PERSUASION MEETS DELEGATION

Anton Kolotilin and Andriy Zapechelnyuk

Abstract. There are two common ways for a principal to influence the decision making of an agent. One is to manipulate the agent's information (persuasion problem). Another is to limit the agent's decisions (delegation problem). We show that, under general assumptions, these two problems are equivalent; so solving one problem solves the other. We illustrate how the methods developed in the persuasion literature can be applied to address unsolved delegation problems by considering monopoly regulation with a participation constraint.

JEL Classification: D82, D83, L43

Keywords: persuasion, delegation, regulation

Date: 7th April 2018.

Kolotilin: School of Economics, UNSW Business School, Sydney, NSW 2052, Australia. *E-mail:* akolotilin@gmail.com.
Zapechelnyuk: School of Economics and Finance, University of St Andrews, Castlecliffe, the Scores, St Andrews KY16 9AR, UK. *E-mail:* az48@st-andrews.ac.uk.

We are grateful to Tymofiy Mylovanov, with whom we are working on related projects. We thank Ricardo Alonso, Benjamin Brooks, Deniz Dizdar, Piotr Dworczak, Alexander Frankel, Gabriele Gratton, Emir Kamenica, Navin Kartik, Ming Li, Hongyi Li, Carlos Oyarzun, Alessandro Pavan, Philip Reny, Joel Sobel, and Thomas Tröger for helpful comments and suggestions. Kolotilin acknowledges support from the Australian Research Council Discovery Early Career Research Award DE160100964. Zapechelnyuk acknowledges support from the Economic and Social Research Council Grant ES/N01829X/1.

1. Introduction

Consider a principal who wants to influence the decision making of an agent. Different instruments of influence have been studied. In the *delegation* literature initiated by Holmström (1984), the principal restricts the set of decisions from which the agent chooses. In the *persuasion* literature set in motion by Kamenica and Gentzkow (2011), the principal restricts the information that the agent has.

The delegation problem has been used to design organizational decision processes (Dessein 2002), monopoly regulation policies (Alonso and Matouschek 2008), and international trade agreements (Amador and Bagwell 2013). The persuasion problem has been used to design school grading policies (Ostrovsky and Schwarz 2010), internet advertising strategies (Rayo and Segal 2010), and media censorship rules (Kolotilin, Mylovanov, Zapechelnyuk, and Li 2017).

This paper shows that, under general assumptions, the delegation and persuasion problems are equivalent, thereby bridging the two strands of literature. The implication is that the existing insights and results in one problem can be used to understand and solve the other problem.

Both delegation and persuasion problems have a principal and an agent whose payoffs depend on the agent's decision and a state of the world. The sets of decisions and states are intervals of the real line. The agent's payoff function satisfies standard concavity and sorting conditions. In a delegation problem, the agent privately knows the state and the principal commits to a set of decisions from which the agent chooses. In a persuasion problem, the principal designs the agent's information structure and the agent freely chooses a decision. The principal's tradeoff is that giving more discretion to the agent in the delegation problem and disclosing more information to the agent in the persuasion problem allows for a better use of information about the state, but limits control over the biased agent's decision.

We consider *balanced delegation* and *monotone persuasion* problems. In the balanced delegation problem, the principal may not be able to exclude certain indispensable decisions of the agent. This problem nests the standard delegation problem and includes, in particular, a novel delegation problem with an agent's participation constraint.

In the monotone persuasion problem, the principal chooses a monotone partitional information structure that either reveals the state or pools it with adjacent states. This problem incorporates constraints faced by information designers in practice. For example, a non-monotone grading policy that gives better grades to worse performing students will be perceived as unfair and will be manipulated by strategic students. Moreover, under regularity conditions, optimal information structures are monotone partitions.

The main result of the paper is that each primitive of the balanced delegation problem is strategically equivalent to a primitive of the monotone persuasion problem, and vice versa, under a direct one-to-one mapping between the principal's strategies in the two

problems. Moreover, for each primitive of one problem, we can find an equivalent primitive of the other problem by solving a pair of ordinary differential equations. It is worth noting that this equivalence result is fundamentally different from the revelation principle. Specifically, given the principal's and agent's payoff functions, the sets of implementable (and, therefore, optimal) decision outcomes generally differ in the delegation and persuasion problems.

To prove the equivalence result, we show that the balanced delegation and monotone persuasion problems are equivalent to a *discriminatory disclosure* problem. In this problem, the principal's and agent's payoffs depend on the agent's binary action, the state of the world, and the agent's private type. The sets of states and types are intervals of the real line. The agent's payoff function is monotone in the state and type. The principal designs a menu of cutoff tests, where a cutoff test reveals whether the state is below or above a cutoff. The agent selects a test from the menu and chooses between *action* and *inaction* depending on his private type and the information revealed by the test.

To see why the discriminatory disclosure problem is equivalent to the balanced delegation problem, observe that the agent's essential decision is the selection of a cutoff test from the menu. Because the agent's payoff function is monotone, the agent optimally chooses action/inaction if the selected test reveals that the state is above/below the cutoff. Thus, this problem can be interpreted as a delegation problem in which a delegation set is identified with a menu of cutoffs, and the agent's decision with his selection of a cutoff from the menu.

To see why the discriminatory disclosure problem is equivalent to the monotone persuasion problem, observe that each menu of cutoff tests defines a monotone partition of the state space. Because the agent's payoff function is monotone, the privately informed agent optimally chooses action/inaction on the same interval of states, whether he observes the result of his preferred cutoff test or the partition element that contains the state. Thus, this problem can be interpreted as a persuasion problem in which a monotone partition is identified with a menu of cutoffs, and the agent's decision with a threshold type above/below which the agent chooses action/inaction.

We use our equivalence result to solve a monopoly regulation problem in which a welfare-maximizing regulator (principal) restricts the set of prices available to a monopolist (agent) who privately knows his cost. This problem was first studied by Baron and Myerson (1982) as a mechanism design problem with transfers. Alonso and Matouschek (2008) pointed out that transfers between the regulator and monopolist are often forbidden, and thus, the monopoly regulation problem can be formulated as a delegation problem. However, Alonso and Matouschek (2008) omitted the monopolist's participation constraint, so under their optimal regulation policy, the monopolist sometimes operates at a loss.

In contrast, we take the monopolist's participation constraint into account. This monopoly regulation problem is thus a *balanced* delegation problem, which has been

previously neglected by the literature. We provide an elegant method of solving this problem, as well as its counterpart without the participation constraint, by recasting it as a monotone persuasion problem and using a single result from the persuasion literature. Remarkably, when the demand function is linear and the cost distribution is unimodal, the optimal regulation policy takes a simple form that is often used in practice. The regulator imposes a price cap and allows the monopolist to choose any price not exceeding the cap. We show that the optimal price cap is higher when the participation constraint is present, because a higher price cap increases the probability that the monopolist operates.

