect **revealed** if it induces a message that causes the follower to assign probability one to this prospect. Call two prospects **pooled** if they sometimes induce the seller to send the same message. Graphically, this shared message can be represented by a **pooling link**, a line segment that connects two pooled prospects in the graph Γ^n. If the prospect set Γ^n is obtained from Γ that satisfies Condition 2, the following partial characterization of a disclosure rule that is optimal for $\mathcal{P}_n$ obtains.

Fact 1. *By RS's Lemma [1], no two prospects are pooled if (i) both lie in* $\Gamma^n \cap \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in [\underline{\theta}, \theta^*]\}$ *(in Figure 3, the segment of* Γ *connecting prospects* $(\pi(\underline{\theta}), \alpha(\underline{\theta}))$ *and* $(\pi(\theta^*), \alpha(\theta^*))$*) or (ii) both lie in* $\Gamma^n \cap \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in [\bar{\theta}, 1]\}$ *(in Figure 3, the segment of* Γ *connecting prospects* $(\pi(\bar{\theta}), \alpha(\bar{\theta}))$ *and* $(\pi(1), \alpha(1))$*).*

Fact 2. *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 line, by Condition 2.*[38]

Fact 3. *By RS's Lemma [4], no two pooling links intersect.*

Fact 4. *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 an optimal disclosure rule in $\mathcal{P}_n$ is further refined by Lemma 4, which exploits the special structure of the seller's problem. The lemma identifies an s_* and an

[38] Lemma [3] does not rule out the situation in which a prospect probabilistically invokes one of multiple messages.

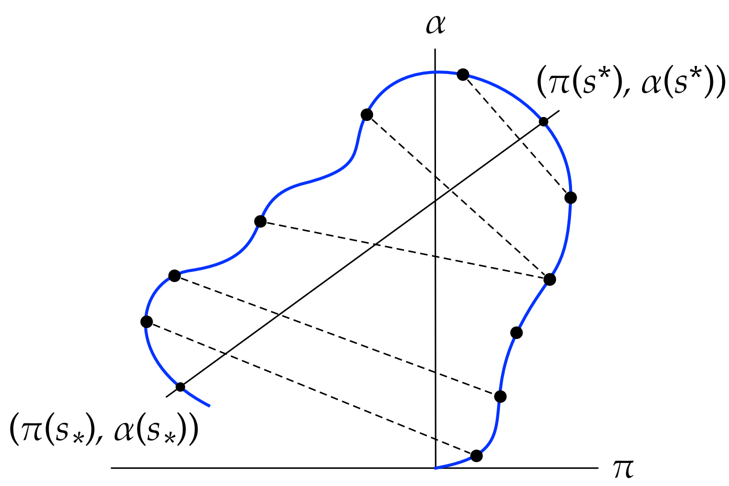

Figure 4: The discretized prospect set Γ^n is the union of solid dots. Without loss of generality, the seller can restrict attention to disclosure rules such as the one illustrated here. Each dashed link denotes a message that pools two prospects. These links never intersect and are oriented so that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

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. Moreover, the existence of such an s^* implies that the optimal effort schedule is single-peaked in θ_1, with the peak that is to the right of the peak of the first-best effort schedule.

Lemma 4. *Suppose that Condition 2 holds. Any discrete disclosure problem $\mathcal{P}_n$ has an optimal disclosure rule that is partially characterized by an $s_* \in [0, \theta^*]$ 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 $([0, s_*] \cup [s^*, 1]) \cap \Theta_1^n$, and, 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 the pooling links never intersect.

(ii) The optimal effort is single-peaked and is maximized at type s^* if $s^* \in \Theta_1^n$ (and, in general, is maximized at either type $\max\{[0, s^*] \cap \Theta_1^n\}$ *or type* $\min\{[s^*, 1] \cap \Theta_1^n\}$*).*

Proof. See Appendix A.3. □

By Lemma 4, a disclosure rule that is optimal in $\mathcal{P}_n$ can be represented by a matrix $\mathbf{p}_n \equiv [p_{ij}]_{i,j \in \{1,..,2^n\}}$, where 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 revealed is denoted by p_{ii}. The prospect indices that correspond to threshold types s_* and s^* defined in Lemma 4 are denoted by

$$n_* \equiv \min \{i : y_i \geq s_*\} \tag{17}$$

and

$$n^* \equiv \max \{i : y_i \leq s^*\} . \tag{18}$$

In this notation, the seller's maximal payoff in $\mathcal{P}_n$ is

$$V_n \equiv \sum_{i=1}^{2^n} p_{ii}\pi(y_i)\alpha(y_i) + \sum_{i=n_*}^{n^*} \sum_{j\in\{1,..,n_*-1\}\cup\{n^*+1,..,2^n\}} (p_{ij}+p_{ji}) \frac{p_{ij}\pi(y_i)+p_{ji}\pi(y_j)}{p_{ij}+p_{ji}} \frac{p_{ij}\alpha(y_i)+p_{ji}\alpha(y_j)}{p_{ij}+p_{ji}}, \tag{19}$$

where the first term is the payoff from revealed prospects, and the second term is the payoff from pooled prospects.

Implications for Optimal Disclosure with the Continuous Prospect Set

Lemma 4, which describes optimal disclosure in the discrete problem, will be used to suggest an optimal disclosure in the continuous problem, denoted by $\mathcal{P}$. The argument, which is technical, will justify the seller's focus on conjugate disclosure rules, each of which is induced by a matching function:

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

Case (i) in Definition 2 prevails in some natural examples, such as Example 1, discussed in Section 6. Case (ii) reduces to case (i) when $s_* = s' = 0$. 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 $(s^*, 1)$.

A matching function τ induces a conjugate disclosure rule:

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 $\theta_1 = s$. Intuitively, when $\tau'(s) = 0$, the seller pools a continuum of types near s with a single type $\tau(s)$. Finally, when the range of τ excludes some types, the seller reveals each of these types, as is explicitly specified in Definition 3.

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

Lemma 5. *A conjugate disclosure rule exists that delivers to the seller payoff $V_n + O(2^{-n})$ in the continuous disclosure problem $\mathcal{P}$, where V_n, defined in (19), is the seller's optimal payoff in the discrete disclosure problem $\mathcal{P}_n$.*

Proof. See Appendix A.4. □

Lemma 5 does not rule out the possibility that, in $\mathcal{P}$, the seller can improve upon the conjugate disclosure rule. Lemma 6 addresses this issues and shows 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 a value denoted by V^*. The discrete disclosure problems in the sequence $\{\mathcal{P}_n\}$ have solutions, which induce a corresponding sequence of values denoted by $\{V_n\}$. Then, $V^* \leq \liminf_{n\to\infty} V_n$.*

Proof. See Appendix A.5. □

Proof of the Main Result (Theorem 4)

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

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 10), $a^*(s^*) \leq \max_{\theta_1} \alpha(\theta_1) = \alpha(\theta^*)$, 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 $\mathbb{E}[a^*(m)] = \mathbb{E}\left[\mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]\right] = \mathbb{E}[\alpha(\theta_1)]$, where the last equality is by the Law of Iterated Expectations, hence part (ii) of the theorem.

