

Relational communication

ANTON KOLOTILIN

School of Economics, UNSW Business School

HONGYI LI

School of Economics, UNSW Business School

We study a communication game between an informed sender and an uninformed receiver with repeated interactions and voluntary transfers. Transfers motivate the receiver's decision-making and signal the sender's information. Although full separation can always be supported in equilibrium, partial or complete pooling is optimal if the receiver's decision-making is highly responsive to information. In this case, the receiver's decision-making is disciplined by pooling states where she is most tempted to defect.

KEYWORDS. Strategic communication, relational contracts.

JEL CLASSIFICATION. C73, D82, D83.

1. INTRODUCTION

Decision-makers and informed parties often develop relationships in which communication and decision-making are governed by informal agreements. We study how such interactions can be disciplined using relational contracts: discretionary compensation schemes that are self-enforcing in a repeated game. We characterize communication and decision-making patterns in optimal equilibria.

As an example of such relational communication, consider the relationships between political advocates and politicians. Political advocates seek to influence politicians, often over a broad range of policy decisions; think, for instance, of the Koch brothers' extensive lobbying activities. They ply politicians with information about the electoral consequences of various policy choices, such as voter attitudes toward clean-energy legislation or gun control. They also make transfers to politicians in the form of political contributions. Such transfers serve as contingent payments for favorable policy decisions (Grossman and Helpman (1994, 1996)) and credible signals of advocates'

Anton Kolotilin: akolotilin@gmail.com

Hongyi Li: hongyi@hongyi.li

We thank Luis Zermeno for important early contributions to this paper. We also thank the anonymous referees, and Ricardo Alonso, David Austen-Smith, Wouter Dessein, Sven Feldmann, Yuk-fai Fong, Robert Gibbons, Richard Holden, Johannes Hörner, Navin Kartik, Jin Li, Niko Matouschek, Tymofiy Mylovanov, Marco Ottaviani, Michael Powell, Heikki Rantakari, Joel Sobel, Kathryn Spier, and participants at various seminars and conferences for helpful comments and suggestions. Kolotilin acknowledges support from the Australian Research Council Discovery Early Career Research Award DE160100964 and from MIT Sloan's Program on Innovation in Markets and Organizations.

information (Austen-Smith (1995) and Lohmann (1995)). Because pay-to-play in this context is illegal, agreements between politicians and advocates are largely implicit and sustained within longstanding relationships.

Another example of relational communication is in repeated principal-agent relationships where an agent implements a series of projects and a principal has relevant information about these projects. Think, for instance, of a movie director and a studio, such as Akira Kurosawa and Toho Studios. The director has to decide whether to position each movie project as more mainstream (in which case the movie will be a sure-fire box-office draw) or as more art house (in which case profitability is uncertain, but potentially large). The studio may provide informed advice or recommendations about potential box-office revenues, but the director retains creative control. The studio may also make payments to the director—in the form of bonuses, perks, or additional funding—to reward the director's decision-making or to bolster the credibility of the studio head's advice.¹

To study relational communication, we add repeated interactions and voluntary transfers, as in Levin (2003), to the Crawford and Sobel (1982) model of strategic communication. In each period, the sender privately observes an independent draw of the state and sends a message to the receiver, who then makes a decision. The players' preferred decisions are increasing in the state, but the magnitude and sign of the difference between preferred decisions may depend on the state. The players can make voluntary transfers to each other at any point in the game.

In relational communication, transfers allow the sender not only to reward the receiver for compliant decision-making, but also to credibly signal his private information. In particular, full separation can be supported in equilibrium, even when the players are impatient. Therefore, the essential incentive constraint is that the receiver is tempted to make decisions that benefit herself but hurt the sender.

We show that a message rule can be supported in equilibrium if and only if it is *monotone*: it induces a monotone partition of the set of states. In any (Pareto) optimal equilibrium, the decision rule simply maximizes, subject to the receiver's incentive constraint, the joint payoff for each message. Therefore, given this decision rule, the optimal message rule solves the monotone persuasion problem: it maximizes the expected joint payoff over all monotone message rules. This is the Bayesian persuasion problem (Rayo and Segal (2010) and Kamenica and Gentzkow (2011)) with an additional constraint that message rules must be monotone.

We completely characterize the optimal (second-best) equilibrium when the players' payoffs are quadratic. Our key insights are about how information should be selectively hidden and revealed to manage decision-making within an optimal relationship. Consider *high-conflict* states where conflict of interest is so large that the first-best decision is not self-enforcing. At these states, self-enforcement requires that second-best decision-making be distorted in favor of the receiver. If the sender's and receiver's preferred decisions respond similarly to information, then first-best and second-best

¹Hermalin (1998) and Benabou and Tirole (2003) discuss how payments from principal to agent can be used not only to reward performance, but also to credibly signal the principal's private information.

decision-making also respond similarly to information, so full separation is optimal. In contrast, if the receiver is highly responsive to information relative to the sender, then second-best decision-making is too responsive to information, so high-conflict states are optimally pooled to moderate second-best decision-making. In sum, optimal relational communication involves pooling if and only if high-conflict states exist and the receiver is highly responsive.

The result that the sender hides information only when conflict of interest is sufficiently large seems to be a natural pattern of communication in relationships. Advocates often discuss in detail the costs and benefits of potential legislation with politicians, but may hide their private information in cases that are particularly controversial or consequential. Similarly, principals provide honest advice and agents comply when their preferences are largely aligned, but principals may hide information when agents are most tempted to dissent or disobey.

The result that pooling of high-conflict states occurs in relationships only if the receiver is highly responsive can be illustrated by fleshing out and comparing our political advocacy and movie-making examples. Starting with our advocacy example, suppose that the decision concerns how much to deregulate gun control and that the state is the degree of deregulation that is most popular with voters. Further suppose that the sender/advocate is an ideologue who puts little weight on voter popularity and always prefers more deregulation than does the receiver/politician, while the politician is electorally motivated and simply prefers the most popular policy. This corresponds to a receiver who is highly responsive (that is, his preferred policy is much more responsive to the state than that of the sender). Our model then predicts that the advocate will pool states where the conflict of interest is large, that is, will withhold information from the politician when deregulation is most unpopular.

Turning to our movie-making example, suppose that the decision is where to position a movie project's type along the art-house–mainstream spectrum and that the state is the most profitable movie type. Further suppose that the director/receiver puts little weight on profitability and always prefers a more art-house type than does the studio head/sender, while the studio simply prefers the most profitable type. This corresponds to a receiver who is not highly responsive. Our model then predicts that no pooling takes place: the studio always fully informs the director, even in states where art-house projects would be very unprofitable and so the conflict of interest is large—unlike in the political advocacy example.

In our model, pooling does not only occur at high-conflict states, when the receiver is highly responsive. Suppose the players are neither too patient nor too impatient, so that high-conflict and low-conflict states coexist. Then over-pooling occurs: high-conflict states are optimally pooled with some adjacent low-conflict states to further ease self-enforcement at those high-conflict states. In other words, optimal relationships hide information about some states where full separation and first-best decision-making could be supported in equilibrium.

We also show that relational communication becomes more informative as the discount factor increases. As the players become more patient, second-best decision-making more closely approximates first-best decision-making and, thus, makes better

use of information. Consequently, the sender optimally reveals more information to the receiver.

An implication of our analysis is that in settings where voluntary transfers are available, incomplete information transmission does not imply a failure to motivate communication, but instead is a tool to discipline decision-making. In other words, the Pareto frontier cannot be expanded simply by introducing a technology for credible (monotone) communication.² Indeed, we show that adding public information generally worsens the relationship.

1.1 *Related literature*

In our model, transfers from the sender to the receiver are used to signal information. [Austen-Smith and Banks \(2000\)](#) and [Kartik \(2007\)](#) consider a related (albeit static) setting where the sender burns money to signal information.³ Unlike burning money, signaling information with transfers incurs no welfare cost. This leads to a clean characterization of the set of optimal equilibria; in particular, all optimal equilibria in our model produce identical communication outcomes. In the setting with burned money, equilibrium communication outcomes differ along the Pareto frontier because there is a trade-off between the efficiency of informed decision-making and the costs of burning money.⁴ As a by-product, we establish a general characterization of equilibria in games of cheap talk with burned money: a message rule is implementable if and only if it is monotone.⁵

Our analysis builds on an extensive literature on repeated interactions with transfers. The seminal papers by [Bull \(1987\)](#) and [Macleod et al. \(1989\)](#) focus on settings with symmetric information. [Levin \(2003\)](#) characterizes the optimal relational contract in two important settings with asymmetric information: adverse selection and moral hazard. In these settings, only the decision-maker (agent) has private information. In contrast, our setting involves an informed sender and an uninformed decision-maker (receiver), in the vein of [Crawford and Sobel \(1982\)](#). In such relational communication, pooling affects the receiver's beliefs and, thus, directly influences her decision-making. In contrast, the decision-maker (agent) in [Levin \(2003\)](#) is fully informed, so pooling has no such Bayesian persuasion effect. Instead, in [Levin's \(2003\)](#) analysis of relational adverse selection, pooling of agent's types reduces the variability of transfers to satisfy the self-enforcement constraint.

²This is in contrast with the existing literature on cheap talk and delegation, where the receiver's expected payoff (which is the standard welfare criterion) unambiguously improves if credible communication can be costlessly achieved.

³[Kartik et al. \(2007\)](#) and [Kartik \(2009\)](#) consider related models with lying costs instead of money burning.

⁴The receiver's optimal equilibrium clearly involves full separation; [Karamychev and Visser \(2017\)](#) characterize the sender's optimal equilibrium.

⁵[Ottaviani \(2000\)](#), [Krishna and Morgan \(2008\)](#), and [Ambrus and Egorov \(2017\)](#) consider communication games where contractible transfers from the receiver to the sender are used to elicit the sender's information, as in mechanism design. Due to a limited liability constraint, the receiver has to leave information rents to the sender and thus trades off the efficiency of informed decision-making with the corresponding information rents, leading to information pooling. In contrast, in our setting, information rents do not arise and information pooling occurs for Bayesian persuasion purposes.

Alonso and Matouschek (2007) also consider repeated communication. In contrast to us, they disallow transfers and consider a sequence of short-lived senders rather than a single long-lived sender.⁶ In their setting, the sender may hide information when the receiver cannot credibly commit to take appropriate informed decisions. Repeated interactions allow the receiver to credibly sustain decisions favorable for the sender, so as to motivate the sender to communicate more information. In contrast, in our setting, credible communication can be costlessly achieved for any decision rule. Repeated interactions allow the receiver to credibly sustain decisions that improve joint surplus. If repeated interactions cannot sustain first-best decision-making, then communication is an additional tool to further improve decision outcome via Bayesian persuasion.⁷

In our model, optimal equilibria are supported by carrot-and-stick strategies (Abreu (1986) and Goldlücke and Kranz (2012)) in which a deviator is punished as harshly as possible but only for a single period. We show that the receiver is punished by complete pooling of information and the sender is punished by an extreme incentive compatible decision. These punishments also characterize the receiver's and sender's worst equilibria in games of cheap talk with burned money.

Our paper also contributes to the rapidly growing literature on Bayesian persuasion with transferable utility (Bergemann and Pesendorfer (2007), Eső and Szentes (2007), Li and Shi (2017), Bergemann et al. (2018), and Dworczak (2020)). Similar to these papers, we use tools from mechanism design and Bayesian persuasion. Unlike these papers, commitment power in our model is endogenous and, thus, imperfect.

2. MODEL

2.1 Setup

A sender (S) and a receiver (R) play an infinitely repeated communication game with perfect monitoring and with voluntary transfers. Time is discrete and the players have a common discount factor $\delta \in [0, 1]$. In each period, the same stage game is played. The sender privately observes a state $\theta \in [0, 1]$ and sends a message $m \in [0, 1]$ to the receiver, who then makes a decision $d \in \mathbb{R}$. The state θ is independently drawn each period from a prior distribution $F(\theta)$ with a strictly positive density $f(\theta)$ for all $\theta \in [0, 1]$. The sender's and receiver's payoffs, $u_S(d, \theta)$ and $u_R(d, \theta)$, satisfy Crawford and Sobel's (1982) assumptions:

ASSUMPTION 1. *For each player $i \in \{S, R\}$,*

- (i) *$u_i(d, \theta)$ is twice differentiable in d and θ for all $d \in \mathbb{R}$ and $\theta \in [0, 1]$,*

⁶Baker et al. (2011) consider a model of repeated decision-making with transfers between long-lived players, but assume symmetric information, so communication plays no role.

⁷A model of repeated Bayesian persuasion would reproduce many of the insights from our model of relational communication. There is a literature on dynamic Bayesian persuasion, albeit with persistent information (Kremer et al. (2014), Au and Hung (2015), Ely et al. (2015), Hörner and Skrzypacz (2016), Ely (2017), Che and Horner (2018), Ely and Szydlowski (2020), Bizzotto et al. (2021), Orlov et al. (2020), Best and Quigley (2020), and Smolin (forthcoming)).




FIGURE 1. Timing of the stage game.

- (ii) $\frac{\partial^2 u_i}{\partial d^2}(d, \theta) < 0$ for all $d \in \mathbb{R}$ and $\theta \in [0, 1]$,
- (iii) $\frac{\partial u_i}{\partial d}(\rho_i(\theta), \theta) = 0$ for some function $\rho_i(\theta)$ and for all $\theta \in [0, 1]$,
- (iv) $\frac{\partial^2 u_i}{\partial d \partial \theta}(d, \theta) > 0$ for all $d \in \mathbb{R}$ and $\theta \in [0, 1]$.

Parts (ii) and (iii) of Assumption 1 require that each player's payoff is strictly concave in the decision and each player $i \in \{S, R\}$ has a unique *preferred* decision $\rho_i(\theta)$ for each state $\theta \in [0, 1]$. Similarly, there is a unique first-best decision $\rho_{FB}(\theta)$ that maximizes the joint payoff $u(d, \theta) = u_S(d, \theta) + u_R(d, \theta)$. Part (iv) is a sorting condition that ensures that $\rho_S(\theta)$, $\rho_R(\theta)$, and $\rho_{FB}(\theta)$ are strictly increasing in θ .

The players can make voluntary (noncontractible) transfers at any point in the game. Specifically, we enrich the stage game with three rounds of transfers: (a) an *ex ante* transfer $\tau \in \mathbb{R}$ before the sender observes the state, (b) an *interim* transfer $t \in \mathbb{R}$ after the sender observes the state and sends the message but before the receiver chooses a decision, and (c) an *ex post* transfer $T \in \mathbb{R}$ after the decision is chosen.⁸ At each round, a positive transfer represents a payment from sender to receiver; conversely, a negative transfer represents a payment from receiver to sender. Thus, the stage game payoff of the sender is $u_S(d, \theta) - \tau - t - T$ and the stage game payoff of the receiver is $u_R(d, \theta) + \tau + t + T$. Since transfers are voluntary, the sender can reject a positive transfer and the receiver can reject a negative transfer.

The game has perfect monitoring in that all actions (message, decision, and transfers) are immediately publicly observed, but the state is only observed by the sender. That is, the receiver never observes the state or her payoff.⁹ Figure 1 summarizes the timing of each stage game.

We study pure-strategy perfect Bayesian *equilibria*; later, in footnote 17, we comment on mixed-strategy equilibria. For each period and each history, an equilibrium specifies (on-path) a message rule $\mu(\theta)$ for the sender, a decision rule $\rho(m)$ for the receiver, and transfer rules τ , $t(m)$, and $T(m)$.¹⁰

Conventions. A (pure-strategy) message rule deterministically maps states to the messages they induce. Without loss of generality, we identify each message with the

⁸Because next period's ex ante transfer can substitute for this period's ex post transfer, the set of equilibrium payoffs would not change if ex post transfers were removed. But the analysis is simpler if ex post transfers are allowed.

⁹This assumption is common in the literature on repeated games with incomplete information (Aumann et al. (1995)) and is ubiquitous in models of repeated communication (Renault et al. (2013), Frankel (2016), Margaria and Smolin (2018), and Lipnowski and Ramos (2020)). In Section 4.2, we briefly discuss the case where the state is publicly observed at the end of each period.

¹⁰The functions μ , ρ , τ , t , and T are required to be measurable.

set of states that induce this message: $m = \{\theta : \mu(\theta) = m\}$. Thus, the range $\mu([0, 1])$ of a message rule μ is a partition of the set of states. A message rule μ is *monotone* if each $m \in \mu([0, 1])$ is a convex set (either a singleton or an interval).¹¹

We can now extend the definition of payoffs and preferred decisions from being state dependent to being message dependent. Specifically, $u_i(d, m) = \mathbb{E}_F[u_i(d, \theta)|m]$ and $\rho_i(m) = \arg \max_{d \in \mathbb{R}} u_i(d, m)$ for each player $i \in \{S, R\}$. Similarly, $u(d, m) = u_S(d, m) + u_R(d, m)$ and $\rho_{FB}(m) = \arg \max_{d \in \mathbb{R}} u(d, m)$. Assumption 1 ensures that $\rho_S(m)$, $\rho_R(m)$, and $\rho_{FB}(m)$ are well defined and are strictly increasing in m in the strong set order.

