



BIROn - Birkbeck Institutional Research Online

Nikandrova, Arina and Pancs, R. (2017) Conjugate information disclosure in an auction with learning. *Journal of Economic Theory* 171 , pp. 174-212. ISSN 0022-0531.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/19020/>

Usage Guidelines:

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

or alternatively

Conjugate Information Disclosure in an Auction with Learning

Arina Nikandrova and Romans Pancs*

June 2017

Abstract

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

Keywords: Information Disclosure, Conjugate Disclosure, Bayesian Persuasion

JEL codes: D82, D83

1 Introduction

In the U.K., the government franchises rail passenger services to train-operating companies for a limited time. An auction determines the award of the franchise to run passenger services in a

*Nikandrova (a.nikandrova@gmail.com) is at Birkbeck; Pancs (rpancs@gmail.com) is at ITAM (Av. Camino a Santa Teresa 930, Magdalena Contreras, Heroes de Padierna, 10700 Ciudad de México, CDMX, México). For helpful discussions, we thank Manuel Amador, Paulo Barelli, Simon Board, Arupratan Daripa, Hari Govindan, Hugo Hopenhayn, Sandeep Kapur, Paul Klemperer, Moritz Meyer-Ter-Vehn, Vladimir Parail, Anne Roesler, Joel Sobel, and Juuso Toikka. We also thank the editor for guidance and anonymous referees for detailed comments and suggestions.

certain region. The pool of bidders typically includes an incumbent operator and potential entrants. The bidders' valuations of the franchise vary depending on the bidder's operating costs. The incumbent is likely to know its costs from past experience, whereas the entrants are poorly informed and must perform due diligence to evaluate the purchase opportunity. This paper is a step towards understanding what a profit-maximizing government should reveal to potential entrants about the well-informed incumbent's intended bid in order to influence their due diligence.

The paper formulates a model that captures the essential features of the government's information-disclosure problem. A seller sells an item by sequentially bargaining with two bidders: the leader, who knows its valuation, and the follower, who must exert costly effort to better estimate his valuation. The bidders' valuations are private and statistically independent.

The seller maximizes his revenue by designing, announcing, and committing to a mechanism. He is restricted to choosing a mechanism with an ex-post efficient allocation rule, which must assign the item to the bidder with the higher expected valuation (conditional on the bidders' information). The focus on ex-post efficient allocation rules is restrictive and is motivated primarily by tractability, as well as by the desire to isolate the distortions due to information disclosure from the distortions due to the allocation rule.¹ In addition, in some applications, the violation of ex-post efficiency is politically costly or outright infeasible.² We also restrict attention to mechanisms that rule out the sale of information about already collected bids; only in such mechanisms is there room for strategic disclosure.

The paper's main result is the rule according to which the seller partially obfuscates the bid he receives from the leader before passing this information on to the follower, who then decides how much to learn.³ In particular, the seller optimally partitions the leader's possible types (i.e., valuations, reported as bids) into singletons and pairs of so-called conjugate types. Thus, the seller may disclose to the follower just a pair to which the leader's type belongs, without revealing that type. Figure 1 illustrates.

The seller's strategic disclosure distorts the follower's effort away from the effort in the first-

¹In the working-paper version (Nikandrova and Pans, 2015), we show that our analysis extends to a class of allocation rules that are fixed (as opposed to being optimized over).

²For instance, the government may face disgruntled voters if it allocates a procurement contract or a franchise to a bidder whom everyone knows not to be the best choice. The working-paper version of Gershkov and Szentes (2009) motivates ex-post efficiency by appealing to the legal ramifications of ex-post inefficient decisions.

³Due to the option value of influencing the follower's learning by the information about the leader's bid, it is optimal for the seller to approach the leader first.

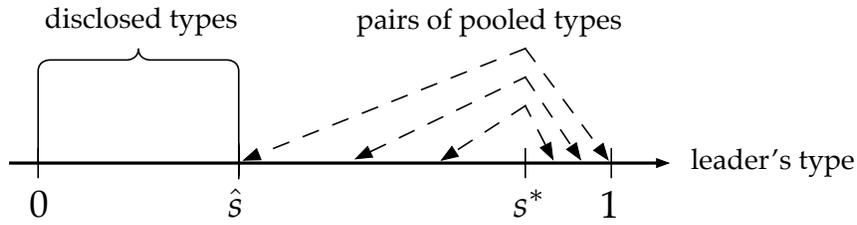


Figure 1: The seller’s optimal disclosure rule. The seller reveals to the follower the leader’s type if it is less than \hat{s} and pools any type that exceeds \hat{s} with a corresponding conjugate type to form a pair that straddles s^* .

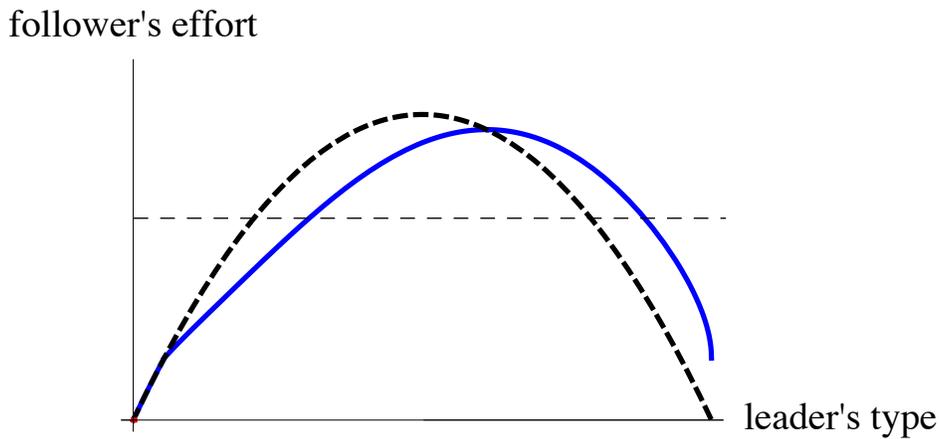


Figure 2: The optimal effort schedule is the solid hump-shaped curve; the first-best effort schedule is the dashed hump-shaped curve. The horizontal dashed line highlights the correspondence between a recommended effort and the equivalent message that pools two leader types together.

best mechanism, which maximizes the ex-ante expected surplus. Figure 2 illustrates both optimal and first-best effort schedules. The first-best effort is hump-shaped in the leader’s valuation; the follower learns more when the uncertainty about the identity of the higher-valuation bidder is the greatest, which is when extra information is needed most. The optimal effort is also hump-shaped but is “shifted to the right” relative to the first-best, meaning that, when the leader’s valuation is low, the follower learns inefficiently little, whereas when the leader’s valuation is high, the follower learns inefficiently much.

This effort distortion arises because the seller seeks to make the follower win as often as possible. To understand why doing so is profitable, note that ex-post efficiency requires the seller to charge the follower exactly θ_1 —the leader’s type—and the leader less than θ_1 . Indeed, ex-post efficiency requires that the follower be charged θ_1 and, thus, buys if and only if his expected valuation, denoted by θ_2 , satisfies $\theta_2 \geq \theta_1$. By contrast, the leader must be charged less than θ_1 . If he

were charged θ_1 , then he would profitably pretend to be a type that would be charged less than that. As a result, the seller prefers selling to the follower.

Whether less or more learning increases the probability that the follower is the efficient winner depends on the leader's type. Learning induces a mean-preserving spread in the probability distribution of θ_2 .⁴ Therefore, event $\{\theta_2 \geq \theta_1\}$ (ex-post efficient sale to the follower) is more likely either when θ_1 is low and the follower learns little, or when θ_1 is high and the follower learns a lot. Thus, the seller nudges the follower to learn more than is first-best efficient when θ_1 is sufficiently high and less than first-best efficient otherwise. It turns out that this nudge can be accomplished by pooling the leader's true type with some other type.

Figure 2 suggests a correspondence between the optimal disclosure rule and the optimal effort schedule. All that the follower needs to be told about the leader's type is summarized in the effort that the seller would like him to exert.⁵ Thus, an optimal disclosure rule can be read off an optimal effort schedule by intersecting the latter with the horizontal line corresponding to a recommended effort. Given its hump shape, the optimal effort schedule in Figure 2 prescribes the same effort for (at most) two leader types; the seller optimally pools conjugate types. Analytically, however, it is more convenient to derive an optimal disclosure rule first and then to recover the corresponding effort schedule; the paper's analysis proceeds in this order.

The seller's profit-maximizing outcome can be implemented in a sequential second-price auction with a tax (or subsidy) for the leader. The tax motivates the leader to bid truthfully. Without the tax, the leader would be tempted to bias his bid away from θ_1 and towards some intermediate value. In response, the follower would learn a lot, thereby introducing greater dispersion into the probability distribution of his type. The leader likes this dispersion; he wins and pays little if the follower's type is low, while loses and pays nothing if the follower's type is high. The tax countervails the leader's incentive to manipulate the follower's learning.

The rest of the paper is structured as follows. This section concludes with a literature review. Section 2 describes the environment. Section 3 derives the first-best outcome and an auction that

⁴The identification of the informativeness of a signal with the induced dispersion of the probability distribution of θ_2 is a standard modeling device (Johnson and Myatt, 2006; Ganuza and Penalva, 2010; Shi, 2012; and Roesler, 2014) and, in the presence of risk neutrality, entails no loss of generality. Indeed, higher information-acquisition effort can be interpreted as delivering a more informative signal about the underlying valuation and thereby induce a more dispersed posterior probability distribution of the underlying valuation. When a bidder is risk-neutral, the expected underlying valuation, denoted by θ_2 , is the only aspect of the posterior probability distribution that he cares about.

⁵This is the Revelation Principle for games with private actions (Myerson, 1982, 1986).

implements it. Section 4 establishes the suboptimality of full disclosure and non-disclosure and then, under additional conditions, partially characterizes optimal disclosure by solving the seller's relaxed problem. Section 5 shows that, for sufficiently costly learning, the focus on the relaxed problem is justified. Section 6 illustrates some of the paper's results in a numerical example. Section 7 concludes. The proofs are in Appendix A, with more-technical arguments relegated to Supplementary Appendix B.

Related Literature

Our paper contributes to the literature on auctions in which a monopolistic seller directly or indirectly influences bidders' information structure and, more broadly, to the literature on Bayesian persuasion.

The existing literature in which a seller, through his choice of a selling mechanism, affects bidders' information-acquisition effort (e.g., Bergemann and Välimäki, 2002; Persico, 2003; Compte and Jehiel, 2007; Crémer et al., 2009; and Shi, 2012) focuses primarily on simultaneous auctions, in which the issue of optimal bid disclosure does not arise. Crémer et al. (2009) examine sequential auctions and design a revenue-maximizing one. Because Crémer et al. (2009) assume that the seller can charge bidders for information, optimal information disclosure turns out to be trivial (full disclosure) and is not their focus. Information disclosure is the focus of Eso and Szentes (2007). In their model, the seller directly designs the signals observed by the bidders instead of motivating bidders to choose signals themselves. Just like Crémer et al. (2009), Eso and Szentes (2007) allow the seller to charge bidders for information and find full disclosure to be optimal. Eso and Szentes's (2007) seller reveals maximal information to maximize the total surplus, which he then taxes away by cleverly charging for the signals he reveals.⁶ In our paper, the critical assumption that rules out selling information and, with it, the optimality of full disclosure is a particularly demanding interim participation constraint. Under this constraint, the follower, upon observing his type, must expect his total payoff to be nonnegative. Without this participation constraint, the logic of Eso and Szentes (2007) would imply the optimality of full disclosure in our model also. Such stringent participation constraints are also imposed by Ganuza (2004), in a second-price auc-

⁶The seller must charge cleverly because the bidders of Eso and Szentes (2007), in contrast to the bidders of Crémer et al. (2009), already have some private information before accepting the seller's mechanism. So the charges are not simple participation fees.

tion, and by [Bergemann and Pesendorfer \(2007\)](#), in an optimally designed auction, and rule out the optimality of full disclosure in their settings.

Another strand of literature to which our paper contributes is on Bayesian persuasion, or sender-receiver games with commitment. The two main papers in this literature are [Rayo and Segal \(2010\)](#)—henceforth RS—and [Kamenica and Gentzkow \(2011\)](#). RS assume additional structure that makes their paper especially pertinent to our problem. We establish the relevance of RS’s results in a novel environment—an auction with costly learning and, crucially, with a continuum of types. A limit argument establishes a formal connection between RS’s model and ours, thereby paving the way for the proof of the optimality of conjugate disclosure. Our sharp characterization of optimal disclosure has no direct counterpart in RS’s discrete model.

Any auction design or agency design in which information disclosure affects some player’s unenforceable (by the seller or the principal) action features Bayesian persuasion. Examples of such a design are a two-player contest of [Zhang and Zhou \(2015\)](#) and auctions with resale. While early resale models ([Bikhchandani and Huang, 1989](#); [Gale et al., 2000](#); [Haile, 2003](#); [Gupta and Lebrun, 1999](#)) fix an auction format and study the informational linkage between the primary market and the resale market, later work ([Calzolari and Pavan, 2006a](#); [Zheng, 2002](#)) adopts the mechanism-design approach. The work of [Calzolari and Pavan \(2006a\)](#) is especially related to ours.

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

In [Calzolari and Pavan’s \(2006a\)](#) model, as in ours, optimality proscribes full disclosure because they assume, as we do, a participation constraint that precludes the seller from expropri-

ating the traders' rents in the resale market.⁷ Calzolari and Pavan's (2006a) assumption that the leader's type is binary delivers tractability and allows them to characterize both an optimal allocation rule and an optimal disclosure rule. By contrast, our assumption of the continuum of the leader's types enables us to study richer disclosure rules, but at the cost of fixing the allocation rule.

2 Model

Environment

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

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. The leader's payoff is

$$\theta_1 x_1 - t_1.$$

The follower's payoff is

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

Both bidders are expected-utility maximizers.

Learning Technology

Interpret θ_2 as the expectation of the follower's (unmodeled) underlying valuation, conditional on the privately observed signal generated by learning. For any $a \in A$, θ_2 is drawn according to a

⁷In a similar spirit, in a sequential common agency model, Calzolari and Pavan (2006b) identify conditions under which the upstream principal may find it strictly optimal to disclose a noisy signal about the agent's type to the downstream principal.

c.d.f that is linear in a :⁸

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

Whenever the p.d.f.s corresponding to the c.d.f.s F , F_H , and F_L exist, they are denoted by f , f_H , and f_L . The c.d.f. F may have mass points on $\{0, 1\}$, but not on $(0, 1)$, and is smooth, with bounded derivatives. Conditional on a , θ_1 and θ_2 are independent.

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

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

According to Condition 1, the follower who exerts effort a , with probability a , receives a more informative signal about his underlying valuation, and, with probability $1 - a$, receives a less informative signal.⁹ Part (i) requires the follower's effort not to affect his expected type.¹⁰ In particular, part (i) rules out the situations in which the follower's effort is a value-enhancing investment. Parts (i) and (ii) taken together imply that a higher effort induces a mean-preserving spread of the distribution of the follower's types.

This mean-preserving spread requirement is implied by Blackwell's informativeness criterion (Blackwell, 1951, 1953). According to this criterion, a more informative signal about the follower's underlying valuation induces a greater dispersion of the probability distribution over the conditional expectation, which is the interpretation of θ_2 in our model.¹¹ Condition 1 generalizes the truth-or-noise information-acquisition technology introduced by Lewis and Sappington (1994) and used by Bergemann and Valimaki (2006, Section 2.2), Johnson and Myatt (2006, Section III.B), and Shi (2012, Example 2), among others.

⁸The linearity condition (1), known as the Linear Distribution Function Condition in the principal-agent literature, is essential for reducing the seller's problem to the information-disclosure problem of RS in Section 4. Supplementary Appendix B.2 discusses what, exactly, linearity rules out.

⁹A signal structure that delivers the probability distribution of types in Condition 1 is given in Supplementary Appendix B.1.

¹⁰For any c.d.f. H on $[0, 1]$, the expectation is $\int x dH(x) = \int (1 - H(x)) dx$.

¹¹Blackwell's informativeness criterion implies Lehmann's accuracy condition (Lehmann, 1988; Persico, 2003), which implies the mean-preserving-spread order on the conditional expectations (Mizuno, 2006, Proposition 1). Directly modeling a signal's informativeness by the induced dispersion of the conditional expectation is standard; see, for instance, Johnson and Myatt (2006), Ganuza and Penalva (2010), Shi (2012), and Roesler (2014).

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

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

Example 1 can be interpreted in this way: with probability a , the follower observes a perfectly informative signal that reveals his underlying valuation, which is either 0 or 1, equiprobably; and with probability $1 - a$, the follower observes a partially informative signal about the underlying valuation.

To guarantee interior solutions for a , we henceforth assume a sufficiently large cost of effort:¹²

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

The Seller's Problem

The seller chooses, publicly announces, and commits to a **mechanism**, which comprises an extensive game-form, a strategy to which the seller commits, and a communication device. The seller chooses the communication device. The game-form is given. Its timing is such that (i) each bidder may leave the mechanism without payment; (ii) the follower exerts his effort and observes his realized type; (iii) each bidder may once again leave the mechanism without payment; and (iv) the seller enforces a trade that is **ex-post efficient**, meaning that the higher-type bidder wins the item. The communication device allows for arbitrary communication between the game-form's stages as long as the seller collects enough information to compute an ex-post efficient allocation.

The logic of the Revelation Principle in environments with private information and private actions applies (Myerson, 1982, 1986): the seller can do no better than to minimize the information revealed to bidders and to maximize the information collected from them. Consequently, no generality is lost by restricting attention to a parsimonious class of direct mechanisms:

Lemma 1. *Without loss of generality, the seller can restrict attention to direct mechanisms in which*

1. *having observed θ_1 , the leader confidentially reports $\hat{\theta}_1$ in Θ_1 to the seller;*

¹²This condition is derived by requiring that the first-best effort in Theorem 1 be less than 1 for every θ_1 .