4.3 The Seller's Optimal-Control Problem

By Theorem 4, the seller can restrict attention to disclosure rules induced by matching functions. Given this restriction, the seller's problem reduces to an optimal-control problem that takes the form of a consumption-saving problem (with the variables reinterpreted appropriately). This problem can be analyzed numerically.

The focus will be on the examples in which, merely by inspecting the prospect set, one can conclude that $s_* = 0$. (Such examples are studied in Section 6.) Then, the matching function is of type (i) of Definition 2, and the seller's objective function, (13), can be written as

$$\int_0^{s^*} \pi(s)\, a^*(s)\, g(s)\, \mathrm{d}s + \int_0^{s^*} \pi(\tau(s))\, a^*(s)\, g(\tau(s))\, \beta(s)\, \mathrm{d}s + \int_{[s^*,1]\setminus \mathrm{range}(\tau)} \alpha(s)\, \pi(s)\, \mathrm{d}s, \quad (20)$$

where $\beta \equiv -\tau'$ is the derivative of a matching function. The first integral in (20) 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 prospects in $[s^*, 1]$ that are pooled with the prospects in $[0, s^*]$ by the matching function. The third integral is the seller's payoff from the prospects in $[s^*, 1]$ that are not pooled with the prospects in $[0, s^*]$; such prospects are revealed.

The second integral in (20) uses the change-of-variables formula. It says that for any $s \in (0, s^*)$ and a "small" $\mathrm{d}s > 0$, interval $(s, s + \mathrm{d}s)$ (whose probability is approximately $g(s)\,\mathrm{d}s$) is pooled with interval $(\tau(s + \mathrm{d}s), \tau(s))$ (whose probability is approximately $g(\tau(s))\,\beta(s)\,\mathrm{d}s$).[39] The change-of-variables formula and Bayes's rule imply that the follower's equilibrium effort (10) takes the special form:

$$a^*(s) \quad = \quad \frac{g(s)\,\alpha(s) + g(\tau(s))\,\beta(s)\,\alpha(\tau(s))}{g(s) + g(\tau(s))\,\beta(s)}, \qquad s \in [0, s^*]. \quad (21)$$

[39]Indeed, the Taylor expansion implies $G(\tau(s) + \mathrm{d}s) \approx G(\tau(s)) + G'(\tau(s))\,\tau'(s)\,\mathrm{d}s$.

One can now formulate the seller's optimal-control problem:

Definition 4. When $s_* = 0$, the **seller's optimal-control problem** is defined as the maximization of (20) subject to (21) with respect to s^* and a piecewise-continuous function $\beta : [0, s^*] \to \mathbb{R}_+$ such that a piecewise-differentiable function $\tau : [0, s^*] \to [s^*, 1]$, where $\tau'(s) = -\beta(s)$ for almost all $s \in (0, s^*)$, satisfies boundary conditions $\tau(0) = 1$ and $\tau(s^*) = s^*$.

5 Generalizations

For the clarity of exposition, the analysis has focused on the case in which the follower's effort translates into information acquisition as described by Condition 1. Some of the conclusions continue to hold for effort technologies that are either more general than or alternative to the one described in Condition 1. In particular, the following generalization permits additional information-acquisition technologies, which can also be combined with (or replaced by) valuation-enhancing investments.

Condition 3 (Productive Effort)**.** For any efforts a' and a with $a' > a$,

$$\int_{\theta_2}^{1} F\left(s \mid a'\right) \mathrm{d}s \le \int_{\theta_2}^{1} F\left(s \mid a\right) \mathrm{d}s \qquad \text{for all } \theta_2 \in \Theta_2.$$

In the spirit of textbook results on the measures of risk (Rothschild and Stiglitz, 1970; Mas-Colell et al., 1995, Chapter 6D), Condition 3 can be shown to be equivalent to the following claim. When $a' > a$, a type drawn according to $F(\cdot \mid a')$ can be obtained from a type drawn according to $F(\cdot \mid a)$ by adding a mean-preserving spread and a first-order stochastic-dominance shift. The mean-preserving spread corresponds to the acquisition of more precise information about the underlying valuation. The first-order stochastic-dominance shift corresponds to investment in enhancing the underlying valuation. In applications, this investment can be interpreted as the acquisition of a good that is complementary to the auctioned item or as investment in a skill that is necessary for operating (or consuming) the auctioned item.

The effort in Condition 3 deserves the label "productive" because both the follower and the surplus-minded planner value it. The planner values the follower's effort because Condition 3 ensures that the first-best effort in (6) is nonnegative. The follower values the effort because his payoff is weakly convex (by incentive compatibility) and weakly increasing in his type, θ_2, and Condition 3 is the necessary and sufficient restriction on F for which the follower's expected payoff is increasing in a for any payoff function that is weakly convex and weakly increasing.[40] In this sense, Condition 3 is the most general formulation of a productive investment. Economically, both the follower and the planner benefit from more precise information about the follower's valuation and from a stochastically higher follower's valuation.

A special case of Condition 3 generalizes Condition 1 to capture information acquisition without requiring the c.d.f.'s linearity in effort:

Condition 4 (Generalized Information Acquisition)**.** For some $\theta^* \in (0,1)$, for all $a, a' \in (0,1)$, (i) (equal means) $\int_0^1 F(s \mid a)\,\mathrm{d}s = \int_0^1 F(s \mid a')\,\mathrm{d}s$, and (ii) (rotation) $s \in (0,\theta^*) \cup (\theta^*, 1) \implies (\theta^* - \theta)\,\partial F(s \mid a)/\partial a > 0$.[41]

Another special case of Condition 3 corresponds to investment in stochastically increasing the valuation:

Condition 5 (Enhancing Valuation)**.** For any efforts a' and a with $a' > a$,

$$F\left(\theta_2 \mid a'\right) \leq F\left(\theta_2 \mid a\right) \qquad \text{for all } \theta_2 \in \Theta_2.$$

The appropriate generalization of the interiority condition (2) is

$$c > \sup_{a \in A, \theta_1 \in \Theta_1} \left\{ -\int_{\theta_1}^{1} \frac{\partial F(s \mid a)}{\partial a} \mathrm{d}s \right\}, \tag{22}$$

which is henceforth assumed to hold.

Assuming Condition 3, but not assuming F's linearity in effort, one can show that:

- If Condition 1 is replaced by Condition 4, the analysis and most of the results of Section 3 apply unchanged (with references to parts of Condition 1 replaced with references to parts of Condition 4). The only difference is that no explicit expression for the first-best effort is available.

[40] The argument is again analogous to the arguments of Rothschild and Stiglitz (1970, Chapter 6D) and Mas-Colell et al. (1995, Chapter 6D).

[41] Without part (i), Condition 4 is the rotation order in Definition 1 of Johnson and Myatt (2006), who build on earlier notions of "simply intertwined" and "satisfying the single-crossing property" random variables. Shi (2012) uses Condition 4 and gives examples of information-acquisition technologies that induce the order in this condition.