Our equivalence result implies that *separable* balanced delegation, where the marginal payoffs are additively separable, is equivalent to *linear* monotone persuasion, where the payoffs depend only on the posterior mean state. Separable delegation has been studied by Holmström (1984), Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), Goltsman, Hörner, Pavlov, and Squintani (2009), Kováč and Mylovanov (2009), and Amador and Bagwell (2013).[1] Linear persuasion has been studied by Ostrovsky and Schwarz (2010), Gentzkow and Kamenica (2016), Kolotilin et al. (2017), and Dworczak and Martini (2018), and has also been considered as a special case in Kamenica and Gentzkow (2011) and Kolotilin (2018).[2]

Alonso and Matouschek (2008) and Amador and Bagwell (2013) also provide partial results for non-separable delegation. Rayo and Segal (2010), Guo and Shmaya (2018), and Kolotilin (2018) study non-linear persuasion. Our equivalence result applies to these settings as well.

2. Two Problems

2.1. **A Problem.** There are a principal (she) and an agent (he). The agent makes a decision x. The set of decisions is an interval of $\mathbb{R}$, without loss of generality normalized to $[0,1]$. The principal's and agent's payoffs

$$V(\theta, x) \quad \text{and} \quad U(\theta, x)$$

depend on the chosen decision $x \in [0,1]$ and a state $\theta \in [\underline{\theta}, \overline{\theta}]$. The state is a random variable with a commonly known distribution F. We assume that

(A_0) F has a strictly positive density;

[1]Despite the same assumptions on the payoff functions, separable balanced delegation is conceptually different from veto-based delegation of Krishna and Morgan (2001), Dessein (2002), and Mylovanov (2008) and from limited-commitment delegation of Kolotilin, Li, and Li (2013).

[2]Dworczak and Martini (2018) and Kolotilin and Li (2018) study linear monotone persuasion. Dworczak and Martini (2018) provide conditions under which monotone partitional information structures are optimal. Kolotilin and Li (2018) characterize the optimal monotone partitional information structures when they differ from the optimal non-monotone ones. Ivanov (2017) and Mensch (2018) consider non-linear monotone persuasion.

(A_1) $U(\theta, x)$ and $V(\theta, x)$ are continuous in θ and continuously differentiable in x;

(A_2) $\frac{\partial}{\partial x}U(\theta, x)$ is strictly increasing in θ and strictly decreasing in x.[3]

Under these assumptions, without loss of generality, we can redefine the state such that its distribution is uniform on [0,1].[4] Thus, in what follows we shall assume

(A_3) F is uniform on $[0, 1]$.

A pair (U, V) is called a *primitive* of the problem. Denote by $\mathcal{P}$ the set of all primitives that satisfy the above assumptions (A_1)–(A_2).

We now describe two variants of this problem that differ only in how the principal can influence the agent's decision.

2.2. **Balanced Delegation Problem.** In a delegation problem, the principal commits to a set of decisions from which the informed agent chooses.

Consider a primitive $(U_D, V_D) \in \mathcal{P}$, where we use subscript D to refer to the delegation problem. The principal chooses a closed subset Π of $[0, 1]$, referred to as a *delegation set*. The agent privately observes the state θ and chooses a decision from Π,

$$x^*_\Pi(\theta) \in \arg\max_{x\in\Pi} U_D(\theta, x).$$

We focus on *balanced* delegation sets that contain the extreme decisions, $\{0, 1\} \subset \Pi$. Denote by $\mathbf{\Pi}$ the set of all such Π.

As illustrated in Section 4.4, a delegation problem can be formulated as a balanced one. This is done by extending the principal's and agent's payoff functions to a sufficiently large interval of decisions such that, for both parties, the extreme decisions of this expanded interval are strictly inferior to all decisions in the original problem. This construction ensures that the requirement to include the extreme decisions into the delegation set is not binding.

The principal's objective is to choose a balanced delegation set $\Pi \in \mathbf{\Pi}$ to maximize her expected payoff, subject to the agent behaving optimally,

$$\max_{\Pi\in\mathbf{\Pi}} \mathbb{E}\big[V_D(\theta, x^*_\Pi(\theta))\big]. \tag{1}$$

[3]Our main result, Theorem 1, continues to hold if this assumption is relaxed, so that $\frac{\partial}{\partial x}U(\theta, x)$ satisfies strict single crossing from below in θ, strict single crossing from above in x, and signed-ratio monotonicity of Quah and Strulovici (2012).

[4]Define $\tilde{\theta} = F(\theta)$. By ($A_0$), the inverse $F^{-1} : [0, 1] \to [\underline{\theta}, \overline{\theta}]$ is continuous and strictly increasing. By (A_1), $\tilde{U}(\tilde{\theta}, x) = U(F^{-1}(\tilde{\theta}), x)$ and $\tilde{V}(\tilde{\theta}, x) = U(F^{-1}(\tilde{\theta}), x)$ are continuous in $\tilde{\theta}$, and, by (A_2), $\frac{\partial}{\partial x}\tilde{U}(\tilde{\theta}, x) = \frac{\partial}{\partial x}U(F^{-1}(\tilde{\theta}), x)$ is strictly increasing in $\tilde{\theta}$.

2.3. **Monotone Persuasion Problem.** In a monotone persuasion problem, the principal designs a monotone partitional information structure that partitions the state space into convex sets: singletons and intervals. The agent observes which partition element contains the state, where singletons are separating elements and intervals are pooling elements. Because the distribution of the state has a density, without loss of generality, each pooling interval is open. Thus, a monotone partition can be equivalently described by a set Π of singletons. Note that Π is a closed subset of $[0,1]$ that contains 0 and 1. Denote by $\mathbf{\Pi}$ the set of all such Π.

Consider a monotone partition $\Pi \in \mathbf{\Pi}$ and a primitive $(U_P, V_P) \in \mathcal{P}$, where we use subscript P to refer to the persuasion problem. To differentiate from the delegation problem, we denote the state by $\omega \in [0,1]$ and the decision by $y \in [0,1]$.

Denote by $\mu_\Pi(\omega)$ the partition element that contains ω, interpreted as a message that state ω induces. After observing message $\mu_\Pi(\omega)$, the agent chooses a decision that maximizes his expected payoff given the posterior belief induced by that message,

$$y^*_\Pi(\omega) \in \arg\max_{y\in[0,1]} \mathbb{E}[U_P(\omega', y) \mid \omega' \in \mu_\Pi(\omega)].$$

The principal's objective is to choose a monotone partition $\Pi \in \mathbf{\Pi}$ to maximize her expected payoff, subject to the agent behaving optimally,

$$\max_{\Pi\in\mathbf{\Pi}} \mathbb{E}[V_P(\omega, y^*_\Pi(\omega))]. \tag{2}$$

2.4. **Persuasion versus Delegation.** We now illustrate that implementable outcomes differ in the persuasion and delegation problems, given the same primitive $(U, V) \in \mathcal{P}$.