2. the seller confidentially sends a message m from a set M to the follower according to a **disclosure rule** $\mu : \Theta_1 \rightarrow \Delta(M)$, which associates with each report of the leader a probability distribution over messages;¹³
3. the follower exerts effort, denoted by $a_c^*(m) \in A$, then observes θ_2 , and confidentially reports $\hat{\theta}_2 \in \Theta_2$; and
4. bidder i with $\hat{\theta}_i > \hat{\theta}_{-i}$ gets the item; payments $(t_1, t_2) : \Theta_1 \times \Theta_2 \rightarrow \mathbb{R}^2$ (the functions of bidders' reports) are assessed.

Proof. See Appendix A. □

Without loss of generality, the seller can focus on the mechanisms whose (perfect Bayes-Nash) equilibria are **truthful**, meaning that $(\hat{\theta}_1, \hat{\theta}_2) = (\theta_1, \theta_2)$. By the Revelation Principle, one can equivalently identify the seller's message space M with the set of recommended learning efforts A , but a different M will sometimes be analytically convenient.

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

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

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

¹³If μ is such that the seller's message is independent of the leader's report, the described mechanism is strategically equivalent to a mechanism in which the seller asks both bidders to submit their reports simultaneously.

3 The First-Best Benchmark

A **first-best** outcome obtains when a planner maximizes the expected total surplus while observing the leader's type, directly controlling the follower's effort, and then observing the follower's realized type. When the leader's type is θ_1 , the follower's first-best effort, denoted by $\alpha_c(\theta_1)$, maximizes the total surplus:

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

Theorem 1 calculates the first-best outcome.

Theorem 1. *The first-best effort is $\alpha_c(\theta_1) \equiv \alpha(\theta_1)/c$, where $\alpha(\theta_1)$, the normalized first-best effort, satisfies*

$$\alpha(\theta_1) = \int_{\theta_1}^1 (F_L(s) - F_H(s)) ds, \quad \theta_1 \in [0, 1], \quad (5)$$

and $\alpha(0) = \alpha(1) = 0$; α is strictly increasing when $\theta_1 < \theta^*$ and strictly decreasing when $\theta_1 \geq \theta^*$.

Proof. See Appendix A. □

Corollary 1 shows that the first-best outcome can be implemented in an incentive-compatible manner. The mechanism in the corollary uses the (possibly negative) tax

$$T(\theta_1) \equiv \int_0^{\theta_1} (F(s | \alpha_c(\theta_1)) - F(s | \alpha_c(s))) ds, \quad \theta_1 \in \Theta_1. \quad (6)$$

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 $T(b)$, given by (6);
2. discloses b to the follower and invites him to bid in the second-price auction; and
3. allocates the item and assesses the payments according to the rules of the second-price auction.

In equilibrium, each bidder bids his type and enjoys a nonnegative expected payoff. The follower exerts the first-best effort.

Proof. See Appendix A. □

Except for the leader's tax, the mechanism in Corollary 1 is the standard second-price auction but executed sequentially, with the leader's bid being public. Without the tax, the leader would try

to manipulate the follower’s learning by bidding untruthfully. Indeed, if the type- θ_1 leader truthfully bids $b = \theta_1$ in the second-price auction without the tax, his payoff is $\mathbb{E}_{\theta_2} [\max \{0, \theta_1 - \theta_2\}]$. Because $\max \{0, \theta_1 - \theta_2\}$ is convex in θ_2 , Jensen’s inequality implies that the leader gains when the distribution of the follower’s valuation is more dispersed, which occurs when the follower learns more. The leader can induce greater learning by nudging his bid towards θ^* , thereby making the follower more uncertain about his payoff from participating in the auction and desirous of more information. Infinitesimal deviation from truth would have only a second-order (detrimental) effect on the leader’s payoff, whereas the beneficial effect from influencing the follower’s learning would be first-order.

Formally, the tax discourages untruthful bidding by altering the leader’s marginal payoff from raising his bid. The marginal payoff is altered by the amount

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

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

4 A Seller-Optimal Auction

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

Section 4.1 reduces the seller’s (relaxed) revenue-maximization problem to an information-disclosure problem.¹⁴ Section 4.2 shows that this problem is nontrivial and is solved by disclosing neither everything nor nothing. Section 4.3 connects the seller’s continuum-of-types disclosure problem to the finite-types disclosure problem of RS. This connection brings out the qualitative

¹⁴The validity of the focus on the relaxed problem is discussed in Section 5.

features of optimal disclosure and the follower's induced effort schedule. In particular, we demonstrate that the optimal disclosure rule is a deterministic function that maps no more than two types of the leader into the same recommended effort. Section 4.4 formulates the seller's disclosure problem as an optimal-control problem.

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

We begin by formalizing the seller's constraints, which we then substitute into his objective function, thereby reducing his problem to an information-disclosure problem. The constraints are introduced by rolling the game-form backwards.

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

Suppose that, after observing the seller's message m , the follower exerts some effort and observes his type θ_2 . He chooses his report $\hat{\theta}_2$ to maximize his interim expected payoff, thereby attaining the value

$$U_2(\theta_2 | m) \equiv \max_{\hat{\theta}_2 \in \Theta_2} \mathbb{E}_{\theta_1 | m} \left[\theta_2 \mathbf{1}_{\{\hat{\theta}_2 > \theta_1\}} - t_2(\theta_1, \hat{\theta}_2) \right], \quad (7)$$

where the expectation is over the leader's type θ_1 , conditional on the seller's message m and, implicitly, on the disclosure rule. By inspection of (7), the follower's effort does not enter his interim expected payoff, and so he chooses his report independently of this effort. The follower's **truth-telling constraint** requires the local truth-telling constraint (implied by the Envelope Theorem applied to (7)),

$$U_2'(\theta_2 | m) \equiv \frac{dU_2(\theta_2 | m)}{d\theta_2} = \mathbb{E}_{\theta_1 | m} [\mathbf{1}_{\{\theta_2 > \theta_1\}}], \quad (8)$$

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

The **interim participation constraint** holds if, even after observing θ_2 , the follower expects his payoff from the mechanism (gross of the cost of information acquisition) to remain nonnegative.

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

$$U_2(0 | m) + \int_0^{\theta_2} \mathbb{E}_{\theta_1|m} [\mathbf{1}_{\{s > \theta_1\}}] ds \geq 0, \quad (9)$$

where the equivalence holds by [Milgrom's \(2004\) Constraint Simplification Theorem](#), which justifies the application of the fundamental theorem of calculus and rewrites U_2 in terms of its derivative from (8). Because the right-hand side of (8) is nonnegative, $U_2(\theta_2 | m)$ is nondecreasing. Thus, the follower's interim participation constraint holds if and only if

$$U_2(0 | m) \geq 0 \quad \text{for all } m \in M. \quad (10)$$

Suppose that $U_2(0 | m) = 0$ for all $m \in M$ (as will be the case in the optimal mechanism). Then, the interim participation constraint rules out mechanisms that ask the follower to commit to a payment in exchange for the right to participate in the mechanism; it also rules out the mechanisms that offer to sell information about the leader's type before the follower decides which effort to exert.¹⁵

One can now take a step back in the game-form and ask which effort is optimal for the follower who observes message m and knows that he will optimally report truthfully in the future. Immediately after observing message m and deciding to exert effort a , the follower expects his payoff net of the cost of effort to be

$$\int_{\Theta_2} U_2(\theta_2 | m) dF(\theta_2 | a) = U_2(0 | m) + \int_0^1 \mathbb{E}_{\theta_1|m} [\mathbf{1}_{\{\theta_2 > \theta_1\}}] (1 - F(\theta_2 | a)) d\theta_2, \quad (11)$$

where the equality uses (8) and integration by parts.¹⁶ Interchanging the order of integration and expectation (by Fubini's theorem) in the right-hand side of the above display yields the expression for the follower's normalized optimal effort $a_c^*(m)$ as a function of the observed message m :

$$a_c^*(m) \in \arg \max_{a \in A} \left(U_2(0 | m) + \mathbb{E}_{\theta_1|m} \left[\int_{\theta_1}^1 (1 - F(\theta_2 | a)) d\theta_2 \right] - C(a) \right). \quad (12)$$

¹⁵In practice, shareholders may forbid managers to commit to any payments until due diligence has been performed.

¹⁶Integration by parts is valid even if F is discontinuous (as in [Example 1](#)), because $U_2(\cdot | m)$ has a bounded derivative everywhere and, hence, is continuous.

Under the maintained Condition 1, the maximization problem in (12) has the unique solution $a_c^*(m) \equiv a^*(m)/c$, with $a^*(m)$ being the follower's normalized optimal effort

$$a^*(m) = \mathbb{E}_{\theta_1|m} [\alpha(\theta_1)], \quad (13)$$

where $\alpha(\theta_1)$, defined in (5), is the normalized first-best effort level when the leader's type is θ_1 . We refer to (13) as the **obedience constraint**; if the seller's message space is the set of recommended efforts, then each recommended effort m satisfies $m = a_c^*(m)$, meaning that the follower must find it optimal to obey the recommendation.

With the knowledge that the follower will be truthful and obedient, one can take another step back and impose the **ex-ante participation constraint**, which requires that the follower expect a nonnegative payoff from the mechanism right after observing the seller's message but before exerting any effort. Formally, for any $m \in M$, the maximand in (12), evaluated at the optimal effort $a_c^*(m)$, must be nonnegative:

$$U_2(0 | m) + \mathbb{E}_{\theta_1|m} \left[\int_{\theta_1}^1 (1 - F(\theta_2 | a_c^*(m))) d\theta_2 \right] - C(a_c^*(m)) \geq 0. \quad (14)$$

Substituting the functional forms of F and C , and the expressions for a^* and α from (13) and (5), and rearranging gives

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

where the inequality follows by inspection. Moreover, the interim participation constraint (10) requires $U_2(0 | m) \geq 0$, and hence, by the display above, the ex-ante participation constraint in (14) is implied by (10). Therefore, from now on, we focus on the follower's interim participation constraint and refer to it simply as his **participation constraint**.

The Leader's Truth-Telling and Participation Constraints

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

$$U_1(\theta_1) \equiv \max_{\hat{\theta}_1 \in \Theta_1} \mathbb{E}_{\theta_2 | \theta_1} \left[\theta_1 \mathbf{1}_{\{\hat{\theta}_1 > \theta_2\}} - t_1(\theta_1, \theta_2) \right]. \quad (15)$$

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

$$U_1(\theta_1) = U_1(0) + \int_0^{\theta_1} \mathbb{E}_{m | \theta_1} [F(\theta_1 | a_c^*(m))] ds, \quad (16)$$

and the monotonicity condition on the probability of winning,

$$\mathbb{E}_{\theta_2 | \theta_1} [\mathbf{1}_{\{\theta_1 > \theta_2\}}] = \mathbb{E}_{m | \theta_1} [F(\theta_1 | a_c^*(m))], \quad \text{is nondecreasing in } \theta_1. \quad (17)$$

In contrast to the follower's monotonicity condition, the leader's monotonicity condition (17) is not immediately implied at the solution. Instead, we proceed with the analysis *assuming* that this condition holds, and we then investigate (in Section 5) the conditions under which it does hold.

The leader's **interim participation constraint** or, simply, his **participation constraint**, ensures that each type of the leader is at least as well off in the mechanism as he would be if he were to refrain from participation and enjoy the payoff of zero: $U_1(\theta_1) \geq 0$ for all θ_1 . Differentiating (16) gives $U_1'(\theta_1) = \mathbb{E}_{m | \theta_1} [F(\theta_1 | a_c^*(m))] \geq 0$; that is, U_1 is nondecreasing, and so the participation constraint holds for all θ_1 as long as it holds for $\theta_1 = 0$:

$$U_1(0) \geq 0. \quad (18)$$

The Seller's Virtual Surplus

We can now use (16) combined with (15) and (11), both evaluated at $a_c^*(m)$ and combined with (7), to substitute out the bidders' transfers from the seller's objective function (3). From the seller's perspective, it is optimal to set the transfers for the lowest-type leader and the lowest-type follower so that their expected payoffs are zero. Because $U_1(0) = 0$ and, for all m , $U_2(0 | m) = 0$,

the seller's virtual surplus is¹⁷

$$\mathbb{E}_{m, \theta_1, \theta_2} \left[\left(\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} \right) \mathbf{1}_{\{\theta_1 > \theta_2\}} + \left(\theta_2 - \frac{1 - F(\theta_2 | a_c^*(m))}{f(\theta_2 | a_c^*(m))} \right) \mathbf{1}_{\{\theta_2 \geq \theta_1\}} \right]. \quad (19)$$

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

Integrating θ_2 out of (19) and simplifying yields a more compact expression for the virtual surplus:¹⁸

$$\int_{\Theta_1} \mathbb{E}_{m | \theta_1} \left[\theta_1 - \frac{1 - G(\theta_1)}{g(\theta_1)} F(\theta_1 | a_c^*(m)) \right] dG(\theta_1). \quad (20)$$

To see why (20) is equivalent to (19), suppose that the leader's type is θ_1 , and the seller sends a message m . Ex-post efficiency requires that the leader win with probability $F(\theta_1 | a_c^*(m))$, which is the probability of the event $\{\theta_2 < \theta_1\}$. In this case, the seller's gain is the leader's virtual valuation, which equals his true valuation θ_1 less the **information rent** $(1 - G(\theta_1)) / g(\theta_1)$. Analogously, if the follower wins, the seller's gain is the follower's expected virtual valuation conditional on winning, which can be verified to be simply θ_1 .¹⁹ Hence, the seller prefers selling to the follower and designs the information disclosure rule so as to maximize the probability of efficient sale to him.

To maximize the follower's probability of winning, the seller aims to encourage the follower to become "stronger," thereby intensifying the competition that the leader faces. Whether a better or worse informed follower is stronger depends on the leader's type. When the leader's type is high, a more informed follower, who draws his valuation from a more dispersed probability distribution, stands a better chance of outbidding the leader. When the leader's type is low, a less informed follower, who draws his valuation from a less dispersed probability distribution, is

¹⁷If F has no density f , skip to (20), which does not rely on the existence of f .

¹⁸The virtual surplus in (20) is nonnegative, for it is bounded below by $\int_{\Theta_1} [\theta_1 - (1 - G(\theta_1)) / g(\theta_1)] dG(\theta_1) = 0$. That is, the seller's optimal revenue is nonnegative.

¹⁹Formally, the follower's expected virtual valuation conditional on winning when the leader's type is θ_1 is

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

Even though the integrand in the above display assumes that F has a positive density, f , the validity of (20) requires no such assumption.

more likely to outbid the leader. Thus, the seller will try to disclose information so as to nudge the follower to learn more when the leader's type is higher.

4.2 The Suboptimality of Full Disclosure and Non-Disclosure

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

The proof of Theorem 2 and the subsequent arguments use the seller's objective function (20) rewritten in a "product form." To arrive at this form, neglect the additive term in (20) that is independent of the disclosure rule; neglect the positive multiple $1/c$, too, and rewrite (20) as

$$\int_{\Theta_1} \mathbb{E}_{m|\theta_1} [\pi(\theta_1) a^*(m)] dG(\theta_1), \quad (21)$$

where

$$\pi(\theta_1) \equiv \frac{1 - G(\theta_1)}{g(\theta_1)} (F_L(\theta_1) - F_H(\theta_1)) \quad (22)$$

is the seller's **marginal benefit** from an increase in the follower's effort. This marginal benefit equals the leader's information rent times the marginal increase in the probability that the follower wins. Equation (21) is further transformed using the Law of Iterated Expectations and $a^*(m)$ in (13) to yield the **product form**

$$\mathbb{E} [\mathbb{E}_{\theta_1|m} [\pi(\theta_1)] \mathbb{E}_{\theta_1|m} [\alpha(\theta_1)]] . \quad (23)$$

The transformed objective function has the same form as the sender's objective function (equation [2]) in RS's model.²⁰ The product structure of (23) relies on two assumptions: the ex-post efficient allocation rule and the linearity of the c.d.f. $F(\theta_2 | a)$ in the follower's effort a . These two assumptions are crucial for adapting RS's techniques to our setting.²¹

To state our results, we borrow vocabulary from RS. A tuple $(\pi(\theta_1), \alpha(\theta_1))$ is called a **prospect**.

²⁰Henceforth, bracketed indices refer to equations and lemmas in RS's paper.

²¹In particular, if the seller were also to maximize over allocation rules, the product structure would be lost.

The **prospect set** is the graph $\Gamma \equiv \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in \Theta_1\}$, which differs from an RS prospect set only in that RS require their prospect sets to be finite.

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

Proof. See Appendix A. □

4.3 The Optimality of Conjugate Disclosure

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

Stating and proving Theorem 3 requires new definitions and intermediate results.

Restriction to Convex Prospect Sets

The analysis relies on the prospect set Γ being convex in the sense of:

Definition 1. A prospect set Γ is **convex** if it is a strictly convex curve (i.e., it intersects any line at most twice).

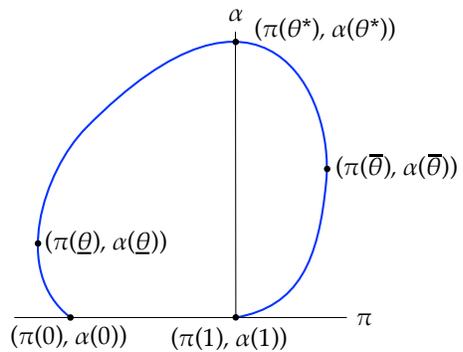
Convex prospect sets are illustrated in Figure 3. The subsequent analysis maintains an additional assumption:

Condition 2. The prospect set Γ is convex.²²

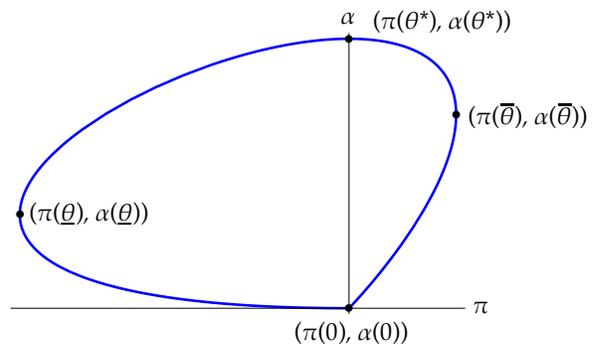
Convexity is satisfied in various "natural" examples (such as those in Figures 3b and 3c), and it renders the optimal-disclosure problem tractable.