- Without Condition 4, the first-best effort α need not be hump shaped. For instance, under Condition 5, α is weakly decreasing.

- The first-best outcome is implementable if and only if the leader's expected allocation, $F(\theta_1 \mid \alpha(\theta_1))$, is weakly increasing in θ_1. Whenever the first-best outcome is implementable, it is implemented by the mechanism of Corollary 1.

- As in Theorem 2, if the seller observes the leader's type, any disclosure rule is optimal. As in Theorem 2, if the seller directly chooses the follower's effort, the seller chooses that effort that minimizes the probability that the leader gets the item. Under Condition 5, this probability is minimized by the largest feasible effort, independently of the leader's type.

Assuming Condition 3 and assuming that the c.d.f. F is linear in effort, one can show that:

- The first-best effort is still given by (6) in Theorem 1. The leader's expected allocation, $F(\theta_1 \mid \alpha(\theta_1))$, is weakly increasing in θ_1, and hence the mechanism described in Corollary 1 implements the first-best outcome.[42]

- The conclusion of Theorem 3, establishing the suboptimality of extreme disclosure policies, may not apply.[43] Nevertheless, Theorem 3 can be strengthened. The suboptimality of full disclosure uses only part (ii) of Condition 1, as the theorem's proof makes clear. Part (ii) combined with Condition 3 is weaker than Condition 1, accommodates a mixture of information acquisition and investment in valuation, and delivers the suboptimality of full disclosure.

  The suboptimality of non-disclosure does not even rely on part (ii) and hence is quite general. Indeed,

$$\frac{\mathrm{d}\pi}{\mathrm{d}\alpha}(\theta_1) = \frac{\pi'(\theta_1)}{\alpha'(\theta_1)} = -c\left[\left(\frac{1-G(\theta_1)}{g(\theta_1)}\right)' + \frac{1-G(\theta_1)}{g(\theta_1)}\frac{f_H(\theta_1)-f_L(\theta_1)}{F_H(\theta_1)-F_L(\theta_1)}\right],$$

  where f_H and f_L are the p.d.f.s corresponding to the c.d.f.s F_H and F_L. By inspection of the above display, the right-hand side is independent of θ_1 only "nongenerically." Hence, non-

[42] When F has a density, the monotonicity of the leader's allocation is established by differentiating:

$$\frac{\mathrm{d}}{\mathrm{d}\theta_1}F(\theta_1 \mid \alpha(\theta_1)) = f(\theta_1 \mid \alpha(\theta_1)) + \frac{1}{c}\left(F_H(\theta_1) - F_L(\theta_1)\right)^2 > 0.$$

[43] Indeed, let $F_L \equiv 1$, let F_H be any c.d.f. (so Condition 5 holds), and let G have a weakly increasing hazard rate. Then, both α and π are weakly decreasing and hence co-monotone. By Lemma 1, full disclosure is optimal.

disclosure (or indeed, merely pooling an open interval of the leader's types under the same message) can be optimal only "nongenerically" because it requires $\mathrm{d}\pi/\mathrm{d}\alpha$ to be constant.

- The conclusion of Theorem 4 still holds; that is, an optimal disclosure rule can be summarized by a matching function. The theorem relies on the prospect set's regularity, which is ensured by a version of Lemma 3 that has Condition 5 instead of Condition 1 in part (ii).

6 Optimal Disclosure in Examples

The focus is on Examples 1 and 2. Example 1 satisfies Condition 1 and models pure information acquisition. Example 2 satisfies Condition 5 and models pure valuation enhancement. In both examples, G is uniform. The disclosure rule is independent of c; the induced effort varies proportionally to $1/c$, which only scales the follower's effort (provided $c \geq 1/6$, which is a maintained assumption, by (22)).

Numerical "solutions" for the seller's problems are identified by applying Hamiltonian techniques to the optimal-control problem in Definition 4. Because the Hamiltonian analysis imposes the additional technical assumptions 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 presented numerical "solutions" are informed guesses, which have been verified to improve upon full disclosure and non-disclosure.[44] All solutions have also been verified to satisfy the monotonicity condition (i.e., that $F\left(\theta_1 \mid a^*\left(\theta_1\right)\right)$ is nondecreasing in θ_1).

For each example, Figure 5 plots a selection from the set of pooling links in an optimal mechanism. These links are derived from the 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 $\tau\left(\theta_1\right)$, and induces the follower's effort $a^*\left(\theta_1\right)$, 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**.

[44]An analogous caveat applies to this section's references to "optimality." In the remainder, the quotation marks are suppressed.

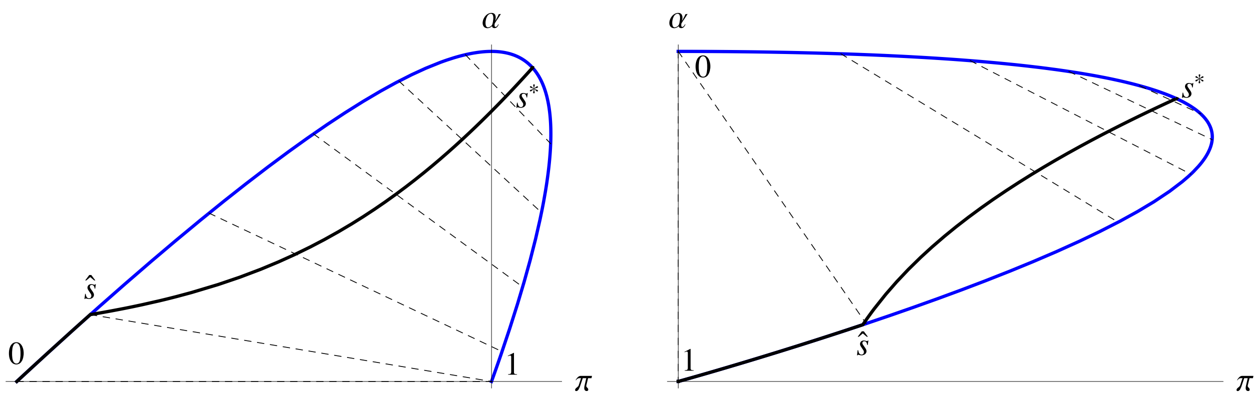

(a) Example 1. The leader's types in $[0, \hat{s}]$ are revealed; the rest are pooled according to a matching function; $\hat{s} = 0.054$ and $s^* = 0.61$.

(b) Example 2. The leader's types in $[\hat{s}, 1]$ are revealed; the rest are pooled according to a matching function; $\hat{s} = 0.74$ and $s^* = 0.24$.

Figure 5: An optimal disclosure rule. The solid blue arc is the prospect set, with four prospects labeled by the leader's types that induce them. Each point on the solid black curve (optimal-prospect path) is a tuple containing the follower's optimal action and the seller's associated expected marginal benefit from that action. The tuple is induced by pooling the prospects at the endpoints of the dashed link that passes through that tuple. The seller's types that are revealed lie at the intersection of the solid black curve and the prospect set.

