Relational Communication with Transfers*

Anton Kolotilin and Hongyi Li†

December 9, 2016

Abstract

We enrich a cheap-talk game between an informed sender and an uninformed receiver by adding repeated interactions and voluntary transfer payments. Transfers play two roles here: they 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 preferences are sufficiently divergent. In this case, extreme information is optimally pooled to discipline the receiver's decision-making by reducing her reneging temptation. As an extension, we consider a partially-informed receiver. As the receiver becomes more informed, welfare strictly decreases because self-enforcing agreements become harder to sustain.

JEL Classification: C73, D82, D83

Keywords: communication, signaling, relational contracts

*We thank Luis Zermeno for important early contributions to this paper. We also thank 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, and participants at various seminars and conferences for helpful comments. Kolotilin acknowledges support from the Australian Research Council.

†UNSW Australia, School of Economics. Emails: akolotilin@gmail.com and hongyi@gmail.com.

1 Introduction

Decision-makers often rely on information supplied by other interested parties. In such interactions, decision and communication outcomes are often difficult to enforce contractually. The existing literature either assumes that compensation contracts are completely unavailable (strategic communication), or ignores the contractibility issue by restricting attention to fully-contractible settings (mechanism design). This paper covers a middle ground. It studies how such interactions can be disciplined using relational contracts: compensation schemes that are discretionary and self-enforced by the players' concerns for the value of their future relationship. We show that optimal relational contracts may involve deliberately incomplete communication, which serves to mitigate decision-making problems, in a way that improves welfare for all players.

Our analysis is based on an infinitely-repeated cheap-talk game, played by a sender and receiver who share a common discount factor. In each period, the sender privately observes an independent draw of the state and sends a cheap-talk message to the receiver, who then makes a decision. The players' payoffs are quadratic in the state and decision, but the magnitude and sign of the relative bias (the difference between the receiver's and sender's preferred decisions) may depend on the state.[1] At the end of the period, the state is publicly revealed, and time moves on to the next period. This framework is particularly tractable and allows us to fully characterize optimal communication and decision outcomes. But we also show that the main insights from our analysis apply more generally.[2]

We allow players to make voluntary transfers to each other at any point in the game. Transfers play two key roles in the relational contract. First, discretionary transfers motivate the receiver to make mutually beneficial decisions. Second, and more novel, transfers allow the sender to credibly signal his private information at no welfare cost.

This second role seems particularly relevant: many relationships rely on (monetary and non-monetary) transfers not only to reward behavior but also to support communication. A lobbyist may double his usual contribution to a politician to emphasize the importance of a piece of legislation to his constituents. An effusive annual employee evaluation may be discarded as cheap talk, but the same evaluation accompanied by a hefty bonus is a credible

[1]This setting generalizes the 'uniform-quadratic' example from Crawford and Sobel (1982), where the relative bias does not depend on the state.

[2]Specifically, in Section 6, we consider the following extensions: non-quadratic payoff functions; different allocations of decision-making authority; imperfect monitoring of the state and of the receiver's decision; and correlation of states across time.

statement that the employee is highly valued. Sweet talk and declarations of love on every Valentine's day may be music to a significant other's ears, but are ultimately non-credible signals of affection unless accompanied by expensive gifts.[3]

In our setting, compensation schemes for decision-making and communication must constitute a self-enforcing relational contract: the value of the future relationship must be large enough that neither party wishes to renege on any promise to compensate the other. This self-enforcement constraint limits incentive provision within the relational contract, and thus the productivity of the relationship.

Because utility is transferable in our setting, optimality in our setting is tantamount to surplus maximization. In fact, to characterize Pareto-optimal equilibria, it suffices to focus on a simple class of stationary contracts that produce identical communication and decision-making outcomes, and differ only in transfers (which determine the division of surplus).

Our analysis hinges on the remarkable power of transfers as a signaling device. In particular, virtually any message rule can be credibly implemented using an appropriate transfer rule, without directly affecting the sender's temptation to renege on the relational contract. In other words, the relational contract is not constrained by the sender's incentive problem. Therefore, the only constraint on the relational contract is the receiver's temptation to make decisions that benefit herself but hurt the sender.

Unsurprisingly, relational contracts can attain the first-best when both players are patient and the shadow of the future looms large: in this case the optimal equilibrium induces full separation of the state, as well as ex-post efficient decision-making.