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

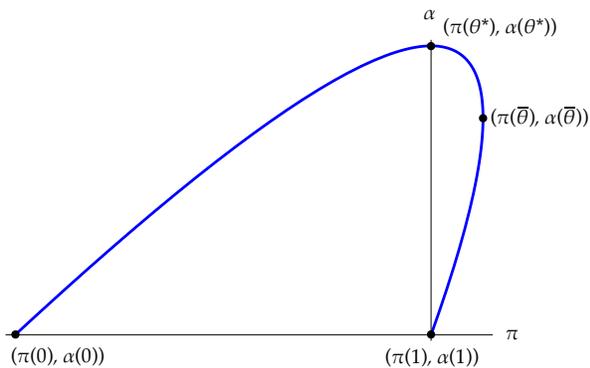
²²An analytical condition that is essentially equivalent to convexity of the prospect set is spelled out in Lemma B.1 in Supplementary Appendix B.3.



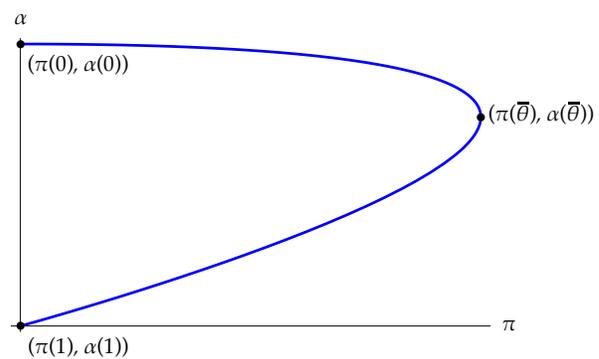
(a) A “typical” convex prospect set that satisfies Condition 1.



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



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



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

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

- On $(0, \underline{\theta})$, Γ is downward-sloping (α is strictly increasing; π is strictly decreasing).
- On $(\underline{\theta}, \theta^*)$, Γ is upward-sloping (both α and π are strictly increasing).
- On $(\theta^*, \bar{\theta})$, Γ is downward-sloping (α is strictly decreasing; π is strictly increasing).
- On $(\bar{\theta}, 1)$, Γ is upward-sloping (both α and π are strictly decreasing).

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

A Discretized Prospect Set

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

Optimal Disclosure with the Discrete Prospect Set

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

RS's results (some of which, for completeness, are reproduced in this paper) imply the following facts about every \mathcal{P}^n -optimal disclosure rule.

Fact 1. *By RS's Lemma [1] (also by this paper's Lemmas A.1 and A.2), no two prospects are pooled if both lie on the upward-sloping regions of—that is, (i) both lie in $\Gamma^n \cap \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in [\underline{\theta}, \theta^*]\}$; or (ii) both lie in $\Gamma^n \cap \{(\pi(\theta_1), \alpha(\theta_1)) : \theta_1 \in [\bar{\theta}, 1]\}$.*

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

Fact 3. *By RS's Lemma [3] (also by this paper's Lemma A.2), at most two prospects can be pooled under the same message because no more than two prospects lie on the same nonincreasing line, by Condition 2.*

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

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

The partial characterization of a \mathcal{P}^n -optimal disclosure rule in Facts 1–5 is refined in Lemma 2, which exploits the special structure of the seller's problem.

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

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

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

Proof. See Appendix A. □

Part (i) of Lemma 2 defines types s_* and s^* , both in Θ_1 (not necessarily in Θ^n), such that each pooling link intersects the line that passes through prospects $(\pi(s_*), \alpha(s_*))$ and $(\pi(s^*), \alpha(s^*))$, as shown in Figure 4. Part (ii) of the lemma shows that the optimal effort schedule is single-peaked in θ_1 , with the peak, at s^* , to the right of the peak of the first-best effort schedule, at θ^* .

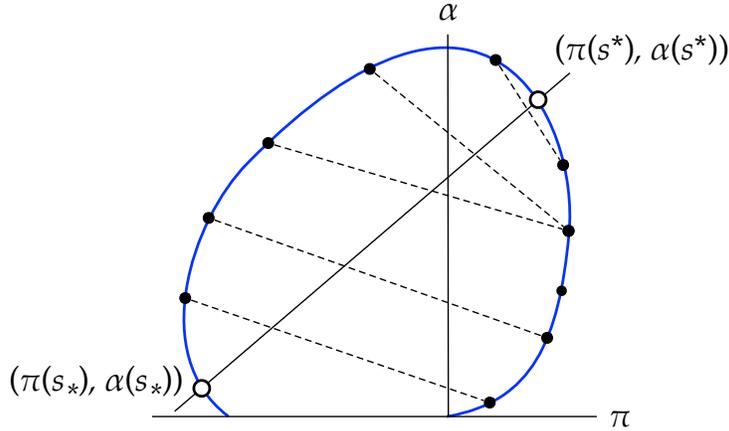


Figure 4: A discretized prospect set Γ^n is the union of solid dots. Without loss of generality, the seller can restrict attention to disclosure rules such as the one illustrated here. Each dashed link denotes a message that pools two prospects. These links never intersect and are oriented so that one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

Optimal Disclosure with the Continuous Prospect Set: The Main Result

The \mathcal{P}^n -optimal disclosure rule of Lemma 2 either reveals a prospect or pools it under the same message with another prospect. Lemma 2 does not rule out situations in which a prospect probabilistically invokes multiple messages (i.e., several pooling links could emanate from a single prospect). However, Theorem 3 shows that, in the continuous problem, denoted by \mathcal{P} , there is no loss of generality in focusing on disclosure rules that deterministically associate each prospect with a unique message.

The formal argument proceeds in two steps. Lemma A.3 in Appendix A shows that, starting from a \mathcal{P}^n -optimal disclosure rule, one can construct a disclosure rule for \mathcal{P} that pools prospects deterministically and delivers a payoff close to the optimal payoff in \mathcal{P}^n . Roughly, when optimality in \mathcal{P}^n calls for probabilistically pooling a prospect under multiple messages, the disclosure rule in \mathcal{P} splits the corresponding “prospect” into nearby prospects; each such nearby prospect is then pooled deterministically. This splitting exploits the continuity of the type space. The disclosure rule, constructed in this way, is then verified to be optimal in \mathcal{P} by using a limit argument of Lemma A.4 in Appendix A.²³ The results of Lemma A.3 and Lemma A.4 combine in Theorem 3,

²³Lemmas A.3 and A.4 do not claim that, with the continuum, splitting prospects must lead to a strict improvement; they show only that whatever can be achieved by pooling a prospect under multiple messages can also be achieved by pooling it under a single message.

which describes \mathcal{P} -optimal disclosure and the associated effort schedule.

To state and prove Theorem 3, we need two more definitions. Definition 2 describes a function that partitions the prospect set into revealed singletons and pooled pairs.

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

In Definition 2, case (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 $(s^*, 1)$. Point s' is the point of discontinuity at which the matching function jumps upward from 0 to 1, thereby switching from taking values in one to taking values in the other interval of its codomain. That is, the matching function takes values in $(0, s_*]$ when $s \in [s_*, s')$ and in $[s^*, 1]$ when $s \in [s', s^*]$. Case (i) in Definition 2 prevails in Example 1, for which the prospect set is depicted in Figure 3c. An example of a prospect set corresponding to case (ii) is depicted in Figure 3b.²⁴

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

Definition 3. Under the **conjugate disclosure rule** induced by a matching function τ , the seller who receives the 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 $\theta_1 = s$ and the complementary probability to $\theta_1 = \tau(s)$, by Bayes' rule. When $\tau'(s) = 0$, the seller's announcement $\{s, \tau(s)\}$ induces the follower to assign probability one to $\theta_1 = s$.²⁵ Finally, when the range of τ omits some type, the seller reveals this type.

²⁴Supplementary Appendix B.4 further contrasts cases (i) and (ii).

²⁵Informally, when $\tau'(s) = 0$, the seller effectively pools a "small positive-measure interval" of types near s with the infinitesimal-measure type $\tau(s)$. The infinitesimal measure of $\tau(s)$ is further spread thinly over the positive measure of types near s , thereby endowing each message $\{s, \tau(s)\}$ with the infinitesimal odds of $\tau(s)$ relative to s .

We are now ready to state the main result:

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

(i) *is a conjugate disclosure rule; and*

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

Proof. See Appendix A. □

According to part (i) of Theorem 3, the optimality of the conjugate disclosure rule defies the pooling pattern common in several models of strategic disclosure, in which all types in a certain interval are pooled (e.g., Crawford and Sobel, 1982). In their footnote [11], RS conjecture that, with a continuum of prospects, one would be unable to dismiss interval pooling as nongeneric. By contrast, our model dismisses interval pooling for a continuous prospect set, which, while “nongeneric,” emerges in an economically interesting setting.²⁶ Nevertheless, RS’s results apply in our setting because, under Conditions 1 and 2, RS’s critical feature is preserved: no three prospects lie on the same line.

According to part (ii) of Theorem 3, the seller’s strategic information disclosure distorts the follower’s effort schedule by shifting its peak to the right of the peak of the first-best effort schedule. The overall (i.e., expected) effort is the same as in the first-best. The peak shifts rightwards because the seller strives to intensify the competition that the leader faces by inducing the follower to learn a lot when the leader’s type is high.

Corollary 2 shows how the optimal outcome can be implemented in an indirect mechanism. The mechanism uses the (possibly negative) tax

$$T^*(\theta_1) \equiv \int_0^{\theta_1} (F(s | a_c^*(m(\theta_1))) - F(s | a_c^*(m(s)))) ds, \quad \theta_1 \in \Theta_1, \quad (24)$$

where $m(\theta_1)$ is the seller’s message induced by type θ_1 , and $a^*(m(\theta_1))$ is the corresponding action of the follower.

²⁶Our prospect set is nongeneric if only because it is a one-dimensional curve in a two-dimensional space. On top of that, this curve is also convex, by Condition 2.

Corollary 2. *In an optimal mechanism, the seller*

1. *asks the leader to submit a bid, denoted by b , and charges him tax $T^*(b)$, defined in (24);*
2. *discloses to the follower the message prescribed by the optimal conjugate disclosure rule of Theorem 3 and asks him to submit a bid; and*
3. *allocates the item and assesses the payments according to the rules of the second-price auction.*

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

Proof. See Appendix A. □

In contrast to the first-best auction of Corollary 1, the optimal auction of Corollary 2 does not fully disclose the leader's bid and specifies a different tax schedule. As in the first-best case, the leader's tax is increasing in his bid if the follower's induced effort is increasing in the leader's bid and is decreasing otherwise. Thus, as in the first-best case, the leader's tax countervails the leader's motive to manipulate his bid so as to induce the follower to learn more.

4.4 The Seller's Optimal-Control Problem and the Euler Equation

By Theorem 3, the seller's optimal conjugate disclosure rule is described by a matching function. We demonstrate shortly that the matching function can be found by solving an optimal-control problem. The optimal-control representation is useful for both formal and numerical analyses.

The Optimal-Control Formulation

For the sake of parsimony, we illustrate the optimal-control representation when the optimal matching function takes the form shown in case (i) of Definition 2; the optimal-control representation for case (ii) is similar.

In case (i), $s_* = 0$, and the seller's objective function (21) can be written as

$$\int_0^{s_*} \pi(s) a^*(s) g(s) ds + \int_0^{s_*} \pi(\tau(s)) a^*(s) g(\tau(s)) \beta(s) ds + \int_{[s^*, 1] \setminus \text{range}(\tau)} \alpha(s) \pi(s) ds, \quad (25)$$

where $\beta \equiv -\tau'$ denotes the derivative of a matching function τ , and $a^*(s)$ denotes the follower's

normalized optimal action when the leader's type is $s \in [0, s^*]$.²⁷ The first integral in (25) 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 those types in $[s^*, 1]$ that the matching function pools with types in $[0, s^*]$. The third integral is the seller's payoff from those types in $[s^*, 1]$ that are fully revealed.

The first and last integrals in (25) copy the corresponding terms from the seller's objective function (21). The second integral is derived using the change-of-variables formula, according to which, for any type $s \in (0, s^*)$ and a "small" $ds > 0$, interval $(s, s + ds)$, whose probability is approximately $g(s) ds$, is pooled with interval $(\tau(s + ds), \tau(s))$, whose probability is approximately $g(\tau(s)) \beta(s) ds$.²⁸ The follower's normalized effort in (25) is derived from (13) by appealing to Bayes' rule and performing the change of variables:

$$a^*(s) = \frac{g(s) \alpha(s) + g(\tau(s)) \beta(s) \alpha(\tau(s))}{g(s) + g(\tau(s)) \beta(s)}, \quad s \in [0, s^*]. \quad (26)$$

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

Definition 4. When $s_* = 0$, the **seller's optimal-control problem** consists in maximizing (25) over s^* , a piecewise-continuous function $\beta : [0, s^*] \rightarrow \mathbb{R}_+$, and the implied piecewise-differentiable function $\tau : [0, s^*] \rightarrow [s^*, 1]$, which, together, induce a^* from (26), subject to $\tau(0) = 1$, $\tau(s^*) = s^*$, and, for almost all $s \in (0, s^*)$, $\tau'(s) = -\beta(s)$.

The Euler Equation

A solution to the seller's optimal-control problem induces an **optimal-prospect path**

$$\Pi \equiv \{(\pi^*(s), a^*(s)) : s \in [0, s^*]\},$$

where $\pi^*(s)$ is the seller's expected marginal benefit from the follower's action when message s is sent:

$$\pi^*(s) = \frac{g(s) \pi(s) + g(\tau(s)) \beta(s) \pi(\tau(s))}{g(s) + g(\tau(s)) \beta(s)}, \quad s \in [0, s^*]. \quad (27)$$

²⁷Equation (25) implicitly normalizes the set of messages to $[0, s^*] \cup ([s^*, 1] \setminus \text{range}(\tau))$. Any type θ_1 in $[0, s^*]$ generates message $m(\theta_1) = \theta_1$. Any fully revealed type θ_1 in $(s^*, 1]$ (i.e., a type in $(s^*, 1] \setminus \text{range}(\tau)$) generates message $m(\theta_1) = \theta_1$. Any pooled type θ_1 in $(s^*, 1]$ generates message $m(\theta_1) \in \tau^{-1}(\theta_1) \subset [0, s^*]$.

²⁸Indeed, the Taylor expansion implies $G(\tau(s + ds)) \approx G(\tau(s)) + G'(\tau(s)) \tau'(s) ds$.

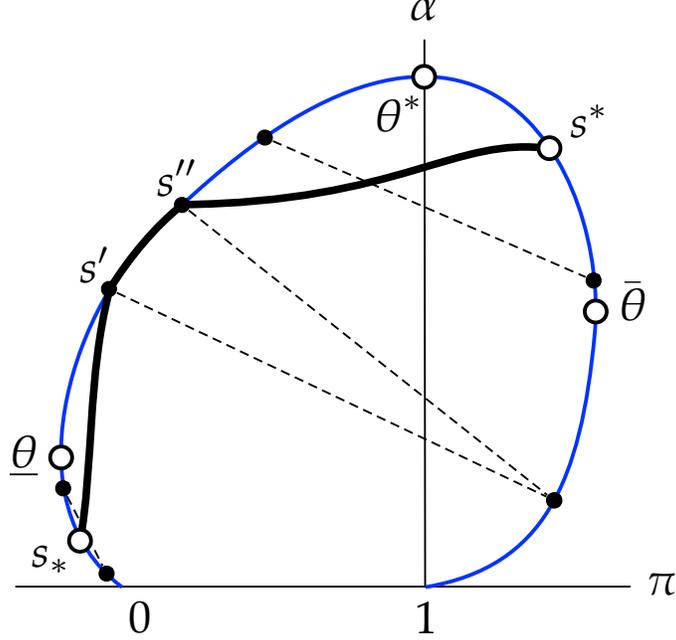


Figure 5: An optimal-prospect path, Π , for a “typical” convex prospect set, Γ , that satisfies Condition 1. The solid thin arc is Γ ; the solid thick curve is Π ; each dashed line segment is a pooling link. Each point on Π is a tuple containing the follower’s normalized optimal action, a^* , and the seller’s expected marginal benefit from that action, π^* . Each tuple is induced by pooling the prospects at the endpoints of the dashed link that passes through that tuple. Here, the leader’s types in $[s', s''] = \Gamma \cap \Pi$ are revealed. As per the Euler equation, where Π and the top pooling link intersect, the absolute values of their slopes are equal.

Figure 5 illustrates Π for a “typical” convex prospect set that satisfies Condition 1. In the figure, each depicted pooling link, for some $s < s^*$, connects the prospects induced by s and by $\tau(s)$, the optimal matching function evaluated at s . The pooling link induces the follower’s normalized effort $a^*(s)$, which is the ordinate of the intersection point of the optimal-prospect path and the pooling link. The corresponding abscissa is the seller’s expected marginal benefit, $\pi^*(s)$.

Any Π is nondecreasing, which is a necessary condition for optimality. Indeed, if Π had a strictly decreasing segment, then, by Lemma A.2 in Appendix A, it would be optimal to pool under a single message some of the messages that induced that segment. Any prospect in the intersection of Π and Γ is fully revealed.

A standard variational argument can be invoked to establish that, almost everywhere in the

interior of the convex hull of Γ , Π satisfies the **Euler equation**:²⁹

$$\frac{da^*}{d\pi^*} = \frac{\alpha - \alpha(\tau)}{\pi(\tau) - \pi'} \quad (28)$$

where the argument θ_1 has been suppressed. In (28), $da^*/d\pi^*$ is the slope of Π , and $(\alpha - \alpha(\tau)) / (\pi - \pi(\tau))$ is the slope of the corresponding pooling link. Thus, graphically, (28) requires that, wherever Π and a pooling link intersect, the absolute values of their slopes are equal (Figure 5 illustrates).