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. This monotonicity of the optimal-prospect path implies, in particular, that in Example 2 (Figure 5b), no full disclosure of prospects induced by $\theta_1 < s^*$ can occur, because the prospect set is strictly decreasing when $\theta_1 < s^*$. Whenever a prospect is revealed, however, the optimal-prospect path must coincide with the prospect set.

For each example, Figure 6 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 4, doing so involves inducing a "rightward shift" in the follower's effort schedule, relative to the first-best effort schedule.

7 Conclusions

This paper's primary purpose is to isolate, in a simplest model, the distortions that the seller's strategic bid-disclosure introduces into an otherwise efficient auction in the presence of moral

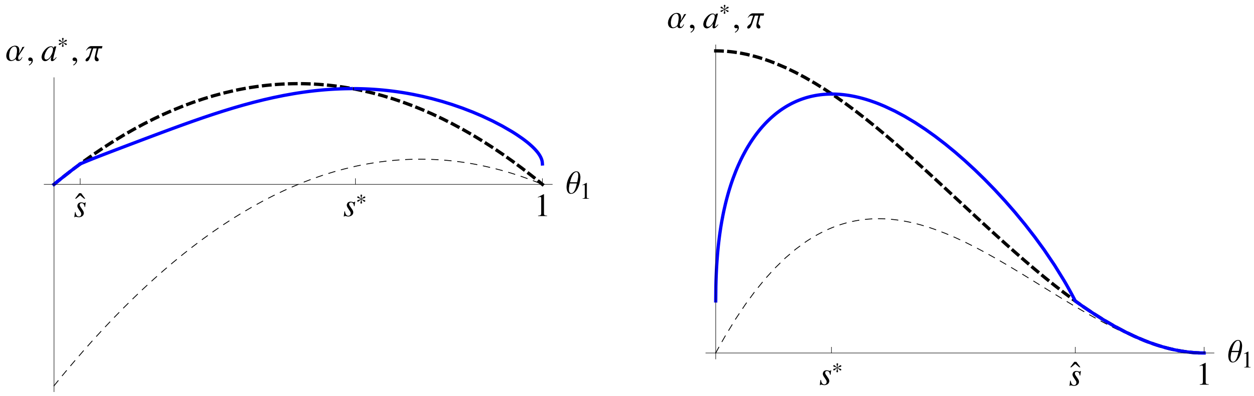

(a) Example 1. (b) Example 2.

Figure 6: An optimal effort schedule. The solid curve is the optimal effort, a^*. 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 align a^* with π. The areas under the solid and the dashed thick curves coincide; that is, $\mathbb{E}\left[a^*\left(m\right)\right] = \mathbb{E}\left[\alpha\left(\theta_1\right)\right]$.

hazard. In particular, the profit-maximizing seller induces the initially uninformed bidder to acquire inefficiently little information when the initially informed bidder's valuation is low, and to acquire inefficiently much information otherwise. Pooling pairs of extreme bids under the same message accomplishes this distortion.

A natural question is whether such disclosure is ever used in practice. As mentioned in the Introduction, auctions of rail passenger service franchises fit the model's assumptions and may feature a disclosure rule similar to the optimal rule derived in this paper. In particular, the incumbent's request for a subsidy before the auction can be interpreted as a signal that pools two extreme types: a productive incumbent who seeks to expand the capacity and an unproductive incumbent who needs cash to remain in business. Even though the public subsidy request is unlikely to have been optimally designed, the practice of publicly revealing all such requests may in part be motivated by the optimality of such a disclosure policy.

Another example of pooling the extremes is the now-defunct practice at the University of Oxford of awarding an "alpha-gamma" grade to a student who is sometimes outstanding ("alpha"), sometimes abysmal ("gamma"), but never mediocre ("beta"). This grading scheme can be rationalized by the seller's disclosure problem (in its canonical form) if one interprets (i) α as a graduate's productivity with an employer, (ii) $\mathbb{E}_{|m}\left[\alpha\right]$ as the probability that the employer hires a graduate whose grade is m, (iii) π as an employed graduate's donation (in cash or in service) to the university, which for each ability α, can be either high or low, depending on the amount of credit the graduate gives to Oxford for his success, and (iv) $\mathbb{E}_m\left[\mathbb{E}_{|m}[\pi]\,\mathbb{E}_{|m}[\alpha]\right]$ as the expected donation, which Oxford's grading scheme is designed to maximize. A relationship between α and π that is consistent with the grading-scheme application is in Figure 3c.

Even though strategic information disclosure is pervasive in private transactions mediated by investment banks and realtors, it is rare in government-run procurement auctions, in which as a part of the formal tender process, the government ought to be able to commit to a particular disclosure rule. Our analysis suggests that the government may refrain from strategic disclosure if any of the following conditions hold: (i) the government seeks the first-best outcome, instead of maximizing its profit; (ii) the government cannot commit to an inefficient disclosure rule; (iii) the government knows the incumbent's valuation; or (iv) an optimal disclosure rule is too complex to compute and too hard to explain to bidders.

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. Given this bidding strategy, it is optimal for him to exert the first-best effort because the braced expression in the surplus-maximization problem (4) coincides with the follower's expected payoff in the mechanism in the corollary. The follower's payoff is nonnegative because the payoff in the second-price auction is always nonnegative.

The leader chooses b to maximize

$$\int_{\Theta_2} \mathbf{1}_{\{b>\theta_2\}}\left(\theta_1-\theta_2\right)\mathrm{d}F\left(\theta_2 \mid \alpha\left(b\right)\right)-T\left(b\right).$$

Integration by parts and the substitution of T from (7) transforms the above display into

$$\int_0^{\theta_1} F\left(s \mid \alpha\left(s\right)\right)\mathrm{d}s+\int_{\theta_1}^{b}\left(F\left(s \mid \alpha\left(s\right)\right)-F\left(b \mid \alpha\left(b\right)\right)\right)\mathrm{d}s,$$

which is maximized at $b=\theta_1$ because $F\left(s \mid \alpha\left(s\right)\right)$ is increasing in s. That $F\left(s \mid \alpha\left(s\right)\right)$ is increasing in s can be seen by letting $s' > s$ and writing

$$F\left(s' \mid \alpha\left(s'\right)\right) - F\left(s \mid \alpha\left(s\right)\right) = \left[F\left(s' \mid \alpha\left(s'\right)\right) - F\left(s \mid \alpha\left(s'\right)\right)\right] + \int_{\alpha(s)}^{\alpha(s')} \frac{\partial F\left(s \mid a\right)}{\partial a} \mathrm{d}a > 0,$$

where the inequality follows by observing that the bracketed term is nonnegative because F is a c.d.f.; pertaining to the integral, one of three cases prevails: (i) $s < s' \leq \theta^*$, implying $\alpha\left(s\right) < \alpha\left(s'\right)$ (Theorem 1) and $\partial F\left(s \mid a\right)/\partial a > 0$ (Condition 1), and so the integral is positive, (ii) $\theta^* \leq s < s'$, implying $\alpha\left(s\right) > \alpha\left(s'\right)$ (Theorem 1) and $\partial F\left(s \mid a\right)/\partial a < 0$ (Condition 1), and so the integral is positive, or (iii) $s < \theta^* < s'$, in which case the inequality follows by considering the change in the leader's type from s to θ^* and applying (i) and then considering the change from θ^* to s' and applying (ii).