Our first main contribution is the complete characterization of the optimal relational contract when the first best is unattainable, because players are insufficiently patient. In this case, the receiver's self-enforcement constraint – representing her temptation to choose inefficient decisions – is binding. Our key insight is that this constraint can be relaxed by pooling information. The receiver's reneging temptation is greatest in extreme states where the relative bias is high. By pooling extreme states with less extreme states, the receiver's maximum reneging temptation is moderated. In fact, we show that if the receiver is prone to overreaction (the relative bias is highly sensitive to the state), then (partial or complete) pooling occurs in equilibrium. In this case, extreme states are optimally pooled. In other words, the sender reveals moderate states and conceals extreme states from an

[3]The signalling role of political lobbying is discussed by Bennedsen and Feldmann (2006) and Esteban and Ray (2006), of employee rewards by Benabou and Tirole (2003) and Suvorov and van de Ven (2009), and of gifts in relationships by Camerer (1988).

overreaction-prone receiver. On the other hand, if the receiver is not prone to overreaction, then pooling does not occur. Instead, the optimal second-best relational contract induces full separation, and tolerates inefficient decision-making in extreme states.

The result that the sender reveals (hides) information when conflict of interest is low (high) seems to capture common patterns of communication within relationships. Movie distributors typically share information about an upcoming movie's prospects with exhibitors (who decide how widely and frequently the movie will be shown in theatres), but may be less forthcoming with negative test results for a particularly expensive movie. Lobbyists often discuss in detail the costs and benefits of potential legislation with politicians, but may hide or distort their private information in cases that are particularly controversial or consequential. Technology companies regularly feed information about upcoming product launches to 'beat' reporters in exchange for favorable reporting, with the implicit understanding that the company may withhold key information for products or issues where controlling the media narrative is particularly important.

One important implication of our analysis is that in settings where monetary or non-monetary 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 communication. This 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 achieved.

Our second main contribution is to show that productivity within the relationship is strictly decreasing in the availability of public information about the state. One might naively expect public information to enable more informed decision-making by the receiver. This turns out not to be the case: the availability of transfers as a signaling device implies that improving the receiver's information can always be achieved without imposing additional shadow costs on the relational contract, so public information brings no informational benefits to the relationship. On the flip side, public information makes a self-enforcing relational contract more difficult to sustain in two ways. First, public information improves both players' worst possible equilibrium payoffs, and thus limits the severity of off-path punishments. Second, public information prevents information pooling, the benefits of which we discussed above. Both of these channels tighten the self-enforcement constraint and thus weaken the relationship.

This comparative static result has implications for optimal organizational and institutional design. Within organizations, parties who hold useful information are often not in control of the relevant decisions. Our model suggests that such 'arms-length' organizational structures do not necessarily result in maladapted decisions. On the contrary, they may enable effective informal communication and thus improve decision-making. In fact, formal communication processes that mechanically increase decision-makers' access to information may be counterproductive because they weaken the relational contracts within the organization.

1.1 Related Literature

To our knowledge, our model is the first to combine three key ingredients: communication, repeated interactions, and transfers. Consequently, our paper relates to a number of different literatures.

Our analysis builds on an extensive literature on repeated interactions with transfers. The seminal papers by Bull (1987) and Macleod and Malcomson (1989) focused 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, so there is no role for information transmission between the principal and agent. In contrast, our setting involves an informed sender and an uninformed decision-maker (receiver), in the vein of Crawford and Sobel (1982). Consequently, relational contracts are used to manage both communication and decision-making.[4]

Similar to us, Alonso and Matouschek (2007) consider repeated communication.[5] In contrast to us, they do not allow for transfers, and they consider a sequence of short-lived senders rather than a single long-lived sender. In their setting, repeated interaction disciplines decision-making, in order to sustain more informative communication. In contrast, in our setting, credible communication is easy to achieve; so repeated interaction is used to improve decision-making which in turn determines the informativeness of optimal communication.

In our model, transfers from sender to receiver are used to signal information.[6] Austen-

[4]Baker, Gibbons and Murphy (2011) consider a model of repeated decision-making with transfers, but assume symmetric information, so communication plays no role.

[5]Relatedly, Ottaviani and Sørensen (2006a,b) study communication where the sender has reputational concerns.

[6]Krishna and Morgan (2008) consider a communication game with transfers. However, in their setting, transfers are from receiver to sender, and thus cannot be used to signal information. Bester and Krähmer

Smith and Banks (2000) and Kartik (2007) consider a related (albeit static) setting where the sender burns money to signal information.[7] They show that the ability to burn money dramatically expands the set of equilibrium communication outcomes. In contrast to their setting, 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.[8]

By signaling with transfers, the sender can endogenously commit to almost any message rule as long as payoff functions satisfy standard sorting conditions (as in Crawford and Sobel 1982). In fact, this commitment ability is the basic premise of the literature on Bayesian Persuasion (Rayo and Segal 2010 and Kamenica and Gentzkow 2011).[9] Our paper thus provides a rationale for how a privately informed sender can endogenously commit to a message rule. From a technical perspective, our proofs rely on results from Kamenica and Gentzkow (2011) and Kolotilin (2016).