5 The Monotonicity Condition

So far, the analysis has been performed under the hypothesis that the identified solution to the seller's relaxed problem, which ignores the leader's monotonicity constraint, also solves the seller's full problem. Theorem 4 shows that this is indeed so if c is sufficiently large.

The probability that a type- θ_1 leader wins is the probability that $\theta_2 \leq \theta_1$ and equals

$$F(\theta_1 | a_c^*(\theta_1)) = a_c^*(\theta_1) F_H(\theta_1) + (1 - a_c^*(\theta_1)) F_L(\theta_1). \quad (29)$$

This probability depends on θ_1 both directly and indirectly, through $a_c^*(\theta_1)$. The leader's monotonicity condition requires $F(\theta_1 | a_c^*(\theta_1))$ to be weakly increasing in θ_1 . If a_c^* were fixed, the leader's monotonicity condition would be satisfied automatically, because F_H and F_L are c.d.f.s and, hence, weakly increasing. However, the dependence of a_c^* on θ_1 may threaten monotonicity on (θ^*, s^*) , where $F_H(\theta_1) < F_L(\theta_1)$, and a_c^* is increasing in θ_1 . In this case, a small increase in θ_1 , while (weakly) increasing the values of both F_H and F_L , shifts the weight in F towards F_H , the smaller of the two constituent c.d.f.s., making the overall direction of change in $F(\theta_1 | a_c^*(\theta_1))$ ambiguous.

Intuitively, on (θ^*, s^*) , the fact that the leader becomes stronger as θ_1 increases is at least partially offset by the fact that the follower also becomes stronger as the seller instructs him to learn more. Even though a more informed follower is not a stronger bidder in the first-order stochastic dominance sense, he has a higher chance of an extremely high type realization, which he needs to

²⁹The argument leading up to Lemma B.4 in Supplementary Appendix B.6 illustrates the derivation.

outbid a high-type leader.³⁰

To show that monotonicity is guaranteed to hold for a sufficiently large c , differentiate (29):

$$\frac{dF(\theta_1 | a_c^*(\theta_1))}{d\theta_1} = f_L(\theta_1) + \frac{a^{*'}(\theta_1)(F_H(\theta_1) - F_L(\theta_1)) + a^*(\theta_1)(f_H(\theta_1) - f_L(\theta_1))}{c}. \quad (30)$$

By inspection, for a sufficiently large c , (30) is nonnegative if f_L is bounded away from zero and from above, and if $a^{*'}$ is bounded from above on (θ^*, s^*) .³¹ Ascertaining the requisite boundedness of $a^{*'}$ is a delicate procedure because a^* depends on optimal disclosure, which is not an explicit function of the primitives. Condition 3 helps.

Condition 3. The functions g and f_L are bounded below away from zero; the functions g' , f_L , and f_H are bounded above; and

$$\lim_{\theta_1 \rightarrow 1} \frac{\alpha'(\theta_1)}{\pi'(\theta_1)} > 0.$$

Condition 3 is satisfied in the paper's leading example: Example 1 with the uniform G .

Theorem 4. *Under Condition 3, the leader's monotonicity condition holds if c is sufficiently large. Moreover, the leader's monotonicity condition satisfies a cut-off property: if monotonicity holds for some c , then it holds for any larger c .*

Proof. See Appendix A. □

Intuitively, in Theorem 4, a larger c attenuates the follower's effort, thereby making it less sensitive to θ_1 . As a result, when c is large, the effect of a higher θ_1 on the leader's probability of winning is dominated by the direct effect of his becoming a stronger bidder; the indirect effect, due to the change in the distribution of the follower's types associated with learning, is insignificant by comparison.

³⁰The seller escalates bidder competition so much because he is motivated by profit; first-best efficiency rules out such an escalation and guarantees monotonicity.

³¹There are no other boundedness conditions to attend to. Function $a^{*'}$ is guaranteed to be bounded from below by zero because a^* is guaranteed to be nondecreasing on $(0, s^*)$. Function a^* is bounded because $a^*(m) \equiv \mathbb{E}_{\theta_1|m} \left[\int_{\theta_1}^1 (F_L(s) - F_H(s)) ds \right]$ has values in the bounded interval $[0, 1 - \theta^*]$.

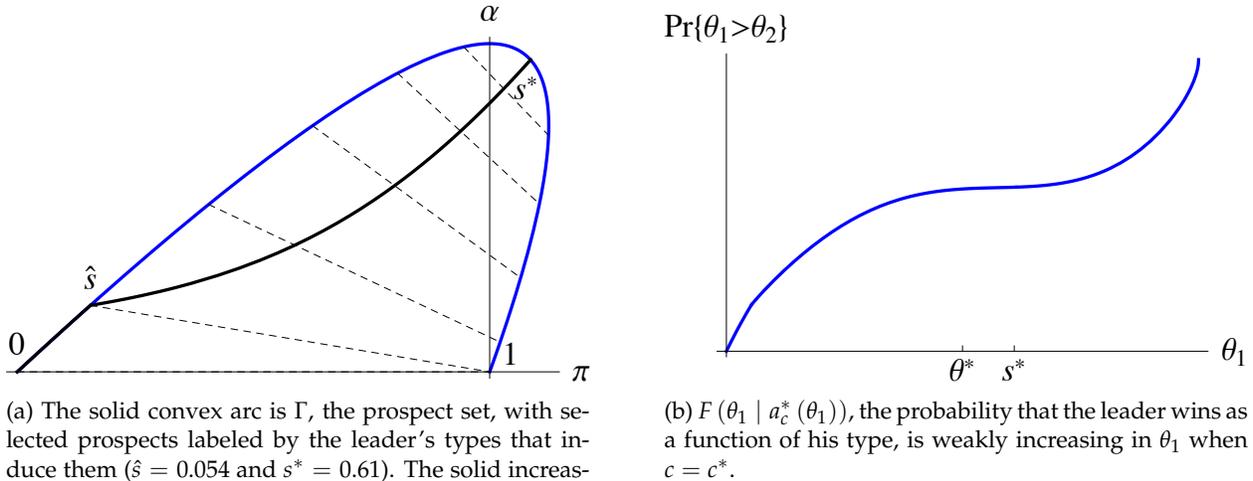


Figure 6: Optimal disclosure in Example 1 with the uniform c.d.f. G .

6 A Numerical Example

In this section, to recover the exact optimal disclosure rule, we solve the seller's optimal-control problem from Section 4.4 for Example 1 with the uniform c.d.f. G . To do so, we apply Hamiltonian techniques.³² The identified solution also solves the full problem, as we shall show.

Figures 2 and 6 report a numerical solution. The first-best and optimal effort schedules are in Figure 2. Consistent with Theorem 3, the areas below the two effort schedules coincide (i.e., $\mathbb{E}[a^*(m)] = \mathbb{E}[a(\theta_1)]$), and the optimal effort schedule is shifted rightwards relative to the first-best schedule.

Figure 6a plots selected pooling links for the optimal disclosure policy. The solid increasing curve is Π . When Π is in the interior of the convex hull of Γ , the Euler equation (28) holds; thus, wherever Π and a pooling link intersect, the absolute values of their slopes are equal.³³

Finally, at the solution to the optimal-control problem, the leader's monotonicity condition holds for any c . Indeed, Figure 6b confirms that, when $c = c^*$, the leader's probability of winning is nondecreasing in his type. By Theorem 4, because the leader's monotonicity condition holds for $c = c^*$, it also holds for any larger c —that is, for any c admissible according to (2).

³²Because the Hamiltonian analysis imposes the additional assumption of piecewise *continuous* differentiability of the matching function, and because we have been unable to show that the problem in Definition 4 is convex, the numerical "solution" we report is an informed guess, which has been verified to improve upon full disclosure and non-disclosure.

³³Consistent with Lemma B.6 in Supplementary Appendix B.6, $s^* < \hat{\theta}$, and, on $(s^*, 1]$, Π has no points in common with Γ .

7 Concluding Remarks

This paper reinterprets and extends the techniques of the Bayesian persuasion model of RS to study the distortions that the seller's strategic bid-disclosure introduces into an otherwise efficient sequential auction. The mapping of the seller's optimal-auction problem into the optimal-disclosure problem of RS relies on three assumptions: (i) the seller chooses the ex-post efficient allocation rule; (ii) the follower's learning effort is the probability with which he gains access to a more precise signal; and (iii) in the seller's relaxed problem, the leader's monotonicity constraint holds. Relaxing any of these assumptions would call for techniques that differ substantially from those that we use in our analysis.

A Appendix: Omitted Proofs

Proof of Lemma 1

The leader takes no private action, and so, without loss of generality, the seller contacts him only once, to elicit his type. The follower exerts effort just once, and so the seller contacts him first, to inform the effort choice, and then contacts him again, to elicit his type. Because of the option value of influencing the follower's effort by disclosing something to him about the leader's type, without loss of generality, the seller contacts the follower after receiving the leader's report.

Proof of Theorem 1

Integrating (4) by parts gives:

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

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

$$B(\theta_1, a) \equiv R(\theta_1) - C'(a), \tag{A.1}$$

where

$$R(\theta_1) \equiv - \int_{\theta_1}^1 \frac{\partial F(\theta_2 | a)}{\partial a} d\theta_2 \tag{A.2}$$

is the return to information acquisition (independent of a because F is linear in a), and $C'(a)$ is the marginal cost of the follower's information-acquisition effort. For any $a \in A$, $B(1, a) = -C'(a) < 0$, and, by part (i) of Condition 1, $B(0, a) = -C'(a) < 0$. Hence, when $\theta_1 = 1$ or $\theta_1 = 0$, the follower exerts zero effort at the first-best.

Part (ii) of Condition 1 implies that, for any $a \in A$, $B(\theta_1, a)$ is strictly increasing in θ_1 for $\theta_1 < \theta^*$ and is strictly decreasing in θ_1 for $\theta_1 \geq \theta^*$. Hence, by the Monotone Selection Theorem (Milgrom and Shannon, 1994), the first-best effort, $\alpha_c(\theta_1)$, is weakly increasing in θ_1 for $\theta_1 < \theta^*$ and is weakly decreasing in θ_1 for $\theta_1 \geq \theta^*$, independently of the exact functional form of C . Moreover, the dependence is strict when $\alpha_c(\theta_1) \in (0, 1)$ (Edlin and Shannon, 1998), which can be shown to be the case for $\theta_1 \in (0, 1)$ as long as the convex C has $C'(0) = 0$ and $C'(1)$ sufficiently large, as the quadratic C and condition (2) indeed imply. From the first-order condition $R(\theta_1) = C'(\alpha_c(\theta_1))$, the quadratic C delivers an expression for the first-best effort in (5).

Proof of Corollary 1

For the follower, it is a weakly dominant strategy to bid his type in the second-price auction. If he bids his type, he also finds it optimal to exert the first-best effort; his expected payoff in the mechanism described in the corollary has been constructed to coincide with the planner's maximand in the surplus-maximization problem (4). The follower's payoff is nonnegative because he can obtain a nonnegative payoff by bidding in the second-price auction without having exerted any effort.

The leader chooses his bid, b , to maximize

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

Integration by parts and the substitution of T from (6) transforms the display above into

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

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

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

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

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

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

Proof of Theorem 2

Showing that neither full disclosure, nor non-disclosure is optimal requires two lemmas: Lemma A.1 and Lemma A.2.

For Lemma A.1, call functions α and π **ordered** on a subset S of the prospect set Γ if

$$(\alpha(s') - \alpha(s)) (\pi(s') - \pi(s)) \geq 0$$

for almost all $(\pi(s), \alpha(s)), (\pi(s'), \alpha(s')) \in S$.

Lemma A.1. *Full disclosure on $S \subseteq \Gamma$ is optimal if and only if α and π are ordered on S .*

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

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

In the display above, multiplying the parentheses, defining $|I| \equiv \int_I ds$, and rearranging yields³⁴

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

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

*Sufficiency:*³⁵ Suppose that α and π are ordered on S . By contradiction, suppose that, with probability-density p , prospect $(\pi(s), \alpha(s))$ occurs and induces a message m . Suppose, also, that, with probability-density p' , prospect $(\pi(s'), \alpha(s'))$ with $s' \neq s$ occurs and induces the same message m .³⁶ The seller's gain from pooling the two prospects under message m relative to revealing each of them is

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

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

Taking $S = \Gamma$ in Lemma A.1 immediately yields

³⁴The obtained inequality is a continuous version of Chebyshev's sum inequality.

³⁵The sufficiency argument is Lemma [1] of RS and is included for completeness.

³⁶The message m may also be induced by some other prospect. Either prospect may also induce some other message.

Corollary 3. *Full disclosure is optimal if and only if α and π are ordered on Γ .*

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

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

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

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

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

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

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

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

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

Corollary 4. *Non-disclosure is optimal if and only if the prospect set Γ lies on a nonincreasing line.*

One can now prove Theorem 2.

Recall that

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

By Theorem 1, α is uniquely maximized at $\theta_1 = \theta^* \in (0, 1)$. By part (ii) of Condition 1 and by the above display, $\theta_1 < \theta^* \implies \pi(\theta_1) < 0$ and $\theta_1 > \theta^* \implies \pi(\theta_1) > 0$. Thus, α and π are not ordered, and Corollary 3 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; the reason is that, 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 4 implies that non-disclosure is suboptimal.

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

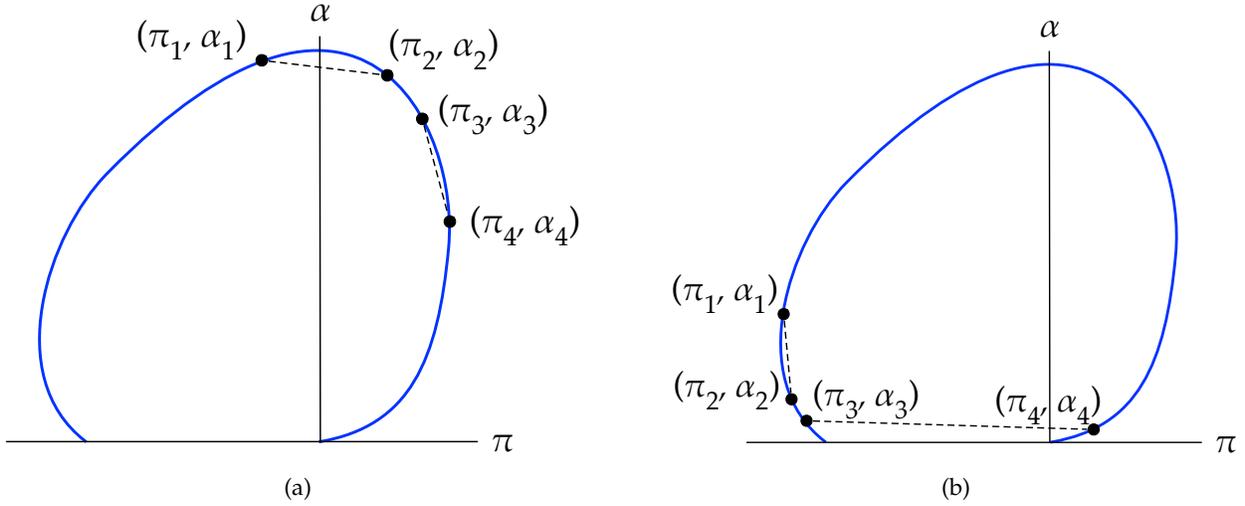


Figure A.1: Contradiction hypotheses for Step 1 in the proof of Lemma 2. Each dashed link denotes a pair of prospects that are pooled under the same message. In neither panel can the two links be crossed by an upward-sloping line.

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

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

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

By Lemma A.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.

Proof of Lemma 2

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

Step 1: Take any pooling link that has a northwest prospect, denoted by (π_3, α_3) .³⁷ Then, optimality rules out the existence of a link that pools two prospects, say (π_1, α_1) and (π_2, α_2) , such that each of these prospects lies to the northwest of (π_3, α_3) . (See both panels of Figure A.1.)

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

³⁷We define the cardinal directions in the (π, α) -space in the obvious manner. For instance, a point (π_3, α_3) is **northwest** of point (π_4, α_4) if $\pi_3 < \pi_4$ and $\alpha_3 > \alpha_4$.

By contradiction, suppose that one can pick four prospects $\{(\pi_i, \alpha_i)\}_{i=1,2,3,4}$ such that prospects (π_1, α_1) and (π_2, α_2) are optimally pooled under some message—say, m ; prospects (π_3, α_3) and (π_4, α_4) are optimally pooled under another message—say m' ; and either (a) $\pi_4 \geq \pi_3 \geq \pi_2 > \pi_1$ and $\alpha_1 \geq \alpha_2 \geq \alpha_3 > \alpha_4$ or (b) $\pi_4 > \pi_3 \geq \pi_2 \geq \pi_1$ and $\alpha_1 > \alpha_2 \geq \alpha_3 \geq \alpha_4$ holds.³⁸ For each $i = 1, 2, 3, 4$, let p_i denote the joint probability that (i) type x_i —which is assumed to invoke prospect (π_i, α_i) —is realized; and (ii) type x_i induces either message m or m' (whichever is appropriate).