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 his payoff from trading is always nonnegative.

The leader chooses his bid by solving

$$\max_{b} \left[\int_{\Theta_2} \mathbf{1}_{\{b>\theta_2\}} \left(\theta_1 - \theta_2\right) \mathrm{d}F\left(\theta_2 \mid a^*\left(m\left(b\right)\right)\right) - T^*\left(b\right)\right],$$

where $a^*\left(m\left(b\right)\right)$ is the follower's optimal action conditional on the message $m\left(b\right)$, which the seller sends when the leader bids b. Integration by parts and the substitution of T^* from (16) transforms the above maximand into

$$\int_0^{\theta_1} F\left(s \mid a^*\left(m\left(s\right)\right)\right) \mathrm{d}s + \int_{\theta_1}^{b} \left(F\left(s \mid a^*\left(m\left(s\right)\right)\right) - F\left(b \mid a^*\left(m\left(b\right)\right)\right)\right) \mathrm{d}s,$$

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

A.3 Proof of Lemma 4

To conclude part (i) of the lemma's statement from Facts 1–4, it remains to rule out the pooling patterns depicted in Figure 7. A single argument rules out both patterns.

By contradiction, suppose that an s_* and s^* with the sought properties do not exist; that is, one can pick four prospects (π_i, α_i), $i = 1, 2, 3, 4$, such that prospects (π_1, α_1) and (π_2, α_2) are pooled under some message, say, m, whereas prospects (π_3, α_3) and (π_4, α_4) are pooled under another message, say m', and these four prospects satisfy 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 $\alpha_1 > \alpha_2 \geq \alpha_3 \geq \alpha_4$.[45] Let p_i, $i = 1, 2, 3, 4$, denote the joint probability that type x_i (which induces prospect (π_i, α_i)) is realized and induces either message m or m' (whichever is appropriate).

The seller's expected gain from using distinct messages m and m' (as in Figure 7) instead of pooling all four prospects under a single message is

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

which can be rearranged to give

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

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 Figure 7 (both panels) are suboptimal. Part (i) of the lemma follows.

For part (ii), take any four leader's types $x_1, x_2, x_3, x_4 \in \Theta_1^n$ 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. For such pooling to be optimal, it must be that, in particular, 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

[45] Cases (a) and (b) differ in the placement of the strict and weak inequalities. Case (a) prevails if leader's types $x_1, x_2, x_3, x_4 \in \Theta_1^n$ satisfy $0 < x_1 < x_2 \leq x_3 < x_4 < 1$, $x_2 > \theta^*$, and $x_3 < \bar{\theta}$. Case (b) prevails if 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 $x_2 < \underline{\theta}$.

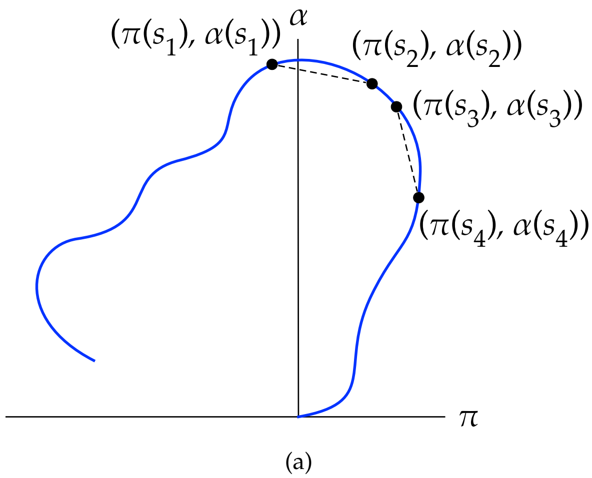

Figure 7: Contradiction hypotheses for the proof of Lemma 4. Each dashed link denotes a pair of prospects that are pooled under the same signal. In both panels, the two links cannot be crossed by an upward-sloping line.

reversed, or $\Delta \geq 0$.

Without loss of generality, let the segment connecting the prospects induced by points $\{x_1, x_2\}$ lie southwest of (or "under") the segment connecting the prospects induced by points $\{x_3, x_4\}$. Because the two segments cannot intersect, $\Delta \geq 0$ immediately implies

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

Consequently, under the optimal disclosure rule, the effort induced by a message monotonically increases as one moves in the northeast direction across the pooling links toward s^*, which establishes part (ii).

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. The value of problem $\mathcal{P}_n$, V_n, is achieved by a solution that is given by a matrix $\mathbf{p}_n$, 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 disclosure rule $\mathbf{p}_n$ will be used to construct an approximately optimal disclosure rule for $\mathcal{P}$. Roughly, if in $\mathcal{P}_n$, prospects i and j are pooled only with each other (i.e., $\Sigma_k p_{ik} = p_{ij}$

and $\Sigma_k p_{jk} = p_{ji}$), then in $\mathcal{P}$, the intervals $(y_{i-1}, y_i]$ and $\left(y_{j-1}, y_j\right]$ will be "linked" pointwise, by pooling every element in $(y_{i-1}, y_i]$ with a corresponding element in $\left(y_{j-1}, y_j\right]$ according to some matching function.[46] If in $\mathcal{P}_n$, prospects i and j are only sometimes pooled with each other (i.e., $\Sigma_k \left(p_{ik} + p_{jk}\right) > p_{ij} + p_{ji}$), then in $\mathcal{P}$, the intervals $(y_{i-1}, y_i]$ and $\left(y_{j-1}, y_j\right]$ are divided into subintervals and only one subinterval in $(y_{i-1}, y_i]$ is linked pointwise with a subinterval in $\left(y_{j-1}, y_j\right]$.

To make the linking procedure precise, define $P_i \equiv \left\{j : p_{ij} > 0\right\}$ to be the set of prospects that are optimally pooled with prospect i, $i \in \{1, .., 2^n\}$, in $\mathcal{P}_n$. Note that if $P_i = \{i\}$, prospect i is revealed in $\mathcal{P}_n$. Partition interval $(y_{i-1}, y_i]$ into a collection of $|P_i|$ subintervals[47]

$$\mathcal{C}_i \equiv \left\{ \left(\underline{b}_{ij}, \bar{b}_{ij}\right] : j \in P_i \right\} \tag{A.2}$$

so that $G\left(\bar{b}_{ij}\right) - G\left(\underline{b}_{ij}\right) = p_{ij}$, and so that whenever prospects i and j are pooled in $\mathcal{P}_n$, one can draw a link between element $\left(\underline{b}_{ij}, \bar{b}_{ij}\right]$ in $\mathcal{C}_i$ and element $\left(\underline{b}_{ji}, \bar{b}_{ji}\right]$ 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 links between (sub)intervals is illustrated in Figure 8. Thus linked subintervals form the basis for pooling prospects in $\mathcal{P}$.