In the persuasion problem with $U(\omega, y) = -(\omega - y)^2$, consider a monotone partition that reveals whether the state is above or below 1/3. The induced decision of the agent is

$$y^*(\omega) = \begin{cases} \frac{1}{6}, & \omega < \frac{1}{3}, \\ \frac{2}{3}, & \omega \geq \frac{1}{3}, \end{cases}$$

where 1/6 is the midpoint between 0 and 1/3, and 2/3 is the midpoint between 1/3 and 1.

This outcome cannot be implemented in the delegation problem with the same primitive $U(\theta, x) = -(\theta - x)^2$. To see this, consider a delegation set that permits only two decisions, 1/6 and 2/3. The induced decision of the agent is

$$x^*(\theta) = \begin{cases} \frac{1}{6}, & \theta < \frac{5}{12}, \\ \frac{2}{3}, & \theta \geq \frac{5}{12}, \end{cases}$$

where 5/12 is the midpoint between 1/6 and 2/3. Thus, the induced decisions in the persuasion and delegation problems differ on the interval of intermediate states between 1/3 and 5/12.[5]

Note that the outcome y^* is the first best for the principal whose payoff function V is maximized at $y = 1/6$ for states below 1/3 and is maximized at $y = 2/3$ for states above 1/3. This first best is not implementable in the delegation problem with the same primitive. Conversely, the outcome x^* is the first best for the principal whose payoff function V is maximized at $x = 1/6$ for states below 5/12 and is maximized at $x = 2/3$ for states above 5/12. This first best is not implementable in the persuasion problem with the same primitive.[6]

3. Equivalence

3.1. **Main Result.** We use von Neumann and Morgenstern's (1944) notion of strategic equivalence. Primitives $(U_D, V_D) \in \mathcal{P}$ and $(U_P, V_P) \in \mathcal{P}$ of the balanced delegation and monotone persuasion problems are *equivalent* if there exists a constant C such that

$$\mathbb{E}\big[V_D(\theta, x^*_\Pi(\theta))\big] = \mathbb{E}\big[V_P(\omega, y^*_\Pi(\omega))\big] + C \quad \text{for all } \Pi \in \mathbf{\Pi}. \tag{3}$$

That is, if (U_D, V_D) and (U_P, V_P) are equivalent, then, for each strategy, the principal gets the same expected payoff, up to a constant, in both problems; consequently, the principal's optimal solution is also the same.

Theorem 1. *Consider primitives $(U_D, V_D) \in \mathcal{P}$ and $(U_P, V_P) \in \mathcal{P}$ of the balanced delegation and monotone persuasion problems. If, for all $(\theta, \omega) \in [0,1]^2$,*

$$\frac{\partial}{\partial \omega} U_D(\theta, \omega) + \frac{\partial}{\partial \theta} U_P(\omega, \theta) = 0 \quad \textit{and} \quad \frac{\partial}{\partial \omega} V_D(\theta, \omega) + \frac{\partial}{\partial \theta} V_P(\omega, \theta) = 0, \tag{4}$$

then (U_D, V_D) and (U_P, V_P) are equivalent.

The proof of Theorem 1 is deferred to Section 5.

For a given primitive of one problem, we can use Theorem 1 to construct an equivalent primitive of the other problem.

Corollary 1. *For each $(U_D, V_D) \in \mathcal{P}$, there exists an equivalent $(U_P, V_P) \in \mathcal{P}$, and vice versa.*

Proof. Given $(U_D, V_D) \in \mathcal{P}$, define (U_P, V_P) as

$$U_P(\omega, y) = -\int_0^y \frac{\partial}{\partial \omega} U_D(\theta, \omega) \mathrm{d}\theta \quad \text{and} \quad V_P(\omega, y) = -\int_0^y \frac{\partial}{\partial \omega} V_D(\theta, \omega) \mathrm{d}\theta. \tag{5}$$

[5]Since the outcome y^* cannot be implemented in the delegation problem, it cannot be implemented in the balanced delegation problem.

[6]Similarly, if the principal's payoff V is maximized at decision 0 for states below 1/2 and is maximized at decision 1 for states above 1/2, then the first best outcome is implementable by the balanced delegation set $\{0, 1\}$, but is not implementable by any monotone partition.

It is easy to see that (U_D, V_D) and (U_P, V_P) satisfy (4), and that (U_P, V_P) is in $\mathcal{P}$. Conversely, given $(U_P, V_P) \in \mathcal{P}$, define (U_D, V_D) as

$$U_D(\theta, x) = -\int_0^x \frac{\partial}{\partial \theta} U_P(\omega, \theta) \mathrm{d}\omega \quad \text{and} \quad V_D(\theta, x) = -\int_0^x \frac{\partial}{\partial \theta} V_P(\omega, \theta) \mathrm{d}\omega. \tag{6}$$

It is easy to see that (U_P, V_P) and (U_D, V_D) satisfy (4), and that (U_D, V_D) is in $\mathcal{P}$. □

3.2. **Linear Persuasion and Separable Delegation.** We now use Theorem 1 to illustrate the equivalence of linear monotone persuasion, as in Kamenica and Gentzkow (2011, p. 2601), and separable balanced delegation, as in Amador and Bagwell (2013, p. 1551).

Definition 1. A primitive (U_P, V_P) of a monotone persuasion problem is *linear* if

$$U_P(\omega, y) = -(\psi(\omega) - y)^2 \quad \text{and} \quad V_P(\omega, y) = \nu(y). \tag{7}$$

where $\omega \in [0, 1]$ is the uniformly distributed state, $y \in [0, 1]$ is the agent's decision, ψ is a continuous strictly increasing function, and ν is a continuously differentiable function.

Definition 2. A primitive (U_D, V_D) of a balanced delegation problem is *separable* if

$$\frac{\partial}{\partial x} U_D(\theta, x) = b(x) + c(\theta) \quad \text{and} \quad \frac{\partial}{\partial x} V_D(\theta, x) = Ab(x) + d(\theta), \tag{8}$$

where $\theta \in [0, 1]$ is the uniformly distributed state, $x \in [0, 1]$ is the agent's decision, $A \in \mathbb{R}$ is a constant, and b, c, and d are continuous functions such that b is strictly decreasing and c is strictly increasing.

It is easy to verify that (U_P, V_P) defined by (7) and (U_D, V_D) defined by (8) are in $\mathcal{P}$.

Consider a primitive (U_D, V_D) of a separable balanced delegation problem. By (5), the equivalent primitive (U_P, V_P) of the monotone persuasion problem is given by

$$\begin{aligned} U_P(\omega, y) &= -\int_0^y (b(\omega) + c(\theta)) \mathrm{d}\theta = \psi(\omega) y + \alpha(y), \\ V_P(\omega, y) &= -\int_0^y (Ab(\omega) + d(\theta)) \mathrm{d}\theta = A\psi(\omega) y + \beta(y), \end{aligned} \tag{9}$$

where $\psi(\omega) = -b(\omega)$, $\alpha(y) = -\int_0^y c(\theta) \mathrm{d}\theta$, and $\beta(y) = -\int_0^y d(\theta) \mathrm{d}\theta$.