2.2 Stationarity

We focus on stationary equilibria. An equilibrium is *stationary* if on the equilibrium path, the message rule μ , the decision rule ρ , and the transfer rules τ , t , and T are identical in every period. An equilibrium is *optimal* if it is not Pareto dominated by any other equilibrium. An equilibrium is *sequentially optimal* if the continuation equilibrium following any history on the equilibrium path is optimal.

LEMMA 1. *There exist $\underline{v}_S$, $\underline{v}_R$, and $\bar{v}$ such that the set of optimal equilibrium payoffs is the line segment*

$$\bar{V} = \{(v_S, v_R) \in \mathbb{R}^2 : v_S \geq \underline{v}_S, v_R \geq \underline{v}_R, v_S + v_R = \bar{v}\}. \quad (1)$$

Any optimal equilibrium is sequentially optimal. Further, there exists a stationary optimal equilibrium σ_* such that any optimal equilibrium payoff vector (v_S, v_R) can be supported by an equilibrium that differs from σ_* only in the first-period ex ante transfer.

Lemma 1 extends some of Levin's (2003) and Goldlücke and Kranz's (2012) results to our setting, with an extensive-form stage game of incomplete information. Because players' payoffs are quasi-linear in money, payoffs are fully transferable, and contingent transfers can substitute for contingent continuation payoffs. Consequently, we can restrict attention to stationary equilibria, and all optimal equilibria induce the message and decision rules that maximize joint payoff $v = v_S + v_R$.

3. EQUILIBRIUM

3.1 Implementability

We now show that the presence of interim and ex post voluntary transfers enables separation of the sender's and receiver's incentive constraints. The sender's incentive constraint requires that the decision outcome be monotone. The receiver's incentive constraint requires that induced decisions be close to the receiver's preferred decisions.

¹¹For example, a message rule that separately pools low states $[0, 1/3]$ and high states $(2/3, 1]$ into two distinct messages while separating intermediate states $[1/3, 2/3]$ is monotone. However, a message rule that pools low and high states into a single non-convex message $[0, 1/3] \cup (2/3, 1]$ while separating intermediate states is not monotone.

Define the receiver's temptation to deviate from decision d given message m as

$$w(d, m) = u_R(\rho_R(m), m) - u_R(d, m)$$

and define the *net discounted surplus* given joint payoff v as

$$L(v) = \frac{\delta}{1 - \delta}(v - \underline{v}_S - \underline{v}_R).$$

PROPOSITION 1. *A message rule μ and a decision rule ρ that produce a joint payoff v can be supported in a stationary equilibrium if and only if the decision outcome is monotone,*

$$\rho(\mu(\theta)) \text{ is nondecreasing in } \theta, \quad (2)$$

and the receiver's temptation to deviate never exceeds the net discounted surplus,

$$w(\rho(m), m) \leq L(v) \quad \text{for all } m \in \mu([0, 1]). \quad (3)$$

We first argue that (2) and (3) are necessary. In any equilibrium, the message rule $\mu(\theta)$ must be incentive compatible for the sender. Since the sender's payoff is quasi-linear in money and satisfies a sorting condition, a standard characterization of incentive compatibility in mechanism design (see, for example, [Rochet \(1987\)](#)) implies that $\rho(\mu(\theta))$ must be nondecreasing in θ .

Also, in any equilibrium, the decision rule ρ must be incentive compatible for the receiver. Therefore, given a message m , the receiver's one-period payoff gain from choosing her preferred decision $\rho_R(m)$ instead of equilibrium decision $\rho(m)$ must be less than the maximum available punishment equal to the discounted surplus.

We now argue that (2) and (3) are sufficient. Ignoring the sender's incentive compatibility constraint, any decision rule ρ that satisfies (3) can be made incentive compatible for the receiver by giving all surplus to the receiver ($v_R = v - \underline{v}_S$) and threatening her with her worst equilibrium payoff ($v_R = \underline{v}_R$) following any deviation from $\rho(m)$.

In such a construction, the sender receives his worst equilibrium payoff $\underline{v}_S$ and, thus, cannot be punished for deviating. However, for any message rule μ that satisfies (2), we can separately construct a (voluntary) interim transfer rule that makes μ incentive compatible for the sender.

The envelope theorem (see, for example, [Milgrom and Segal \(2002\)](#)) implies that there exists a unique (up to a constant C) interim transfer rule t such that the sender prefers to induce $\rho(\mu(\theta))$ and pay $t(\mu(\theta))$ rather than to induce $\rho(\mu(\hat{\theta}))$ and pay $t(\mu(\hat{\theta}))$ for all $\hat{\theta} \neq \theta$,

$$t(m) = u_S(\rho(m), \theta(m)) - \int_0^{\theta(m)} \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta} + C, \quad (4)$$

where $\theta(m)$ is an arbitrary state $\theta \in m$.¹² The constant C can be chosen in such a way that the sender does not want to deviate to any out-of-equilibrium message-transfer

¹²Since $\rho(\mu(\theta))$ is nondecreasing in θ , $t(m)$ is independent of the choice of a representative state $\theta \in m$.

pair $(\hat{m}, \hat{t})$. Specifically, choose C such that the minimum transfer is equal to zero and is achieved for some punishment message m^P ,¹³

$$t(m) \geq 0 \text{ for all } m \in \mu([0, 1]), \text{ with equality for some } m^P \in \mu([0, 1]). \quad (5)$$

If, following any out-of-equilibrium pair $(\hat{m}, \hat{t})$, the receiver believes that the state is in m^P and chooses the punishment decision $d^P = \rho(m^P)$, then the sender prefers to report m^P and pay $t(m^P) = 0$ rather than to report $\hat{m}$ and pay $\hat{t}$. Thus, the sender's incentive constraint is satisfied.¹⁴

This argument implies that voluntary interim transfers are powerful in signaling information, even if the players are myopic.¹⁵

COROLLARY 1. *Suppose $\delta = 0$. A message rule μ and a decision rule ρ can be supported in an equilibrium if and only if μ is monotone and $\rho(m) = \rho_R(m)$ for all $m \in \mu([0, 1])$.*

Corollary 1 is closely connected to existing results from the literature on cheap talk and burned money (Austen-Smith and Banks (2000), Kartik (2007), and Karamychev and Visser (2017)). In the myopic setting, interim transfers serve the same signaling role as burned money. In fact, the set of implementable message and decision rules does not depend on whether the sender transfers money to the receiver or whether the sender burns money.¹⁶

In contrast to burned money, interim transfers are not wasteful: the sender's loss is the receiver's gain. Further, since ex ante transfers are available, the use of interim transfers does not create a distributional imbalance. Any surplus obtained by the receiver from interim transfers can be redistributed to the sender using ex ante transfers. Such ex ante transfers can be supported by the threat of complete pooling of information. Consequently, the sender can commit at no welfare cost to any monotone message rule.

3.2 Optimality

An optimal equilibrium solves a monotone persuasion problem: it maximizes the expected joint payoff over monotone message rules, subject to the second-best decision rule.

¹³In the proof, we allow for the possibility that $\inf t(m)$ is not attained by any m^P .

¹⁴Note that self-enforcement does not impose any limit on the variability of interim transfers and, thus, does not create a shadow cost to screening, which is a key driving force in Levin's (2003) analysis of relational adverse selection.

¹⁵Although interim transfers are powerful, messages are still used to convey information. For example, suppose the players' preferred decision rules intersect at some state. Then in any fully separating equilibrium, the interim transfer function is non-monotone and takes the same value for multiple state realizations. Messages are thus used to distinguish between these realizations.

¹⁶Karamychev and Visser's (2017) Proposition 1 characterizes implementable outcomes with money burning. Our mechanism design approach to characterization provides a much simpler proof of the result and removes the assumptions that the bias $\rho_R - \rho_{FB}$ has constant sign and that the receiver's payoff satisfies a sorting condition. Indeed, if the receiver's payoff did not satisfy part (iv) of Assumption 1, our Proposition 1 and its proof would still hold, but Corollary 1 would require that $\rho_R(\mu(\theta))$ be nondecreasing in θ , rather than that μ be monotone.

Define the second-best decision given message m as

$$\begin{aligned}\rho_*(m) &= \arg \max_d u(d, m) \\ \text{subject to } w(d, m) &\leq L(\bar{v})\end{aligned}\tag{6}$$

and define the joint payoff under the second-best decision as

$$u_*(m) = u(\rho_*(m), m) \quad \text{for all } m \in [0, 1].\tag{7}$$

PROPOSITION 2. *In a stationary optimal equilibrium, the message rule satisfies*

$$\begin{aligned}\mu_* &\in \arg \max_{\mu} \mathbb{E}[u_*(\mu(\theta))] \\ \text{subject to } \mu &\text{ is monotone,}\end{aligned}\tag{8}$$

and the second-best decision $\rho_*(m)$ is taken for each on-path message $m \in \mu_*([0, 1])$.

The intuition for Proposition 2 is as follows. By Proposition 1, an optimal equilibrium maximizes v jointly over message and decision rules that satisfy the sender's and receiver's incentive constraints, (2) and (3). By constraint (2) and a revelation principle argument, we can restrict attention to monotone message rules. Consider a relaxed problem in which the constraint (2) is replaced with the constraint that the message rule is monotone. It is easy to see that ρ_* given by (6) and μ_* given by (8) solve this relaxed problem. Further, we show that $\rho_*(m)$ is nondecreasing in m because the sender's and receiver's payoffs satisfy the sorting condition (part (iv) of Assumption 1). Therefore, $\rho_*(\mu_*(\theta))$ is nondecreasing in θ , the constraint (2) is automatically satisfied, and ρ_* and μ_* constitute an optimal equilibrium.

Proposition 2 shows that the decision rule and message rule in any optimal equilibrium can be calculated in two steps. First, the decision rule is characterized without reference to the message rule. The decision rule is point-wise equal to the second-best decision $\rho_*(m)$ given by (6). For each message m , the second-best decision $\rho_*(m)$ can be found as follows. If $d = \rho_{FB}(m)$ satisfies the constraint of (6), then $\rho_*(m) = \rho_{FB}(m)$. Otherwise, $\rho_*(m)$ is the unique decision d that lies between $\rho_R(m)$ and $\rho_{FB}(m)$ and satisfies the constraint of (6) with equality. Second, given ρ_* and thus u_* , the message rule μ_* solves the *monotone persuasion* problem (8): it maximizes the expected joint payoff $\mathbb{E}[u_*(\mu(\theta))]$ over all monotone message rules μ .¹⁷

¹⁷We have restricted attention to pure-strategy equilibria. Proposition 2 would continue to hold if we instead restricted attention to equilibria where the receiver uses a pure strategy. Indeed, since u_S is super-modular, in any such equilibrium, the sender's incentive constraint implies that the decision outcome must be monotone, with randomizations only at a countable set of states, which—since θ has a density—do not affect expected payoffs. Moreover, Proposition 2 would continue to hold without restriction to pure strategies for either player if we additionally assume that $\partial^2 u_S(d, \theta) / \partial d \partial \theta = \beta'(\theta) \gamma'(d)$ for some increasing functions β and γ , for example, as in Assumption 2. Under this additional assumption, in any mixed-strategy equilibrium, the sender's incentive constraint implies that a higher state induces a higher expectation of $\gamma(d)$, with randomizations by the sender only at a countable set of states. Thus, as far as payoffs are concerned, we can restrict attention to equilibria with a pure monotone message rule. Given this restriction and the assumption that payoffs are concave in the decision, the second-best decision rule is pure and so is the receiver's strategy in any optimal equilibrium.

3.3 High-conflict states

We say that a state θ is *high conflict* if the first-best decision is not enforceable at this state, $w(\rho_{FB}(\theta), \theta) > L(\bar{v})$; otherwise, the state is *low conflict*.

COROLLARY 2. *If full separation is suboptimal, then some states are high conflict, and each nonsingleton message $m \in \mu_*([0, 1])$ contains some high-conflict states.*

PROOF. Consider the second-best decision rule and a monotone message rule where some nonsingleton message m consists only of low-conflict states. Then the expected joint payoff can be increased by separating all states in m , thus implementing the first-best decision $\rho_{FB}(\theta)$ for each $\theta \in m$, while keeping the other messages unchanged. $\square$

Further, we can specify sufficient local conditions for pooling to be optimal at some high-conflict states.

PROPOSITION 3. *Suppose $L(\bar{v}) > 0$, and $u_S(d, \theta)$ and $u_R(d, \theta)$ are thrice differentiable in d and θ . Full separation is suboptimal if there exists a high-conflict state $\hat{\theta} \in [0, 1]$ such that*

$$\frac{1}{2}u''_{dd}(\rho_*(\hat{\theta}), \hat{\theta})(\rho'_*(\hat{\theta}) - 2\tilde{\rho}'_{FB}(\hat{\theta})) + u'_d(\rho_*(\hat{\theta}), \hat{\theta})\left(\frac{\rho''_*(\hat{\theta})}{2\rho'_*(\hat{\theta})} - \frac{\varphi''(\hat{\theta})}{2\varphi'(\hat{\theta})}\right) < 0, \quad (9)$$

$$\text{where } \tilde{\rho}'_{FB}(\hat{\theta}) = -\frac{u''_{d\theta}(\rho_*(\hat{\theta}), \hat{\theta})}{u''_{dd}(\rho_*(\hat{\theta}), \hat{\theta})} \quad \text{and} \quad \varphi(\theta) = u_R(\rho_R(\hat{\theta}), \theta) - u_R(\rho_*(\hat{\theta}), \theta).$$

Separating a high-conflict state $\hat{\theta}$ is suboptimal if for a small neighborhood m around $\hat{\theta}$, pooling is better than separation, $\mathbb{E}[u(\rho^*(m), \theta)|m] > \mathbb{E}[u(\rho^*(\theta), \theta)|m]$. Equation (9) shows that the trade-off between separation and pooling is affected by the slope of ρ_* and the curvatures of ρ_* and φ .

Importantly, as it turns out, the second-best decision rule can be approximated as

$$\rho_*(m) \approx \rho_*(\varphi^{-1}(\mathbb{E}[\varphi(\theta)|m])). \quad (10)$$

To see why, notice that the receiver's binding incentive constraint $w(\rho_*(m), m) = L(\bar{v})$ implies

$$\rho_*(m) - \rho_*(\hat{\theta}) \approx \frac{\mathbb{E}[u_R(\rho_R(\hat{\theta}), \theta) - u_R(\rho_*(\hat{\theta}), \theta)|m]}{u_{Rd}'(\rho_*(\hat{\theta}), \hat{\theta})} = \frac{\mathbb{E}[\varphi(\theta)|m]}{u_{Rd}'(\rho_*(\hat{\theta}), \hat{\theta})}.$$

That is, locally, $\rho_*(m)$ depends on m only through $\mathbb{E}[\varphi(\theta)|m]$, so we can approximately express $\rho_*(m) \approx \zeta(\mathbb{E}[\varphi(\theta)|m])$ for some function ζ . Expression (10) then follows from the observation that $\zeta(\cdot) = \rho_*(\varphi^{-1}(\cdot))$, which in turn follows from $\rho_*(\theta) = \zeta(\varphi(\theta))$.

The curvatures of ρ^* and φ jointly determine the expected second-best decision under pooling, $\rho_*(\varphi^{-1}(\mathbb{E}[\varphi(\theta)|m]))$, and separation, $\mathbb{E}[\rho_*(\varphi^{-1}(\varphi(\theta))|m]$. Comparing the expected decision under pooling and separation is equivalent to comparing the expected utility of a consumer with a Bernoulli utility function $\rho_*(\varphi^{-1}(\cdot))$ under two fair

lotteries over consumption quantities. Pooling corresponds to a degenerate lottery that yields a certain quantity $\mathbb{E}[\varphi(\theta)|m]$ and separation corresponds to a risky lottery that yields a random quantity $\varphi(\theta)|m$. Consequently, as ρ_* and φ^{-1} become more concave, the consumer becomes more risk-averse, which makes the degenerate lottery more attractive, so the expected decision becomes relatively higher under pooling. Thus, pooling is more favorable if the first-best decision is higher than the second-best decision, $\rho_{FB}(\hat{\theta}) > \rho_*(\hat{\theta})$ or, equivalently, $u'_d(\rho_*(\hat{\theta}), \hat{\theta}) > 0$.

The slope of ρ_* does not affect the expected decisions under pooling and separation, but instead determines whether pooling or separation better approximates the slope of the “locally first-best” decision rule.¹⁸ The decision outcome has slope $\rho'_*(\hat{\theta})$ under separation and slope 0 under pooling. Pooling is thus more favorable if $\tilde{\rho}'_{FB}(\hat{\theta})$ is closer to 0 than to $\rho'_*(\hat{\theta})$ or, equivalently, $\rho'_*(\hat{\theta}) > 2\tilde{\rho}'_{FB}(\hat{\theta})$. We develop this intuition more fully in Section 4.1, where the payoffs are quadratic.