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 $i = n_*$, where n_* is defined in (17):

1. If no subinterval in $\mathcal{C}_i$ is linked to any other subinterval, set $\tau\left(\theta_1\right) = \tau\left(y_{i-1}\right)$ for all $\theta_1 \in (y_{i-1}, y_i]$, with the convention that $\tau\left(y_{n_*-1}\right) = 1$ if $s_* = 0$ and $\tau\left(y_{n_*-1}\right) = y_{n_*-1}$ if $s_* > 0$.

2. If a subinterval $\left(\underline{b}_{ij}, \bar{b}_{ij}\right]$ in $\mathcal{C}_i$ is linked to some subinterval $\left(\underline{b}_{ji}, \bar{b}_{ji}\right]$ in $\mathcal{C}_j$, define a strictly decreasing τ so that for all $s \in \left[\underline{b}_{ij}, \bar{b}_{ij}\right]$, $g\left(s\right) / \left(\beta\left(s\right) g\left(\tau\left(s\right)\right)\right) = p_{ij} / p_{ji}$, and $\tau\left(\underline{b}_{ij}\right) = \bar{b}_{ji}$ and

[46] Recall 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.

[47] Here, $|P_i|$ denotes the number of elements in the set P_i.

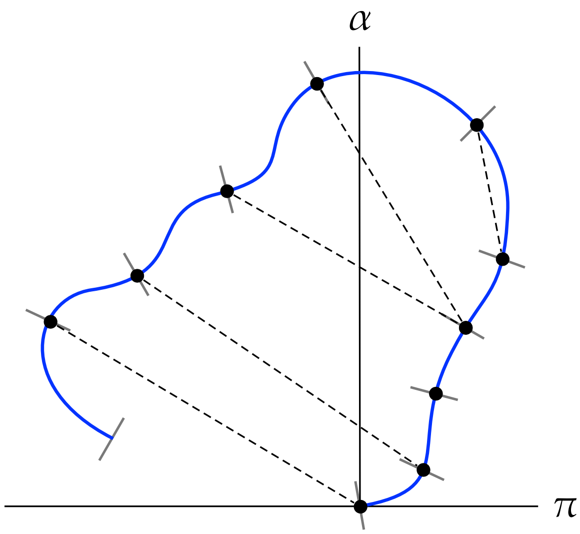

(a) Optimal disclosure in $\mathcal{P}_n$. The solid dots are prospects. The dashed links pool these prospects. The prospect that is not pooled is revealed.

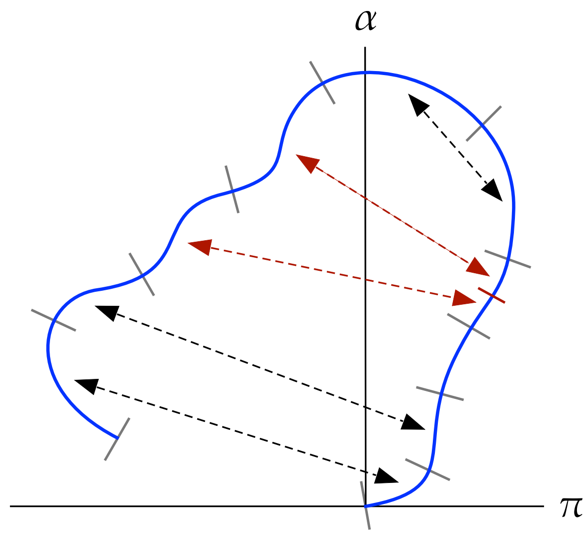

(b) Disclosure in $\mathcal{P}$ derived from disclosure in $\mathcal{P}_n$. Arrow-headed dashed segments indicate subintervals whose prospects are pooled pointwise. Each prospect in the interval that is not linked with any other interval is revealed.

Figure 8: An optimal disclosure rule in the discrete problem $\mathcal{P}_n$ is used to construct a disclosure rule in the continuous problem $\mathcal{P}$.

$\tau\left(\bar{b}_{ij}\right) = \underline{b}_{ji}$.[48] Otherwise, go to Step 3.

3. If $i < n^*$, where n^* is defined in (18), increment i by 1 and go to Step 1; otherwise, terminate.[49]

[48]When the c.d.f. G is uniform, the sought τ is linear: $\tau(s) = \bar{b}_{ji} - \left(s - \underline{b}_{ij}\right) p_{ji}/p_{ij}$, 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 $\tau' = -p_{ji}g(s)/\left(p_{ij}g(\tau)\right)$ on $\left[\underline{b}_{ij}, \bar{b}_{ij}\right]$ subject to $\tau\left(\underline{b}_{ij}\right) = \bar{b}_{ji}$. 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 $\tau\left(\bar{b}_{ij}\right) = \underline{b}_{ji}$, rewrite the ODE as $-g(\tau(s))\tau'(s) = g(s)p_{ji}/p_{ij}$ and integrate to obtain

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

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

$$\int_{\tau\left(\bar{b}_{ij}\right)}^{\bar{b}_{ji}} g(z)\,\mathrm{d}z = p_{ji}.$$

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

[49]The 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 intervals indexed by $i \in \{n_*, n_* + 1, .., n^*\}$ are considered.

Any interval $(y_{i-1}, y_i]$ whose elements are revealed contributes to the seller's payoff (15) amount

$$\int_{y_{i-1}}^{y_i} \pi(s)\,\alpha(s)\,g(s)\,\mathrm{d}s \equiv p_{ii}\pi(z_i)\,\alpha(z_i)\,, \tag{A.3}$$

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.[50]

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

$$\begin{aligned}
&\int_{\underline{b}_{ij}}^{\bar{b}_{ij}} \frac{g(s)\,\pi(s) + \beta(s)\,g(\tau(s))\,\pi(\tau(s))}{g(s) + \beta(s)\,g(\tau(s))} \frac{g(s)\,\alpha(s) + \beta(s)\,\alpha(\tau(s))}{g(s) + \beta(s)\,g(\tau(s))} \left(g(s) + \beta(s)\,g(\tau(s))\right) \mathrm{d}s \\
&\qquad \equiv \left(p_{ij} + p_{ji}\right) \frac{g(z_{ij})\,\pi(z_{ij}) + \beta(z_{ij})\,g(z_{ji})\,\pi(z_{ji})}{g(z_{ij}) + \beta(z_{ij})\,g(z_{ji})} \frac{g(z_{ij})\,\alpha(z_{ij}) + \beta(z_{ij})\,\alpha(z_{ji})}{g(z_{ij}) + \beta(z_{ij})\,g(z_{ji})} \\
&\qquad\qquad = \left(p_{ij} + p_{ji}\right) \frac{p_{ij}\pi(z_{ij}) + p_{ji}\pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij}\alpha(z_{ij}) + p_{ji}\alpha(z_{ji})}{p_{ij} + p_{ji}},
\end{aligned} \tag{A.4}$$

where the identity uses $p_{ij} = G(\bar{b}_{ij}) - G\left(\underline{b}_{ij}\right)$ and $p_{ji} = G(\bar{b}_{ji}) - G\left(\underline{b}_{ji}\right)$, and implicitly (and not necessarily uniquely) defines $z_{ij} \in \left(\underline{b}_{ij}, \bar{b}_{ij}\right)$ 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, $z_{ji} \in \left(\underline{b}_{ji}, \bar{b}_{ji}\right)$.