The seller's expected gain from using the distinct messages m and m' (as in Figure A.1) relative to pooling all four prospects under a single message is

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

which can be rearranged to give

$$\Delta = \frac{(p_1 + p_2)(p_3 + p_4)}{p_1 + p_2 + p_3 + p_4} \left(\frac{p_3\pi_3 + p_4\pi_4}{p_3 + p_4} - \frac{p_1\pi_1 + p_2\pi_2}{p_1 + p_2} \right) \left(\frac{p_3\alpha_3 + p_4\alpha_4}{p_3 + p_4} - \frac{p_1\alpha_1 + p_2\alpha_2}{p_1 + p_2} \right) < 0, \quad (\text{A.3})$$

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

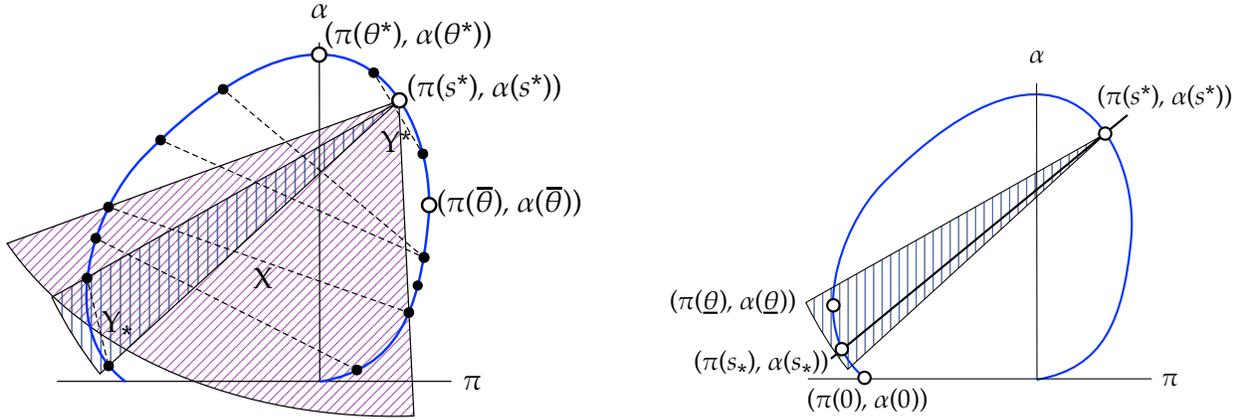
Step 2: *There exist an $s_* \in [0, \theta]$ and an $s^* \in [\theta^*, \bar{\theta}]$ that satisfy part (i) of the lemma.*

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

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

We now turn to the construction of s_* . Draw a sequence of pairs of secants of Γ . Each secant in the pair passes through the prospect $(\pi(s^*), \alpha(s^*))$ and either endpoint of a pooling link. Each

³⁸Cases (a) and (b) differ in the placement of the strict and weak inequalities. Case (a) prevails if the leader's types x_1, x_2, x_3 , and x_4 in Θ_1^n satisfy $0 < x_1 < x_2 \leq x_3 < x_4 < 1$, $x_2 > \theta^*$, and $x_3 < \bar{\theta}$. Case (b) prevails if the leader's types satisfy either $\theta^* > x_1 > x_2 \geq x_3 > x_4 > 0$ or $x_4 > \theta^* > x_1 > x_2 \geq x_3 > 0$ and $x_2 < \bar{\theta}$.



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

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

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

secant pair delimits a cone in the (π, α) -space, as depicted in Figure A.2a. By Step 1, and because Γ is a convex curve, \prec induces an inclusion order on the cones generated by the pooling links. In particular, the cones associated with \prec -smaller pooling links are smaller in the inclusion sense. As a result, the intersection of all the cones (depicted in Figure A.2b) is nonempty and is the cone induced by the \prec -minimal pooling link, denoted by Y_* .³⁹ Consequently, the intersection of all cones contains the arc of Γ that lies in the lower half-space of Y_* . Because Y_* is nonincreasing, this arc contains at least one prospect induced by some type in $[0, \underline{\theta}]$. Define s_* to be an arbitrary such type.

The line that passes through the prospects induced by s^* and s_* constructed above (Figure A.2b) is nondecreasing and traverses all pooling links. This line partitions the prospects in those induced by types in $[s_*, s^*] \cap \Theta_1^n$ and those induced by types in $([0, s_*] \cup [s^*, 1]) \cap \Theta_1^n$. By construction of s^* and s_* , every prospect in Γ^n is either pooled with a prospect in the other element of this partition or is fully revealed. Hence, part (i) of the lemma follows.

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

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

³⁹Just as Y^*, Y_* exists because Γ^n , and, thus, the number of pooling links, is finite.

\prec -smaller than the link connecting the prospects induced by $\{x_3, x_4\}$.⁴⁰ For the described pooling to be optimal, it must be, in particular, that the seller does not gain from pooling $\{x_1, x_2, x_3, x_4\}$ all under the same message. That is, the inequality in (A.3) must be reversed to yield $\Delta \geq 0$. The only way to satisfy $\Delta \geq 0$ is to have⁴¹

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

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

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

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

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

Proof of Theorem 3

The proof of Theorem 3 relies on two preliminary results: Lemma A.3 and Lemma A.4.

Lemma A.3 shows that, for \mathcal{P} , one can construct a conjugate disclosure rule that delivers to the seller a payoff approximately equal to the seller's optimal payoff in \mathcal{P}^n when n is large. The lemma's statement uses the big-O notation, in which O stands for a function that satisfies $\limsup_{n \rightarrow \infty} |O(2^{-n})/2^{-n}| < \infty$.

Lemma A.3. *Suppose that V^n is the seller's optimal payoff in the discrete disclosure problem \mathcal{P}^n . Then, there exists a conjugate disclosure rule that delivers to the seller payoff $V^n + O(2^{-n})$ in the continuous disclosure problem \mathcal{P} .*

⁴⁰If $x_3 = x_4$, extend \prec (to compare a link and a fully revealed prospect) so that the $\{x_1, x_2\}$ -pooling link is \prec -smaller than the $\{x_3, x_4\}$ -prospect if x_3 is in the upper half-space of the $\{x_1, x_2\}$ -pooling link. Similarly, if $x_1 = x_2$, extend \prec so that the $\{x_1, x_2\}$ -prospect is \prec -smaller than the $\{x_3, x_4\}$ -pooling link if x_1 is in the lower half-space of the $\{x_3, x_4\}$ -pooling link.

⁴¹Indeed, $\left(\frac{p_i\pi_i + p_{i+1}\pi_{i+1}}{p_i + p_{i+1}}, \frac{p_i\alpha_i + p_{i+1}\alpha_{i+1}}{p_i + p_{i+1}}\right)$ is an average point on the link connecting prospects (π_i, α_i) and $(\pi_{i+1}, \alpha_{i+1})$, $i = 1, 3$. By $\Delta \geq 0$, the two averages must be product-ordered. Because, by normalization, the link connecting (π_1, α_1) and (π_2, α_2) is \prec -smaller than the link connecting (π_3, α_3) and (π_4, α_4) , the only way the averages can be ordered (and a picture makes this clear) is if the average for $i = 1$ is weakly smaller than the average for $i = 3$, which is what the displayed pair of inequalities advocates.

Proof. Let prospect $(\pi(y_i), \alpha(y_i))$ be referred to as prospect i and denote it by (π_i, α_i) . By Lemma 2 and Facts 1–5, a \mathcal{P}^n -optimal disclosure rule can be represented by a matrix $\mathbf{p}^n \equiv [p_{ij}]_{i,j \in \{1, \dots, 2^n\}}$, whose typical element p_{ij} is the joint probability that prospect i arises and that it induces the message that pools prospects i and j . The probability that prospect i arises and is fully revealed is denoted by p_{ii} . The probability that prospect i arises is denoted by p_i and equals $\sum_{j \in \{1, \dots, 2^n\}} p_{ij}$, which is the joint probability that prospect i arises and either is pooled with any other prospect or is fully revealed. Because, by Fact 5, a prospect cannot be fully revealed sometimes and pooled at other times, $p_{ii} > 0$ implies that $p_{ii} = p_i$ (i.e., $p_{ij} = 0$ for every $j \neq i$). A prospect can be pooled with more than one other prospect (depending on the realized message); that is, $p_{ij} > 0$ does not imply that $p_{ij} = p_i$.

The value of problem \mathcal{P}^n is denoted by V^n . To define this value, first define n_* and n^* , the prospect indices that correspond to the threshold types s_* and s^* defined in Lemma 2:

$$n_* \equiv \min \{i : y_i \geq s_*\} \quad (\text{A.4})$$

and

$$n^* \equiv \max \{i : y_i \leq s^*\}. \quad (\text{A.5})$$

With this notation,

$$V^n \equiv \max_{\mathbf{p}^n} \left\{ \sum_{i=1}^{2^n} p_{ii} \pi(y_i) \alpha(y_i) + \sum_{i=n_*}^{n^*} \sum_{j \in \{1, \dots, n_*-1\} \cup \{n^*+1, \dots, 2^n\}} (p_{ij} + p_{ji}) \frac{p_{ij} \pi(y_i) + p_{ji} \pi(y_j)}{p_{ij} + p_{ji}} \frac{p_{ij} \alpha(y_i) + p_{ji} \alpha(y_j)}{p_{ij} + p_{ji}} \right\}, \quad (\text{A.6})$$

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

Henceforth, let \mathbf{p}^n be the \mathcal{P}^n -optimal disclosure rule, which attains the maximum in (A.6). This rule will be used to construct an approximately \mathcal{P} -optimal disclosure rule. Roughly, if in \mathcal{P}^n , \mathbf{p}^n pools prospects i and j only with each other (i.e., $\sum_k p_{ik} = p_{ij}$ and $\sum_k p_{jk} = p_{ji}$), then in \mathcal{P} , the intervals $(y_{i-1}, y_i]$ and $(y_{j-1}, y_j]$ will be “linked” pointwise, by pooling every element in $(y_{i-1}, y_i]$ with a corresponding element in $(y_{j-1}, y_j]$ according to some matching function.⁴² If in \mathcal{P}^n , \mathbf{p}^n only sometimes pools prospects i and j with each other (i.e., $\sum_k (p_{ik} + p_{jk}) > p_{ij} + p_{ji}$), then in \mathcal{P} , the intervals $(y_{i-1}, y_i]$ and $(y_{j-1}, y_j]$ are divided into subintervals, and only one subinterval in $(y_{i-1}, y_i]$ is linked pointwise with a subinterval in $(y_{j-1}, y_j]$.

To make the linking procedure precise, define $P_i \equiv \{j : p_{ij} > 0\}$ to be the set of prospects that \mathbf{p}^n pools with prospect i , $i \in \{1, \dots, 2^n\}$. If $P_i = \{i\}$, prospect i is revealed in \mathcal{P}^n . Partition interval $(y_{i-1}, y_i]$ into a collection of $|P_i|$ subintervals⁴³

$$\mathcal{C}_i \equiv \left\{ \left(\underline{b}_{ij}, \bar{b}_{ij} \right) : j \in P_i \right\} \quad (\text{A.7})$$

⁴²Recall 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 .

⁴³Here, $|P_i|$ denotes the number of elements in the set P_i .

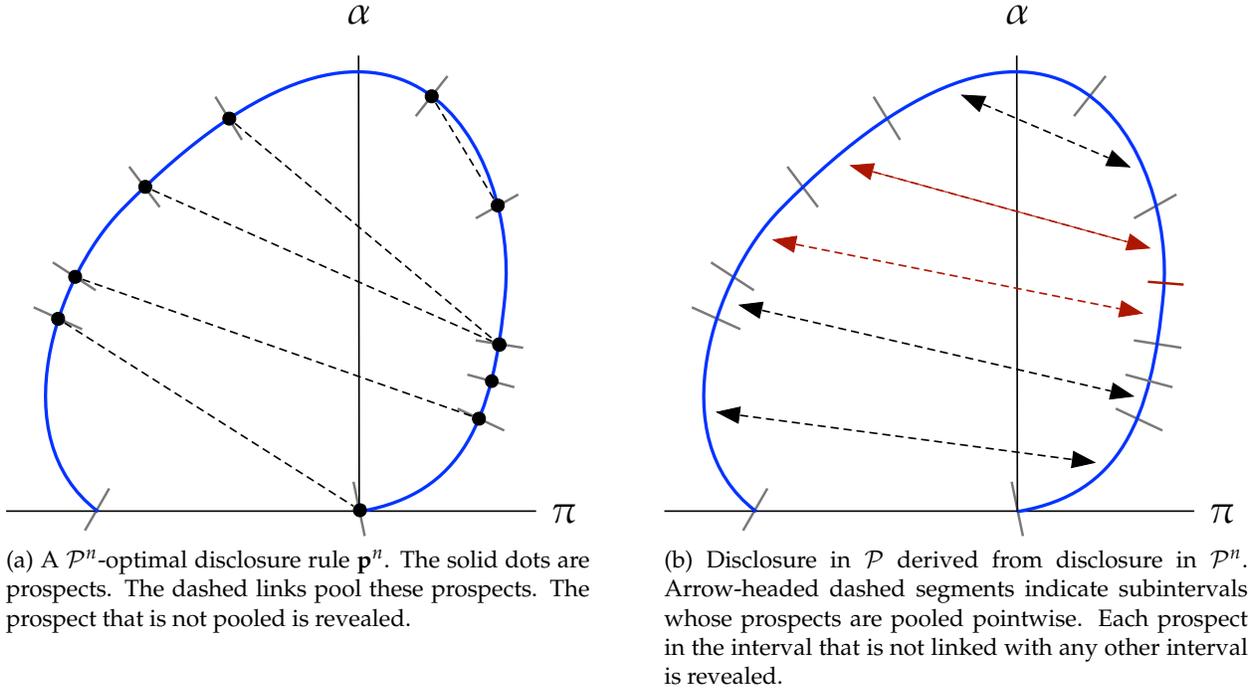


Figure A.3: An optimal disclosure rule in the discrete problem \mathcal{P}^n is used to construct a disclosure rule in the continuous problem \mathcal{P} .

so that $G(\bar{b}_{ij}) - G(\underline{b}_{ij}) = p_{ij}$, and so that whenever \mathbf{p}^n pools prospects i and j , one can draw a link between element $(\underline{b}_{ij}, \bar{b}_{ij}]$ in \mathcal{C}_i and element $(\underline{b}_{ji}, \bar{b}_{ji}]$ in \mathcal{C}_j in such a manner that no two links intersect. If $|P_i| = 1$, the only subinterval is the interval $(y_{i-1}, y_i]$ itself, which either is linked to some (sub)interval or remains unlinked. The construction of the links between (sub)intervals is illustrated in Figure A.3.

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 (A.4)):

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

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

p_{ij}/p_{ji} , and $\tau(\underline{b}_{ij}) = \bar{b}_{ji}$ and $\tau(\bar{b}_{ij}) = \underline{b}_{ji}$.⁴⁴ Otherwise, go to Step 3.

3. If $i < n^*$ (where n^* is defined in (A.5)), increment i by 1 and go to Step 1; otherwise, terminate.⁴⁵

Any interval $(y_{i-1}, y_i]$ whose elements are revealed contributes to the seller's payoff (23) amount

$$\int_{y_{i-1}}^{y_i} \pi(s) \alpha(s) g(s) ds \equiv p_{ii} \pi(z_i) \alpha(z_i), \quad (\text{A.8})$$

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.⁴⁶

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

$$\begin{aligned} & \int_{\underline{b}_{ij}}^{\bar{b}_{ij}} \frac{g(s) \pi(s) + \beta(s) g(\tau(s)) \pi(\tau(s))}{g(s) + \beta(s) g(\tau(s))} \frac{g(s) \alpha(s) + \beta(s) \alpha(\tau(s))}{g(s) + \beta(s) g(\tau(s))} (g(s) + \beta(s) g(\tau(s))) ds \\ & \equiv (p_{ij} + p_{ji}) \frac{g(z_{ij}) \pi(z_{ij}) + \beta(z_{ij}) g(z_{ji}) \pi(z_{ji})}{g(z_{ij}) + \beta(z_{ij}) g(z_{ji})} \frac{g(z_{ij}) \alpha(z_{ij}) + \beta(z_{ij}) \alpha(z_{ji})}{g(z_{ij}) + \beta(z_{ij}) g(z_{ji})} \\ & = (p_{ij} + p_{ji}) \frac{p_{ij} \pi(z_{ij}) + p_{ji} \pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij} \alpha(z_{ij}) + p_{ji} \alpha(z_{ji})}{p_{ij} + p_{ji}}, \quad (\text{A.9}) \end{aligned}$$

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

Assembling the contributions (A.8) and (A.9) gives the value of the seller's objective function

⁴⁴When the c.d.f. G is uniform, the sought τ is linear: $\tau(s) = \bar{b}_{ji} - (s - \underline{b}_{ij}) p_{ji}/p_{ij}$, where $p_{ij} = \bar{b}_{ij} - \underline{b}_{ij}$ and $p_{ji} = \bar{b}_{ji} - \underline{b}_{ji}$. For a general G , set up the initial-value problem $\tau' = -p_{ji}g(s)/(p_{ij}g(\tau))$ on $[\underline{b}_{ij}, \bar{b}_{ij}]$ subject to $\tau(\underline{b}_{ij}) = \bar{b}_{ji}$. Because the right-hand side of the problem's ordinary differential equation (ODE) is continuous in (s, τ) , the Peano existence theorem implies the existence of a solution. The solution is strictly decreasing because the right-hand side of the ODE is negative. To see that the solution satisfies $\tau(\bar{b}_{ij}) = \underline{b}_{ji}$, rewrite the ODE as $-g(\tau(s)) \tau'(s) = g(s) p_{ji}/p_{ij}$ and integrate to obtain

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

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

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

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

⁴⁵The non-linked intervals indexed by $i > n^*$ or $i < n_*$ do not affect the matching function; they automatically translate into discontinuities. The linked intervals indexed by $i > n^*$ or $i < n_*$ are accounted for when the intervals indexed by $i \in \{n_*, n_* + 1, \dots, n^*\}$ are considered.

⁴⁶See http://en.wikipedia.org/wiki/Mean_value_theorem#First_mean_value_theorem_for_integration

(23), for problem \mathcal{P} , under the disclosure rule induced by τ , constructed form \mathbf{p}^n :

$$\hat{V}^n \equiv \sum_{i=1}^{2^n} p_{ii} \pi(z_i) \alpha(z_i) + \sum_{i=n^*}^{n^*} \sum_{j \in \{1, \dots, n^*-1\} \cup \{n^*+1, \dots, 2^n\}} (p_{ij} + p_{ji}) \frac{p_{ij} \pi(z_{ij}) + p_{ji} \pi(z_{ji})}{p_{ij} + p_{ji}} \frac{p_{ij} \alpha(z_{ij}) + p_{ji} \alpha(z_{ji})}{p_{ij} + p_{ji}}. \quad (\text{A.10})$$

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 $z \in (0, 1)$:

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

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:

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

which are substituted into (A.10) to obtain

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

as desired. \square

Lemma A.3 does not rule out a discontinuity: the possibility that, in \mathcal{P} , the seller can improve upon the conjugate disclosure rule that is the limit of \mathcal{P}^n -optimal disclosure rules as n goes to infinity. Lemma A.4 rules out this discontinuity by showing that the value in \mathcal{P} is no greater than the limit of the values in \mathcal{P}^n as n increases.