One of our insights is that optimal relational contracts may involve pooling of information in states of extreme conflict. The idea that optimal communication may involve a combination of pooling and separation has been discussed elsewhere (e.g., Dye 1985, Krishna and Morgan 2008, and Kartik 2009), albeit driven by very different economic mechanisms. In Dye (1985), the sender's information is verifiable but some senders may be uninformed. In this case, low-quality types will pool with each other by pretending to be uninformed but high-quality types will fully separate. In Kartik (2009), the upwardly-biased sender incurs a lying cost from misreporting his type, which lies in a bounded interval. In equilibrium, low types separate, and high types pool by pretending to be the highest possible type. In these papers, full separation can never be achieved in equilibrium. In contrast, full separation is always feasible in our model, even in a static setting. In Krishna and

(2016) consider a related setting with contractible transfer schemes and study the optimal allocation of authority, similar to Dessein (2002).

[7]Relatedly, a number of papers enrich the Crawford and Sobel (1982) setting with other means of costly signaling such as lying costs (Kartik, Ottaviani and Squintani 2007 and Kartik 2009) and costly information production (Li and Li 2013).

[8]In the setting with burned money, equilibrium communication outcomes differ along the Pareto frontier because there is a tradeoff between the informativeness of communication and the costs of burning money. The receiver's optimal equilibrium clearly involves full separation; Karamychev and Visser (2016) characterize the sender's optimal equilibrium.

[9]A model of repeated Bayesian persuasion would reproduce many of the insights from our model of relational communication. The existing literature has studied dynamic Bayesian persuasion with persistent information (Au 2015, Ely, Frankel and Kamenica 2015, Bizzotto, Rüdiger and Vigier 2016, Ely 2016, Hörner and Skrzypacz 2016, and Orlov, Skrzypacz and Zryumov 2016).

Morgan (2008), discussed in Footnote 6, full separation is also always feasible, but it turns out to be never optimal, because incentivizing high types of the upwardly-biased sender to separate is very costly for the receiver. In contrast to Krishna and Morgan (2008) and Kartik (2009) where the relative bias is constant, pooling can be optimal in our setting only if the relative bias varies with the state. In a repeated setting, Levin (2003) shows that under adverse selection the optimal relational contract involves pooling of low-cost agent types at the same effort level. In our setting, pooling serves to affect the receiver's beliefs and thus directly improves her decision-making. In contrast, the decision-maker (agent) in Levin (2003) is fully informed, so pooling has no such effect.

Finally, our 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) clarifies 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.[10] This relates to the logic of our model, where the availability of public information makes it impossible to pool self-enforcement 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 relational surplus and thus tightens the self-enforcement constraint.[11]

2 Model

2.1 Setup

A *sender* (S) and a *receiver* (R) play an infinitely repeated communication game with perfect monitoring and with voluntary transfer payments. Time is discrete and the players have common discount factor $\delta \in [0, 1)$. In each period, the same stage game is played. The sender privately observes the state $\theta \in [0, 1)$ and sends a cheap-talk message $m \in [0, 1)$ to the receiver who then makes a decision $d \in \mathbb{R}$; the state is subsequently observed by the receiver. The state θ is independently drawn each period from a distribution with strictly

[10] Crémer (1995), Fong and Li (2016) and Kolotilin (2015) discuss other settings where public information may be detrimental.

[11] This point relates to the insight from the relational contracting literature that improvements in the players' outside options make cooperation within the relationship more difficult to sustain (e.g., Baker, Gibbons and Murphy 1994).

positive density f for all $\theta \in [0,1)$. The players' payoffs are $u_R(d,\theta) = -\alpha_R (d-\theta)^2$ and $u_S(d,\theta) = -\alpha_S (d - a\theta - b)^2$, where $a > 0$, $b \in \mathbb{R}$, $\alpha_R > 0$, $\alpha_S > 0$, and $\alpha_R + \alpha_S = 1$ (normalization). That is, the players' preferences over d differ: the receiver's preferred decision is $d_R(\theta) = \theta$, the sender's preferred decision is $d_S(\theta) = a\theta + b$, and the first-best decision is $d_{FB}(\theta) = (\alpha_R + \alpha_S a)\theta + \alpha_S b$.[12]