Since (U_P, V_P) is linear in $\psi(\omega)$, the principal's expected payoff from a message $\mu_\Pi(\omega)$ of a monotone partition Π is a function ν of the posterior mean m of $\psi(\omega)$ induced by that message, $m = \mathbb{E}[\psi(\omega') | \omega' \in \mu_\Pi(\omega)]$.[7] Thus, the primitive (U_P, V_P) given by (9) and that given by (7) are equivalent.

[7]It is easy to show that the resulting function ν is differentiable almost everywhere, and our analysis in Section 3.1 can be appropriately extended to this case.

Conversely, consider a primitive (U_P, V_P) of a linear monotone persuasion problem. By (4), the equivalent primitive (U_D, V_D) of the balanced delegation problem satisfies

$$\frac{\partial}{\partial x}U_D(\theta, x) = \theta - \psi(x),$$
$$\frac{\partial}{\partial x}V_D(\theta, x) = -\nu'(\theta),$$

which is a special case of (8).

4. Application to Monopoly Regulation

We consider the classical problem of monopoly regulation as in Baron and Myerson (1982). The monopolist privately knows his cost and chooses a price to maximize profit. The welfare-maximizing regulator can restrict the set of prices the monopolist can choose from, for example, by imposing a price cap. Following Alonso and Matouschek (2008), we assume that the demand function is linear and the marginal cost has a unimodal distribution. Importantly, unlike Baron and Myerson (1982), transfers between the monopolist and regulator are prohibited.

We study two versions of this problem: (i) with the monopolist's participation constraint, as in Baron and Myerson (1982), and (ii) without any participation constraint, as in Alonso and Matouschek (2008). We formulate both versions as balanced delegation problems and address them by solving the equivalent monotone persuasion problems. We show that, in both versions, the welfare-maximizing regulator imposes a price cap, which is higher when the participation constraint is present.

Version (ii) was originally formulated as a standard (unbalanced) delegation problem and solved in Alonso and Matouschek (2008). Version (i), however, is novel and cannot be formulated as an unbalanced delegation problem. Moreover, it is not straightforward how to modify the existing techniques in Alonso and Matouschek (2008) and Amador and Bagwell (2013) to solve a balanced delegation problem. In contrast, it is easy to cast the balanced and unbalanced delegation problems as monotone persuasion problems. We use a single result from the persuasion literature to solve both problems.

4.1. Setup.

The demand function is $q = 1 - x$ where x is the price and q is the quantity demanded at this price. The cost of producing quantity q is γq. The marginal cost $\gamma \in [0, 1]$ is a random variable with a *unimodal* distribution F that has a strictly positive density f.

The monopolist's (agent's) profit is

$$U_D(\gamma, x) = (x - \gamma)(1 - x). \tag{10}$$

Welfare is the sum of the profit and consumer's surplus,

$$V_D(\gamma, x) = U_D(\gamma, x) + \tfrac{1}{2}(1 - x)^2. \tag{11}$$

The welfare-maximizing regulator (principal) chooses a set of prices $\Pi \subset [0,1]$ available for the monopolist. The monopolist privately observes the marginal cost γ and chooses a price x from Π to maximize profit $U_D(\gamma, x)$.

4.2. **Participation Constraint.** We first assume that the monopolist cannot be forced to operate at a loss. Formally, the monopolist can always choose to produce zero quantity, or equivalently, set price $x = 1$; so $1 \in \Pi$.

Notice that selling at zero price gives a lower profit than not producing at all, regardless of the value of the marginal cost. Thus, allowing the price $x = 0$ does not affect the monopolist's behavior; so, without loss of generality, $0 \in \Pi$.

To sum up, the regulator chooses a closed set of prices $\Pi \subset [0,1]$ that contains 0 and 1; so $\Pi \in \mathbf{\Pi}$.

To interpret this problem as a balanced delegation problem defined in Section 2.2, we change the variable $\theta = F(\gamma)$, so that θ is uniformly distributed on $[0,1]$. Denote $c(\theta) = F^{-1}(\theta)$. The monopolist's and regulator's payoffs are now given by

$$U_D(\theta, x) = (x - c(\theta))(1 - x) \quad \text{and} \quad V_D(\theta, x) = U_D(\theta, x) + \tfrac{1}{2}(1 - x)^2. \tag{12}$$

This is a separable balanced delegation problem, because U_D and V_D satisfy (8). The equivalent monotone persuasion problem is linear, as follows from Section 3.2.

4.3. **Translation to Persuasion Problem.** By (5), the equivalent primitive (U_P, V_P) of the monotone persuasion problem is given by

$$U_P(\omega, y) = -\int_0^y (1 + c(\theta) - 2\omega)\mathrm{d}\theta \quad \text{and} \quad V_P(\omega, y) = -\int_0^y (c(\theta) - \omega)\mathrm{d}\theta,$$

where $y \in [0,1]$ is a decision and $\omega \in [0,1]$ is a state. Now we find the equivalent (U_P, V_P) that satisfies (7).

Let $m = \mathbb{E}[\omega' | \omega' \in \mu_\Pi(\omega)]$ be the posterior mean state induced by a message $\mu_\Pi(\omega)$ of a monotone partition Π. Since U_P is linear in ω, the agent's optimal decision is a function $\bar{y}^*$ that depends on $\mu_\Pi(\omega)$ only through m. Thus, $\bar{y}^*(m)$ satisfies the first-order condition

$$\frac{\partial}{\partial y}\mathbb{E}\big[U_P(\omega', y)\big|\omega' \in \mu_\Pi(\omega)\big] = -(1 + c(y) - 2m) = 0$$

if $2m - 1 \geq 0$, and $y^*(m) = 0$ if $2m - 1 < 0$. Recalling that $c = F^{-1}$, we have

$$\bar{y}^*(m) = F(2m - 1),$$

where, by convention, $F(2m - 1) = 0$ if $2m - 1 < 0$.

Since V_P is linear in ω, the principal's expected payoff from a message $\mu_\Pi(\omega)$ is a function ν that depends on $\mu_\Pi(\omega)$ only through m:

$$\nu(m) = \mathbb{E}\big[V_P(\omega', \bar{y}^*(m))\big|\omega' \in \mu_\Pi(\omega)\big] = -\int_0^{\bar{y}^*(m)} (c(\theta) - m)\mathrm{d}\theta$$
$$= \int_0^{2m-1} (m - \gamma)\mathrm{d}F(\gamma). \tag{13}$$

Thus, the principal's objective is to choose a monotone partition that maximizes the expectation of ν,

$$\max_{\Pi\in\mathbf{\Pi}} \mathbb{E}[\nu(\mathbb{E}[\omega'|\omega' \in \mu_\Pi(\omega)])]. \tag{14}$$

This problem, albeit without the restriction to monotone partitions, has been formulated by Kamenica and Gentzkow (2011), who argue that the curvature of ν determines the form of the optimal information structure.