In the myopic case, where $L(\bar{v}) = 0$ and, correspondingly, $\varphi(\theta) = u_{R,d}'(\rho_R(\hat{\theta}), \theta)$, our setting reduces to that of the monotone persuasion problem in which the decision and the state are both one dimensional. Proposition 3 can thus be interpreted as specifying local sufficient conditions for optimal persuasion to involve some pooling. Closest to this analysis is Jehiel (2015), who also considers local sufficient conditions for optimal persuasion to involve pooling. His main result is that full separation is suboptimal if there exist two distinct states at which the receiver’s preferred decision is the same. This result has no bite in our setting, where the receiver’s preferred decision is strictly increasing in the state.

To summarize, Corollary 2 provides necessary conditions for pooling to be optimal, whereas Proposition 3 provides sufficient conditions. The next section specializes to the case of quadratic payoffs, so as to provide conditions that are both necessary and sufficient for pooling to be optimal, and to fully characterize optimal relational communication.

4. QUADRATIC PAYOFFS

ASSUMPTION 2. *The sender’s and receiver’s payoffs are given by $u_S(d, \theta) = \lambda_S(\rho_S(\theta)d - d^2/2)$ and $u_R(d, \theta) = \lambda_R(\rho_R(\theta)d - d^2/2)$ for all $d \in \mathbb{R}$ and $\theta \in [0, 1]$, where preferred decisions $\rho_S(\theta)$ and $\rho_R(\theta)$ are increasing and linear in θ , and weights $\lambda_S > 0$ and $\lambda_R > 0$ satisfy $\lambda_S + \lambda_R = 1$.*

Assumption 2 implies Assumption 1, so all our previous results apply. Given the normalization $\lambda_S + \lambda_R = 1$, the first-best decision is $\rho_{FB}(\theta) = \lambda_R\rho_R(\theta) + \lambda_S\rho_S(\theta)$ and the joint payoff is $u(d, \theta) = \rho_{FB}(\theta)d - d^2/2$.¹⁹

¹⁸The locally first-best decision rule $\tilde{\rho}_{FB}(\theta)$ maximizes $\mathbb{E}[u(\rho(\theta), \theta)|m]$ among all linear rules $\rho(\theta)$ that satisfy $\rho(\hat{\theta}) = \rho_*(\hat{\theta})$.

¹⁹These payoff functions nest two special cases. First, in Crawford and Sobel’s (1982) example, the sender has constant upward bias, so that $\rho_S(\theta) = \rho_R(\theta) + b$ with $b > 0$. Second, in Kamenica and Gentzkow’s (2011) lobbying example, the sender is biased toward a specific decision $d_c \in \mathbb{R}$, so that $\rho_S(\theta) = a\rho_R(\theta) + (1-a)d_c$ with $a \in (0, 1)$.

Assumption 2 ensures that $u_*(m)$ depends on a message only through the induced posterior mean state: $u_*(m) = u_*(\mathbb{E}[\theta|m])$ for all $m \in [0, 1]$.²⁰ Therefore, without loss of generality, we identify each message m with the induced posterior mean state, $m = \mathbb{E}[\theta|m]$. This simplifies the previous convention that identified each message with the set of states that induce it.

4.1 Pooling versus separation

We first show that it is optimal to pool some states whenever the players are not patient enough to enforce the first-best outcome and the slope of the receiver's preferred decision rule is sufficiently high.

Under Assumption 2, the second-best decision (6) given message m pushes d as close to $\rho_{FB}(m)$ as possible, while still keeping d within distance ℓ from the receiver's preferred decision $\rho_R(m)$,

$$\begin{aligned} \rho_*(m) &= \arg \max_d u(d, m) \\ \text{subject to } |d - \rho_R(m)| &\leq \ell = \sqrt{\frac{L(\bar{v})}{\lambda_R}}, \end{aligned}$$

where we call ℓ the relational *leeway*. The second-best decision rule ρ_* is parallel to ρ_R at high-conflict states, where $|\rho_{FB}(\theta) - \rho_R(\theta)| > \ell$, and coincides with ρ_{FB} at low-conflict states.

The trade-off between pooling and separation is tightly linked to the curvature of the joint payoff under the second-best decision

$$u_*(\theta) = \rho_{FB}(\theta)\rho_*(\theta) - \frac{\rho_*^2(\theta)}{2},$$

which is in turn determined by how the second-best decision responds to the state. Define the receiver to be *highly responsive* if

$$\rho'_R(\theta) > 2\rho'_{FB}(\theta). \quad (11)$$

We can strengthen Corollary 2 and Proposition 3 as follows.

COROLLARY 3. *Full separation is suboptimal if and only if some states are high conflict and the receiver is highly responsive. In that case, u_* is concave for high-conflict states and convex for low-conflict states.*

PROOF. As shown in Kamenica and Gentzkow (2011), if u_* is convex, then full separation is optimal. Since each message rule μ is less informative than the fully separating

²⁰Bayesian persuasion in settings where payoffs depend only on the induced posterior mean state is studied, for example, in Gentzkow and Kamenica (2016), Kolotilin et al. (2017), Kolotilin (2018), and Dworczak and Martini (2019).

message rule, the prior distribution F is a mean-preserving spread of the distribution of posterior means induced by μ . Thus, if u_* is convex, then

$$\mathbb{E}[u_*(\theta)] \geq \mathbb{E}[u_*(\mu(\theta))],$$

showing that full separation is optimal.

Conversely, if u_* is strictly concave on some nonempty interval (θ_1, θ_2) , then full separation is suboptimal. Indeed, if u_* is strictly concave on (θ_1, θ_2) , then

$$\mathbb{E}[u_*(\theta)|\theta \in (\theta_1, \theta_2)] < u_*(\mathbb{E}[\theta]|\theta \in (\theta_1, \theta_2)),$$

showing that full separation can be improved upon by pooling states (θ_1, θ_2) .

To complete the proof of the corollary, it suffices to show that u_* is continuously differentiable, and that $u_*''(\theta) < 0$ if and only if the state θ is high conflict and the receiver is highly responsive. This is established in Lemma 3 in Appendix C, with

$$\begin{aligned} u_*''(\theta) &= (2\rho'_{FB}(\theta) - \rho'_*(\theta))\rho'_*(\theta) \\ &= \begin{cases} \rho'^2_{FB}(\theta) & \text{if } \theta \text{ is low conflict,} \\ (2\rho'_{FB}(\theta) - \rho'_R(\theta))\rho'_R(\theta) & \text{if } \theta \text{ is high conflict.} \end{cases} \quad \square \end{aligned}$$

Corollary 3 shows that at high-conflict states, condition (9) simplifies to (11), and further becomes necessary and sufficient for pooling to be optimal.²¹ We pause to explain in detail how the receiver's responsiveness drives the separation–pooling trade-off. Suppose that all states are high conflict and the receiver is downward biased, so that the second-best decision rule is $\rho_*(m) = \rho_R(m) + \ell$. The optimal message rule should induce a decision outcome that approximates the first-best decision outcome as closely as possible.

To understand how the message rule shapes decision-making, first notice that any message rule must induce receiver's beliefs that are correct in expectation, so that every message rule always induces the same expected decision outcome. Thus, the expected decision outcome cannot be moved closer to the expected first-best decision by changing the message rule. Optimizing the message rule thus involves making the “slope” of the induced decision outcome as close as possible to that of the first-best decision outcome.

Let us compare complete pooling and full separation. Complete pooling induces a completely “flat” decision outcome that is completely unresponsive to the state. Full separation induces a decision outcome that runs parallel to the receiver's preferred decision, so $\rho'_*(\theta) = \rho'_R(\theta)$. This decision outcome is clearly more responsive to the state than complete pooling and, further, is more responsive than the first-best outcome if $\rho'_R(\theta) > \rho'_{FB}(\theta)$. Thus, the first-best slope is better approximated by full separation than by complete pooling if and only if $\rho'_{FB}(\theta)$ is closer to $\rho'_R(\theta)$ than to 0 or, equivalently, $\rho'_R(\theta) < 2\rho'_{FB}(\theta)$; that is, the receiver is not highly responsive.

²¹The curvatures of φ and ρ_* play no role in the separation–pooling trade-off because φ is linear and ρ_* is piecewise linear.

This logic extends to the more general case where not all states are high conflict. It is optimal to fully separate states if the receiver is not highly responsive and to pool the high-conflict states if the receiver is highly responsive. We next show how the high-conflict states should be optimally pooled with low-conflict states when the receiver is highly responsive.

4.2 Optimal communication

We characterize optimal equilibria in the case where the receiver is highly responsive and is downward biased, $\rho_R(\theta) < \rho_{FB}(\theta)$ for all $\theta \in [0, 1]$. Here, the set of high-conflict states is a (possibly empty) interval $[0, \theta_c]$ with $\theta_c \in [0, 1]$. At the end of this section, we briefly discuss the case where the receiver is upward biased for some states and downward biased for others.

PROPOSITION 4. *Suppose the receiver is highly responsive and downward biased. There exists a cutoff state $\theta_* \in [0, 1]$ such that the optimal message rule pools the states below θ_* into one message and separates the states above θ_* . Furthermore, the pooling interval includes all high-conflict states and some adjacent low-conflict states, $\theta_* > \theta_c$, if high-conflict and low-conflict states coexist, $\theta_c \in (0, 1)$. Otherwise, $\theta_c \in \{0, 1\}$, the pooling interval is equal to the set of high-conflict states, $\theta_* = \theta_c$.*

PROOF. By Corollary 3, if $\theta_c = 0$, then u_* is convex and full separation is optimal. Similarly, if $\theta_c = 1$, then u_* is concave and complete pooling is optimal. Moreover, if $\theta_c \in (0, 1)$, then u_* is concave on $[0, \theta_c]$ and convex on $[\theta_c, 1]$. Then Proposition 4 follows from Remark 1 in Kolotilin et al. (forthcoming).²² For completeness, we outline the proof here. Consider an auxiliary payoff function

$$u_A(\theta) = \begin{cases} u_*(m_*) + u'_*(m_*)(\theta - m_*) & \text{if } \theta < \theta_*, \\ u_*(\theta) & \text{if } \theta \geq \theta_*, \end{cases}$$

where $m_* = \mathbb{E}[\theta | \theta < \theta_*]$ and $\theta_* \geq \theta_c$ is such that $u_A(\theta)$ is continuous at θ_* (Figure 2). Notice that u_A is convex and $u_A \geq u_*$. Let us solve the auxiliary problem of choosing a message rule μ to maximize $\mathbb{E}[u_A(\mu(\theta))]$. Since u_A is convex, full separation is optimal in the auxiliary problem, and the maximum expected auxiliary payoff is

$$\begin{aligned} \bar{v}_A &= F(\theta_*)\mathbb{E}[u_A(\theta) | \theta < \theta_*] + (1 - F(\theta_*))\mathbb{E}[u_A(\theta) | \theta \geq \theta_*] \\ &= F(\theta_*)u_*(m_*) + (1 - F(\theta_*))\mathbb{E}[u_*(\theta) | \theta \geq \theta_*]. \end{aligned}$$

Since $u_A \geq u_*$, $\bar{v}_A$ is an upper bound on the maximum expected joint payoff $\bar{v}$ in problem (8). But the message rule that pools the states below θ_* and separates the rest achieves this upper bound. Since this message rule is monotone, it solves the monotone persuasion problem (8). $\square$

²²Related results also appear in Kolotilin et al. (2017), Kolotilin (2018), and Dworczak and Martini (2019).




FIGURE 2. Joint and auxiliary payoff functions.

Proposition 4 highlights what we call *over-pooling*: all high-conflict states are optimally pooled with some adjacent low-conflict states (Figure 3). To see why over-pooling occurs, consider the effect of marginally expanding the pooling interval from $[0, \theta_*)$ to $[0, \theta_* + d\theta)$. The cost of this expansion is that marginal states switch from the first-best decision to the second-best decision, for a loss of $(u_*(\theta_*) - u_*(m_*))f(\theta_*)d\theta$. The benefit of this expansion is that decision-making for the pooled states marginally improves toward the first-best decision (Figure 3(b)), for a gain of $u'_*(m_*)(\theta_* - m_*)f(\theta_*)d\theta$. The cutoff state, if interior, is such that the benefit equals the cost:

$$u'_*(m_*)(\theta_* - m_*) = (u_*(\theta_*) - u_*(m_*)). \quad (12)$$

If the receiver is highly responsive, then u_* is concave on $[0, \theta_c)$ (Figure 2). In this case, evaluation of (12) at $\theta_* = \theta_c$ indicates that the benefit of the marginal expansion outweighs the cost, leading to over-pooling: the cutoff state θ_* is greater than θ_c . Intuitively, with a highly responsive receiver, the second-best decision rule—which runs parallel to the receiver's preferred decision at high-conflict state—is also highly responsive to the message, so expanding the pool results in a large shift in the pooled decision toward the first-best and, thus, a large benefit (relative to the cost) from over-pooling.

As the players become more patient, the interval of high-conflict states $[0, \theta_c)$ shrinks²³ and the pooling interval shrinks with it.

COROLLARY 4. *Suppose the receiver is highly responsive and downward biased. As the players become more patient, the pooling interval shrinks, $d\theta_*/d\delta \leq 0$, and strictly so if and only if $\theta_* \in (0, 1)$.*

Corollary 4 states that more information is optimally revealed in the sense of Blackwell (1953) as the players become more patient. With more patient players, the second-best decision rule more closely approximates the first-best decision rule and, thus, makes better use of information. Consequently, the sender optimally reveals more information.²⁴

²³This is because the relational leeway ℓ increases with δ , as shown in Lemma 4 in Appendix C.

²⁴As Corollary 5 in Appendix C shows, more information is optimally revealed as the highly responsive receiver becomes less responsive to the state (that is, as $\rho'_R(\theta)/\rho'_{FB}(\theta)$ decreases).




FIGURE 3. Optimal decision rules and decision outcomes given a highly responsive, downward-biased receiver. Figure (a) shows the high- δ case, with full separation and $\theta_c = \theta_* = 0$; Figure (b) shows the intermediate- δ case, with partial pooling and $0 < \theta_c < \theta_* < 1$; Figure (c) shows the low- δ case, with complete pooling and $0 < \theta_c < \theta_* = 1$. The last two cases illustrate over-pooling: $\theta_c < \theta_*$.

As an aside, we briefly discuss how optimal relational communication depends on the monitoring structure. We have assumed that the decision is publicly observed, but the state is only observed by the sender. Alternatively, we might assume that the state is publicly observed at the end of each period. In this case, non-monotone message rules could be enforced by conditioning continuation play on the state. However, by the proof of Proposition 4, the second-best unconstrained message rule turns out to be monotone, so Proposition 4 and Corollary 4 would continue to hold. Alternatively still, we might assume that the receiver never observes the state and the sender never observes the decision.²⁵ Then the receiver would always take her preferred decision; thus, complete pooling would be optimal if the receiver is highly responsive and full separation would be optimal otherwise.

The results of this section extend to the case where the receiver is highly responsive and is upward biased for some states and downward biased for others.²⁶ Both high and low states may be high conflict, with intermediate states being low conflict. In this case, optimal communication takes one of three forms (Figure 4). In all forms, over-pooling occurs: high-conflict states are pooled with adjacent low-conflict states. Some intermediate interval of low-conflict states may be separated. As the players become more impatient, the high-conflict-state intervals expand, as do the corresponding pooling intervals. Eventually, the interval of separated states vanishes, so that the state space is

²⁵Kuvalekar et al. (2020) study a repeated communication game with such a monitoring structure, but without transfers. Hörner et al. (2015) consider a much richer setting with communication and general monitoring structures, but focus on the case of sufficiently patient players.

²⁶This case is technically challenging because, unlike the case where the receiver is downward biased, the optimal unconstrained message rule is not necessarily monotone, so existing results from the literature on Bayesian persuasion no longer apply. A working-paper version of this paper (Kolotilin and Li (2019)) derives new results on monotone persuasion (Section 4) and applies them to characterize optimal equilibria in this case (Section 5.3). The working paper also derives a closed-form solution for the case where the state is uniformly distributed (Section 5.4).




FIGURE 4. Optimal decision rules and decision outcomes given a highly responsive receiver and two intervals of high-conflict states, $[0, \underline{\theta}_x)$ and $(\bar{\theta}_x, 1]$. Figure (a) shows the case of relatively high δ , with two pooling intervals $[0, \underline{\theta}_*)$ and $(\bar{\theta}_*, 1]$ and separation of the middle states $\theta \in (\underline{\theta}_*, \bar{\theta}_*)$. Figure (b) shows the intermediate- δ case, with two adjacent pooling intervals $[0, \theta_{**})$ and $(\theta_{**}, 1]$. Figure (c) shows the low- δ case with complete pooling. Over-pooling occurs in all cases.

partitioned into two pooling intervals. The two-pooling-interval structure remains optimal as the players become yet more impatient—until, at some sufficiently small discount factor, the two pooling intervals suddenly coalesce into one and complete pooling becomes optimal.