Assembling the contributions (A.3) and (A.4) gives the value of the seller's objective function (15) under the disclosure rule induced by τ:

$$\hat{V}_n \equiv \sum_{i=1}^{2^n} p_{ii}\pi(z_i)\,\alpha(z_i) + \sum_{i=n_*}^{n^*} \sum_{j \in \{1,..,n_*-1\} \cup \{n^*+1,..,2^n\}} \left(p_{ij} + p_{ji}\right) \frac{p_{ij}\pi(z_{ij}) + p_{ji}\pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij}\alpha(z_{ij}) + p_{ji}\alpha(z_{ji})}{p_{ij} + p_{ji}}. \tag{A.5}$$

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 $z \in (0,1)$:

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

[50]See http://en.wikipedia.org/wiki/Mean_value_theorem#First_mean_value_theorem_for_integration

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

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

which are substituted into (A.5) to obtain

$$\hat{V}_n = V_n + O\left(2^{-n}\right)$$

as desired.

A.5 Proof of Lemma 6

Preliminary Definitions

Normalize the set of messages to equal 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.[51] 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:

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

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).

[51]The topology of weak convergence is metrizable under the Prohorov metric.

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

$$V^* \equiv \max_{\hat{\nu} \in \Delta(\Delta\Theta_1)} \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\hat{\nu} \qquad \text{s.t.} \qquad \int_{\Delta\Theta_1} P \mathrm{d}\hat{\nu} = P^0, \tag{A.6}$$

where

$$\bar{\pi}(P) = \int_{\Theta_1} \pi(\theta_1)\, \mathrm{d}P \qquad \text{and} \qquad a^*(P) = \int_{\Theta_1} \alpha(\theta_1)\, \mathrm{d}P.$$

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 discontinuities 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.6) exists. Let ν^* denote this solution.

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

$$V_n \equiv \max_{\hat{\nu}_n \in \Delta\left(\Delta\Theta_1^n\right)} \int_{\Delta\Theta_1^n} \bar{\pi}(P_n)\, a^*(P_n)\, \mathrm{d}\hat{\nu}_n \qquad \text{s.t.} \qquad \int_{\Delta\Theta_1^n} P_n \mathrm{d}\hat{\nu}_n = P_n^0, \tag{A.7}$$

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:

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

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.7). 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 $\nu_n \in \Delta(\Delta\Theta_1^n) \subseteq \Delta(\Delta\Theta_1)$.

<u>The Proof of the Lemma</u>

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

$n' \to \infty$, $\nu_{n'}$ 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,

$$\int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu_{n'} \to \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu. \tag{A.8}$$

By contradiction, suppose that ν, the limit of $\{\nu_{n'}\}$, does not solve the continuous problem (A.6) and that:

$$\epsilon \equiv \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^* - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu > 0. \tag{A.9}$$

Let $\mathcal{N} \equiv \{n'_1, n'_2, ...\}$ be the set indexing the convergent sequence $\{\nu_{n'}\}$. Because of convergence (A.8), one can choose an $N \in \mathcal{N}$ such that for all $n' \geq N$, $n' \in \mathcal{N}$:

$$\left| \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu_{n'} - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu \right| < \frac{\epsilon}{2}. \tag{A.10}$$

Because space $\Delta(\Delta\Theta_1)$ is separable and because ν^* solves the seller's maximization problem (A.6) 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 $\nu^*_{n'}$ to ν^* with a support on $\Theta_1^{n'}$:

$$0 \leq \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^* - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^*_{n'} < \frac{\epsilon}{2}. \tag{A.11}$$

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

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

$$\begin{aligned}
&\int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu_{\bar{N}} - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^*_{\bar{N}} \\
&= \left(\int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu_{\bar{N}} - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu \right) - \left(\int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^* - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu \right) \\
&\qquad + \left(\int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^* - \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^*_{\bar{N}} \right) < \frac{\epsilon}{2} - \epsilon + \frac{\epsilon}{2} = 0,
\end{aligned}$$

where the term in the first parenthesis is less than $\epsilon/2$ by (A.10), the term in the second parenthesis equals ϵ by the contradiction hypothesis (A.9), and the term in the third parenthesis is less than $\epsilon/2$ by (A.11). The inequality is a contradiction, however, because $\nu_{\bar{N}}$ solves the seller's discrete maximization problem with the leader's type space $\Theta_1^{\bar{N}}$. Hence,

$$\int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu \geq \int_{\Delta\Theta_1} \bar{\pi}(P)\, a^*(P)\, \mathrm{d}\nu^* \equiv V^*.$$

The argument above shows that every convergent subsequence $\{\nu_{n'}\}$ converges weakly to a limit that delivers a payoff at least as high as V^*. Consequently, it must be the case that $\liminf_{n\to\infty} V_n \geq V^*$. If not, there exists an $\epsilon > 0$ such that infinitely many ν_n deliver payoff $V_n < V^* - \epsilon$. Then it is possible to pick a subsequence $\{\nu_{n_k}\}$ for which no term delivers payoff $V_{n_k} \geq V^* - \epsilon$. By sequential compactness, $\{\nu_{n_k}\}$ has a convergent subsequence, which is a subsequence of the original $\{\nu_n\}$, 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 $\{\nu_n\}$ must converge weakly to a limit that delivers at least V^*.

References

Amador, Manuel and Kyle Bagwell, "The Theory of Optimal Delegation With an Application to Tariff Caps," *Econometrica*, 2013, *81* (4), 1541–1599.

Bergemann, Dirk and Juuso Välimäki, "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 2002, *70* (3), 1007–1033.

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.

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.

Fudenberg, Drew and Jean Tirole, *Game Theory*, The MIT Press, 1991.

Ganuza, Juan-José and José Penalva, "Signal orderings based on dispersion and the supply of private information in auctions," *Econometrica*, 2010, *78* (3), 1007–1030.

Golosov, Mikhail, 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.

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.

Kovac, Eugen and Tymofiy Mylovanov, "Stochastic mechanisms in settings without monetary transfers: The regular case," *Journal of Economic Theory*, 2009, *144* (4), 1373–1395.

Krishna, Vijay and John Morgan, "Contracting for information under imperfect commitment," *The RAND Journal of Economics*, 2008, *39* (4), 905–925.

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

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.

Myerson, Roger B, "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

— , "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.

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.

Rothschild, Michael and Joseph E. Stiglitz, "Increasing risk: I. A definition," *Journal of Economic theory*, 1970, *2* (3), 225–243.

Shi, Xianwen, "Optimal Auctions with Information Acquisition," *Games and Economic Behavior*, 2012, *74* (2), 666–686.