The players can make voluntary (non-contractible) transfers at any point in the game. Specifically, we enrich the stage game with three rounds of transfers: (i) an *ex-ante* round before the sender observes the state, (ii) an *interim* round after the sender observes the state and sends the message but before the receiver chooses a decision, and (iii) an *ex-post* round after the decision is chosen and the state is publicly observed. In each round, transfers are made sequentially, first by the sender and then by the receiver. Each player chooses a non-negative *gross* transfer to the other player and a non-negative amount of money to burn. The players' transfer choices in each round determine their *net* transfers in that round. Specifically, the sender's net transfer equals his gross transfer, minus the receiver's gross transfer, plus the sender's money burned (and similarly for the receiver). The net transfers by player $i \in \{S, R\}$ in the ex-ante, interim, and ex-post rounds are denoted by τ_i, t_i, and T_i; so the stage game payoff of player i is $u_i(d,\theta) - \tau_i - t_i - T_i$. Note that net transfers in each round must satisfy $\tau_S + \tau_R \geq 0$, $t_S + t_R \geq 0$, and $T_S + T_R \geq 0$, with strict inequality in the case of burned money.[13] Although we allow for both ex-ante and ex-post transfers, ex-ante transfers can substitute for ex-post transfers (and vice versa).[14]

The game has perfect monitoring in that all actions (message, decision, and transfers) are immediately publicly observed and the state is publicly observed immediately after the decision is chosen. Figure 1 summarizes the timing of each stage game.

We focus on pure-strategy perfect Bayesian equilibria, called *equilibria* henceforth. Given a message rule $m(\theta)$, we abuse notation and refer to a realization of this message rule as m. We restrict attention to direct message rules such that $m = \mathbb{E}[\theta|m]$ for any realization m; this restriction is without loss because payoffs are quadratic in the state and

[12]These payoff functions nest two special cases. First, Crawford and Sobel (1982)'s example has a constant upward bias of the sender, so that $a = 1$ and $b > 0$. Second, Kamenica and Gentzkow (2011)'s lobbying example has a bias of the sender towards a specific decision in the sense that $d_S(\theta) = \beta\theta + (1-\beta)d^*$ with $\beta \in (0,1)$ and $d^* > 1$, so that $a = \beta$ and $b = (1-\beta)d^*$.

[13]Conversely, for any net transfers that satisfy these three constraints, we can construct gross transfers and burned money amounts that correspond to these net transfers.

[14]Thus we may, e.g., restrict attention to equilibria where the ex-ante transfers (τ_S, τ_R) are zero in every period except the first period. In this case, we may think of the first period's ex-ante transfers as 'up-front' payments that determine the division of surplus in the relationship.

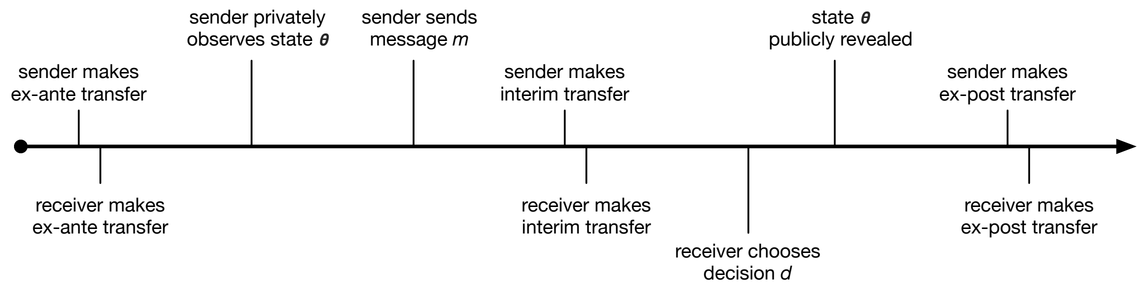

Figure 1: Timing of stage game

decision.[15]

2.2 Optimality and Stationarity

An equilibrium is *stationary* if on the equilibrium path, the message rule $m(\theta)$, the decision rule $d(m)$, and the ex-ante, interim, and ex-post transfer rules τ_i, $t_i(m)$ and $T_i(\theta)$ for $i \in \{S, R\}$ 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.

The following proposition extends some of Levin (2003)'s and Goldlücke and Kranz (2012)'s results to our setting.

Proposition 1 *There exist $\underline{v}_S$, $\underline{v}_R$ and $\overline{v}$ such that the set of equilibrium payoffs $V \subset \mathbb{R}^2$ is a simplex of the form*

$$V = \left\{(v_S, v_R) : v_S \geq \underline{v}_S, v_R \geq \underline{v}_R, v_S + v_R \leq \overline{v}\right\}. \tag{1}$$

Any optimal equilibrium is sequentially optimal and involves no burned money. Further, there exists a stationary optimal equilibrium σ^* such that any $(v_S, v_R) \in V$ can be supported by an equilibrium that differs from σ^* only in the first period's ex-ante transfers.