Lemma A.4. *The continuous disclosure problem \mathcal{P} has a solution, which induces the value denoted by V^* . The discrete disclosure problems in the sequence $\{\mathcal{P}^n : n \geq 1\}$ have solutions, which induce the corresponding sequence of values denoted by $\{V^n\}$. Furthermore, $V^* \leq \liminf_{n \rightarrow \infty} V^n$.*

Proof. Preliminary Definitions

Normalize the set of the seller's messages by setting it equal to the set of the follower's posterior probability distributions: $M \equiv \Delta\Theta_1$, where $\Delta\Theta_1$ denotes the set of Borel probabilities on the space of leader's types $\Theta_1 = [0, 1]$. The space $\Delta\Theta_1$ is a compact metric space when endowed with the topology of weak convergence.⁴⁷ Let Δ ($\Delta\Theta_1$) denote the space of probability measures on the

⁴⁷The topology of weak convergence is metrizable under the Prohorov metric.

subsets of $\Delta\Theta_1$. Like $\Delta\Theta_1$, the space $\Delta(\Delta\Theta_1)$ is also a compact metric space when endowed with the topology of weak convergence.

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

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

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

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

$$V^* \equiv \max_{\hat{\nu} \in \Delta(\Delta\Theta_1)} \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) d\hat{\nu} \quad \text{s.t.} \quad \int_{\Delta\Theta_1} P d\hat{\nu} = P^0, \quad (\text{A.11})$$

where

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

Note that both $\bar{\pi}(P)$ and $a^*(P)$ are continuous in P . To see this, take an arbitrary sequence $\{P_k\}$ of probability measures on Θ_1 that converge weakly to P . Because Lebesgue measurable functions π and α can have at most a countable number of discontinuity points, the set of 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.11) exists. Let ν^* denote this solution.

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

$$V^n \equiv \max_{\hat{\nu}_n \in \Delta(\Delta\Theta_1^n)} \int_{\Delta\Theta_1^n} \bar{\pi}(P_n) a^*(P_n) d\hat{\nu}_n \quad \text{s.t.} \quad \int_{\Delta\Theta_1^n} P_n d\hat{\nu}_n = P_n^0, \quad (\text{A.12})$$

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

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

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.12). 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 $\Delta(\Delta\Theta_1)$ because, for each n , $P_n \in \Delta\Theta_1^n \subseteq \Delta\Theta_1$ and thus $\nu_n \in \Delta(\Delta\Theta_1^n) \subseteq \Delta(\Delta\Theta_1)$.

The Proof of the Lemma

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

$$\int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{n'} \rightarrow \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv. \quad (\text{A.13})$$

By contradiction, suppose that v , the limit of $\{v_{n'}\}$, does not solve the continuous problem (A.11) and that:

$$\epsilon \equiv \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv^* - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv > 0. \quad (\text{A.14})$$

Let $\mathcal{N} \equiv \{n'_1, n'_2, \dots\}$ be the set indexing the convergent sequence $\{v_{n'}\}$. Because of convergence (A.13), one can choose an $N \in \mathcal{N}$ such that for all $n' \geq N$, $n' \in \mathcal{N}$:

$$\left| \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{n'} - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv \right| < \frac{\epsilon}{2}. \quad (\text{A.15})$$

Because space $\Delta (\Delta\Theta_1)$ is separable and because v^* solves the seller's maximization problem (A.11) with the leader's type space Θ_1 , one can choose an $\hat{N} \in \mathcal{N}$ such that for any $n' \geq \hat{N}$, $n' \in \mathcal{N}$, there exists an approximation $v_{n'}^*$ to v^* with a support on $\Theta_1^{n'}$:

$$0 \leq \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv^* - \int_{\Delta\Theta_1} \bar{\pi}(P) a^*(P) dv_{n'}^* < \frac{\epsilon}{2}. \quad (\text{A.16})$$

Proof that such an approximation exists is in Supplementary Appendix B.5.

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

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

where the term in the first parenthesis is less than $\epsilon/2$ by (A.15); the term in the second parenthesis equals ϵ by the contradiction hypothesis (A.14); and the term in the third parenthesis is less than $\epsilon/2$ by (A.16). The inequality is a contradiction, however, because $v_{\bar{N}}$ solves the seller's discrete maximization problem with the leader's type space $\Theta_1^{\bar{N}}$. Hence,

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

The argument above shows that every convergent subsequence $\{v_{n'}\}$ converges weakly to a limit that delivers a payoff at least as high as V^* . Consequently, it must be the case that $\liminf_{n \rightarrow \infty} V^n \geq$

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

We are now ready to prove Theorem 3. Lemmas A.3 and A.4 imply the optimality of the conjugate disclosure rule and, with it, part (i) of the theorem.

For part (ii) of the theorem, note that Lemma 2's part (ii) and Lemma A.4 imply that the induced normalized 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 (13)), $a^*(s^*) \leq \max_{\theta_1} \alpha(\theta_1) = \alpha(\theta^*)$, 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 that $\mathbb{E}[a^*(m)] = \mathbb{E}[\mathbb{E}_{\theta_1|m}[\alpha(\theta_1)]] = \mathbb{E}[\alpha(\theta_1)]$, where the last equality is by the Law of Iterated Expectations.

Proof of Corollary 2

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

The leader chooses his bid by solving

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

where $a_c^*(m(b))$ is the follower's optimal action conditional on the message $m(b)$, which the seller sends when the leader bids b . Integration by parts and the substitution of T^* from (24) transforms the maximand above into

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

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

Proof of Theorem 4

As the discussion that precedes the statement of Theorem 4 establishes, to show that monotonicity holds for a sufficiently large c , it suffices to show that $a^{s'}$ is bounded above on (θ^*, s^*) . We will show that $a^{s'}$ is indeed bounded because (i) Π is connected, by Lemma B.3 in Supplementary Appendix B.6; and (ii) Π and Γ are disjoint on $(s^*, 1]$, by Lemma B.6 in Supplementary Appendix B.6.

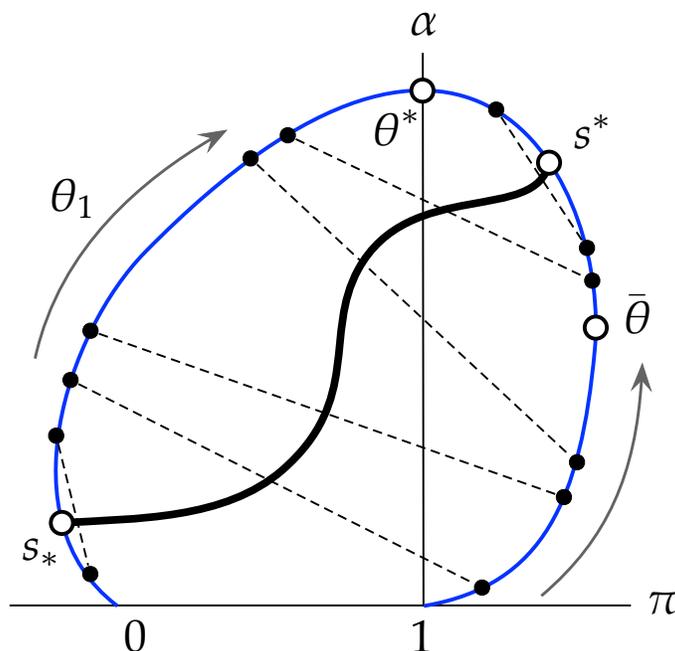


Figure A.4: When Π (the solid thick curve) and Γ (the convex curve) are disjoint, as θ_1 increases on $(0, s^*)$ at a bounded rate, the corresponding pooling segment's endpoints move at a bounded rate, and, hence, so does a^* , the ordinate of the segment's intersection with Π .

To establish the sufficiency of (i) and (ii) for the boundedness, consider Figure A.4, which depicts a connected Π that is disjoint from Γ on $(s^*, 1]$. As θ_1 increases on (s_*, s^*) at a bounded rate, the northwestern end of the corresponding pooling link moves at a bounded rate (because α and π have bounded derivatives, guaranteed by Condition 3). This movement must translate into a bounded-rate movement on $(s^*, 1]$ of the southeastern end of the pooling link. Indeed, if it did not, the southeastern end would have to trace out Π that would have to lie on Γ —the configuration that is ruled out by the hypothesis that Π is disjoint from Γ on $(s^*, 1]$. Because both ends of the pooling segment move at a finite rate, so must the pooling segment's intersection point with Π . In particular, this intersection's ordinate, a^* on (θ^*, s^*) , moves at a finite rate. In other words, a^* is bounded on (θ^*, s^*) .

References

- Bergemann, Dirk and Juuso Välimäki**, "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 2002, 70 (3), 1007–1033.
- **and Juuso Valimäki**, "Information in Mechanism Design," in Richard Blundell, Whitney K. Newey, and Torsten Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Vol. I, Cambridge University Press, 2006, chapter 5, pp. 186–221.

— **and Martin Pesendorfer**, “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 2007, 137 (1), 580–609.

Bikhchandani, Sushil and Chi-Fu Huang, “Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets,” *The Review of Financial Studies*, 1989, 2 (3), 311–339.

Billingsley, Patrick, *Convergence of probability measures*, Wiley, 1968.

Blackwell, David, “Comparisons of Experiments,” *Proceedings of the Second Berkeley Symposium in Mathematical Statistics*, 1951, pp. 93–102.

—, “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 1953, 24 (2), 265–272.

Calzolari, Giacomo and Alessandro Pavan, “Monopoly with Resale,” *RAND Journal of Economics*, 2006, 32 (2), 362–375.

— **and —**, “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, 2006, 130 (1), 168–204.

Compte, Olivier and Philippe Jehiel, “Auctions and information acquisition: sealed bid or dynamic formats?,” *The RAND Journal of Economics*, 2007, 38 (2), 355–372.

Crawford, Vincent P. and Joel Sobel, “Strategic Information Transmission,” *Econometrica*, 1982, 50 (6), 1431–1451.

Crémer, Jacques, Yossi Spiegel, and Charles Z. Zheng, “Auctions with Costly Information Acquisition,” *Economic Theory*, 2009, 38 (1), 41–72.

Edlin, Aaron S. and Chris Shannon, “Strict Monotonicity in Comparative Statics,” *Journal of Economic Theory*, 1998, 81 (1), 201–219.

Eso, Peter and Balázs Szentes, “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 2007, 74 (3), 705–31.

Gale, Ian L., Donald B. Hausch, and Mark Stegeman, “Sequential Procurement with Subcontracting,” *International Economic Review*, 2000, 41 (4), 989–1020.

Ganuzza, Juan-José, “Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations,” *RAND Journal of Economics*, 2004, 35 (3), 583–598.

— **and José S. Penalva**, “Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions,” *Econometrica*, 2010, 78 (3), 1007–1030.

Gershkov, Alex and Balázs Szentes, “Optimal Voting Scheme with Costly Information Acquisition,” *Journal of Economic Theory*, 2009, 144 (1), 36–68.

- Gupta, Madhurima and Bernard Lebrun**, "First Price Auctions with Resale," *Economics Letters*, 1999, 64 (2), 181–185.
- Haile, Philip A.**, "Auctions with Private Uncertainty and Resale Opportunities," *Journal of Economic Theory*, 2003, 108 (1), 72–110.
- 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.
- Lehmann, E. L.**, "Comparing Location Experiments," *The Annals of Statistics*, 1988, 16 (2), 521–533.
- Lewis, Tracy R. and David E. M. Sappington**, "Supplying information to facilitate price discrimination," *International Economic Review*, 1994, 35 (2), 309–327.
- Milgrom, Paul**, *Putting Auction Theory to Work*, Cambridge University Press, 2004.
- **and Chris Shannon**, "Monotone Comparative Statics," *Econometrica*, 1994, 62 (1), 157–180.
- Mizuno, Toshihide**, "A relation between positive dependence of signal and the variability of conditional expectation given signal," *Journal of applied probability*, 2006, 43 (4), 1181–1185.
- Myerson, Roger B.**, "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 1982, 10 (1), 67–81.
- Myerson, Roger B.**, "Multistage Games with Communication," *Econometrica*, 1986, 54 (2), 323–358.
- Nikandrova, Arina and Romans Pans**, "An Optimal Auction with Moral Hazard," *Birkbeck Working Papers in Economics and Finance*, 2015, (1504).
- Persico, Nicola**, "Information Acquisition in Auctions," *Econometrica*, 2003, 68 (1), 135–148.
- Rayo, Luis and Ilya Segal**, "Optimal Information Disclosure," *Journal of Political Economy*, 2010, 118 (5), 949–987.
- Roesler, Anne-Katrin**, "Information Disclosure in Market Design: Auctions, Contests and Matching Markets," 2014.
- Shi, Xianwen**, "Optimal Auctions with Information Acquisition," *Games and Economic Behavior*, 2012, 74 (2), 666–686.
- Zhang, Jun and Junjie Zhou**, "Information Disclosure in Contests: A Bayesian Persuasion Approach," *The Economic Journal*, 2015, (*forthcoming*).
- Zheng, Charles Zhoucheng**, "Optimal Auctions with Resale," *Econometrica*, 2002, 70 (6), 2197–2224.

Supplement to “Conjugate Information Disclosure in an Auction with Learning”

Arina Nikandrova and Romans Pancs

Birkbeck and ITAM

B Appendix

B.1 A Signal Structure that Rationalizes the Information-Acquisition Technology in Section 2

The model’s description, in Section 2, *could* have been specified as follows. The follower exerts 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 $p \equiv \int_0^1 s dF_H(s) = \int_0^1 s dF_L(s)$. Then, by construction, $\Pr \{v = 1\} = \mathbb{E} [\theta_2 | a]$ for all $a \in A$, meaning that the probability that the follower assigns to $v = 1$ before observing z equals his expectation of the conditional (on z) probability that $v = 1$, which is also his conditional expectation of v , denoted by θ_2 . This Bayesian consistency condition is necessary and sufficient for θ_2 to represent the follower’s conditional expectation of his underlying valuation (Kamenica and Gentzkow, 2011).

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

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

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

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

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

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

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

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

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

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

B.2 What the Information-Acquisition Technology in Section 2 Rules Out

The linear specification (1) and Condition 1 are restrictive. The linearity in (1) rules out information-acquisition technologies that let the follower choose among three or more signals, as Example B.1 clarifies.

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

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

Condition 1 remains restrictive even conditional on the linear specification (1), as Example B.2 illustrates.

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

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

B.3 An Analytical Equivalent of Condition 2

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

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

denote the inverse hazard rate of the leader's c.d.f. As is standard, $r(\theta_1)$ is interpreted as the profit that the seller forgoes—equivalently, the information rent that the leader reaps—when the seller commits to sell to a type- θ_1 leader.¹ In addition, recall that $R(\theta_1)$, defined in (A.2), denotes the planner's return to the follower's information acquisition in the first-best benchmark when the leader's type is θ_1 . This return is closely related to the follower's information-acquisition technology (in particular, $R'(\theta_1) = F_H(\theta_1) - F_L(\theta_1)$ and $R''(\theta_1) = f_H(\theta_1) - f_L(\theta_1)$) and so can be treated as a primitive.

Lemma B.1. *Suppose that Condition 1 holds and $f_L(\theta^*) \neq f_H(\theta^*)$.² Then, a prospect set is convex if and only if*

$$r''(\theta_1) + \left(r(\theta_1) \frac{R''(\theta_1)}{R'(\theta_1)} \right)' < 0 \quad \text{for all } \theta_1 \in (0, \theta^*) \cup (\theta^*, 1). \quad (\text{B.2})$$

Proof. A prospect set, $\Gamma \equiv \{(\pi(\theta_1), \alpha(\theta_1)) \mid \theta_1 \in \Theta_1\}$, is a parametrically given plane curve. Its signed curvature at θ_1 is given by³

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

where primes refer to derivatives with respect to θ_1 . Because Γ is simple⁴ and regular,⁵ it is strictly convex if and only if κ is either always positive or always negative. Because the denominator in (B.3) is always positive, requiring that κ does not change the sign is equivalent to requiring that the numerator in (B.3) not change the sign.

When $\theta_1 = \theta^*$, the numerator in (B.3) is negative, or $-r(\theta^*) (f_H(\theta^*) - f_L(\theta^*))^2 / c < 0$, because $f_H(\theta^*) \neq f_L(\theta^*)$ by the lemma's hypothesis and $F_L(\theta^*) = F_H(\theta^*)$ by part (ii) of Condition 1. Thus, the strict convexity of Γ is equivalent to the numerator in (B.3) being always negative.⁶

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

¹When the seller commits to sell to type θ_1 at some price, all types higher than θ_1 may be tempted to imitate type θ_1 and buy at the same price, thereby constraining the seller in how much he can charge these higher types.

²Condition $f_L(\theta^*) \neq f_H(\theta^*)$, which can be interpreted to hold "generically," simplifies the analytical characterization in the lemma but is not required for the convexity of the prospect set.

³The curvature of Γ at a point is the reciprocal of the radius of the circle osculating Γ at that point; see the Wikipedia entry on curvature: <https://en.wikipedia.org/wiki/Curvature>.

⁴A curve is simple if it does not intersect itself.

⁵A curve Γ is regular if its derivative $(\alpha', \pi') \neq (0, 0)$ for all $\theta_1 \in \Theta_1$, which holds in our model.

⁶In general, the sign of the curvature κ indicates the direction in which the unit tangent vector rotates as a function of the parameter along the curve. If the unit tangent rotates counterclockwise, then $\kappa > 0$. If it rotates clockwise, then $\kappa < 0$. In our model, as θ_1 increases, the unit tangent vector of Γ rotates clockwise, and, thus, κ must be negative everywhere.

Substituting the definitions of α and π into the above display, dividing by $(R'(\theta_1))^2$, which is positive when $\theta_1 \neq \theta^*$, and rearranging gives the sought inequality (B.2) of Lemma B.1.