5. PUNISHMENT

In this section, we characterize each player's worst equilibrium: the communication and decision-making outcomes that are used, as part of an optimal equilibrium, to punish the player as harshly as possible following a deviation. Unsurprisingly, the receiver is punished by complete pooling of information, so that she takes an uninformed decision.

We show in the general setting of Assumption 1 that the sender is punished by the highest or lowest incentive compatible decision. The sender's worst message rule may combine pooling and separation, which hurt the sender in distinct ways. Pooling misadapts decisions to the state, while separation extracts signaling transfers from the sender.

The worst equilibrium payoffs $\underline{v}_R$ and $\underline{v}_S$ can be supported by single-period punishment strategy profiles in which a deviator is punished as harshly as possible but only for a single period.²⁷ A *single-period punishment* strategy profile is described by: normal as well as penal ex ante transfers, $\underline{\tau}_0$ as well as $\underline{\tau}_R$ and $\underline{\tau}_S$; normal as well as penal message rules, $\underline{\mu}_0$ as well as $\underline{\mu}_R$ and $\underline{\mu}_S$; normal as well as penal decision rules, $\underline{\rho}_0$ as well as $\underline{\rho}_R$ and $\underline{\rho}_S$.

²⁷Baker et al. (1994, 2002) restrict attention to trigger-strategy equilibria where off-path punishments correspond to some static equilibria of the stage game; such trigger strategies are suboptimal in our setting because static equilibria are not the harshest possible punishments. Alternatively, Levin (2003) specifies exogenous outside option payoffs. Our results continue to hold in these alternative approaches, with the worst equilibrium payoffs replaced by either static equilibrium payoffs or outside option payoffs.

Play proceeds as follows. The ex ante transfer is $\underline{\tau}_i$ if player $i \in \{0, R, S\}$ deviated last in the previous period, where $i = 0$ denotes that no player deviated. The message rule, interim transfer rule, decision rule, and punishment message are $\underline{\mu}_j$, $\underline{t}_j$, $\underline{\rho}_j$, and $\underline{m}_j^P$ if player $j \in \{0, R, S\}$ deviated from the ex ante transfer in this period, where $\underline{t}_j$ and $\underline{m}_j^P$ are defined by (4) and (5) given $\underline{\mu}_j$ and $\underline{\rho}_j$. The punishment decision is $\underline{d}_j^P = \underline{\rho}_j(\underline{m}_j^P)$ if the sender deviated to some $(\hat{m}, \hat{t}) \notin (\underline{\mu}_j, \underline{t}_j)([0, 1])$ in this period. Ex post transfers are always zero.

PROPOSITION 5. *There exists an optimal equilibrium in single-period punishment strategies where the following statements hold:*

- (i) *The on-path rules are $\underline{\mu}_0 = \mu_*$ and $\underline{\rho}_0 = \rho_*$, so $\bar{v} = \mathbb{E}[u_*(\mu_*(\theta))]$.*
- (ii) *The receiver's penal rules are $\underline{\mu}_R = [0, 1]$ and $\underline{\rho}_R = \rho_R$, so $\underline{v}_R = u_R(\rho_R([0, 1]), [0, 1])$.*
- (iii) *The sender's penal rules $\underline{\mu}_S$ and $\underline{\rho}_S$ solve*

$$\begin{aligned} \underline{v}_S = \min_{\mu, \rho, \theta^P} & \left\{ u_S(\rho(m^P), \theta^P) + \mathbb{E} \left[\int_{\theta^P}^{\theta} \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta} \right] \right\} \\ & \text{subject to } \rho(\mu(\theta)) \text{ is nondecreasing in } \theta, \\ & \rho(m) \begin{cases} = \rho_-(m) & \text{if } m > m^P, \\ \in \{\rho_-(m), \rho_+(m)\} & \text{if } m = m^P, \\ = \rho_+(m) & \text{if } m < m^P, \end{cases} \end{aligned} \quad (13)$$

where $m_P = \mu(\theta_P)$ and $[\rho_-(m), \rho_+(m)] = \{d : w(d, m) \leq L(\bar{v})\}$ for all $m \in \mu([0, 1])$.

Proposition 5 specifies optimal punishments for the receiver and the sender: a deviator is punished as harshly as possible for a single period and then optimal play resumes. The deviator's worst equilibrium payoff equals his or her payoff in the punishment period. Following a deviation from τ by the receiver, the message rule is completely uninformative and no transfers are made. Following a deviation from τ by the sender, the receiver makes either the highest or lowest incentive compatible decision, and the message and interim transfer rules are chosen to minimize the sender's expected payoff.

The Uniform-Quadratic Example. Suppose that the players' payoffs are quadratic and the receiver is downward biased, as in Section 4. Suppose further that the state is uniformly distributed. As shown in Proposition 7 in Appendix D, the (upward-biased) sender's penal decision rule is the lowest incentive compatible decision rule, $\underline{\rho}_S(m) = \rho_R(m) - \ell$ for all m . Moreover, the sender's penal message rule either pools some interval of low states $[0, \underline{\theta}_S)$ or separates all states. As the players become more patient, signaling transfers become relatively more effective as a punishment; thus the optimal penal pooling interval shrinks and full separation eventually becomes optimal.

6. PUBLIC INFORMATION

In this section, we show that increasing transparency by adding public information generally worsens the relationship. The availability of transfers as a signaling device implies that better informed decision-making can always be achieved without tightening incentive constraints, so transparency adds no informational benefits for the relationship, but tightens incentive constraints in two ways. First, it improves both players' worst possible equilibrium payoffs and, thus, limits the severity of off-path punishments. Second, it prevents information pooling and, thus, limits the ability to discipline decision-making in high-conflict states.

We augment our model so that at the start of each period, the receiver observes a realization of a state-dependent signal. We consider the general payoffs of Assumption 1, rather than the quadratic payoffs of Assumption 2. Just as with message rules, we assume that the signal rule $\psi(\theta)$ is deterministic and (without loss) identify each signal realization s with the set of states that induce it, $s = \{\theta : \psi(\theta) = s\}$. We also assume that the signal rule ψ is monotone in the sense that each $s \in \psi([0, 1])$ is a convex set.

Since the signal and message rules are deterministic, we can restrict attention to message rules that are *refinements* of the signal rule in that for each realization s of ψ , there exists a realization m of μ such that $m \subset s$. In particular, this restriction allows us to consider decision rules ρ that depend on the message m but not the signal realization s , because m incorporates all information contained in s .

The set of optimal equilibrium payoffs under signal ψ can be computed by applying Proposition 5 separately to each realization of signal ψ . In particular, the optimal and penal message and decision rules are defined for each signal realization $s \in \psi([0, 1])$ as follows: ρ_* is given by (6) and μ_* solves (8) given that the set of states is s rather than $[0, 1]$; $\underline{\rho}_R = \rho_R$ and $\underline{\mu}_R = \psi$; and $\underline{\rho}_S$ and $\underline{\mu}_S$ solve (13) given that the set of states is s .

We say that ψ is *more informative* than $\hat{\psi}$ if ψ is a refinement of $\hat{\psi}$. For monotone signal rules, this notion coincides with the informativeness criterion of Blackwell (1953).

PROPOSITION 6. *Suppose that signal ψ is more informative than signal $\hat{\psi}$. If the sender's penal decision outcome $\underline{\rho}_S(\underline{\mu}_S(\theta))$ is nondecreasing in θ under ψ , then the best joint payoff is lower under ψ than under $\hat{\psi}$.*

To build intuition for Proposition 6, we start with the myopic case. Consider moving from a fully informative signal ($\psi_f(\theta) = \theta$) to a completely uninformative signal ($\psi_u(\theta) = [0, 1]$). We argue that the worst equilibrium payoffs $\underline{v}_S$ and $\underline{v}_R$ strictly decrease, and the best joint payoff $\bar{v}$ weakly increases.

The receiver's worst equilibrium payoff $\underline{v}_R$ is lower under ψ_u than ψ_f . In the receiver's worst equilibrium, the receiver always chooses her preferred decision $\rho_R(\psi(\theta))$ given the signal ψ and always receives zero transfers. Public information improves the receiver's decision-making and, thus, her worst equilibrium payoff.

The sender's worst equilibrium payoff $\underline{v}_S$ is lower under ψ_u than ψ_f . The basic idea is that any equilibrium decision outcome implemented under ψ_f (and, thus, a fully informed receiver) can also be implemented under ψ_u by inducing the sender to fully reveal the state to the receiver. The sender's payoff $\underline{v}_S$ is strictly smaller under ψ_u because

inducing full separation requires the sender to make positive interim transfers to the receiver.

The best joint payoff $\bar{v}$ is weakly higher under ψ_u than ψ_f , again because any equilibrium under ψ_f can be implemented under ψ_u . In fact, the best joint payoff may be strictly higher under ψ_u than ψ_f . For example, under ψ_u , the joint payoff is maximized under complete pooling of the states if the payoffs are quadratic and the receiver is highly responsive (see Section 4.1). Such pooling, however, is precluded under ψ_f (and, thus, a fully informed receiver).

In the non-myopic case, these effects are preserved, and are further amplified by the shadow of the future. Moving from ψ_f to ψ_u increases the net discounted surplus $L(\bar{v})$ (which increases with $\bar{v}$ and decreases with v_S and v_R). This in turn relaxes constraints on decision-making and increases the best joint payoff.

A subtle issue arises in the non-myopic case. An implementable decision outcome $\rho(\mu(\theta))$ under an informative signal ψ must be nondecreasing in θ on each signal realization, but it is not required to be nondecreasing across signal realizations. Thus, there are non-monotone decision outcomes implementable under ψ but not under a less informative signal $\hat{\psi}$. Nevertheless, to show that the best joint payoff is higher under $\hat{\psi}$ than under ψ , it is sufficient for the second-best and penal decision outcomes under ψ to be monotone, and, thus, implementable under $\hat{\psi}$. By Proposition 5, the second-best and receiver's penal decision outcomes are always monotone, but the sender's penal decision outcome can be non-monotone. Hence, in Proposition 6, we require $\underline{\rho}_S(\underline{\mu}_S(\theta))$ to be nondecreasing in θ under ψ . This requirement is satisfied, for example, in the uniform-quadratic example of Section 5, where $\underline{\rho}_S$ is the lowest incentive compatible decision rule.²⁸

The result that public information hurts the relationship relates to various papers that study the social value of public information. [Hirshleifer \(1971\)](#) argues that welfare may be decreasing in the amount of public information available to agents. [Bergemann and Morris \(2016\)](#) clarify this point: making more information available to an agent may, by increasing the set of incentive constraints she faces, shrink the set of equilibrium outcomes.²⁹ This relates to the logic of our model, where the availability of public information makes it impossible to pool incentive constraints across states and, thus, worsens incentive provision within the relationship. Public information in our model also improves the worst possible equilibrium payoffs for both players; this decreases the surplus and, thus, worsens intertemporal incentives.³⁰

²⁸Moreover, this requirement holds if $\delta = 0$ or if $\psi(\theta) = \theta$ for all $\theta \in [0, 1]$ and $\rho_R(\theta)$ does not cross $\rho_{FB}(\theta)$ from above. In general, however, $\underline{\rho}_S(\underline{\mu}_S(\theta))$ can be non-monotone. For example, suppose that the players' payoffs are quadratic and that $\rho_S(\theta)$ crosses $\rho_R(\theta)$ from below at $\theta_0 \in (0, 1)$. Then $\underline{\rho}_S(\underline{\mu}_S(\theta))$ is not monotone under fully informative ψ_f . Indeed, for states $\theta < \theta_0$, the sender is downward biased, so $\underline{\rho}_S(\underline{\mu}_S(\theta)) = \rho_R(\theta) + \ell$, but for states $\theta > \theta_0$, the sender is upward biased, so $\underline{\rho}_S(\underline{\mu}_S(\theta)) = \rho_R(\theta) - \ell$. Thus, $\underline{\rho}_S(\underline{\mu}_S(\theta))$ jumps down at θ_0 .

²⁹[Cr  mer \(1995\)](#), [Kloosterman \(2015\)](#), and [Fong and Li \(2016\)](#) discuss other settings where public information may be detrimental.

³⁰This point relates to an insight from [Baker et al. \(1994\)](#). There, objective performance measures, rather than transparency, improve the players' outside options and make cooperation within the relationship more difficult to sustain.

7. CONCLUSION

In our model, incomplete information transmission does not reflect a failure to motivate communication, but instead is an instrument for managing decision-making. This finding relies on the capacity of voluntary transfers to credibly support any monotone message rule at no welfare cost. It suggests that when modeling strategic communication in applied settings, it is crucial to understand whether monetary or nonmonetary transfers (such as wages or favors) are available, because our implications differ significantly from those of the standard literature on strategic communication without transfers. In fact, one interpretation of our model is that voluntary transfers endogenously endow the privately informed sender with the ability to commit to any monotone message rule, even with impatient players. Such commitment is the premise of the literature on Bayesian persuasion (Kamenica and Gentzkow (2011)). So, our analysis extends the applicability of the Bayesian persuasion framework to settings without commitment but with transfers.

Our model is remarkably tractable and, thus, allows for a thorough treatment of repeated interactions. This analysis produces a rich and intuitive set of results. In particular, incomplete information transmission is implemented only for high-conflict states and only if the receiver's decision-making is too responsive to information. One implication is that with constant bias, pooling does not occur. In contrast, in the standard constant-bias cheap-talk game (Crawford and Sobel (1982)), information transmission is always incomplete, and this is generally exacerbated in high (low) states if the sender is upward (downward) biased.

Our result that adding public information worsens optimal equilibria highlights the benefits of an "arms length" approach where information and control are separated. For instance, giving the decision right to the sender may worsen the relationship by increasing the players' worst equilibrium payoffs and, thus, reducing the severity of punishments. Similarly, introducing mediators who control the flow of information from the sender to the receiver cannot improve the relationship. This is because it is optimal to give the sender as much control over the release of information as possible.

We hope that future work will use our tractable framework to study other challenging problems in strategic communication. For example, one might examine the case with multiple senders and receivers, possibly connected by a communication network. Another promising avenue would be to allow for costly information acquisition.

APPENDIX A: STATIONARITY

This appendix specifies necessary and sufficient conditions for equilibrium, and proves Lemma 1.

To show that the set of equilibrium payoffs is compact, restrict decisions and transfers to compact sets $d \in [-\bar{d}, \bar{d}]$ and $\tau, t, T \in [-\bar{t}, \bar{t}]$. Under this restriction and under Assumption 1, it can be shown that the set of equilibrium payoffs is compact (see, for example, Mailath and Samuelson (2006)). Now, observe that this restriction is without loss of generality if the bounds $\bar{d}$ and $\bar{t}$ are chosen to be large enough that (in any equilibrium) decisions and transfers are interior. Indeed, under Assumption 1, we can show that such bounds exist.

Let V be the set of equilibrium payoffs. We now show that the set $\bar{V}$ of optimal equilibrium payoffs is given by (1). Consider an optimal equilibrium payoff vector (v_S^*, v_R^*) that maximizes the joint payoff, so that $v_S^* + v_R^* = \bar{v}$, and let σ_* be an equilibrium supporting (v_S^*, v_R^*) . Let (v_S, v_R) be any point in $\bar{V}$. Notice that we can modify σ_* to produce (v_S, v_R) by changing only the ex ante transfer in the first period from τ^* to $\tau = \tau^* + (v_S^* - v_S)/(1 - \delta)$. This modification affects the players' incentives only at the ex ante round of the first period. Each player is willing to make the ex ante transfer τ because $v_S \geq \underline{v}_S$ and $v_R \geq \underline{v}_R$ by definition of $\bar{V}$. Thus, the modified strategy profile is an equilibrium. Moreover, it is an optimal equilibrium, because $\bar{v}$ is the maximum equilibrium joint payoff. Finally, if $v_S < \underline{v}_S$ or $v_R < \underline{v}_R$, then (v_S, v_R) cannot be supported in equilibrium. We conclude that $\bar{V}$ is the set of optimal equilibrium payoffs.

A message rule $\mu(\theta)$, a decision rule $\rho(m)$, transfer rules τ , $t(m)$, $T(m)$, continuation payoff function $v_i(m)$ for each $i \in \{S, R\}$, and punishment decision d^P and message m^P constitute an equilibrium if and only if the following six conditions hold (see, for example, [Mailath and Samuelson \(2006\)](#)):

Condition C1. Each player is willing to make ex ante transfer τ :

$$\begin{aligned} v_S &= (1 - \delta)[- \tau + \mathbb{E}[u_S(\rho(\mu(\theta)), \theta) - t(\mu(\theta)) - T(\mu(\theta))]] + \delta \mathbb{E}[v_S(\mu(\theta))] \geq \underline{v}_S \\ v_R &= (1 - \delta)[\tau + \mathbb{E}[u_R(\rho(\mu(\theta)), \theta) + t(\mu(\theta)) + T(\mu(\theta))]] + \delta \mathbb{E}[v_R(\mu(\theta))] \geq \underline{v}_R. \end{aligned}$$

Condition C2. For each state θ , the sender is willing to send message $\mu(\theta)$ and to make interim transfer $t(\mu(\theta))$.