[15] This formulation excludes the possibility of distinct messages m_1 and m_2 with $\mathbb{E}[\theta|m_1] = \mathbb{E}[\theta|m_2]$. To see why this exclusion is innocuous, consider an equilibrium with a message rule $m(\theta)$, a decision rule $d(m)$, and some transfer rule. Suppose that $m(\theta)$ sends m_1 and m_2 with probabilities p_1 and p_2 such that $\mathbb{E}[\theta|m_1] = \mathbb{E}[\theta|m_2]$. Modify $m(\theta)$ so that a new message m_0 is sent in place of both m_1 and m_2, and modify $d(m)$ so that the receiver chooses the 'average' decision $(p_1\, d(m_1) + p_2\, d(m_2))\,/(p_1 + p_2)$ whenever she receives m_0. With these modifications, it is easy to check that both players' expected payoffs (weakly) increase. Now, modify the transfer rule appropriately so that both players burn just enough money to reproduce the expected payoffs from the original equilibrium. These modified message, decision, and transfer rules constitute a payoff-equivalent equilibrium.

Because players' payoffs are quasi-linear in money, surplus is fully transferable and all optimal equilibria induce the message and decision rules that maximize joint surplus v_S+v_R. Further, due to free disposal (both players can burn money), the set of equilibrium payoffs has a simplex structure.

Optimal equilibria do not involve burned money, because burning money would only tighten incentive constraints and reduce joint surplus. Therefore, the Pareto frontier would not change if we modified the model by disallowing money burning.

Notice that the worst equilibrium payoffs are endogenously determined. It is easy to see that the receiver's worst equilibrium payoff is supported by the repetition of a static 'babbling' equilibrium, whereby the receiver chooses the uninformed decision $d = \mathbb{E}[\theta]$ and does not pay or receive any transfers. Consequently, $\underline{v}_R = -\alpha_R \mathrm{Var}(\theta)$. On the other hand, the sender's worst equilibrium payoff $\underline{v}_S$ depends on the discount factor δ.[16]

3 Myopic Benchmark

3.1 Implementable Communication

We start our analysis by considering the myopic benchmark where both players have zero discount factor ($\delta = 0$). This benchmark corresponds to the static version of our model, so the receiver always chooses her preferred decision, $d_R(m) = m$ for all m. In this benchmark, interim transfers are quite powerful and allow players to sustain almost any communication outcome, even though transfers are voluntary and players have no commitment power.

A (direct) message rule $m(\theta)$ is *monotone* if it is non-decreasing in θ. So, any monotone message rule is characterized by a set of disjoint intervals whereby states within each interval I are pooled into a message $\mathbb{E}[\theta \,|\, \theta \in I]$, and all remaining states are fully separated so that $m(\theta) = \theta$.

Lemma 1 *A message rule is supported in some equilibrium if and only if it is monotone.*

In particular, Lemma 1 implies that the sender can always credibly induce full separation in equilibrium by using interim transfers to signal information.[17] To illustrate this point,

[16] When players are patient, the equilibrium can support decisions that are distorted away from the receiver's ex-post preferred decision in a way that hurts the sender.

[17] Although 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.

consider the Crawford and Sobel (1982) example where the sender's bias is a positive constant, $b = d_S(\theta) - d_R(\theta) > 0$. Let us show that the message rule $m(\theta) = \theta$ together with the interim transfer rule $t_S(m) = 2bm = -t_R(m)$ can credibly induce full separation. Given that the (myopic) receiver always chooses her preferred decision $d_R(m) = m$, the sender always prefers to truthfully report the state ($m = \theta$) and make the associated interim transfer $t_S(\theta)$ because $m = \theta$ uniquely maximizes the sender's payoff[18]

$$u_S(d_R(m), \theta) - t_S(m) = -(m - \theta - b)^2 - 2bm.$$

Lemma 1 is closely connected to existing results from the literature on cheap talk and burned money (e.g., Austen-Smith and Banks 2002, Kartik 2007, and Karamychev and Visser 2016). 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 ($t_R(m) = -t_S(m)$) or whether the sender burns money ($t_R(m) = 0$).

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 are supported by the threat of playing a babbling equilibrium.

In other words, the sender can effectively commit at no welfare cost to any monotone message rule. Thus, to characterize the Pareto frontier, we may reformulate the problem as that of a social planner who wants to maximize joint surplus and can choose any monotone message rule:

$$\max_{m(\theta)} \mathbb{E}\left[\sum\nolimits_{i \in \{S,R\}} u_i\big(d_R\left(m(\theta)\right), \theta\big)\right]$$

subject to $m(\theta)$ is direct and monotone.

3.2 Optimal Communication

The optimal myopic equilibrium involves an extreme communication outcome: either complete pooling or full separation.

[18] In this example, if the sender chooses a message-transfer pair that is not observed on the equilibrium path, then the receiver believes that $\theta = 0$. Therefore, any out-of-equilibrium pair (m', t') is weakly dominated by reporting $m = 0$ and paying $t(0) = 0$.