Lemma 1. *Let $\gamma_m \in (0,1)$ be the mode of the distribution F. Then, ν is convex on $[0,(1+\gamma_m)/2]$ (strictly so on $[1/2,(1+\gamma_m)/2]$) and strictly concave on $[(1+\gamma_m)/2,1]$.*

Proof. The derivative of ν is

$$\nu'(m) = 2(1-m)f(2m-1) + \int_0^{2m-1} f(\gamma)\mathrm{d}\gamma.$$

For any $m_1, m_2 \in [0,1]$,

$$\begin{aligned}\nu'(m_2) - \nu'(m_1) &= 2(1-m_2)[f(2m_2-1) - f(2m_1-1)] \\ &\quad + \int_{2m_1-1}^{2m_2-1} [f(\gamma) - f(2m_1-1)]\mathrm{d}\gamma.\end{aligned}$$

Thus, if $0 \le m_1 < m_2 \le (1+\gamma_m)/2$, then $\nu'(m_2) \ge \nu'(m_1)$, because $f(2m-1)$ is increasing for $m \in [0,(1+\gamma_m)/2]$. Moreover, if $1/2 \le m_1 < m_2 \le (1+\gamma_m)/2$, then $\nu'(m_2) > \nu'(m_1)$, because $f(2m-1)$ is strictly increasing for $m \in [1/2,(1+\gamma_m)/2]$. Finally, if $(1+\gamma_m)/2 \le m_1 < m_2 \le 1$, then $\nu'(m_2) < \nu'(m_1)$, because $f(2m-1)$ is strictly decreasing for $m \in [(1+\gamma_m)/2,1]$. □

Because ν is S-shaped (see Figure 1), the optimal information structure is an *upper-censorship*: the states below a threshold ω^* are separated, and the states above ω^* are pooled and induce the posterior mean state equal to $(1+\omega^*)/2$.[8]

Proposition 1. *Let $\gamma_m \in (0,1)$ be the mode of the distribution F. The monotone partition $\Pi^* = [0,\omega^*] \cup \{1\}$ is optimal, where $\omega^* \in (\gamma_m,(1+\gamma_m)/2)$ is the unique solution to*

$$\nu\left(\frac{1+\omega^*}{2}\right) - \nu(\omega^*) = \left(\frac{1+\omega^*}{2} - \omega^*\right)\nu'\left(\frac{1+\omega^*}{2}\right). \tag{15}$$

[8]Proposition 1 below follows from Proposition 3 in Kolotilin (2018). However, to illustrate the simplicity of the persuasion approach to this problem, we provide a complete proof, which is an adaptation of Dworczak and Martini's (2018) proof of Theorem 1.

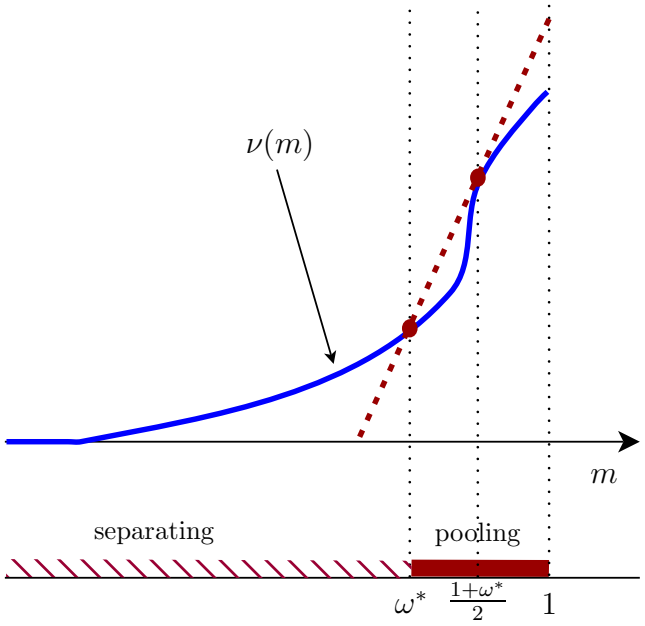


FIGURE 1. Optimal information structure with the participation constraint.

Proof. An information structure $\Pi \in \mathbf{\Pi}$ induces a distribution G_Π of the posterior mean state, $m = \mathbb{E}[\omega'|\omega' \in \mu_\Pi(\omega)]$. Since any such Π is a refinement of the fully informative partition $\overline{\Pi} = [0,1]$, the distribution $G_{\overline{\Pi}}$ is a mean-preserving spread of G_Π, by Blackwell (1953). Thus, to prove that Π^* is optimal, it is sufficient to show that $\int_0^1 \nu(m)\mathrm{d}G_{\Pi^*}(m) \geq \int_0^1 \nu(m)\mathrm{d}H(m)$ for all H such that $G_{\overline{\Pi}}$ is a mean-preserving spread of H.

By Lemma 1, ν is convex on $[0,(1+\gamma_m)/2]$ and strictly concave on $[(1+\gamma_m)/2,1]$. By assumption, ω is uniformly distributed on $[0,1]$. Thus, there exists a unique $\omega^* \in (\gamma_m,(1+\gamma_m)/2)$ that satisfies (15), as illustrated in Figure 1. Define function p as

$$p(m) = \begin{cases} \nu(m), & m \in [0,\omega^*], \\ \nu\left(\frac{1+\omega^*}{2}\right) + \left(m - \frac{1+\omega^*}{2}\right)\nu'\left(\frac{1+\omega^*}{2}\right), & m \in (\omega^*,1]. \end{cases} \tag{16}$$

It is easy to verify (see Figure 1) that

$$p(m) \geq \nu(m) \quad \text{for all } m \in [0,1], \tag{17}$$

$$p(m) \text{ is convex on } [0,1]. \tag{18}$$

We have

$$\int_0^1 (p(m)-\nu(m))\mathrm{d}H(m) \geq 0 = \int_0^1 (p(m)-\nu(m))\mathrm{d}G_{\Pi^*}(m), \tag{19}$$

where the inequality holds by (17) and the equality by (16) and the support of G_{Π^*} being $[0,\omega^*]\cup\{(1+\omega^*)/2\}$. Next,

$$\int_0^1 \nu(m)\mathrm{d}G_{\Pi^*}(m) - \int_0^1 \nu(m)\mathrm{d}H(m) \geq \int_0^1 p(m)\mathrm{d}G_{\Pi^*}(m) - \int_0^1 p(m)\mathrm{d}H(m)$$
$$= \int_0^1 p(m)\mathrm{d}G_{\overline{\Pi}}(m) - \int_0^1 p(m)\mathrm{d}H(m) \geq 0,$$

where the first inequality holds by (19), the equality by the linearity of $p(m)$ on $[\omega^*,1]$, and the second inequality by (18) and $G_{\overline{\Pi}}$ being a mean-preserving spread of H. □

Since the upper-censorship information structure $\Pi^* = [0,\omega^*]\cup\{1\}$ is optimal in the monotone persuasion problem, the same delegation set is optimal in the original balanced delegation problem (10)–(11). That is, the welfare-maximizing regulator imposes the price cap ω^*, thus implementing the price function

$$x^*(\gamma) = \begin{cases} \frac{1+\gamma}{2}, & \gamma < 2\omega^* - 1, \\ \omega^*, & 2\omega^* - 1 \leq \gamma < \omega^*, \\ 1, & \gamma \geq \omega^*. \end{cases}$$

In words, the monopolist chooses not to participate if his marginal cost γ is above the price cap ω^*. The participating monopolist chooses his preferred price $(1+\gamma)/2$ if it is below the cap, and he chooses the cap otherwise.