- (a) There is no profitable deviation to another message–interim-transfer pair $(\mu(\hat{\theta}), t(\mu(\hat{\theta})))$ that is observed on the equilibrium path:

$$\begin{aligned} &(1 - \delta)[u_S(\rho(\mu(\theta)), \theta) - t(\mu(\theta)) - T(\mu(\theta))] + \delta v_S(\mu(\theta)) \\ &\geq (1 - \delta)[u_S(\rho(\mu(\hat{\theta})), \theta) - t(\mu(\hat{\theta})) - T(\mu(\hat{\theta}))] \\ &\quad + \delta v_S(\mu(\hat{\theta})) \quad \text{for all } \theta, \hat{\theta} \in [0, 1]. \end{aligned}$$

- (b) There is no profitable deviation to some pair $(\hat{m}, \hat{t})$ that is never observed on the equilibrium path:

$$\begin{aligned} &(1 - \delta)[u_S(\rho(\mu(\theta)), \theta) - t(\mu(\theta)) - T(\mu(\theta))] + \delta v_S(\mu(\theta)) \\ &\geq (1 - \delta)u_S(d^P, \theta) + \delta \underline{v}_S \quad \text{for all } \theta \in [0, 1]. \end{aligned}$$

Here, we specify that following any such deviation, the receiver chooses punishment decision d^P .

Condition C3. The receiver is willing to make interim transfer $t(m)$:

$$\begin{aligned} &(1 - \delta)[u_R(\rho(m), m) + t(m) + T(m)] + \delta v_R(m) \\ &\geq (1 - \delta)u_R(\rho(m), m) + \delta \underline{v}_R \quad \text{for all } m \in \mu([0, 1]). \end{aligned}$$

Condition C4. The receiver is willing to choose decision $\rho(m)$ on path and d^P off path.

- (a) After an on-path message–interim-transfer pair, the receiver is willing to choose decision $\rho(m)$:

$$\begin{aligned} (1 - \delta)[u_R(\rho(m), m) + T(m)] + \delta v_R(m) \\ \geq (1 - \delta)u_R(\rho_R(m), m) + \delta \underline{v}_R \quad \text{for all } m \in \mu([0, 1]). \end{aligned}$$

- (b) After an off-path message–interim-transfer pair, the receiver is willing to choose decision d^P :

$$(1 - \delta)u_R(d^P, m^P) + \delta(\bar{v} - \underline{v}_S) \geq (1 - \delta)u_R(\rho_R(m^P), m^P) + \delta \underline{v}_R.$$

Here, we specify that following any deviation by the sender, the receiver believes that the state is in $m^P \subset [0, 1]$.

Condition C5. Each player is willing to make ex post transfer $T(m)$:

$$\begin{aligned} -(1 - \delta)T(m) + \delta v_S(m) &\geq \delta \underline{v}_S \quad \text{for all } m \in \mu([0, 1]), \\ (1 - \delta)T(m) + \delta v_R(m) &\geq \delta \underline{v}_R \quad \text{for all } m \in \mu([0, 1]). \end{aligned}$$

Condition C6. The continuation payoffs can be supported in equilibrium:

$$(v_S(m), v_R(m)) \in V \quad \text{for all } m \in \mu([0, 1]).$$

PROOF OF LEMMA 1. We have already shown that the set of optimal equilibrium payoffs is $\bar{V}$ given by (1). In any optimal equilibrium, continuation is optimal: $v_S(m) + v_R(m) = \bar{v}$ for all $m \in \mu([0, 1])$. Otherwise, one could increase $v_S(m)$ and $v_R(m)$ without violating Condition C6, thereby relaxing the constraints of Conditions C1–C5 and increasing joint payoff $v_S + v_R$.

An optimal equilibrium σ with zero first-period ex ante transfers clearly exists. Let $(v_S, v_R) \in \bar{V}$ be the payoff vector under σ . We modify σ to construct an optimal stationary equilibrium with the same payoff vector. Let $\mu(\theta)$, $\rho(m)$, $t(m)$, $T(m)$, and $v_i(m)$ for each $i \in \{S, R\}$ be the message rule, decision rule, transfer rules, and continuation payoff function in the first period on the equilibrium path of σ . Define $T^*(m)$ by

$$-(1 - \delta)T^*(m) + \delta v_S = -(1 - \delta)T(m) + \delta v_S(m).$$

Since $v_S + v_R = v_S(m) + v_R(m) = \bar{v}$ by optimality of σ , we also have

$$(1 - \delta)T^*(m) + \delta v_R = (1 - \delta)T(m) + \delta v_R(m).$$

Consider the following stationary strategy profile σ_* . On the equilibrium path, $\mu(\theta)$, $\rho(m)$, $\tau = 0$, $t(m)$, and $T^*(m)$ are played in each period. Following any deviation, expect for an undetectable deviation by the sender as in Condition C2(a), play proceeds

according to σ . By construction, the sender's and receiver's expected payoffs under σ_* are the same as under σ .

We now show that σ_* constitutes an equilibrium. In each period, the constraints of Conditions C1–C5 continue to hold under σ_* because they are identical to the first-period constraints under σ . Condition C6 holds because (v_S, v_R) belongs to $\bar{V}$ by supposition.

Finally, by modifying the first-period ex ante transfer in σ_* from 0 to $\tau = (v_S - \hat{v}_S)/(1 - \delta)$, we can support any optimal equilibrium payoff vector $(\hat{v}_S, \hat{v}_R) \in \bar{V}$. $\square$

LEMMA 2. *If $0 \leq \delta < \delta' < 1$, then the corresponding best joint payoffs satisfy $\bar{v} \leq \bar{v}'$, and the corresponding worst equilibrium payoffs satisfy $\underline{v}_S \geq \underline{v}'_S$ and $\underline{v}_R \geq \underline{v}'_R$.*

PROOF. Given $\delta \in [0, 1)$, consider a stationary optimal equilibrium σ_* with zero ex ante transfers. Let this equilibrium produce an equilibrium payoff vector (v_S^*, v_R^*) , with $v_S^* + v_R^* = \bar{v}$. We can support any optimal equilibrium payoff vector $(v_S, v_R) \in \bar{V}$ under δ by modifying the first-period ex ante transfer in σ_* from 0 to $\tau = (v_S^* - v_S)/(1 - \delta)$. Notice that Conditions C1–C6 continue to hold under $\delta' \in (\delta, 1)$, after replacing $\tau = (v_S^* - v_S)/(1 - \delta)$ with $\tau' = (v_S^* - v_S)/(1 - \delta')$, because

$$\frac{\delta'}{1 - \delta'}(v_i^* - \underline{v}_i) \geq \frac{\delta}{1 - \delta}(v_i^* - \underline{v}_i) \quad \text{for each } i \in \{S, R\}.$$

Therefore, the set $\bar{V}$ is self-generating under δ' , which proves the lemma (see, for example, Mailath and Samuelson (2006)). $\square$

APPENDIX B: EQUILIBRIUM

PROOF OF PROPOSITION 1. Consider a stationary equilibrium σ that produces a joint payoff v . Let $\mu(\theta)$, $\rho(m)$, τ , $t(m)$, and $T(m)$ be the message rule, decision rule, and transfer rules on the equilibrium path of σ . Define $U_S(\theta)$ as the one-period payoff of the sender if the state is θ ,

$$U_S(\theta) = u_S(\rho(\mu(\theta)), \theta) - p(\theta),$$

where $p(\theta)$ is the net one-period transfer of the sender if the state is θ ,

$$p(\theta) = \tau + t(\mu(\theta)) + T(\mu(\theta)).$$

Condition C2(a) requires that

$$U_S(\theta) \geq u_S(\rho(\mu(\hat{\theta})), \theta) - p(\hat{\theta}) \quad \text{for all } \theta, \hat{\theta} \in [0, 1].$$

Since $\partial^2 u_S(d, \theta)/\partial d \partial \theta > 0$ by Assumption 1, this inequality holds if and only if $\rho(\mu(\theta))$ is nondecreasing in θ and

$$U_S(\theta) = U_S(0) + \int_0^\theta \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta} \quad \text{for all } \theta \in [0, 1], \quad (14)$$

by Proposition 1 of Rochet (1987) and Corollary 1 of Milgrom and Segal (2002).

Adding the constraint of Condition C4(a) and the sender's constraint of Condition C5, and taking into account that $v_S + v_R = v$, gives (3).

Conversely, suppose that $\mu(\theta)$ and $\rho(m)$ are such that $\rho(\mu(\theta))$ is nondecreasing in θ and (3) holds. We construct transfer rules and punishment variables that satisfy Conditions C1–C6 and, thus, constitute a stationary equilibrium. We consider the case $\delta > 0$; the case $\delta = 0$ is simpler but slightly different. Let $T(m) = 0$ and $\tau = \mathbb{E}[u_S(\rho(\mu(\theta)), \theta) - t(\mu(\theta))] - \underline{v}_S$. Moreover, let $t(m)$ and m^P be defined by (4) and (5), and let $d^P = \rho(m^P)$.

Equation (5) assumes that $t(m^P) = \inf_{m \in \mu([0, 1])} t(m)$ for some message $m^P \in \mu([0, 1])$. If this assumption does not hold, then we specify m^P and d^P as follows. By the Bolzano–Weierstrass theorem, there exists a sequence $\{m_k\} \in \mu([0, 1])$ such that as $k \rightarrow \infty$, $t(m_k) \rightarrow \inf_{m \in \mu([0, 1])} t(m)$, $\theta(m_k) \rightarrow \theta_*$, and $\rho(m_k) \rightarrow d_*$ for some $\theta(m_k) \in m_k$, $\theta_* \in [0, 1]$, and $d_* \in \mathbb{R}$. Set $m^P = \theta_*$ and $d^P = d_*$. Since $u_R(d, \theta)$ is continuous and (3) holds for all $(m_k, \rho(m_k))$, it also holds for (m^P, d^P) .

Notice that the left-hand side of (3) is nonnegative, so $v \geq \underline{v}_S + \underline{v}_R$. Condition C6 holds because the continuation payoffs are $v_S = \underline{v}_S$ and $v_R = v - \underline{v}_S$. Condition C5 holds because Condition C6 holds and $T(m) = 0$. The sender's constraint of Condition C1 holds with equality. The receiver's constraint of Condition C1 holds because it can be simplified to $v \geq \underline{v}_S + \underline{v}_R$. Condition C2(a) holds because $\rho(\mu(\theta))$ is nondecreasing in θ and (14) holds. Condition C2(b) holds because by deviating to a message-transfer pair $(\hat{m}, \hat{t})$ that is not observed on the equilibrium path, the sender induces $d^P = \rho(m^P)$, which he can induce more cheaply on the equilibrium path with message m^P and zero interim transfer $t(m^P) = 0$. This argument assumes that there exists m^P such that $t(m^P) = \inf_{m \in \mu([0, 1])} t_S(m)$. Condition C2(b) still holds even if such m^P does not exist. This is because Condition C2(a) holds for each $\hat{\theta} = \theta(m_k)$ and, thus, in the limit $k \rightarrow \infty$. But in this limit, Condition C2(a) coincides with Condition C2(b). Condition C4 is a restatement of (3). Note that, as for Condition C2(b), a limiting argument needs to be made for Condition C4(b) if $\inf_{m \in \mu([0, 1])} t(m)$ is not attained by any m^P . Condition C3 holds because Condition C4 holds and $t(m)$ is nonnegative. $\square$

PROOF OF PROPOSITION 2. By Lemma 1 and Proposition 1, in an optimal equilibrium, the decision and message rules solve

$$\bar{v} = \max_{\mu, \rho} \mathbb{E}[u(\rho(\mu(\theta)), \theta)] \quad (15)$$

$$\text{subject to } \rho(\mu(\theta)) \text{ is nondecreasing in } \theta, \quad (16)$$

$$w(\rho(m), m) \leq L(\bar{v}) \text{ for all } m \in \mu([0, 1]). \quad (17)$$

Without loss of generality, we can restrict attention to monotone message rules. The argument is similar to the revelation principle. To this end, consider any μ and ρ that satisfy (16) and (17). Define new rules $\tilde{\mu}$ and $\tilde{\rho}$ as $\tilde{\mu}(\tilde{\theta}) = \{\theta : \rho(\mu(\theta)) = \rho(\mu(\tilde{\theta}))\}$ for all $\tilde{\theta} \in [0, 1]$ and $\tilde{\rho}(\tilde{m}) = \rho(\mu(\tilde{\theta}(\tilde{m})))$ for all $\tilde{m} \in \tilde{\mu}([0, 1])$, where $\tilde{\theta}(\tilde{m})$ is an arbitrary state $\tilde{\theta} \in \tilde{m}$. It is easy to see that $\tilde{\rho}(\tilde{m})$ is independent of the choice of a representative state $\tilde{\theta} \in \tilde{m}$ and that $\tilde{\rho}(\tilde{\mu}(\theta)) = \rho(\mu(\theta))$ for all $\theta \in [0, 1]$. Since $\rho(\mu(\theta))$ is nondecreasing in θ

by (16), $\tilde{\rho}(\tilde{\mu}(\theta))$ is also nondecreasing in θ and $\tilde{\mu}$ is monotone. Moreover, since each set $\tilde{m} \in \tilde{\mu}([0, 1])$ is the union of some disjoint sets $m \in \mu([0, 1])$ and the constraint (17) holds for $\rho(m)$ for each $m \in \mu([0, 1])$, the constraint (17) also holds for $\tilde{\rho}(\tilde{m})$ for each $\tilde{m} \in \tilde{\mu}([0, 1])$.

Consider a relaxed problem

$$\begin{aligned} \bar{v} &= \max_{\mu, \rho} \mathbb{E}[u(\rho(\mu(\theta)), \theta)] \\ \text{subject to } \mu &\text{ is monotone,} \\ w(\rho(m), m) &\leq L(\bar{v}) \text{ for all } m \in \mu([0, 1]). \end{aligned}$$

We can solve this relaxed problem in two steps. First, for a given monotone message rule μ , the optimal decision rule specifies that the second-best decision is taken for each on-path message $m \in \mu([0, 1])$. Second, given this optimal decision rule, the optimal message rule is clearly μ_* defined by (8). To prove that the solution to the relaxed problem also solves the original problem (15), it suffices to show that $\rho_*(m)$ is nondecreasing in m .

We first rewrite the constraint of the problem (6) as $d \in D(m)$, where $D(m)$ is nondecreasing in m in the strong set order. Since $u_R(d, \theta)$ is strictly concave in d and has a unique maximum, $w(d, m)$ is strictly convex in d and has a unique minimum. Taking into account that $w(\rho_R(m), m) = 0$ and $L(\bar{v}) \geq 0$, we have that the set of decisions d that satisfy the constraint of the problem (6) is a nonempty closed convex set and, thus, can be written as $D(m) = [\rho_-(m), \rho_+(m)]$, where $\rho_-(m)$ and $\rho_+(m)$ satisfy the constraint with equality. Moreover, since $u_R(d, \theta)$ is concave in d and is supermodular, $w(d, m)$ is nonincreasing in d and nondecreasing in m if $d < \rho_R(m)$, and $w(d, m)$ is nondecreasing in d and nonincreasing in m if $d > \rho_R(m)$. This implies that $\rho_-(m)$ and $\rho_+(m)$ are nondecreasing in m and, thus, $D(m)$ is nondecreasing in m . Taking into account that $u(d, m)$ is strictly concave and has increasing differences, $\rho_*(m) = \arg \max_{d \in D(m)} u(d, m)$ is nondecreasing in m , as follows, for example, from Theorem 4 of [Milgrom and Shannon \(1994\)](#). $\square$

PROOF OF PROPOSITION 3. Consider an interval $m = (\theta_1, \theta_2)$. Let all expectations condition on m , that is, $\mathbb{E}[x] = \mathbb{E}[x|m]$. Let the expected state given m be $\hat{\theta} = \mathbb{E}[\theta]$. Denote $\hat{\rho}_* = \rho_*(\hat{\theta})$, $\hat{\rho}_R = \rho_R(\hat{\theta})$, and $(\Delta\theta)^2 = \mathbb{E}[(\theta - \hat{\theta})^2]$. In our calculations, we fix $\hat{\theta}$ and consider small m : $\theta_1 \uparrow \hat{\theta}$ and $\theta_2 \downarrow \hat{\theta}$.