This curvature condition captured by (B.2) is local and, alone, does not suffice to conclude that the prospect set is convex (in the sense of Definition 1); a spiral is a counterexample. Condition 1, however, which ensures that $\alpha(0) = \alpha(1) = 0$, thereby ruling out a spiral and ensuring that the curvature condition in (B.2) is equivalent to the convexity of Γ . \square

B.4 Examples that Illustrate Cases in Definition 2

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

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

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

Figure B.1 illustrates how a downward-sloping segment for Γ near $\theta_1 = 0$ is necessary for case (ii) in Definition 2 not to collapse into case (i). The figure also illustrates the role played by s' .

Figure 3b illustrates an example of case (ii); c.d.f.s F_H and F_L are Beta distributions chosen to satisfy Condition 1. Then, $F_H(0) = F_L(0) = 0$. Furthermore, one can (merely to simplify the argument) choose F_H and F_L so that $f_H(0) > f_L(0)$. As a result, $\pi'(0) = r(0)(f_L(0) - f_H(0)) < 0$; π is decreasing near 0. Because α is increasing near 0, the prospect set is downward-sloping near 0.

B.5 Justifying Equation (A.16) in the Proof of Lemma A.4

To justify (A.16), Lemma B.2 demonstrates that one can approximate any $\nu \in \Delta(\Delta\Theta_1)$ by a probability measure that puts some mass only on discrete measures in a countable set. The proof proceeds in two steps. First, it shows that by choosing n sufficiently large, any probability measure in $\Delta\Theta_1$ can be approximated by a probability measure that puts some mass on a countable set $\{\frac{1}{2^n}, \frac{2}{2^n}, \dots, 1\}$ in Θ_1 . Then, a similar argument is repeated to show that if one chooses n sufficiently large, any measure in $\Delta(\Delta\Theta_1)$ can be approximated by a measure that puts positive mass only on discrete measures with support $\{\frac{1}{2^n}, \frac{2}{2^n}, \dots, 1\}$. This second half is slightly trickier because it requires finding a countable set of non-overlapping neighborhoods in $\Delta\Theta_1$ that almost cover space

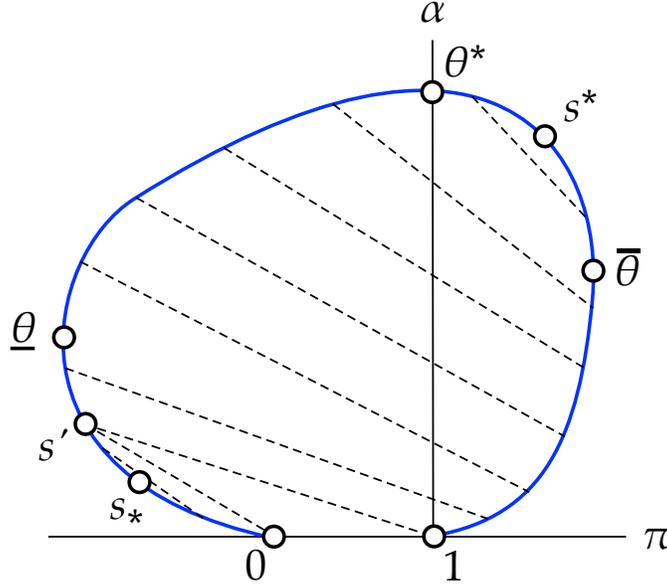


Figure B.1: The convex curve is the prospect set Γ . The circles mark the prospects induced by the leader's types in $0, s_*, s', \underline{\theta}, \theta^*, s^*, \bar{\theta}$, and 1. The dashed links comprise a subset of the links that pool prospects into messages. Type s' demarcates the leader's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

Δ_{Θ_1} .

Lemma B.2. Fix an arbitrary measure $\nu \in \Delta(\Delta_{\Theta_1})$. For every $\epsilon > 0$, there exists N such that for $n \geq N$,

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

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

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

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

Proof. The proof proceeds in two steps.

Step 1: It is possible to approximate any measure μ in Δ_{Θ_1} with a measure in \mathcal{D}_n by choosing n sufficiently high.

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

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

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

by

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

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

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

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

To show that $\mu_n \Rightarrow \mu$, take a uniformly continuous bounded function g on $\Theta_1 = [0, 1]$.⁷ Let $\|g\|_\infty \equiv \sup_{x \in \Theta_1} g(x)$ denote the supremum norm. Then,

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

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

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

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

⁷Statements that $\mu_n \Rightarrow \mu$ and that $\lim \int g d\mu_n = \int g d\mu$ for all uniformly continuous, bounded functions are equivalent because (i) every Lipschitz function between two metric spaces is uniformly continuous; and (ii) 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).

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

Let

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

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

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

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

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 $m \geq 1$,

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

Take N and J such that

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

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

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

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

$$\nu\left(\cup_{k=1}^J B_k^m\right) = \nu\left(\cup_{k=1}^J B\left(\mu_N^k, \frac{1}{m}\right)\right) \geq 1 - 1/m. \quad (\text{B.4})$$

It is possible to approximate

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

by

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

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

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

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

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

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

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

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

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

B.6 Preliminary Results for the Proof of Theorem 4

The proof of Theorem 4 demonstrates that the boundedness of $a^{*'}$ on (θ^*, s^*) follows from Π being connected and from Π and Γ being disjoint on $(s^*, 1]$. This section proves the two required intermediate results in Lemma B.3 and Lemma B.6, respectively.

Lemma B.3. Π is connected.

Proof. By contradiction, suppose that Π is not connected. Any non-connectedness in Π must be caused by some prospects being revealed, as in Figure B.2. The revealed prospects, which lie in Γ , belong to Π . As a result, Π fails to be nondecreasing, thereby contradicting optimality. \square

It remains to show that the optimal-prospect path Π is disjoint from the prospect set Γ on $(s^*, 1]$. Lemma B.6 accomplishes this task. The lemma's proof relies on a number of preliminary results.

It is convenient to cast the analysis in terms of a decreasing function $\lambda \equiv \tau^{-1}$, the inverse of the matching function τ , whenever this inverse exists. Function λ is defined on the interval

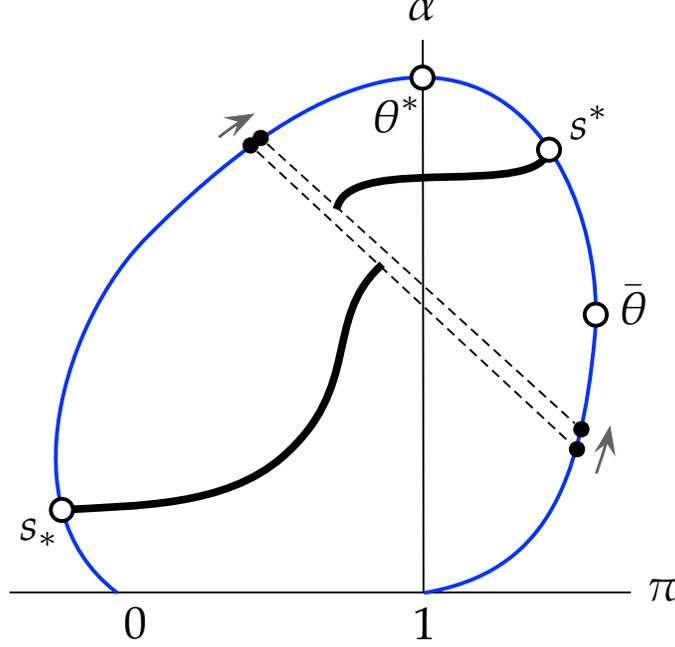


Figure B.2: The broken solid thick curve is a counterfactually non-connected Π . Prospects that lie on Γ 's chords that are demarcated by the two pooling segments are revealed and belong to Π ; Π fails to be nondecreasing, thereby contradicting optimality.

$[s^*, 1]$. Operating under the assumption that λ exists is justified because the goal is to rule out the situation in which λ is flat (and so, by implication, exists). Define $\gamma \equiv -\lambda'$.

For expositional purposes only, assume that the leader's type is distributed uniformly: $G(\theta_1) = \theta_1$.⁸ After a change of variables, the seller's objective function (21) restricted to $[s^*, 1]$ (which is the interval of particular interest in our analysis) can be written in terms of γ and λ as

$$\int_{s^*}^1 \Lambda(s, \lambda, \gamma) ds, \quad (\text{B.5})$$

where

$$\Lambda(s, \lambda, \gamma) \equiv \frac{(\pi(s) + \gamma(s) \pi(\lambda(s))) (\alpha(s) + \gamma(s) \alpha(\lambda(s)))}{1 + \gamma(s)}.$$

Towards optimality, take some λ and perturb it towards $\lambda + \epsilon\eta$ for some $\epsilon \in \mathbb{R}$ and for some $\eta : [s^*, 1] \rightarrow \mathbb{R}$ with $\eta(s^*) = \eta(1) = 0$. The value of the perturbed objective function is denoted by

$$\Phi \equiv \int_{s^*}^1 \Lambda(s, \lambda + \epsilon\eta, \gamma - \epsilon\eta') ds,$$

where the dependence of λ , γ , and η on s is implicit and will remain so as long as no ambiguity

⁸The Euler equation, in (B.7), and the subsequent arguments all hold for a general G . The unwieldy derivation for this general case is available upon request.

arises. The marginal benefit from perturbing λ in the direction of η is denoted by

$$J \equiv \frac{d\Phi}{d\epsilon} \Big|_{\epsilon=0} = \int_{s^*}^1 \eta \left(\frac{\partial \Lambda}{\partial \lambda} - \frac{d}{ds} \frac{\partial \Lambda}{\partial \gamma} \right) ds.$$

If λ is optimal, any (feasible) perturbation η requires $J \leq 0$. If, in addition, an optimal λ is interior, the parenthetical term in the display above must be identically zero. The parenthetical term equated to zero becomes the Euler equation.

Computing $\partial \Lambda / \partial \lambda$, $\partial \Lambda / \partial \gamma$, and $(d/ds)(\partial \Lambda / \partial \gamma)$ and substituting the results into J in the display above gives

$$J = \int_{s^*}^1 \eta \left[\frac{(\pi(\lambda) - \pi)(\alpha' - \gamma^2 \alpha'(\lambda)) - (\pi' - \gamma^2 \pi'(\lambda))(\alpha - \alpha(\lambda))}{(1 + \gamma)^2} - \frac{2\gamma'(\pi(\lambda) - \pi)(\alpha - \alpha(\lambda))}{(1 + \gamma)^3} \right] ds.$$

To interpret the expression for J graphically, in terms of the slope of Π , with some abuse of notation, write $\Pi \equiv \{(\pi^*(s), a^*(s)) \mid s \in [s^*, 1]\}$, where⁹

$$\pi^* = \frac{\pi + \gamma\pi(\lambda)}{1 + \gamma} \quad \text{and} \quad a^* = \frac{\alpha + \gamma\alpha(\lambda)}{1 + \gamma}.$$

Combining

$$\frac{d\pi^*}{ds} = \frac{(\pi' - \gamma^2 \pi'(\lambda))(1 + \gamma) - \gamma'(\pi - \pi(\lambda))}{(1 + \gamma)^2}$$

and

$$\frac{da^*}{ds} = \frac{(\alpha' - \gamma^2 \alpha'(\lambda))(1 + \gamma) - \gamma'(\alpha - \alpha(\lambda))}{(1 + \gamma)^2},$$

one obtains the slope of Π :

$$\frac{da^*}{d\pi^*} = \frac{da^*/ds}{d\pi^*/ds} = \frac{(\alpha' - \gamma^2 \alpha'(\lambda))(1 + \gamma) - \gamma'(\alpha - \alpha(\lambda))}{(\pi' - \gamma^2 \pi'(\lambda))(1 + \gamma) - \gamma'(\pi - \pi(\lambda))}.$$

Rearranging the expression for J and substituting $d\pi^*/ds$ and $da^*/d\pi^*$ gives

$$J = \int_{s^*}^1 \eta \sigma \left(\frac{da^*}{d\pi^*} - \frac{\alpha(\lambda) - \alpha}{\pi - \pi(\lambda)} \right) ds, \quad \text{where} \quad \sigma \equiv \frac{\pi - \pi(\lambda)}{1 + \gamma} \left(-\frac{d\pi^*}{ds} \right). \quad (\text{B.6})$$

Now, from (B.6), we can extract the Euler equation. On $(s^*, 1)$, $\pi - \pi(\lambda) > 0$ and $\gamma \geq 0$ (and so $1 + \gamma > 0$). Because Π cannot be vertical, $d\pi^*/ds < 0$. As a result, $\sigma > 0$. Therefore, whenever Π is interior, it must satisfy the Euler equation:

$$\frac{da^*}{d\pi^*} = \frac{\alpha(\lambda) - \alpha}{\pi - \pi(\lambda)}. \quad (\text{B.7})$$

⁹The abuse of notation here consists in reinterpreting $\pi^*(s)$ and $a^*(s)$ to be indexed by s in $[s^*, 1]$ (not $[s_*, s^*]$), which is consistent with the dummy variable in the integrand of the rewritten objective function (B.5) being defined on $[s^*, 1]$ (not $[s_*, s^*]$).

To summarize,

Lemma B.4. *When Π is interior, its slope, $da^*/d\pi^*$, obeys the Euler equation (B.7).*

Proof. The argument for the uniform G precedes the lemma's statement. The argument for a general, non-uniform, G is available upon request. \square

Lemma B.5 is an auxiliary result that roughly says that, if Π were to touch Γ at a point in $[\bar{\theta}, 1]$, it would do so at an angle, instead of pasting smoothly at that point.

Lemma B.5. *Suppose that Π coincides with Γ at a single point $s' \in [\bar{\theta}, 1]$ and is disjoint from Γ on some subset $S \subset [\bar{\theta}, 1]$ adjacent to s' (i.e., $s' \in cl(S)$). Then, Π cannot be tangent to Γ at s' .*

Proof. By contradiction, suppose that Π coincides with, and is tangent to, Γ at a point $s' \in [\bar{\theta}, 1]$. Define a subset $S \subset (\bar{\theta}, 1)$ so that (i) for some $\varepsilon > 0$, S is either $(s', s' + \varepsilon)$ or $(s' - \varepsilon, s')$; (ii) on S , Π and Γ are disjoint; and (iii) ε is "sufficiently small" in the sense that will be made precise. That is, S is chosen so that, on S , Π is close to, but disjoint from, Γ . The disjointness of Π and Γ implies that, on S , Π satisfies the Euler equation (B.7). Differentiating the Euler equation on S gives

$$\frac{d}{ds} \left(\frac{da^*}{d\pi^*} \right) = \frac{d}{ds} \frac{\alpha(\lambda) - \alpha}{\pi - \pi(\lambda)} = - \frac{\alpha' + \gamma\alpha'(\lambda) + (da^*/d\pi^*)(\pi' + \gamma\pi'(\lambda))}{\pi - \pi(\lambda)}.$$

By requirement (iii) in the construction of S , Π is "close" to Γ on S , and so $a^* \approx \alpha$, which requires $\gamma \approx 0$ (by Bayes' rule, $a^* = (\alpha + \gamma\alpha(\lambda)) / (1 + \gamma)$).¹⁰ With this justification, we neglect the terms multiplied by γ . Then, $\pi - \pi(\lambda) > 0$ combined with $da^*/d\pi^* > 0$ (Π is increasing, by optimality), $\alpha' < 0$, and $\pi' < 0$ (by $s \in (\bar{\theta}, 1)$), imply

$$\frac{d}{ds} \left(\frac{da^*}{d\pi^*} \right) > 0,$$

which, in turn, implies that Π is concave.¹¹

Thus, we have reached a contradiction because, at s' , Π and Γ coincide and are tangent to each other, and yet, on S , Π is concave, whereas Γ is convex (by Condition 2). This situation is a geometric impossibility; it requires Π to exit the convex hull of Γ . Hence, Π and Γ cannot be tangent at s' . \square

We are now ready to formulate and prove

Lemma B.6. *Suppose that Condition 3 holds. Then, the intersection of Π and Γ on $(s^*, 1]$ is empty.*

¹⁰The role of tangency in the contradiction hypothesis is to ensure that Π is "close" to Γ on S , and so $\gamma \approx 0$ indeed holds.

¹¹As s rises, the induced prospect is moving southwest along Γ —or from right to left if one projects this movement on the horizontal axis. Hence, the sign in the criterion for concavity is positive, flipped from the customary. Formally, Π , a parametric curve, is concave if its second derivative is negative; that is, if $d^2a^*/(d\pi^*)^2 = d(da^*/d\pi^*)/d\pi^* = (d(da^*/d\pi^*)/ds) / (d\pi^*/ds) < 0$. Because, on $[\bar{\theta}, 1] \subseteq [s^*, 1]$, $d\pi^*/ds < 0$, Π is concave if $d(da^*/d\pi^*)/ds > 0$.

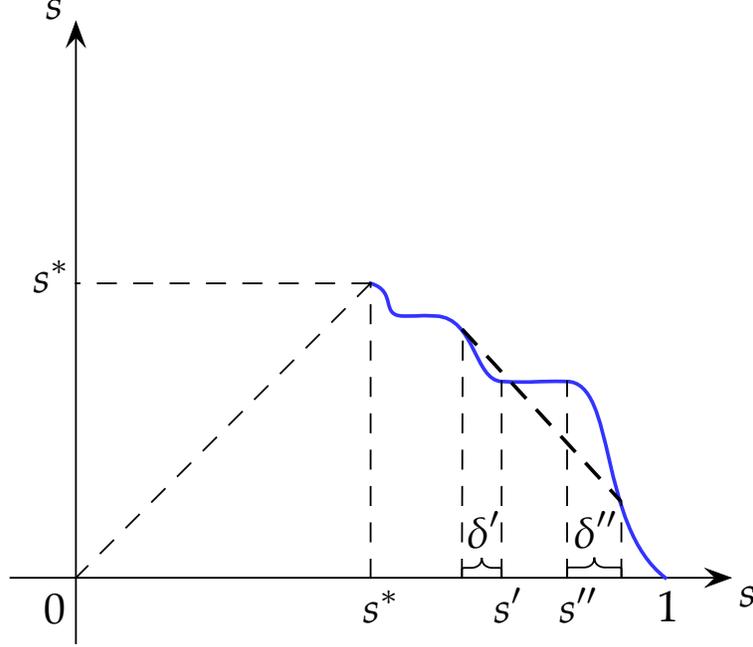


Figure B.3: The thick solid curve is an inverse matching function λ with a (counterfactually) flat segment on $[s', s'']$. The dashed segment is the perturbed (by η) inverse matching function $\hat{\lambda} \equiv \lambda + \eta$.