Proposition 2 *Let $\delta = 0$. In an optimal equilibrium, $d(m) = m$ and*

$$m(\theta) = \begin{cases} \mathbb{E}[\theta] & \text{if } d'_{FB}(\theta) = \alpha_R + \alpha_S a \leq \frac{1}{2}, \\ \theta & \text{otherwise.} \end{cases}$$

In this myopic setting, whenever players' preferences do not coincide ($(a, b) \neq (1, 0)$), the receiver's decision-making is ex-post inefficient, even if she is fully informed: she chooses her own preferred decision over the first-best decision.

The optimal message rule should induce an equilibrium decision outcome that approximates the first-best decision rule as closely as possible. To understand how the message rule shapes decision-making, first notice that any message rule must induce, in equilibrium, receiver's beliefs that are correct in expectation, so that the average decision induced by any message rule is simply the expectation of the state: $\mathbb{E}[d_R(\theta)] = \mathbb{E}[\theta]$. Consequently, every message rule always induces the same expected difference between the equilibrium outcome and the first-best decision rule: $\mathbb{E}[d_R(m(\theta)) - d_{FB}(\theta)] = \mathbb{E}[\theta] - d_{FB}(\mathbb{E}[\theta])$. Given this constraint, optimizing the message rule involves matching the 'slopes' of the equilibrium decision outcome and the first-best decision rule as closely as possible.

Under full separation, the receiver chooses her preferred decision ($d(\theta) = \theta$), which as a function of the state has slope equal to 1. On the other hand, under complete pooling, the receiver always chooses a completely uninformed decision ($d(\theta) = \mathbb{E}[\theta]$) which as a function of the state has slope equal to 0. If the slope $d'_{FB}(\theta)$ of the first-best decision rule is greater than 1, then clearly the decision outcome under full separation is the best possible approximation of the first-best decision rule. If $d'_{FB}(\theta)$ is less than 1, then a tradeoff between separation and pooling arises. In this case, separation (pooling) produces decisions that are too sensitive (insensitive) to the state. When $d'_{FB}(\theta)$ is greater (less) than 1/2, the steep decision outcome induced by full separation is a better (worse) approximation of the first-best decision rule than the flat decision outcome induced by complete pooling.

4 Relational Communication

4.1 First Best

We now consider non-myopic players ($\delta > 0$). In this setting, repeated interactions endow the receiver with endogenous commitment power, so the equilibrium decision rule $d(m)$ can differ from the receiver's preferred decision rule $d_R(m)$.

Leading up to our main result, we first discuss the case where players are sufficiently patient, so that the first-best can be achieved: specifically, full separation occurs and first-best decisions are chosen ($m(\theta) = \theta$ and $d(m) = d_{FB}(m)$). The following proposition shows that there is a simple self-enforcement constraint that gives necessary and sufficient conditions for the first-best outcome to be an equilibrium. Given the set of equilibrium payoffs V, define the relational *leeway* as

$$\Delta \equiv \sqrt{\frac{\delta}{1-\delta}\frac{\overline{v} - \underline{v}_S - \underline{v}_R}{\alpha_R}}. \tag{2}$$

Proposition 3 *An optimal equilibrium achieves the first-best outcome if and only if*

$$|d_{FB}(\theta) - \theta| \leq \Delta \textit{ for all } \theta \in [0, 1). \tag{3}$$

In repeated settings with transferable utility, the set of implementable outcomes can be characterized by a single self-enforcement constraint (Levin 2003). The gist of this self-enforcement constraint is that the total reneging temptation, aggregated over all players, can never exceed the surplus from the relationship.

In the context of first-best outcomes, Proposition 3 shows that the aggregate reneging temptation consists solely of the receiver's short-run temptation to deviate from the first-best decision $d_{FB}(\theta)$ to her preferred decision $d_R(\theta) = \theta$. The self-enforcement constraint can thus be stated as follows: for each state θ, the distance between the first-best decision $d_{FB}(\theta)$ and the receiver's preferred decision $d_R(\theta)$ cannot exceed the relational leeway (see Figure 2a).

Importantly, the sender's temptation to deviate from a fully-separating message rule does not contribute to the self-enforcement constraint. This is not entirely surprising in light of Section 3.2, where we showed that the sender can use interim transfers to endogenously commit to any monotone message rule without relying on the threat of future punishment.

4.2 Second Best

We now present the first main result of the paper: a characterization of the second-best optimal equilibrium, when the first-best outcome is not achievable. It focuses on the case where the sender is upwardly biased, in the sense that his preferred decision is larger than the receiver's preferred decision in every state. (Other cases produce qualitatively similar insights and are briefly discussed at the end of this section.)