4.4. **Analysis without Participation Constraint.** We now assume that the monopolist operates even when making a loss.

To interpret this problem as a balanced delegation problem defined in Section 2.2, we observe that when the price x is sufficiently high or sufficiently low, both the monopolist and regulator would prefer intermediate prices. Thus, the requirement to include sufficiently extreme prices into the delegation set is not binding.

Formally, let us extend the definition of U_D and V_D in (12) to the domain of prices $[0,2]$, and let the regulator choose a closed set of prices $\Pi\subset[0,2]$.

Lemma 2. *If $\Pi\subset[0,2]$ is optimal, then $\Pi\cup\{0,2\}$ is optimal.*

Proof. First, we show that an optimal delegation set Π must contain a price in $[0,1]$. To see this, note that, by (12) and the fact that $c(\theta)\in[0,1]$, the monopolist's and regulator's payoffs are strictly decreasing in x for all $x\in[1,2]$ and all $\theta\in[0,1]$. Thus, the monopolist chooses the smallest price $\tilde{x}$ from Π whenever $\Pi\subset(1,2]$. It follows that the regulator strictly prefers $\tilde{\Pi}=\{1\}$ to any $\Pi\subset(1,2]$.

It remains to show that for any Π that contains a price $x\in[0,1]$, the delegation sets Π and $\Pi\cup\{0,2\}$ implement the same price function. This holds, because, by (12) and the fact that $c(\theta)\in[0,1]$, $U_D(\theta,x)\geq U_D(\theta,0)$ and $U_D(\theta,x)\geq U_D(\theta,2)$ for all $x\in[0,1]$ and all $\theta\in[0,1]$. □

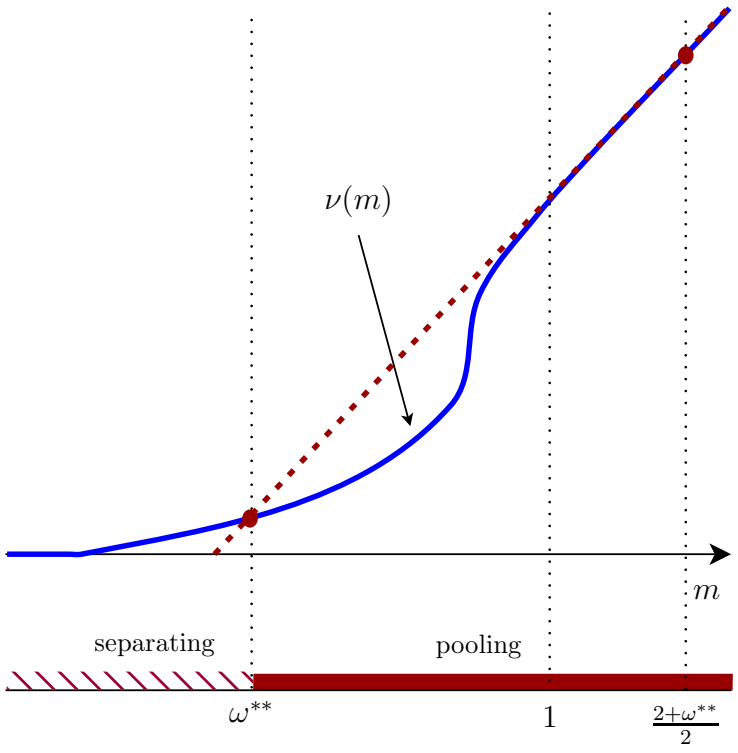


FIGURE 2. Optimal information structure without the participation constraint.

To sum up, the regulator chooses a closed delegation set $\Pi \subset [0,2]$ that contains 0 and 2. Rescaling the monopolist's decision so that it is in $[0,1]$, we obtain a separable balanced delegation problem. We can use Theorem 1 to find the equivalent primitive of the monotone persuasion problem, which qualitatively coincides with the problem in Section 4.3.

For comparability, it is convenient to rescale the state in the monotone persuasion problem, so that it is uniformly distributed on $[0,2]$. It is straightforward to verify that the principal's problem now is to choose a monotone partition $\Pi \subset [0,2]$ such that $\{0,2\} \in \Pi$ to maximize the expectation of ν given by

$$\nu(m) = \int_0^{2m-1} (m-\gamma)\mathrm{d}F(\gamma), \tag{20}$$

where, by convention, $F(2m-1) = 0$ if $2m-1 < 0$ and $F(2m-1) = 1$ if $2m-1 > 1$. Observe that $\nu(m)$ is just a linear extension of (13) to the interval $[0,2]$ (see Figure 2); so ν is convex on $[0,(1+\gamma_m)/2]$ and concave on $[(1+\gamma_m)/2,2]$.

Because ν is S-shaped, the optimal information structure is an upper-censorship: the states below a threshold ω^{**} are separated, and the states above ω^{**} are pooled and induce the posterior mean state equal to $(2+\omega^{**})/2$.

Proposition 2. *Let $\gamma_m \in (0,1)$ be the mode of the distribution F. The monotone partition $\Pi^{**} = [0, \omega^{**}] \cup \{2\}$ is optimal, where $\omega^{**} \in (0, (1+\gamma_m)/2)$ is the unique solution to*

$$\nu\left(\frac{2+\omega^{**}}{2}\right) - \nu\left(\omega^{**}\right) = \left(\frac{2+\omega^{**}}{2} - \omega^{**}\right)\nu'\left(\frac{2+\omega^{**}}{2}\right). \tag{21}$$

The proof of Proposition 2 is the same as the proof of Proposition 1.

Proposition 2, together with Lemma 2, implies that the optimal delegation set is $[0, \omega^{**}]$ in the original problem without the participation constraint. That is, the welfare-maximizing regulator imposes the price cap ω^{**}, thus implementing the price function

$$x^{**}(\gamma) = \begin{cases} \frac{1+\gamma}{2}, & \gamma < 2\omega^{**} - 1, \\ \omega^{**}, & \gamma \geq 2\omega^{**} - 1. \end{cases}$$

In words, the monopolist chooses his preferred price $(1+\gamma)/2$ if it is below the price cap, and he chooses the cap otherwise.

4.5. **Discussion.** Note that the optimal regulation policy takes the form of a price cap in both versions of the problem, with and without the monopolist's participation constraint. However, the optimal price cap is higher when the participation constraint is present, as follows from Propositions 1 and 2. Indeed, since ν is concave on $[(1+\gamma_m)/2, 2]$ and

$$\frac{1+\gamma_m}{2} < \frac{1+\omega^*}{2} < 1 < \frac{2+\omega^{**}}{2} < 2,$$

the slope of ν is higher at $(1+\omega^*)/2$ than at $(2+\omega^{**})/2$; so (15) and (21) imply that $\omega^* > \omega^{**}$ (see Figures 1 and 2).