To start, we approximate $\mathbb{E}[\rho_*(\theta) - \hat{\rho}_*]$. Suppose $\hat{\theta}$ is high conflict, so the receiver's incentive constraint for singleton message $\hat{\theta}$ is strictly binding (and, thus, is also strictly binding for nearby states and intervals),

$$\begin{aligned} u_R(\rho_R(\theta), \theta) - u_R(\rho_*(\theta), \theta) &= L(\bar{v}) \quad \text{for } \theta \in m, \\ \mathbb{E}[u_R(\rho_R(m), \theta) - u_R(\rho_*(m), \theta)] &= L(\bar{v}). \end{aligned} \tag{18}$$

Differentiating (18) and noting that $u_{R_d}'(\rho_R(\theta), \theta) = 0$, $u_{R_{d\theta}}''(d, \theta) > 0$, and $\rho_*(\theta) \neq \rho_R(\theta)$, we obtain

$$u_{R_d}'(\rho_R(\theta), \theta)\rho_R'(\theta) + u_{R_\theta}'(\rho_R(\theta), \theta) - u_{R_d}'(\rho_*(\theta), \theta)\rho_*'(\theta) - u_{R_\theta}'(\rho_*(\theta), \theta) = 0, \text{ so}$$

$$\rho'_*(\theta) = \frac{u_{R\theta}'(\rho_R(\theta), \theta) - u_{R\theta}'(\rho_*(\theta), \theta)}{u_{Rd}'(\rho_*(\theta), \theta)} \in (0, \infty).$$

We can then write

$$\begin{aligned} \mathbb{E}[\rho_*(\theta) - \hat{\rho}_*] &= \mathbb{E}\left[\rho'_*(\hat{\theta})(\theta - \hat{\theta}) + \frac{1}{2}\rho''_*(\hat{\theta})(\theta - \hat{\theta})^2 + o(\theta - \hat{\theta})^2\right] \\ &= \frac{1}{2}\rho''_*(\hat{\theta})(\Delta\theta)^2 + o(\Delta\theta)^2. \end{aligned}$$

Next, we approximate $\mathbb{E}[u(\rho_*(m), \theta) - u(\hat{\rho}_*, \theta)]$. Starting with

$$u_R(\hat{\rho}_R, \hat{\theta}) - u_R(\hat{\rho}_*, \hat{\theta}) = L(\bar{v}) = \mathbb{E}[u_R(\rho_R(m), \theta) - u_R(\rho_*(m), \theta)],$$

we rearrange to obtain

$$\mathbb{E}[u_R(\rho_R(m), \theta) - u_R(\hat{\rho}_R, \hat{\theta})] = \mathbb{E}[u_R(\rho_*(m), \theta) - u_R(\hat{\rho}_*, \hat{\theta})].$$

The left-hand side expands to become

$$\begin{aligned} &\mathbb{E}[u_R(\rho_R(m), \theta) - u_R(\hat{\rho}_R, \hat{\theta})] \\ &= \mathbb{E}\left[\begin{aligned} &u_{Rd}'(\hat{\rho}_R, \hat{\theta})(\rho_R(m) - \hat{\rho}_R) + u_{R\theta}'(\hat{\rho}_R, \hat{\theta})(\theta - \hat{\theta}) \\ &+ u_{Rd\theta}''(\hat{\rho}_R, \hat{\theta})(\rho_R(m) - \hat{\rho}_R)(\theta - \hat{\theta}) \\ &+ \frac{1}{2}u_{Rdd}''(\hat{\rho}_R, \hat{\theta})(\rho_R(m) - \hat{\rho}_R)^2 + \frac{1}{2}u_{R\theta\theta}''(\hat{\rho}_R, \hat{\theta})(\theta - \hat{\theta})^2 \\ &+ o((\rho_R(m) - \hat{\rho}_R)^2 + (\theta - \hat{\theta})^2) \end{aligned}\right] \\ &= \frac{1}{2}u_{Rdd}''(\hat{\rho}_R, \hat{\theta})(\rho_R(m) - \hat{\rho}_R)^2 + \frac{1}{2}u_{R\theta\theta}''(\hat{\rho}_R, \hat{\theta})(\Delta\theta)^2 + o((\rho_R(m) - \hat{\rho}_R)^2 + (\Delta\theta)^2) \\ &= \frac{1}{2}u_{R\theta\theta}''(\hat{\rho}_R, \hat{\theta})(\Delta\theta)^2 + o(\Delta\theta)^2, \end{aligned}$$

while the right-hand side expands to become

$$\mathbb{E}[u_R(\rho_*(m), \theta) - u_R(\hat{\rho}_*, \hat{\theta})] = \mathbb{E}\left[\begin{aligned} &u_{Rd}'(\hat{\rho}_*, \hat{\theta})(\rho_*(m) - \hat{\rho}_*) + \frac{1}{2}u_{R\theta\theta}''(\hat{\rho}_*, \hat{\theta})(\theta - \hat{\theta})^2 \\ &+ o(\theta - \hat{\theta})^2 + o(\rho_*(m) - \hat{\rho}_*) \end{aligned}\right].$$

Equating both sides then yields

$$\rho_*(m) - \hat{\rho}_* = \frac{u_{R\theta\theta}''(\hat{\rho}_R, \hat{\theta}) - u_{R\theta\theta}''(\hat{\rho}_*, \hat{\theta})}{2u_{Rd}'(\hat{\rho}_*, \hat{\theta})}(\Delta\theta)^2 + o(\Delta\theta)^2.$$

Substituting in this last expression yields

$$\begin{aligned} \mathbb{E}[u(\rho_*(m), \theta) - u(\hat{\rho}_*, \theta)] &= \mathbb{E}[u'_d(\hat{\rho}_*, \theta)(\rho_*(m) - \hat{\rho}_*) + o(\rho_*(m) - \hat{\rho}_*)] \\ &= u'_d(\hat{\rho}_*, \theta) \frac{u_{R\theta\theta}''(\hat{\rho}_R, \hat{\theta}) - u_{R\theta\theta}''(\hat{\rho}_*, \hat{\theta})}{2u_{Rd}'(\hat{\rho}_*, \hat{\theta})}(\Delta\theta)^2 + o(\Delta\theta)^2. \end{aligned}$$

It is better to pool than to separate all states in interval m if and only if

$$\begin{aligned}
 0 &> \mathbb{E}[u(\rho_*(\theta), \theta) - u(\rho_*(m), \theta)] \\
 &= \mathbb{E} \left[u''_{d\theta}(\hat{\rho}_*, \hat{\theta})(\theta - \hat{\theta})(\rho_*(\theta) - \hat{\rho}_*) + \frac{1}{2} u''_{dd}(\hat{\rho}_*, \hat{\theta})(\rho_*(\theta) - \hat{\rho}_*)^2 \right. \\
 &\quad \left. + u'_d(\hat{\rho}_*, \hat{\theta})(\rho_*(\theta) - \hat{\rho}_*) + u(\hat{\rho}_*, \hat{\theta}) - u(\rho_*(m), \hat{\theta}) + o(\Delta\theta)^2 \right] \\
 &= \left\{ u''_{d\theta}(\hat{\rho}_*, \hat{\theta})\rho'_*(\hat{\theta}) + \frac{1}{2} u''_{dd}(\hat{\rho}_*, \hat{\theta})(\rho'_*(\hat{\theta}))^2 \right. \\
 &\quad \left. + \frac{1}{2} u'_d(\hat{\rho}_*, \hat{\theta}) \left(\rho''_*(\hat{\theta}) + \frac{u_{R''\theta\theta}(\hat{\rho}_*, \hat{\theta}) - u_{R''\theta\theta}(\hat{\rho}_R, \hat{\theta})}{u_{R'd}(\hat{\rho}_*, \hat{\theta})} \right) \right\} (\Delta\theta)^2 + o(\Delta\theta)^2,
 \end{aligned}$$

so full disclosure can be improved in some neighbourhood of $\hat{\theta}$ if

$$\begin{aligned}
 0 &> u''_{d\theta}(\hat{\rho}_*, \hat{\theta})\rho'_*(\hat{\theta}) + \frac{1}{2} u''_{dd}(\hat{\rho}_*, \hat{\theta})(\rho'_*(\hat{\theta}))^2 \\
 &\quad + \frac{1}{2} u'_d(\hat{\rho}_*, \hat{\theta}) \left(\rho''_*(\hat{\theta}) + \frac{u_{R''\theta\theta}(\hat{\rho}_*, \hat{\theta}) - u_{R''\theta\theta}(\hat{\rho}_R, \hat{\theta})}{u_{R'd}(\hat{\rho}_*, \hat{\theta})} \right) \\
 &= \frac{1}{2} \rho'_*(\hat{\theta}) u''_{dd}(\rho_*(\hat{\theta}), \hat{\theta})(\rho'_*(\hat{\theta}) - 2\tilde{\rho}'_{FB}(\hat{\theta})) + u'_d(\rho_*(\hat{\theta}), \hat{\theta}) \left(\frac{\rho''_*(\hat{\theta})}{2} - \frac{\rho'_*(\hat{\theta})\varphi''(\hat{\theta})}{2\varphi'(\hat{\theta})} \right),
 \end{aligned}$$

where the equality can be verified by substituting $\tilde{\rho}'_{FB}(\theta)$, $\rho'_*(\hat{\theta})$, $\varphi'(\hat{\theta})$, and $\varphi''(\hat{\theta})$. Dividing this expression by $\rho'_*(\hat{\theta}) > 0$, we obtain (9).

Two final observations: first, if $\rho_*(\hat{\theta}) = \rho_{FB}(\hat{\theta})$, then $\rho'_*(\hat{\theta}) = \rho'_{FB}(\hat{\theta}) = \tilde{\rho}'_{FB}(\hat{\theta}) > 0$ and $u'_d(\rho_*(\hat{\theta}), \hat{\theta}) = 0$, in which case (9) must fail; second, in footnote 18, we remark that $\tilde{\rho}'_{FB}(\theta)$ is the slope of the locally first-best decision rule. To justify this interpretation, we approximate $\mathbb{E}[u(\hat{\rho}_* + a(\theta - \hat{\theta}), \theta)]$ as

$$\begin{aligned}
 &\mathbb{E}[u(\hat{\rho}_* + a(\theta - \hat{\theta}), \theta)] \\
 &= u(\hat{\rho}_*, \hat{\theta}) + \left(\frac{1}{2} u''_{dd}(\hat{\rho}_*, \hat{\theta})a^2 + u''_{d\theta}(\hat{\rho}_*, \hat{\theta})a + \frac{1}{2} u''_{\theta\theta}(\hat{\rho}_*, \hat{\theta}) \right) (\Delta\theta)^2 + o(\Delta\theta)^2.
 \end{aligned}$$

This expression is maximized at $a = \tilde{\rho}'_{FB}(\theta)$ when $(\Delta\theta)^2 \rightarrow 0$. □

APPENDIX C: QUADRATIC PAYOFFS

LEMMA 3. *Under Assumption 2,*

- (i) $u_*(\theta)$ is continuously differentiable in θ for all $\theta \in [0, 1]$
- (ii) $u_*(\theta)$ is twice continuously differentiable in θ for almost all $\theta \in [0, 1]$, with

$$u''_*(\theta) = (2\rho'_{FB}(\theta) - \rho'_*(\theta))\rho'_*(\theta), \quad \text{if } |\rho_{FB}(\theta) - \rho_R(\theta)| \neq \ell. \quad (19)$$

PROOF. Since

$$\rho_*(\theta) = \begin{cases} \rho_R(\theta) + \ell & \text{if } \rho_{FB}(\theta) - \rho_R(\theta) > \ell, \\ \rho_{FB}(\theta) & \text{if } |\rho_{FB}(\theta) - \rho_R(\theta)| \leq \ell, \\ \rho_R(\theta) - \ell & \text{if } \rho_{FB}(\theta) - \rho_R(\theta) < -\ell, \end{cases}$$

we have

$$\rho'_*(\theta) = \begin{cases} \rho'_R(\theta) & \text{if } |\rho_{FB}(\theta) - \rho_R(\theta)| > \ell, \\ \rho'_{FB}(\theta) & \text{if } |\rho_{FB}(\theta) - \rho_R(\theta)| < \ell. \end{cases} \quad (20)$$

Further, $u_*(\theta)$ is continuously differentiable in θ for all $\theta \in [0, 1]$, with

$$u'_*(\theta) = \begin{cases} \rho'_{FB}(\theta)\rho_*(\theta) + (\rho_{FB}(\theta) - \rho_*(\theta))\rho'_*(\theta) & \text{if } |\rho_{FB}(\theta) - \rho_R(\theta)| \neq \ell, \\ \rho'_{FB}(\theta)\rho_*(\theta) & \text{if } |\rho_{FB}(\theta) - \rho_R(\theta)| = \ell. \end{cases}$$

Finally, since $\rho_*(\theta)$ is twice continuously differentiable in θ everywhere except at most two states where $|\rho_{FB}(\theta) - \rho_R(\theta)| = \ell$, it follows that $u_*(\theta)$ is twice continuously differentiable everywhere except at most these two states, with $u''_*(\theta)$ given by (19). $\square$

LEMMA 4. *If the receiver is highly responsive, ℓ is strictly increasing in δ .*

PROOF. If $\delta = 0$, then $u''_*(\theta) < 0$ for all $\theta \in [0, 1]$. By Corollary 3, the expected joint payoff is strictly higher under complete pooling than under full separation. Thus, Lemma 2 implies that $\bar{v} - \underline{v}_S - \underline{v}_R > 0$ and ℓ is strictly increasing in δ for all $\delta \in [0, 1)$. $\square$

PROOF OF COROLLARY 4. Since, by Lemma 4, the relational leeway ℓ increases with δ , the set of high-conflict states consists of up to one interval that shrinks and eventually vanishes as δ increases (Figure 3). Specifically, there exist $\delta^A, \delta^B \in (0, 1)$ such that $\delta^A > \delta^B$, and the set of high-conflict states is

$$X = \begin{cases} \emptyset & \text{if } \delta \in (\delta^A, 1), \\ [0, \theta_c) \text{ for some } \theta_c \in (0, 1) & \text{if } \delta \in (\delta^B, \delta^A), \\ [0, 1] & \text{if } \delta \in [0, \delta^B). \end{cases}$$

By Proposition 4, full separation is optimal if $\delta \in (\delta^A, 1)$ and complete pooling is optimal if $\delta \in [0, \delta^B)$. Moreover, it is optimal to pool the states below $\theta_* \in (\theta_c, 1]$ and separate the rest if $\delta \in (\delta^B, \delta^A)$. The expected joint payoff under this message rule is

$$v_* = F(\theta_*)u_*(m_*) + \int_{\theta_*}^1 u_*(\theta) dF(\theta).$$

Thus, taking into account that

$$\frac{dm_*}{d\theta_*} = \frac{f(\theta_*)}{F(\theta_*)}(\theta_* - m_*),$$

we have

$$\begin{aligned}\frac{dv_*}{d\theta_*} &= F(\theta_*)u'_*(m_*)\frac{dm_*}{d\theta_*} + f(\theta_*)u_*(m_*) - f(\theta_*)u_*(\theta_*) \\ &= f(\theta_*)(u'_*(m_*)(\theta_* - m_*) + u_*(m_*) - u_*(\theta_*))\end{aligned}\quad (21)$$

and

$$\begin{aligned}\frac{d^2v_*}{d\ell d\theta_*} &= f(\theta_*)\left(\frac{du'_*(m_*)}{d\ell}(\theta_* - m_*) + \frac{du_*(m_*)}{d\ell}\right) \\ &= f(\theta_*)((\rho'_{FB}(m_*) - \rho'_*(m_*))(\theta_* - m_*) + (\rho_{FB}(m_*) - \rho_*(m_*))) \\ &< f(\theta_*)((\rho'_{FB}(m_*) - \rho'_*(m_*))(\theta_c - m_*) + (\rho_{FB}(m_*) - \rho_*(m_*))) = 0,\end{aligned}$$

where the inequality holds because $\theta_* > \theta_c > m_*$ and $\rho'_{FB}(m_*) < \rho'_*(m_*)$, and the last equality holds because $\rho_{FB}(\theta_c) = \rho_*(\theta_c)$ and $\rho_{FB}(\theta) - \rho_*(\theta)$ is linear in θ for $\theta \in (0, \theta_c)$. So θ_* is nonincreasing in ℓ , and $d\theta_*/d\ell < 0$ if $\theta_* < 1$, as follows, for example, from Theorem 1 of [Edlin and Shannon \(1998\)](#). $\square$

COROLLARY 5. *Suppose the receiver is highly responsive and $\theta_c \in (0, 1)$. Keeping θ_c constant, θ_* is strictly increasing in $\alpha = \rho'_R(\theta)/\rho'_{FB}(\theta)$ if $\theta_* < 1$. Moreover, $\theta_* \rightarrow \theta_c$ as $\alpha \rightarrow 2$.*

PROOF. Denote $\rho_{FB}(\theta_c) = d_c$, $\rho'_{FB}(\theta) = a_{FB}$, and $\rho'_R(\theta) = a_R$. Notice that

$$\rho_{FB}(\theta) = d_c + a_{FB}(\theta - \theta_c) \quad \text{for all } \theta$$

and

$$\rho_*(\theta) = \begin{cases} d_c + a_R(\theta - \theta_c) & \text{if } \theta < \theta_c, \\ d_c + a_{FB}(\theta - \theta_c) & \text{if } \theta \geq \theta_c. \end{cases}$$

Thus, we can rewrite (21) as

$$\begin{aligned}\frac{dv_*}{d\theta_*} &= f(\theta_*)\left([a_{FB}\rho_*(m_*) + a_R(\rho_{FB}(m_*) - \rho_*(m_*))](\theta_* - m_*)\right. \\ &\quad \left.+ \rho_{FB}(m_*)\rho_*(m_*) - \frac{\rho_*^2(m_*)}{2} - \frac{\rho_{FB}^2(\theta_*)}{2}\right) \\ &= f(\theta_*)\left(a_R(a_R - 2a_{FB})(\theta_c - m_*)\left(\theta_* - \frac{\theta_c + m_*}{2}\right) - a_{FB}^2 \frac{(\theta_* - \theta_c)^2}{2}\right).\end{aligned}$$

Hence,

$$\frac{d}{d\alpha}\left(\frac{1}{a_{FB}^2}\frac{dv_*}{d\theta_*}\right) = f(\theta_*)(\alpha - 1)(\theta_c - m_*)(2\theta_* - \theta_c - m_*) > 0.$$

So θ_* is nondecreasing in α , and $d\theta_*/d\alpha > 0$ if $\theta_* < 1$, as follows from Theorem 1 of [Edlin and Shannon \(1998\)](#). Finally, if $\alpha \rightarrow 2$, then $dv_*/d\theta_*|_{\theta_*=\theta_c} \rightarrow 0$, implying that $\theta_* \rightarrow \theta_c$. $\square$

APPENDIX D: PUNISHMENT

PROOF OF PROPOSITION 5. By Proposition 2, $\bar{v} = \mathbb{E}[u_*(\rho_*(\mu_*(\theta)), \theta)]$.