Proposition 4 *Suppose the sender is upwardly biased: $a\theta + b \geq \theta$ for all θ. Suppose also that $\delta \in (0, \delta_{FB})$ where δ_{FB} is the minimum δ that satisfies (3). In an optimal equilibrium, the decision rule is*

$$d(m) = \begin{cases} d_{FB}(m) & \text{if } |d_{FB}(m) - m| \leq \Delta, \\ m + \Delta & \text{otherwise,} \end{cases} \tag{4}$$

where Δ is given by (2). Further, if $d'_{FB}(\theta) = \alpha_R + \alpha_S a \geq 1/2$, then the message rule is $m(\theta) = \theta$; otherwise, the message rule is

$$m(\theta) = \begin{cases} \mathbb{E}[\theta | \theta \leq \theta^*] & \text{if } \theta \leq \theta^*, \\ \theta & \text{if } \theta > \theta^*, \end{cases} \tag{5}$$

for some $\theta^* \geq \widehat{\theta} = \max\{\theta \in [0,1] : d_{FB}(\theta) \geq \theta + \Delta\}$; further, $\theta^* > \widehat{\theta}$ if $\widehat{\theta} < 1$.

Just as in the first-best equilibrium, the equilibrium decisions cannot be too far from the receiver's preferred decisions. In contrast to the first-best equilibrium, however, the receiver may not be fully informed in the second-best equilibrium, so her temptation is to deviate to her preferred decision $d_R(m) = m$ conditional on her available information. The self-enforcement constraint thus requires that

$$|d(m) - m| \leq \Delta \text{ for all messages } m. \tag{6}$$

The decision-making constraint (6) may not be sufficient for self-enforcement if the sender is tempted to deviate from the message rule $m(\theta)$. In the case of the second-best equilibrium, this potential concern is avoided. In the proof of Proposition 4, we consider a relaxed problem in which a social planner can choose any message and decision rules to maximize joint surplus subject only to the decision-making constraint (6). We show that $d(m)$ and $m(\theta)$ given by (4) and (5) solve this problem. Since $d(m)$ is strictly increasing and $m(\theta)$ is monotone, parallel to Lemma 1, the sender can use interim transfers to credibly implement $m(\theta)$ without tightening the self-enforcement constraint. So, the solutions (4) and (5) of the relaxed problem constitute an equilibrium and thus achieve the second-best.

Geometrically, the second-best decision rule (4) pushes $d(m)$ as close to $d_{FB}(m)$ as possible, while keeping it within Δ distance from the receiver's preferred decision rule $d_R(m) = m$.

Just as in the myopic setting, if the first-best decision rule is relatively steep ($d'_{FB}(\theta) > 1/2$), then the second-best message rule involves full separation, so that $m(\theta) = \theta$ (see Figure 2b). For high states, the conflict of interest is within the relational leeway ($|d_{FB}(\theta) -$

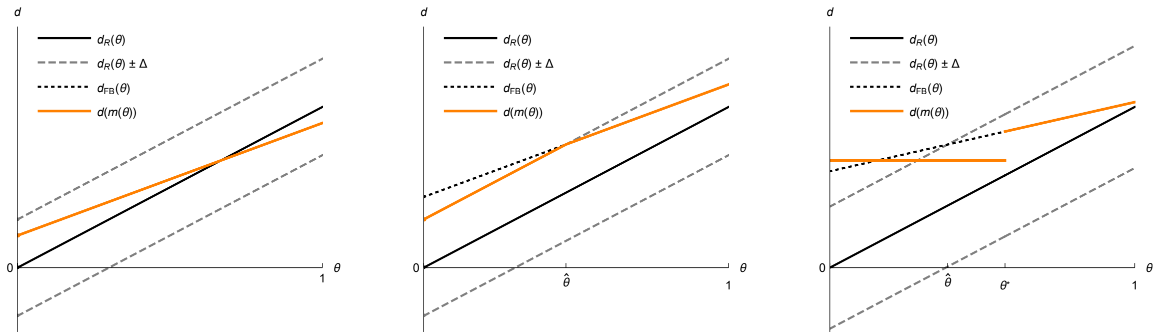

(a) First-best (b) $d'_{FB}(\theta) > 1/2$ (c) $d'_{FB}(\theta) < 1/2$

Figure 2: Second-best outcomes with upwardly-biased sender

$d_R(\theta)| < \Delta$), so the first-best decision is successfully implemented: $d(m(\theta)) = d_{FB}(\theta)$. For low states, the conflict of interest is so large that the first-best decision cannot be implemented. Instead receiver's decision-making is only partially disciplined, so that $d(m(\theta)) = d_R(\theta) + \Delta$.