We now build the intuition for why the optimal price cap is higher when the participation constraint is present. The first-order condition (15) for the optimal price cap ω^* can be written as

$$\int_{2\omega^*-1}^{\omega^*} (\omega^* - \gamma) f(\gamma) \mathrm{d}\gamma = \frac{1}{2}(1-\omega^*)^2 f(\omega^*), \tag{15$'$}$$

where the left-hand side and right-hand side correspond to the regulator's marginal gain and marginal loss of decreasing the price cap by $\mathrm{d}\omega$. The gain is that the monopolist with the cost $\gamma \in (2\omega^* - 1, \omega^*)$ now chooses the decreased price cap $\omega^* - \mathrm{d}\omega$, which is closer to his cost γ. The loss is that the monopolist with the cost $\gamma \in (\omega^* - \mathrm{d}\omega, \omega^*)$ now chooses to exit.

Instead, if the regulator does not take into account that the monopolist with the cost higher than the price cap exits, then the first-order condition (21) for the price cap ω^{**} can be written as

$$\int_{2\omega^{**}-1}^{\omega^{**}} (\omega^{**} - \gamma) f(\gamma) \mathrm{d}\gamma = \int_{\omega^{**}}^{1} (\gamma - \omega^{**}) f(\gamma) \mathrm{d}\gamma \tag{21$'$}$$

The regulator's marginal gain here is the same. But the marginal loss is that the monopolist with the cost $\gamma \in (\omega^{**}, 1)$ chooses the decreased price cap $\omega^{**} - \mathrm{d}\omega$, which is further from his cost γ.

Intuitively, the marginal loss in (15′) is higher than in (21′), because all surplus is lost if the monopolist exits, but only part of surplus is lost if the monopolist sets the price further from his cost.[9] This suggests that the regulator should give more discretion to the monopolist when she is concerned that the monopolist can exit.

5. Proof of Theorem 1

To prove the equivalence of the balanced delegation and monotone persuasion problems, we show that they are equivalent to a discriminatory disclosure problem.

5.1. **Discriminatory Disclosure Problem.** There are a principal and an agent. The agent chooses between *action* ($a = 1$) and *inaction* ($a = 0$). His preferred choice depends on his private type $\theta \in [0, 1]$ and a state $\omega \in [0, 1]$ that are independently and uniformly distributed. We denote by $u(\omega, \theta)$ and $v(\omega, \theta)$ the agent's and principal's payoffs if the agent chooses $a = 1$, and we normalize these payoffs to zero if the agent chooses $a = 0$. We assume that:

(A_1') $u(\omega, \theta)$ and $v(\omega, \theta)$ are continuous in ω and θ;

(A_2') $u(\omega, \theta)$ is strictly increasing in ω and strictly decreasing in θ.

A pair (u, v) that satisfies the above assumptions is a primitive of this problem.

The principal designs a menu Π of cutoff tests, where Π is a closed subset of $[0, 1]$. Each cutoff test $x \in \Pi$ reveals whether the state is below or above x. Note that cutoff tests with extreme cutoffs, 0 and 1, are uninformative, because their results are the same for all states. Thus, without loss of generality, we can assume that each menu Π contains the extreme cutoffs, $\{0, 1\} \in \Pi$; so, $\Pi \in \mathbf{\Pi}$.

The agent who knows his private type θ selects a cutoff test x from the menu Π, observes the result of the selected test, $\omega \geq x$ or $\omega < x$, and then chooses $a = 1$ or $a = 0$. Without loss of generality, assume that the agent chooses $a^*(\omega, x) = \mathbf{1}_{\{\omega \geq x\}}$. Indeed, by ($\mathrm{A}_2'$), the agent gets a higher payoff from $a = 1$ when the state is higher; so either he optimally chooses $a^*(\omega, x) = \mathbf{1}_{\{\omega \geq x\}}$, or he ignores the test result and chooses $a = 1$ or $a = 0$ irrespective of ω. But ignoring the result of a selected test is the same as selecting an uninformative test, 0 or 1, and then choosing $a^*(\omega, 0) = 1$ or $a^*(\omega, 1) = 0$ irrespective of ω.

The agent selects a test from the menu that maximizes his expected payoff,

$$x^*_\Pi(\theta) \in \arg\max_{x \in \Pi} \mathbb{E}_\omega \big[u(\omega, \theta) \big| \omega \geq x \big]. \tag{22}$$

[9] The right-hand side of (15′) can be expressed as $\int_{\omega^*}^1 (\gamma - \omega^*) f(\omega^*) \mathrm{d}\gamma$. This marginal loss is higher than in (21′), because $f(\omega^*) > f(\gamma)$ for $\gamma > \omega^* > \gamma_m$ by the unimodality of f.

The principal's objective is to choose a menu of cutoffs $\Pi \in \mathbf{\Pi}$ to maximize her expected payoff, subject to the agent behaving optimally,

$$\max_{\Pi\in\mathbf{\Pi}} \mathbb{E}_\theta\big[\mathbb{E}_\omega\big[v(\omega,\theta)\big|\omega \geq x^*_\Pi(\theta)\big]\big].$$

5.2. **Equivalence to Balanced Delegation.** For a primitive (u, v) of the discriminatory disclosure problem, which satisfies (A_1')–(A_2'), define

$$U_D(\theta, x) = \mathbb{E}_\omega\big[u(\omega,\theta)\big|\omega \geq x\big] \quad \text{and} \quad V_D(\theta, x) = \mathbb{E}_\omega\big[v(\omega,\theta)\big|\omega \geq x\big].$$

Conversely, for a primitive $(U_D, V_D) \in \mathcal{P}$ of the balanced delegation problem, define

$$u(\omega,\theta) = -\frac{\partial}{\partial\omega}U_D(\theta,\omega) \quad \text{and} \quad v(\omega,\theta) = -\frac{\partial}{\partial\omega}V_D(\theta,\omega).$$

It is easy to verify that (U_D, V_D) satisfies (A_1)–(A_2) if and only if (u, v) satisfies (A_1')–(A_2').

With these notations, (22) is identical to $x^*_\Pi(\theta) \in \arg\max_{x\in\Pi} U_D(\theta, x)$, and the principal's expected payoff is identical in both problems,

$$\mathbb{E}\big[V_D(\theta, x^*_\Pi(\theta))\big] = \mathbb{E}_\theta\big[\mathbb{E}_\omega\big[v(\omega,\theta)\big|\omega \geq x^*_\Pi(\theta)\big] \text{ for all } \Pi \in \mathbf{\Pi}.$$

5.3. **Equivalence to Monotone Persuasion.** Note that a menu $\Pi \in \mathbf{\Pi}$ defines a monotone partition of $[0, 1]$. The key observation is that in the discriminatory disclosure problem, the agent is indifferent between:

(i) optimally selecting a cutoff test x from the menu Π and observing $\omega \geq x$ or $\omega < x$;

(ii) observing which element of the monotone partition Π contains ω.