B Supplementary Appendix: Omitted Technical Details

B.1 A Signal Structure Rationalizing the Information-Acquisition Technology of Section 2

The model's description, in Section 2, *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 2 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 $p \equiv \int_0^1 s \mathrm{d}F_H(s) = \int_0^1 s \mathrm{d}F_L(s)$. Then, by construction, $\Pr\{v = 1\} = \mathbb{E}[\theta_2 \mid a]$ 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 $\sigma_H(z \mid v)$ and $\sigma_L(z \mid v)$, where

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

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

$$\int_{s \le z} \left[p\sigma_j(s \mid 1) + (1-p)\sigma_j(s \mid 0) \right] \mathrm{d}s = F_j(z),$$

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

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

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

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

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

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

B.2 Justifying Equation (A.11) in the Proof of Lemma 6

To justify equation (A.11), 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 $\left\{\frac{1}{2^n}, \frac{2}{2^n}, ..., 1\right\}$ 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 $\left\{\frac{1}{2^n}, \frac{2}{2^n}, ..., 1\right\}$. 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

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

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 $\left\{ \frac{1}{2^n}, \frac{2}{2^n}, ..., 1 \right\}$ 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 $j = 1, 2, ..., 2^n$, so that the family of disjoint sets $\{B_1^n, ..., B_{2^n}^n\}$ completely covers Θ_1. Note that it is possible to approximate a discrete measure

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

by

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

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

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

Such choice of $\left\{\alpha_j^n\right\}$ is possible because rationals are dense in reals. Then for each n, $\mu_n \in \mathcal{D}_n$. Moreover, as $n \to \infty$, $\mu_n \Rightarrow \mu$, 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 $\Theta_1 = [0,1]$.[52] Let

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

$\|g\|_\infty \equiv \sup_{x\in\Theta_1} g(x)$ denote the supremum norm. Then

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

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

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

Because g is bounded, the second term on the right-hand side can be made arbitrarily small by choosing n sufficiently large, whereas ϵ is arbitrary. Hence, $\int g\mathrm{d}\mu_n \to \int g\mathrm{d}\mu$ as $n \to \infty$, 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

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

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

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.

$\Delta(\Delta\Theta_1)$, 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

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

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

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

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

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

Then $B_1^m, \ldots, B_J^m$ 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,

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

It is possible to approximate

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

by

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

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

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

Since rationals are dense in reals, such choice of $\left\{\beta_j^m\right\}$ 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,

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

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

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

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, $\int f \mathrm{d}\nu_m \to \int f \mathrm{d}\nu$ as $m \to \infty$, which implies that $\nu_m \Rightarrow \nu$. □

B.3 Proof of Lemma 3

The proof proceeds in steps. Step 1 gives a differential characterization of a class of regular prospect sets. Step 2 shows that the conditions in that characterization are implied by the conditions of Lemma 3 or its modification discussed in Section 5.

Step 1: Differential Conditions for Regularity

Restrict attention to a class of so-called orderly prospect sets. A prospect set is **orderly** if it satisfies $\alpha(0) = \alpha(1) = 0$ with $\alpha \geq 0$, or $\pi(0) = \pi(1) = 0$ with $\pi \geq 0$, or both, and if it is induced by α and π that are twice differentiable. Because of the differentiability of α and π, an orderly prospect set is a continuous curve whose endpoints lie on at least one of the axes, which the curve does not intersect.

An orderly prospect set is regular if its every decreasing segment is convex toward the "interior" of the prospect set. (That is, there are at most one decreasing segment, at most one increasing segment, and the two segments are "convex away" from each other. This property can be ascertained graphically.) Thus, to establish regularity, one would like to know something about the behavior of α as a function of π—in particular, $\mathrm{d}^2\alpha/\mathrm{d}\pi^2$. Because the curve $\alpha(\pi)$ is given parametrically, the sought derivative must be found in a roundabout way, applying the chain rule of differentiation twice. Then:

$$\frac{\mathrm{d}^2\alpha}{\mathrm{d}\pi^2} = \frac{\alpha''\pi' - \alpha'\pi''}{(\pi')^3}, \tag{B.3}$$

where α' denotes $\mathrm{d}\alpha(\theta_1)/\mathrm{d}\theta_1$, etc. The appropriate convexity assumptions require

$$\begin{aligned} \pi' &> 0 \text{ and } \alpha' < 0 \implies \frac{\mathrm{d}^2\alpha}{\mathrm{d}\pi^2} < 0 \iff \left(\frac{\pi'}{\alpha'}\right)' > 0 \\ \pi' &< 0 \text{ and } \alpha' > 0 \implies \frac{\mathrm{d}^2\alpha}{\mathrm{d}\pi^2} > 0 \iff \left(\frac{\pi'}{\alpha'}\right)' > 0, \end{aligned}$$

where the equivalences are established using (B.3). To summarize, an orderly prospect set is regular if, for all $s \in \Theta_1$

$$\pi'(s)\,\alpha'(s) < 0 \implies \left(\frac{\pi'(s)}{\alpha'(s)}\right)' > 0. \tag{B.4}$$

<u>Step 2: Sufficiency of the Conditions in Lemma 3</u>

The twice-differentiability of α and π and part (ii) of Lemma 3 (or Condition 5 and the linearity of F in a) ensure that $\alpha(0) = \alpha(1) = 0$ and $\alpha \geq 0$ (respectively, $\pi(0) = \pi(1)$ and $\pi \geq 0$), and hence that the induced prospect set is orderly.

It remains to ascertain that (B.4) holds. Define $r \equiv (1-G)/g$. The following equivalence holds:

$$\left(\frac{\pi'}{\alpha'}\right)' > 0 \iff D \equiv \left(\frac{(r(F_H - F_L))'}{F_H - F_L}\right)' < 0.$$

By the chain rule of differentiation,

$$\begin{aligned} D &= r'' + r'\frac{f_H - f_L}{F_H - F_L} + r\left(\frac{f_H - f_L}{F_H - F_L}\right)'. \\ &= r'' + r'\frac{f_H - f_L}{F_H - F_L} + r\left(\frac{\alpha''}{\alpha'}\right)', \end{aligned}$$

with the last equality follows by the differentiation of α, given in (6), Theorem 1.

To conclude (B.4), one must consider two cases: $\alpha' < 0$ and $\pi' > 0$, and $\alpha' > 0$ and $\pi' < 0$. Consider the case of $\pi' > 0$ and $\alpha' < 0$ first. In this case,

$$\alpha' = \frac{1}{c}(F_H - F_L) < 0 \qquad \text{and} \qquad \pi' = r'(F_L - F_H) + r(f_L - f_H) > 0,$$

which imply $f_L - f_H > 0$ as long as $r' \leq 0$, as part (i) of Lemma 3 supposes. As a result, $D < 0$ as long as $r'' \leq 0$ (as part (i) of Lemma 3 supposes) and $(\alpha''/\alpha')' < 0$ (as part (iii) of Lemma 3 supposes).

Consider the remaining case in which $\alpha' > 0$ and $\pi' < 0$. In this case,

$$\alpha' = \frac{1}{c}(F_H - F_L) > 0 \qquad \text{and} \qquad \pi' = r'(F_L - F_H) + r(f_L - f_H) < 0,$$

imply $f_L - f_H < 0$, by $r' \leq 0$. As a result, $D < 0$, by $r'' \leq 0$ and $(\alpha''/\alpha')' < 0$. Hence, (B.4) holds, and the orderly prospect set is regular.