Proof. By contradiction, first suppose that Π and Γ have a nonempty intersection on $(s^*, \bar{\theta})$. Then, because Π and Γ coincide at s^* , and because Γ is decreasing on $(s^*, \bar{\theta})$, Π must have a decreasing segment, thereby contradicting optimality, which requires that Π be nondecreasing.

The remainder of the proof is concerned with showing that Π cannot intersect Γ on $[\bar{\theta}, 1]$. By way of contradiction, let $[s', s'']$ with $s'' \geq s'$ be the largest (lengthwise, $s'' - s'$) interval on which Π and Γ coincide in $[\bar{\theta}, 1]$. The argument goes through a list of cases with all the possible values for s' and s'' relative to $\bar{\theta}$ and 1.

Case 1: $\bar{\theta} < s' < s'' < 1$.

Given λ , define an additive perturbation η parametrized by positive scalars δ' and δ'' . This perturbation is illustrated in Figure B.3. The scalars δ' and δ'' , and the perturbation η are chosen so that (i) on $[s^*, s' - \delta'] \cup [s'' + \delta'', 1]$, $\eta = 0$; (ii) on $s \in (s' - \delta', s'' + \delta'') \subset (s^*, 1)$, η induces a $\hat{\lambda}$ that linearly interpolates between $\lambda(s' - \delta')$ and $\lambda(s'' + \delta'')$:

$$\hat{\lambda}(s) \equiv \eta(s) + \lambda(s) = \lambda(s'' + \delta'') + \frac{s'' + \delta'' - s}{s'' + \delta'' - (s' - \delta')} (\lambda(s' - \delta') - \lambda(s'' + \delta'')); \quad (\text{B.8})$$

and (iii) δ' and δ'' are such that $s^* < s' - \delta'$, $s'' + \delta'' < 1$, Π and Γ are disjoint on $(s^*, s' - \delta') \cup (s'' + \delta'', 1)$,

$$\int_{s' - \delta'}^{s'' + \delta''} \eta \sigma ds = 0, \quad (\text{B.9})$$

and η crosses zero only once (at a single point), from above.¹² Because σ , defined in (B.6), is positive, the single-crossing property of η implies that $\eta\sigma$, too, crosses zero only once (at a single point), from above.

The construction of η implies the inequality

$$\begin{aligned} J &= \int_{s^*}^1 \eta\sigma \left(\frac{da^*}{d\pi^*} - \frac{\alpha(\lambda) - \alpha}{\pi - \pi(\lambda)} \right) ds \\ &= \int_{s'}^{s''} \eta\sigma \left(\frac{da^*}{d\pi^*} - \frac{\alpha(\lambda) - \alpha}{\pi - \pi(\lambda)} \right) ds > 0. \end{aligned} \quad (\text{B.10})$$

Indeed, in the display above, the equality in the second line follows because, by construction, $\eta \equiv 0$ outside the interval $(s' - \delta', s'' + \delta'')$, and because the Euler equation holds on the set $(s' - \delta', s') \cup (s'', s'' + \delta'')$, so that the parenthetical term is zero on this set. To conclude the inequality in (B.10), note that, on (s', s'') , the parenthetical term is decreasing in s because $da^*/d\pi^*$ is decreasing in s (because, by hypothesis, Π coincides with Γ on $[s', s'']$ and, by Condition 2, Γ is convex), and because the term $(\alpha(\lambda) - \alpha) / (\pi - \pi(\lambda))$ is increasing in s (because, by hypothesis, λ does not vary with s ; α and π both decrease in s ; and, since pooling links are non-increasing, $\pi - \pi(\lambda)$ and $\alpha(\lambda) - \alpha$ are both positive). Then, because, by construction of η , $\eta\sigma$ crosses zero once and from above, (B.9) implies the positive sign in (B.10); indeed, the left-hand side of (B.10) puts larger weights on positive $\eta\sigma$'s, and smaller weights on negative $\eta\sigma$'s, than (B.9) does. Thus, when $s'' > s'$, the constructed η induces a profitable perturbation away from λ , and so the coincidence of Π and Γ on $[s', s'']$ is suboptimal.

Case 2: $\bar{\theta} < s' = s'' < 1$.

The Euler equation (B.7) describes the slope of Π arbitrarily close to s' . This slope is the same from whichever direction s' is approached. Thus, Π must be tangent to Γ at s' . Lemma B.5, however, rules out such a tangency; a contradiction is reached. Thus, the coincidence of Π and Γ at a single point s' is suboptimal.

Case 3: $\bar{\theta} \leq s' < s'' = 1$.

Figure B.4 illustrates that if Π coincides with Γ on $[s', 1]$, then Π coincides with Γ also on a nonempty interval $[0, \theta_1]$, for some $\theta_1 > 0$. Such a Π fails to be nondecreasing at θ_1 , thereby contradicting optimality. Therefore, the coincidence of Π and Γ on $[s', 1]$ is suboptimal.

Case 4: $\bar{\theta} < s' = s'' = 1$.

The Euler equation applies in a sufficiently small neighborhood of s' . Because Π lies in the convex hull of Γ , near s' , Π must be at least as steep as Γ ; that is, the slope of Π must be at least $\lim_{s \rightarrow 1} \alpha'(s) / \pi'(s)$. If this slope is nonzero, the pooling pattern implied by the Euler equation requires Π to coincide with Γ on $[0, \theta_1]$ for some $\theta_1 > 0$, thereby leading to a Π that fails to be nondecreasing and, thus, contradicting optimality. Figure B.5 illustrates the contradiction.

¹²Because $\sigma > 0$, by the Second Mean Value Theorem for Integrals, $\int_{s' - \delta'}^{s'' + \delta''} \eta\sigma ds = \eta(x) \int_{s' - \delta'}^{s'' + \delta''} \sigma ds$ for some $x \in (s' - \delta', s'' + \delta'')$. Because η is positive at first and then negative, it must cross zero at some point. Consequently, it is possible to find δ' and δ'' to satisfy $\eta(x) = 0$ and, hence, also (B.9).

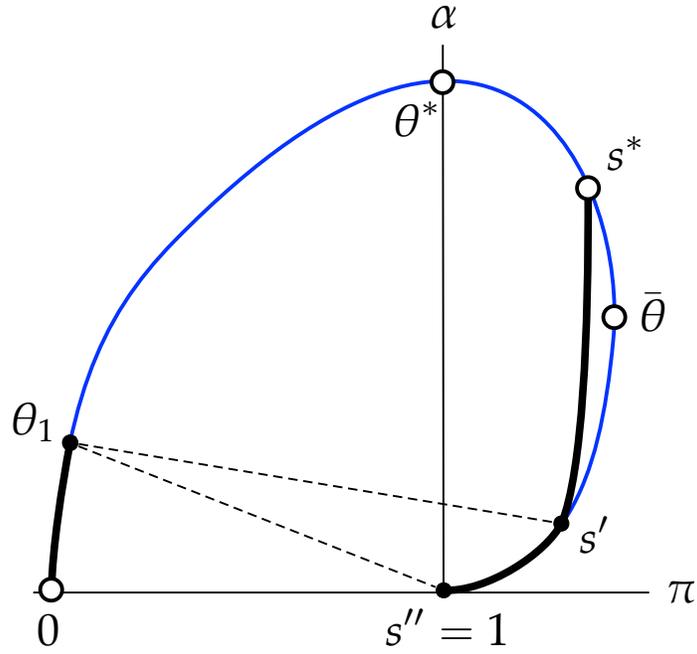


Figure B.4: If Π lies on Γ for $s \in [s', 1]$, then the pooling links must be nonincreasing as indicated, thereby implying that Π lies on Γ for $s \in [0, \theta_1]$, for some $\theta_1 > 0$. As a result, the implied Π fails to be nondecreasing, which contradicts optimality.

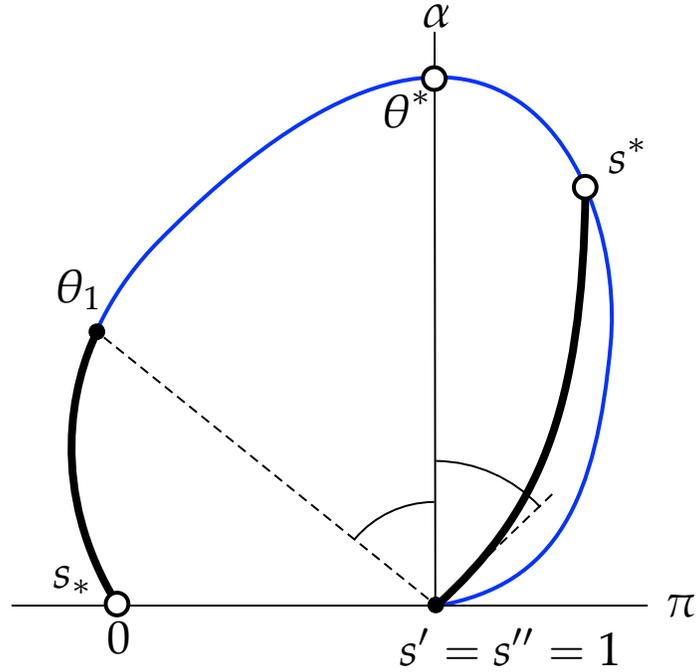


Figure B.5: The thick curves constitute Π . The two angles emanating from s' and marked by arcs are equal by the Euler equation. The dashed pooling link that connects θ_1 and s' suggests two possibilities: (i) $s_* = 0$ and Π coinciding with Γ on $[0, \theta_1]$ (shown); and (ii) $s_* > 0$ (not shown). In either case, Π fails to be nondecreasing, thereby contradicting optimality.

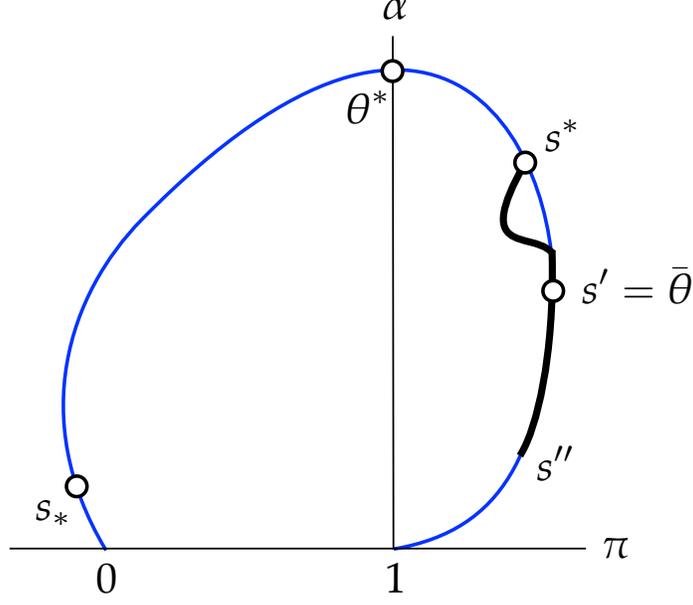


Figure B.6: The thick curve denotes the segment of Π invoked in the argument. If $s^* < s' = \bar{\theta}$, then Π must bend backwards to reach s^* .

The slope is nonzero if

$$\lim_{s \rightarrow 1} \frac{\alpha'(s)}{\pi'(s)} > 0,$$

which holds by Condition 3, in the lemma's hypothesis. Thus, the coincidence of Π and Γ at 1 is suboptimal.

Case 5: $\bar{\theta} = s' < s'' \leq 1$.

Note that $s' = \bar{\theta}$ implies that $s^* = \bar{\theta}$. Indeed, if $s' = \bar{\theta}$ and, by contradiction, $s^* < \bar{\theta}$, then Π must necessarily be decreasing somewhere on $[s^*, \bar{\theta}]$ (see Figure B.6), thereby contradicting optimality.

However, it will be shown that $s^* = \bar{\theta}$ is not possible either. By contradiction, suppose that $s^* = \bar{\theta}$, as in Figure B.7. If $s^* = \bar{\theta} = s'$ and Π coincides with Γ on $[s', s'']$, then Π must also coincide with Γ on an interval $[\theta_1, s^*]$ for some $\theta_1 < s^*$. This geometric arrangement contradicts Π being nondecreasing because Γ is decreasing on $[\theta^*, \bar{\theta}]$. Thus, the coincidence of Π and Γ on $[s', s'']$ is suboptimal.

Case 6: $\bar{\theta} = s' = s'' < 1$.

As in the preceding case, $s' = \bar{\theta}$ requires $s^* = \bar{\theta}$.

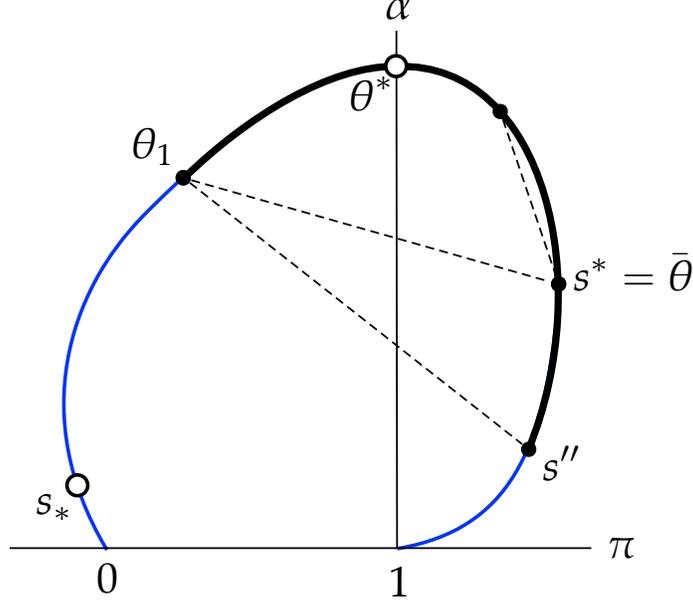


Figure B.7: The thick curve denotes the segment of Π invoked in the argument. If $s' = s^* = \bar{\theta} < s''$, then all prospects in $[s^*, s'']$ must be connected to the same prospect θ_1 by nonincreasing pooling segments. Then, by $s^* = \bar{\theta}$, Π must coincide with Γ on $(\theta_1, \bar{\theta})$. In particular, a segment of Π , on $(\theta^*, \bar{\theta})$, must be downward-sloping, which contradicts optimality.

By the Euler equation, (B.7),

$$\begin{aligned}
\lim_{s \downarrow s^*} \frac{da^*}{d\pi^*} &= \lim_{s \downarrow s^*} \frac{\alpha(\lambda(s)) - \alpha(s)}{\pi(s) - \pi(\lambda(s))} \\
&= - \lim_{s \downarrow s^*} \frac{\gamma(s) \alpha'(\lambda(s)) + \alpha'(s)}{\pi'(s) + \gamma(s) \pi'(\lambda(s))} \\
&= - \lim_{s \downarrow s^*} \frac{\alpha'(s) (1 + \gamma(s))}{\pi'(s) (1 + \gamma(s))} \\
&= - \lim_{s \downarrow s^*} \frac{\alpha'(s)}{\pi'(s)}, \tag{B.11}
\end{aligned}$$

where the second equality is by L'Hôpital's rule, the third one uses the smoothness of the prospect set (i.e., $\lim_{s \downarrow s^*} (\alpha'(s) - \alpha'(\lambda(s))) = 0$) and the continuity of λ , and the fourth one uses $1 + \gamma(s^*) > 0$. When $s^* = \bar{\theta}$, equation (B.11) implies that

$$\lim_{s \downarrow s^* = \bar{\theta}} \frac{da^*}{d\pi^*} = - \lim_{s \downarrow s^* = \bar{\theta}} \frac{\alpha'}{\pi'} = \infty;$$

that is, at $s^* = \bar{\theta}$, Π is tangent to Γ . Lemma B.5 rules out such a tangency; a contradiction is reached. Figure B.8 illustrates the contradiction. Thus, the coincidence of Π and Γ at $\bar{\theta}$ is suboptimal. \square

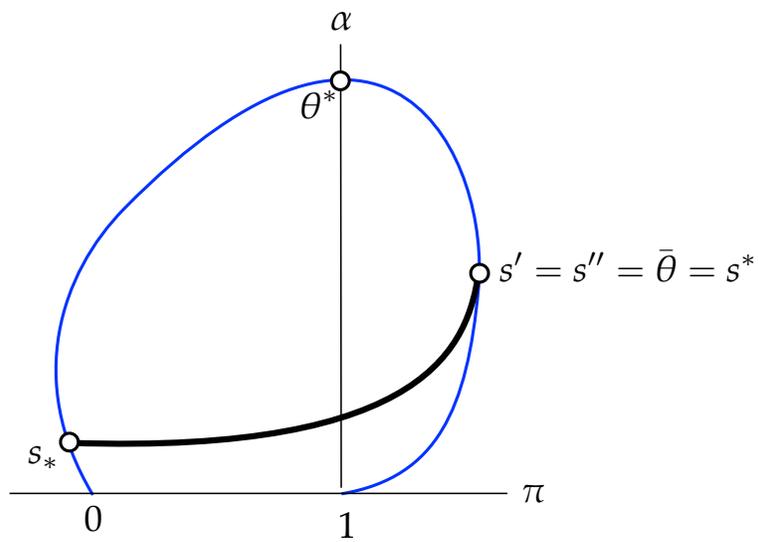


Figure B.8: The thick curve is Π . If $s' = s'' = \bar{\theta}$, it must be that $s^* = \bar{\theta}$, and that Π has an infinite slope at s^* , which contradicts optimality.