This is because in (ii), the agent divides all elements of the monotone partition into two groups, those inducing optimal choice $a = 1$ and those inducing optimal choice $a = 0$. Because the agent's payoff function $u(\omega, \theta)$ is strictly increasing in ω by (A_2'), these two groups form two intervals separated by a cutoff $x \in \Pi$, so the optimal choice is $a = 1$ if $\omega \geq x$ and $a = 0$ otherwise.

Consider the agent's normal-form strategy that maps each element of monotone partition Π and each type θ into a choice between $a = 1$ and $a = 0$. Because the agent's payoff function $u(\omega, \theta)$ is strictly decreasing in θ by (A_2'), for each $\omega \in [0, 1]$ there exists a threshold type $y \in [0, 1]$ such that

$$\mathbb{E}\big[u(\omega',\theta)\big|\omega' \in \mu_\Pi(\omega)\big] \geq 0 \text{ if and only if } \theta \leq y.$$

Consequently, a normal-form strategy of the agent reduces to a choice of threshold type y such that $a = 1$ when $\theta \leq y$ and $a = 0$ when $\theta > y$. Thus, the agent optimally chooses

$$y^*_\Pi(\omega) \in \arg\max_{y\in[0,1]} \mathbb{E}\left[\int_0^y u(\omega',\theta)\mathrm{d}\theta \middle| \omega' \in \mu_\Pi(\omega)\right]. \tag{23}$$

The principal's objective is to choose a monotone partition $\Pi \in \mathbf{\Pi}$ to maximize her expected payoff, subject to the agent behaving optimally,

$$\max_{\Pi \in \mathbf{\Pi}} \mathbb{E}_\omega \big[\mathbb{E}_\theta \big[v(\omega, \theta) \big| \theta \leq y^*_\Pi(\omega) \big] \big].$$

To complete the proof, for a primitive (u, v) of the discriminatory disclosure problem, which satisfies (A_1')–(A_2'), define

$$U_P(\omega, y) = \int_0^y u(\omega, \theta) \mathrm{d}\theta \quad \text{and} \quad V_P(\omega, y) = \int_0^y v(\omega, \theta) \mathrm{d}\theta.$$

Conversely, for a primitive $(U_P, V_P) \in \mathcal{P}$ of the monotone persuasion problem, define

$$u(\omega, \theta) = \frac{\partial}{\partial \theta} U_P(\omega, \theta) \quad \text{and} \quad v(\omega, \theta) = \frac{\partial}{\partial \theta} V_P(\omega, \theta).$$

It is easy to verify that (U_P, V_P) satisfies (A_1)–(A_2) if and only if (u, v) satisfies (A_1')–(A_2').

With these notations, (23) is identical to $y^*_\Pi(\omega) \in \arg\max_{y \in [0,1]} \mathbb{E}\big[U_P(\omega', \theta) \big| \omega' \in \mu_\Pi(\omega) \big]$, and the principal's expected payoff is identical in both problems,

$$\mathbb{E}\big[V_P(\omega, y^*_\Pi(\omega)) \big] = \mathbb{E}_\omega \big[\mathbb{E}_\theta \big[v(\omega, \theta) \big| \theta \leq y^*_\Pi(\omega) \big] \quad \text{for all } \Pi \in \mathbf{\Pi}.$$

6. Conclusion

We have shown the equivalence between balanced delegation and monotone persuasion, with the upshot that insights in delegation can be used for better understanding of persuasion, and vice versa. For instance, persuasion as the design of a distribution of posterior beliefs is notoriously hard to explain to a non-specialized audience. The connection to delegation can thus be instrumental in relaying technical results from the persuasion literature to practitioners and policy makers.

The classical delegation and persuasion problems have numerous extensions, which include a privately informed principal, competing principals, multiple agents, repeated interactions, and multidimensional state and decision spaces. We hope that our equivalence result will be a starting point for studying the connection between delegation and persuasion in these extensions.

Naturally, a principal may wish to influence an agent's decision by a combination of persuasion and delegation tools. How to optimally control both information and decisions of the agent, how these tools interact, and whether they are substitutes or complements are important questions that are left for future research.

References

Alonso, R., and N. Matouschek (2008): "Optimal Delegation," *Review of Economic Studies*, 75, 259–293.

Amador, M., and K. Bagwell (2013): "The Theory of Optimal Delegation with an Application to Tariff Caps," *Econometrica*, 81(4), 1541–1599.

Baron, D. P., and R. B. Myerson (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–930.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 24, 265–272.

Dessein, W. (2002): "Authority and Communication in Organizations," *Review of Economic Studies*, 69, 811–838.

Dworczak, P., and G. Martini (2018): "The Simple Economics of Optimal Persuasion," Working Paper, University of Chicago.

Gentzkow, M., and E. Kamenica (2016): "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review, Papers & Proceedings*, 106, 597–601.

Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani (2009): "Mediation, Arbitration and Negotiation," *Journal of Economic Theory*, 144, 1397–1420.

Guo, Y., and E. Shmaya (2018): "The Interval Structure of Optimal Disclosure," Working Paper, Northwestern University.

Holmström, B. (1984): "On the Theory of Delegation," in *Bayesian Models in Economic Theory*, ed. by M. Boyer, and R. Kihlstrom, pp. 115–141. New York: North-Holland.

Ivanov, M. (2017): "Optimal Monotone Signals in Bayesian Persuasion Mechanisms," Working Paper, McMaster University.

Kamenica, E., and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Kolotilin, A. (2018): "Optimal Information Disclosure: A Linear Programming Approach," *Theoretical Economics*, forthcoming.

Kolotilin, A., and H. Li (2018): "Relational Communication with Transfers," Working Paper, UNSW.

Kolotilin, A., H. Li, and W. Li (2013): "Optimal Limited Authority for Principal," *Journal of Economic Theory*, 148, 2344–2382.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85, 1949–1964.

Kováč, E., and T. Mylovanov (2009): "Stochastic Mechanisms in Settings without Monetary Transfers: The Regular Case," *Journal of Economic Theory*, 144, 1373–1395.

Krishna, V., and J. Morgan (2001): "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 95, 435–452.

Martimort, D., and A. Semenov (2006): "Continuity in Mechanism Design without Transfers," *Economics Letters*, 93(2), 182–189.

Melumad, N. D., and T. Shibano (1991): "Communication in Settings with No Transfers," *RAND Journal of Economics*, 22, 173–198.

Mensch, J. (2018): "Monotone Persuasion," Working Paper, Hebrew University.

Mylovanov, T. (2008): "Veto-based delegation," *Journal of Economic Theory*, 138, 297–307.

Ostrovsky, M., and M. Schwarz (2010): "Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*, 2, 34–63.

Quah, J., and B. Strulovici (2012): "Aggregating the Single Crossing Property," *Econometrica*, 80, 2333–2348.

Rayo, L., and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949 – 987.

von Neumann, J., and O. Morgenstern (1944): *Theory of Games and Economic Behavior*. Princeton University Press.