Because the receiver can guarantee the payoff $u_R(\rho_R([0, 1]), [0, 1])$ by rejecting all transfers and choosing $\rho_R([0, 1])$ in all periods, we have

$$\underline{v}_R \geq u_R(\rho_R([0, 1]), [0, 1]).$$

A sender's worst equilibrium with zero first-period ex ante transfers ($\tau = 0$) clearly exists. Let $\mu(\theta)$, $t(m)$, $\rho(m)$, $T(m)$, $v_S(m)$, d^P , and m^P be used in the first period of such an equilibrium. Define $V_S(\theta)$ as the expected payoff of the sender if the first-period state is θ :

$$\begin{aligned} V_S(\theta) &= (1 - \delta)u_S(\rho(\mu(\theta)), \theta) - p(\theta) \\ \text{where } p(\theta) &= (1 - \delta)[t(\mu(\theta)) + T(\mu(\theta))] - \delta v_S(\mu(\theta)). \end{aligned}$$

Condition C2(a) requires that

$$V_S(\theta) \geq (1 - \delta)u_S(\rho(\mu(\hat{\theta})), \theta) - p(\hat{\theta}) \quad \text{for all } \theta, \hat{\theta} \in [0, 1]. \quad (22)$$

As explained in the proof of Proposition 1, this inequality holds if and only if

$$\rho(\mu(\theta)) \text{ is nondecreasing in } \theta, \quad (23)$$

$$V_S(\theta) = V_S(0) + (1 - \delta) \int_0^\theta \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta} \quad \text{for all } \theta \in [0, 1]. \quad (24)$$

Condition C2(b), the constraint of Condition C4(a) and the sender's constraint of Condition C5, and the constraint of Condition C4(b), respectively, imply that

$$V_S(\theta) \geq (1 - \delta)u_S(d^P, \theta) + \delta \underline{v}_S \quad \text{for all } \theta \in [0, 1], \quad (25)$$

$$w(\rho(m), m) \leq L(\bar{v}) \quad \text{for all } m \in \mu([0, 1]), \quad (26)$$

$$w(d^P, m^P) \leq L(\bar{v}). \quad (27)$$

Thus, $\underline{v}_S$ is greater than or equal to the value of the problem

$$\begin{aligned} &\min_{\mu, \rho, V_S, m^P, d^P} \mathbb{E}[V_S(\theta)] \\ &\text{subject to (23)–(27)}. \end{aligned} \quad (28)$$

CLAIM 1. *There exists an optimal solution to the problem (28) such that $m^P \in \mu([0, 1])$, $d^P = \rho(m^P)$, and (25) holds with equality for $\theta \in m^P$.*

PROOF. Given ρ and μ that satisfy (23) and (26), define the function

$$h(m) = u_S(\rho(m), \theta(m)) - \int_0^{\theta(m)} \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta}, \quad (29)$$

where $\theta(m) \in m$. Define

$$m_\star \in \arg \min_{m \in \mu([0,1])} h(m) \quad \text{and} \quad \theta_\star \in m_\star.$$

Hereafter, we assume that the infimum of h is attained. If the infimum is not attained by any $m_\star$, a limiting argument, as in the proof of Proposition 1, needs to be made. It is easy to see that $\mu, \rho, m^P = m_\star, \theta^P = \theta_\star, d^P = \rho(m^P)$, and

$$V_S(\theta) = (1 - \delta) \left(u_S(\rho(m^P), \theta^P) + \int_{\theta^P}^{\theta} \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta} \right) + \delta \underline{v}_S \quad (30)$$

constitute a feasible solution to the problem (28). In particular, (30) clearly satisfies (24), and (25) holds because

$$\begin{aligned} V_S(\theta) &= (1 - \delta) (u_S(\rho(\mu(\theta)), \theta) - (h(\mu(\theta)) - h(m^P))) + \delta \underline{v}_S \\ &\geq (1 - \delta) u_S(\rho(m^P), \theta) + \delta \underline{v}_S, \end{aligned}$$

where the equality follows from (29) and (30), and the inequality follows from (22) evaluated at $\hat{\theta} = \theta^P$, where (22) holds because (23) and (24) hold.

Suppose to the contrary that there does not exist an optimal solution to (28) with the stated properties. Thus, in an optimal solution, $d^P \notin \rho(\mu([0, 1]))$ and

$$u_S(\rho(\mu(\theta)), \theta) - (h(\mu(\theta)) - h(m_\star)) > u_S(d^P, \theta) \quad \text{for all } \theta \in [0, 1]. \quad (31)$$

There are two cases to consider: $d^P \in (\rho(\mu(0)), \rho(\mu(1)) \setminus \rho(\mu([0, 1]))$ and $d^P < \rho(\mu(0))$ (the case $d^P > \rho(\mu(1))$ is analogous).

Suppose that $d^P \in (\rho(\mu(0)), \rho(\mu(1)) \setminus \rho(\mu([0, 1]))$. Then there exists $\hat{\theta} \in (0, 1)$ such that $d^P \in (\rho(\mu(\hat{\theta}-)), \rho(\mu(\hat{\theta}+)))$. By continuity of u_S and V_S , we have

$$u_S(\rho(\mu(\hat{\theta}-)), \hat{\theta}) - (h(\mu(\hat{\theta}-)) - h(m_\star)) = u_S(\rho(\mu(\hat{\theta}+)), \hat{\theta}) - (h(\mu(\hat{\theta}+)) - h(m_\star)).$$

Since $u_S(d, \theta)$ is concave in d by Assumption 1 and $h(m)$ is minimized at $m_\star$, this equality is incompatible with (31) evaluated at $\hat{\theta}-$, leading to a contradiction.

Suppose that $d^P < \rho(\mu(0))$. The optimal V_S is such that (25) holds with equality for some θ ,

$$\min_{\theta \in [0, 1]} (V_S(\theta) - (1 - \delta) u_S(d^P, \theta)) = \delta \underline{v}_S,$$

which can be rewritten using (24) as

$$(1 - \delta) \min_{\theta \in [0, 1]} \int_0^{\theta} \left(\frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) - \frac{\partial u_S}{\partial \theta}(d^P, \tilde{\theta}) \right) d\tilde{\theta} = (1 - \delta) u_S(d^P, 0) + \delta \underline{v}_S - V_S(0).$$

Since $\partial^2 u_S(d, \theta) / \partial d \partial \theta > 0$ and $\rho(\mu(\theta)) > d^P$, the minimum is achieved at $\theta = 0$. Moreover, (31) implies that $u_S(d^P, 0) < u_S(\rho(\mu(0)), 0)$. Therefore, $u_S(\rho_-(0), 0) \leq u_S(d^P, 0)$ because $\rho_-(0) \leq d^P$ by (27), $d^P < \rho(\mu(0))$ by supposition, and u_S is concave in d . So an optimal $d^P < \rho(\mu(0))$ must be given by $\rho_-(0)$ to minimize $V_S(0)$ and, thus, V_S . But then we can modify μ and ρ only in that μ separates $\theta = 0$ and $\rho(\mu(0))$ is replaced

with $\rho_-(0)$. Under this modification, we can support the same V_S given by (24) with $V_S(0) = (1 - \delta)u_S(\rho_-(0), 0) + \delta \underline{v}_S$, leading to a contradiction. $\square$

Claim 1, together with (30), implies that $\underline{v}_S$ is greater than or equal to

$$\begin{aligned} \min_{\mu, \rho, \theta^P} & \left\{ u_S(\rho(\mu(\theta^P)), \theta^P) + \mathbb{E} \left[\int_{\theta^P}^{\theta} \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta} \right] \right\} \\ & \text{subject to } \rho(\mu(\theta)) \text{ is nondecreasing in } \theta, \\ & w(\rho(m), m) \leq L(\bar{v}) \text{ for all } m \in \mu([0, 1]). \end{aligned} \quad (32)$$

CLAIM 2. *There exists an optimal solution to the problem (32) that solves the problem (13).*

PROOF. Consider an optimal solution (μ, ρ, θ^P) to (32). Without loss of generality,

$$m^P = \mu(\theta^P) = \{\theta : \rho(\mu(\theta)) = \rho(\mu(\theta^P))\};$$

otherwise, we can modify the message and decision rules such that all states in $\{\theta : \rho(\mu(\theta)) = \rho(\mu(\theta^P))\}$ are pooled, the same decision $\rho(\mu(\theta))$ is induced for all θ , the constraints of (32) hold, and the value of (32) remains the same. Moreover,

$$\rho(\mu(\theta)) = \begin{cases} \rho_-(\mu(\theta)) & \text{if } \mu(\theta) > m^P, \\ \rho_+(\mu(\theta)) & \text{if } \mu(\theta) < m^P; \end{cases}$$

otherwise, we can decrease the value of (32) without violating the constraints either by decreasing $\rho(\mu(\theta))$ for $\mu(\theta) > m^P$ or by increasing $\rho(\mu(\theta))$ for $\mu(\theta) < m^P$.

Suppose to the contrary that there does not exist an optimal solution to (32) with $\rho(m^P) \in \{\rho_-(m^P), \rho_+(m^P)\}$. Consider an optimal solution such that no other optimal solution has a strictly larger m^P in the set order. If $\bar{\theta}^P = \sup m^P < 1$, then some states adjacent to m^P from above, say $(\bar{\theta}^P, \bar{\theta}^P + \varepsilon)$, must be pooled; otherwise, we can decrease the value of (32) by pooling states $(\bar{\theta}^P, \bar{\theta}^P + \varepsilon)$ and m^P , and inducing the same decision $\rho(m^P)$. Similarly, if $\underline{\theta}^P = \inf m^P > 0$, then some states adjacent to m^P from below, say $(\underline{\theta}^P - \varepsilon, \underline{\theta}^P)$, must be pooled. Notice that the objective function in (32) is concave in $\rho(m^P)$, so we can decrease the value of (32) without violating the constraints by changing $\rho(m^P)$ to at least one of the four values $\rho(\mu(\bar{\theta}^P +))$, $\rho(\mu(\underline{\theta}^P -))$, $\rho_+(m^P)$, or $\rho_-(m^P)$, leading to a contradiction. $\square$

It remains to show that a single-period punishment strategy profile from Proposition 5 can be supported in an equilibrium using the ex ante transfers $\underline{t}_0$, $\underline{t}_S$, and $\underline{t}_R$ given by

$$\begin{aligned} \underline{t}_0 &= \underline{t}_S = \mathbb{E}[u_S(\rho_*(\mu_*(\theta)), \theta) - \underline{t}_0(\mu_*(\theta))] - \underline{v}_S, \\ (1 - \delta)[\underline{t}_R &+ \mathbb{E}[u_R(\rho_*(\mu_*(\theta)), \theta) + \underline{t}_0(\mu_*(\theta))]] + \delta(\bar{v} - \underline{v}_S) = \underline{v}_R. \end{aligned}$$

Condition C6 holds because the continuation payoffs are $v_S(m) = \underline{v}_S$ and $v_R(m) = \bar{v} - \underline{v}_S$. Condition C5 holds because Condition C6 holds and $T(m) = 0$. The sender's (re-

ceiver's) constraint of Condition C1 holds with equality for $\tau_0 = \tau_S$ (for τ_R). The receiver's (sender's) constraint of Condition C1 holds for $\tau_0 = \tau_S$ (for τ_R) because it can be simplified to $\bar{v} \geq v_S + v_R$. Condition C2(a) holds because $\underline{\rho}_j(\underline{\mu}_j(\theta))$ is nondecreasing in θ and $\underline{t}_j$ satisfies (4). Condition C2(b) holds because by deviating to a message-transfer pair $(\hat{m}, \hat{t})$ that is not observed on the equilibrium path, the sender induces $\underline{d}_j^P = \underline{\rho}_j(\underline{m}_j^P)$, which he can induce more cheaply on the equilibrium path with message $\underline{m}_j^P$ and zero interim transfer $\underline{t}_j(\underline{m}_j^P) = 0$, as required by (5). Condition C4(a) holds because $w(\underline{\rho}_j(m), m) \leq L(\bar{v})$ for all $m \in \underline{\mu}_j([0, 1])$. Condition C4(b) holds because Condition C4(a) holds and $\underline{d}_j^P = \underline{\rho}_j(\underline{m}_j^P)$. Condition C3 holds because Condition C4 holds and $\underline{t}_j(m)$ is nonpositive. $\square$

PROPOSITION 7. *Suppose Assumption 1 holds and the state is uniformly distributed. Suppose also that $\rho_S(1/2) > \rho_R(1/2)$. Denote $\rho_S(\theta) = a_S\theta + b_S$ and $\rho_R(\theta) = a_R\theta + b_R$. Then there exists an optimal equilibrium in single-period punishment strategies where the sender's penal decision rule is $\underline{\rho}_S(m) = \rho_R(m) - \ell$ for all m and the sender's penal message rule $\underline{\mu}_S$ pools the states below $\underline{\theta}_S$ and separates the states above $\underline{\theta}_S$, where*

$$\underline{\theta}_S = \begin{cases} 0 & \text{if } \ell > \frac{3a_R^2}{32a_S} - b_S + b_R, \\ \frac{a_R + \sqrt{a_R^2 - 8a_S(b_S - b_R + \ell)}}{2a_S} & \text{if } \ell < \frac{3a_R^2}{32a_S} - b_S + b_R. \end{cases}$$

PROOF. Define $\bar{\theta}_* = \sup m^P$ and $\underline{\theta}_* = \inf m^P$, where m^P solves (13). We first show by contradiction in each case that if $\bar{\theta}_* < 1$, then $\underline{\rho}_S(m^P) = \rho_R(m^P) - \ell$ and all states $\theta > \bar{\theta}_*$ are separated by $\underline{\mu}_S$. Suppose that $\underline{\rho}_S(m^P) \in (\rho_R(m^P) - \ell, \rho_R(\bar{\theta}_*) - \ell]$. If some states above $\bar{\theta}_*$ are pooled, say (θ_1, θ_2) , we can decrease the value of (13) by separating these states, as follows from

$$\int_{\theta_1}^{\theta} (\rho_R(\theta_1 + \theta_2)/2 - \ell) d\tilde{\theta} > \int_{\theta_1}^{\theta} (\rho_R(\tilde{\theta}) - \ell) d\tilde{\theta} \quad \text{for } \theta \in (\theta_1, \theta_2).$$

If all states above $\bar{\theta}_*$ are separated, we can decrease the value of (13) by pooling states $[\bar{\theta}_*, \bar{\theta}_* + \varepsilon)$ together with the states in m^P , and inducing the same decision $\rho_R(m^P)$ for all these states, leading to a contradiction. Next suppose that $\underline{\rho}_S(m^P) > \rho_R(\bar{\theta}_*) - \ell$. Then some states adjacent to $\bar{\theta}_*$ from above, say $(\bar{\theta}_*, \hat{\theta})$, must be pooled, such that $\underline{\rho}_S(m^P) < \rho_R((\bar{\theta}_* + \hat{\theta})/2) - \ell$. But then we can decrease the value of (13) by separating states $(\hat{\theta} - \varepsilon, \hat{\theta})$, as follows from

$$\begin{aligned} \int_{\bar{\theta}_*}^{\theta} (\rho_R(\bar{\theta}_* + \hat{\theta})/2 - \ell) d\tilde{\theta} &> \int_{\bar{\theta}_*}^{\theta} (\rho_R(\bar{\theta}_* + \hat{\theta} - \varepsilon)/2 - \ell) d\tilde{\theta} \quad \text{for } \theta \in (\bar{\theta}_*, \hat{\theta} - \varepsilon), \\ \int_{\bar{\theta}_*}^{\theta} (\rho_R(\bar{\theta}_* + \hat{\theta})/2 - \ell) d\tilde{\theta} &> \int_{\bar{\theta}_*}^{\hat{\theta} - \varepsilon} (\rho_R(\bar{\theta}_* + \hat{\theta} - \varepsilon)/2 - \ell) d\tilde{\theta} + \int_{\hat{\theta} - \varepsilon}^{\theta} (\rho_R(\tilde{\theta}) - \ell) d\tilde{\theta} \\ &\quad \text{for } \theta \in (\hat{\theta} - \varepsilon, \hat{\theta}). \end{aligned}$$

Analogously, we can show that if $\underline{\theta}_* > 0$, then $\underline{\rho}_S(m^P) = \rho_R(m^P) + \ell$ and all states $\theta < \underline{\theta}_*$ are separated by $\underline{\mu}_S$. This implies that either $\underline{\theta}_* = 0$ or $\bar{\theta}_* = 1$.

Thus, the sender's worst equilibrium payoff $\underline{v}_S$ is achieved either by a message rule that pools the states below $\underline{\theta}^L$ (and separates the rest) and decision rule $\underline{\rho}(m) = \rho_R(m) - \ell$, or by a message rule that pools the states above $\underline{\theta}^H$ (and separates the rest) and decision rule $\underline{\rho}(m) = \rho_R(m) + \ell$. Computation reveals that the value of (13) under pooling interval $[0, \underline{\theta}]$ and decision rule $\underline{\rho}(m) = \rho_R(m) - \ell$ is smaller than the value of (13) under pooling interval $(1 - \underline{\theta}, 1]$ and decision rule $\underline{\rho}(m) = \rho_R(m) + \ell$ for all $\underline{\theta} \in [0, 1]$ if $a_S/2 + b_S > a_R/2 + b_R$. Moreover, the value of (13) is minimized for $\underline{\theta}^L \in [0, 1]$ at either $\underline{\theta}^L = 0$ or

$$\underline{\theta}^L = \frac{a_R + \sqrt{a_R^2 - 8a_S(b_S - b_R + \ell)}}{2a_S} < 1,$$

where the inequality follows from the assumption $a_S/2 + b_S > a_R/2 + b_R$. Further computation then produces $\underline{\theta}_S$, as defined in Proposition 7. $\square$

APPENDIX E: PUBLIC INFORMATION

PROOF OF PROPOSITION 6. Suppose, for the sake of argument, that $L(\bar{v})$, as defined in Section 3.1, takes the same value under ψ and $\hat{\psi}$. We show that the best equilibrium joint payoff is higher and the worst monotone equilibrium payoffs are smaller under $\hat{\psi}$ than under ψ . Specifically, $\hat{v} \geq \bar{v}$, $\hat{v}_R \leq \underline{v}_R$, and $\hat{v}_S \leq \underline{v}_S$. This implies that $\hat{L}(\hat{v}) \geq L(\bar{v})$. The proposition follows easily from this observation.

The best equilibrium joint payoff $\bar{v}$ under ψ can be supported by an equilibrium in single-period punishment strategies such that $\rho_*(\mu_*(\theta))$ is induced in each period on the equilibrium path by application of Proposition 5 to each realization of signal ψ . Since $\rho_*(\mu_*(\theta))$ is nondecreasing in θ on $[0, 1]$, it can be supported in an equilibrium under less informative signal $\hat{\psi}$ by application of an analogue of Proposition 1 to each realization of signal ψ , so $\hat{v} \geq \bar{v}$.

By Proposition 5, the receiver's worst equilibrium payoffs under $\hat{\psi}$ and ψ are

$$\hat{v}_R = \mathbb{E}[u_R(\rho_R(\hat{\psi}(\theta)), \theta)] \leq \mathbb{E}[u_R(\rho_R(\psi(\theta)), \theta)] = \underline{v}_R,$$

where the inequality holds because ψ is more informative than $\hat{\psi}$.

By an argument similar to the proof of Proposition 5, the sender's worst equilibrium payoff under ψ can be supported by $\tau = 0$, $T(m) = 0$, and $v_S(m) = \underline{v}_S$; that is, the sender may refuse to make any ex ante or ex post transfers, and the worst punishment for him would involve zero transfers from the receiver and the worst continuation payoff. Let $\mu(\theta)$ and $\rho(m)$ be penal message and decision rules that support this equilibrium. By assumption, $\rho(\mu(\theta))$ is nondecreasing in θ . Then the interim transfer $t(\mu(\theta))$ is defined by (4) and (5) given that the set of states is $\psi(\theta) \subset [0, 1]$ rather than $[0, 1]$,

$$\begin{aligned} t(m) &= h(m) - \min_{m \in \mu(\psi(\theta))} h(m), \\ h(m) &= u_S(\rho(m), \theta(m)) - \int_0^{\theta(m)} \frac{\partial u_S}{\partial \theta}(\rho(\mu(\tilde{\theta})), \tilde{\theta}) d\tilde{\theta}, \end{aligned} \tag{33}$$

where $\theta(m) \in m$. The message and decision rules $\mu(\theta)$ and $\rho(m)$ such that $\rho(\mu(\theta))$ is nondecreasing in θ can be supported in equilibrium under $\hat{\psi}$ using the interim transfer rule $\hat{t}(m)$ that differs from $t(m)$ given by (33) only in that the minimum of h is taken over $m \in \mu(\hat{\psi}(\theta))$ rather than over $m \in \mu(\psi(\theta))$. Since $\psi(\theta) \subset \hat{\psi}(\theta)$ for all $\theta \in [0, 1]$ by the definition of more informative signals, we have $\hat{t}(\mu(\theta)) \geq t(\mu(\theta))$ for all $\theta \in [0, 1]$ and, thus,

$$\underline{v}_S \leq \mathbb{E}[u_S(\rho(\mu(\theta)), \theta) - \hat{t}(\mu(\theta))] \leq \mathbb{E}[u_S(\rho(\mu(\theta)), \theta) - t(\mu(\theta))] = \underline{v}_S. \quad \square$$

REFERENCES

- Abreu, Dilip (1986), “Extremal equilibria of oligopolistic supergames.” *Journal of Economic Theory*, 39, 191–225. [1395]
- Alonso, Ricardo and Niko Matouschek (2007), “Relational delegation.” *RAND Journal of Economics*, 38, 1070–1089. [1394, 1395]
- Ambrus, Attila and Georgy Egorov (2017), “Delegation and nonmonetary incentives.” *Journal of Economic Theory*, 171, 101–135. [1394]
- Au and Pak Hung (2015), “Dynamic information disclosure.” *RAND Journal of Economics*, 46, 791–823. [1395]
- Aumann, Robert J., Michael Maschler, and Richard E. Stearns (1995), “Repeated games with incomplete information.” MIT Press. [1396]
- Austen-Smith, David (1995), “Campaign contributions and access.” *American Political Science Review*, 89, 566–581. [1392]
- Austen-Smith, David and Jeffrey S. Banks (2000), “Cheap talk and burned money.” *Journal of Economic Theory*, 91, 1–16. [1394, 1399]
- Baker, George, Robert Gibbons, and Kevin J. Murphy (1994), “Subjective performance measures in optimal incentive contracts.” *Quarterly Journal of Economics*, 109, 1125–1156. [1408, 1411]
- Baker, George, Robert Gibbons, and Kevin J. Murphy (2002), “Relational contracts and the theory of the firm.” *Quarterly Journal of Economics*, 117, 39–84. [1408]
- Baker, George, Robert Gibbons, and Kevin J. Murphy (2011), “Relational adaptation.” Working Paper, MIT. [1395]
- Benabou, Roland and Jean Tirole (2003), “Intrinsic and extrinsic motivation.” *Review of Economic Studies*, 70, 489–520. [1392]
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2018), “The design and price of information.” *American Economic Review*, 108, 1–48. [1395]
- Bergemann, Dirk and Stephen Morris (2016), “Bayes correlated equilibrium and the comparison of information structures in games.” *Theoretical Economics*, 11, 487–522. [1411]
- Bergemann, Dirk and Martin Pesendorfer (2007), “Information structures in optimal auctions.” *Journal of Economic Theory*, 137, 580–609. [1395]

- Best, James and Daniel Quigley (2020), “Persuasion for the long run.” Working Paper, University of Oxford. [1395]
- Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier (2021), “Dynamic persuasion with outside information.” *American Economic Journal: Microeconomics*, 13, 179–194. [1395]
- Blackwell, David (1953), “Equivalent comparisons of experiments.” *Annals of Mathematical Statistics*, 24, 265–272. [1406, 1410]
- Bull, Clive (1987), “The existence of self-enforcing implicit contracts.” *Quarterly Journal of Economics*, 102, 147–159. [1394]
- Che, Yeon-Koo and Johannes Horner (2018), “Optimal design for social learning.” *Quarterly Journal of Economics*, 133, 871–925. [1395]
- Crawford, Vincent P and Joel Sobel (1982), “Strategic information transmission.” *Econometrica*, 50, 1431–1451. [1392, 1394, 1395, 1402, 1412]
- Crémer, Jacques (1995), “Arm’s length relationships.” *Quarterly Journal of Economics*, 110, 275–295. [1411]
- Dworczak, Piotr (2020), “Mechanism design with aftermarkets: Cutoff mechanisms.” *Econometrica*, 88, 2629–2661. [1395]
- Dworczak, Piotr and Giorgio Martini (2019), “The simple economics of optimal persuasion.” *Journal of Political Economy*, 127, 1993–2048. [1403, 1405]
- Edlin, Aaron S. and Chris Shannon (1998), “Strict monotonicity in comparative statics.” *Journal of Economic Theory*, 81, 201–219. [1421]
- Ely, Jeffrey, Alexander Frankel, and Emir Kamenica (2015), “Suspense and surprise.” *Journal of Political Economy*, 123, 215–260. [1395]
- Ely, Jeffrey C. (2017), “Beeps.” *American Economic Review*, 107, 31–53. [1395]
- Ely, Jeffrey C. and Martin Szydlowski (2020), “Moving the goalposts.” *Journal of Political Economy*, 128, 468–506. [1395]
- Eső, Péter and Balázs Szentes (2007), “Optimal information disclosure in auctions and the handicap auction.” *Review of Economic Studies*, 74, 705–731. [1395]
- Fong, Yuk-Fai and Jin Li (2016), “Information revelation in relational contracts.” *Review of Economic Studies*, 84, 277–299. [1411]
- Frankel, Alexander (2016), “Discounted quotas.” *Journal of Economic Theory*, 166, 396–444. [1396]
- Gentzkow, Matthew and Emir Kamenica (2016), “A Rothschild–Stiglitz approach to Bayesian persuasion.” *American Economic Review*, 106, 597–601. [1403]
- Goldlücke, Susanne and Sebastian Kranz (2012), “Infinitely repeated games with public monitoring and monetary transfers.” *Journal of Economic Theory*, 147, 1191–1221. [1395, 1397]

- Grossman, Gene M. and Elhanan Helpman (1994), “Protection for sale.” *American Economic Review*, 84, 833–850. [1391]
- Grossman, Gene M. and Elhanan Helpman (1996), “Electoral competition and special interest politics.” *Review of Economic Studies*, 63, 265–286. [1391]
- Hermalin, Benjamin E. (1998), “Toward an economic theory of leadership: Leading by example.” *American Economic Review*, 88, 1188–1206. [1392]
- Hirshleifer, Jack (1971), “The private and social value of information and the reward to inventive activity.” *American Economic Review*, 61, 561–574. [1411]
- Hörner, Johannes and Andrzej Skrzypacz (2016), “Selling information.” *Journal of Political Economy*, 124, 1515–1562. [1395]
- Hörner, Johannes, Satoru Takahashi, and Nicolas Vieille (2015), “Truthful equilibria in dynamic Bayesian games.” *Econometrica*, 83, 1795–1848. [1407]
- Jehiel, Philippe (2015), “On transparency in organizations.” *Review of Economic Studies*, 82, 736–761. [1402]
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101, 2590–2615. [1392, 1402, 1403, 1412]
- Karamychev, Vladimir and Bauke Visser (2017), “Optimal signaling with cheap talk and money burning.” *International Journal of Game Theory*, 46, 813–850. [1394, 1399]
- Kartik, Navin (2007), “A note on cheap talk and burned money.” *Journal of Economic Theory*, 136, 749–758. [1394, 1399]
- Kartik, Navin (2009), “Strategic communication with lying costs.” *Review of Economic Studies*, 76, 1359–1395. [1394]
- Kartik, Navin, Marco Ottaviani, and Francesco Squintani (2007), “Credulity, lies, and costly talk.” *Journal of Economic Theory*, 134, 93–116. [1394]
- Kloosterman, Andrew (2015), “Public information in Markov games.” *Journal of Economic Theory*, 157, 28–48. [1411]
- Kolotilin, Anton (2018), “Optimal information disclosure: A linear programming approach.” *Theoretical Economics*, 13, 607–636. [1403, 1405]
- Kolotilin, Anton and Hongyi Li (2019), “Relational communication.” arXiv preprint arXiv:1901.05645v1. [1407]
- Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li (2017), “Persuasion of a privately informed receiver.” *Econometrica*, 85, 1949–1964. [1403, 1405]
- Kolotilin, Anton, Tymofiy Mylovanov, and Andriy Zapechelnyuk (forthcoming). “Censorship as optimal persuasion.” *Theoretical Economics*. [1405]
- Kremer, Ilan, Yishay Mansour, and Motty Perry (2014), “Implementing the ‘wisdom of the crowd’.” *Journal of Political Economy*, 122, 988–1012. [1395]

Krishna, Vijay and John Morgan (2008), “Contracting for information under imperfect commitment.” *RAND Journal of Economics*, 39, 905–925. [1394]

Kuvalekar, Aditya, Elliot Lipnowski, and Joao Ramos (2020), “Goodwill in communication.” Working Paper, Columbia University. [1407]

Levin, Jonathan (2003), “Relational incentive contracts.” *American Economic Review*, 93, 835–857. [1392, 1394, 1397, 1399, 1408]

Li, Hao and Xianwen Shi (2017), “Discriminatory information disclosure.” *American Economic Review*, 107, 3363–3385. [1395]

Lipnowski, Elliot and Joao Ramos (2020), “Repeated delegation.” *Journal of Economic Theory*, 188. [1396]

Lohmann, Susanne (1995), “Information, access, and contributions: A signaling model of lobbying.” *Public Choice*, 85, 267–284. [1392]

Macleod, W. Bentley, and James M. Malcomson (1989), “Implicit contracts, incentive compatibility, and involuntary unemployment.” *Econometrica*, 57, 447–480. [1394]

Mailath, George J. and Larry Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press. [1412, 1413, 1415]

Margaria, Chiara and Alex Smolin (2018), “Dynamic communication with biased senders.” *Games and Economic Behavior*, 110, 330–339. [1396]

Milgrom, Paul and Ilya Segal (2002), “Envelope theorems for arbitrary choice sets.” *Econometrica*, 70, 583–601. [1398, 1415]

Milgrom, Paul and Chris Shannon (1994), “Monotone comparative statics.” *Econometrica*, 62, 157–180. [1417]

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2020), “Persuading the principal to wait.” *Journal of Political Economy*, 128, 2542–2578. [1395]

Ottaviani, Marco (2000), “The economics of advice.” Working Paper, Bocconi University. [1394]

Rayo, Luis and Ilya Segal (2010), “Optimal information disclosure.” *Journal of Political Economy*, 118, 949–987. [1392]

Renault, Jérôme, Eilon Solan, and Nicolas Vieille (2013), “Dynamic sender–receiver games.” *Journal of Economic Theory*, 148, 502–534. [1396]

Rochet, Jean-Charles (1987), “A necessary and sufficient condition for rationalizability in a quasi-linear context.” *Journal of Mathematical Economics*, 16, 191–200. [1398, 1415]

Smolin, Alex (forthcoming), “Dynamic evaluation design.” *American Economic Journal: Microeconomics*. [1395]

Co-editor Simon Board handled this manuscript.

Manuscript received 24 April, 2019; final version accepted 21 October, 2020; available online 16 December, 2020.