Just as in the myopic setting, if the first-best decision rule is relatively flat ($d'_{FB}(\theta) < 1/2$), then the second-best message rule involves pooling. But, pooling may be incomplete, in which case it takes the following form (see Figure 2c). In states above some threshold θ^*, the state is fully separated and the first-best rule is implemented. All states below θ^* are pooled.

To understand the benefits of pooling low states when $d'_{FB}(\theta) < 1/2$, consider (as in the case $d'_{FB}(\theta) > 1/2$) a fully separating equilibrium that obeys the self-enforcement constraint (4). Decision-making would be first-best efficient for high states ($\theta > \hat{\theta}$), but constrained by (4) for low states ($\theta < \hat{\theta}$). In particular, for $\theta < \hat{\theta}$ the decision outcome $d(\theta) = \theta + \Delta$ would run parallel to the receiver's preferred decision outcome $d(\theta) = \theta$) and thus would have slope of 1; this is 'too steep' compared to the first-best decision rule. Then, as in the myopic setting, decision-making can be improved by pooling states for $\theta < \hat{\theta}$ so as to 'flatten' the decision outcome.

Interestingly, the optimal message rule pools a larger set of states than $[0, \hat{\theta}]$ to relax the (otherwise binding) self-enforcement constraint in those states. Consider the effects of expanding the pooling interval from $[0, \hat{\theta}]$ to $[0, \hat{\theta} + d\hat{\theta}]$. The benefit of such an expansion is that decision-making on $[0, \hat{\theta}]$ improves. Because the receiver's preferred decision is increasing in the expected state and $\mathbb{E}[\theta \,|\, \theta \leq \hat{\theta}] + \Delta$ is increasing in $\hat{\theta}$, adding higher states to the pool allows an increase in the constrained decision for the original pooled states towards the first-best decision (see Figure 2c). The cost of such an expansion is that the

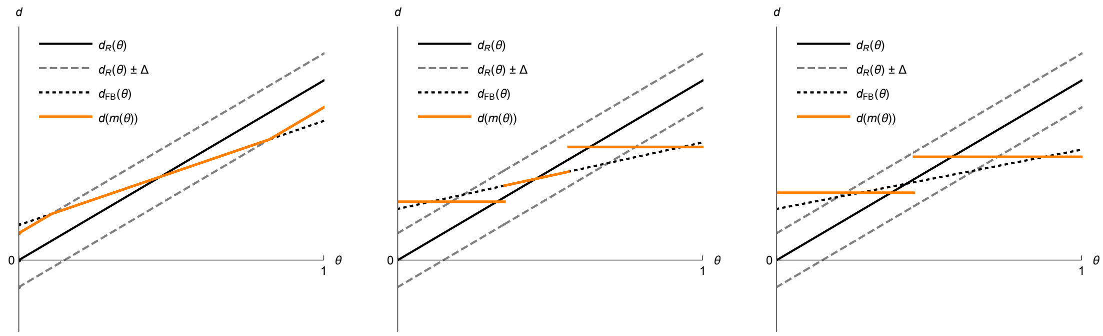

(a) Full separation (b) Pooling at extreme states (c) No separation in middle

Figure 3: Second-best outcomes when bias switches sign

newly-added states $[\hat{\theta}, \hat{\theta}+d\hat{\theta}]$ switch from the first-best decision to the constrained decision. Both the benefit and cost are first-order in $d\hat{\theta}$; nonetheless it turns out that the marginal benefit outweighs the marginal cost at $\hat{\theta}$ whenever $d'_{FB}(\theta) < 1/2$. The optimal threshold θ^*, where the marginal benefit equals the marginal cost, is thus greater than $\hat{\theta}$.

The analysis also admits some intuitive comparative statics results. For example, holding $d_R(\theta)$ and $d_{FB}(\theta)$ fixed, as the relational leeway Δ increases, the pooling interval decreases because it becomes easier to directly discipline the receiver's decision-making. Also, the benefits of pooling decrease in the sensitivity of the relative bias to the state $(d'_{FB}(\theta) - d'_R(\theta))$. In particular, if we vary $\beta \equiv d'_{FB}(\theta)$ by rotating $d_{FB}(\theta)$ around $\hat{\theta}$, while holding Δ and $d_R(\theta)$ fixed, we get $\theta^* \to \hat{\theta}$ as $\beta \to 1/2$, and $\partial\theta^*/\partial\beta < 0$.

The key insight from Proposition 4 is that pooling occurs in states where the conflict of interest $|d_{FB}(\theta) - d_R(\theta)|$ is large and the receiver is prone to overreaction ($d'_{FB}(\theta) < 1/2$, or equivalently $d'_R(\theta) > 2d'_{FB}(\theta)$). Under the assumption that the sender is upwardly-biased, if the receiver is prone to overreaction, then extreme conflict occurs in low states (see Figure 2c).

This insight does not rely on the upward bias of the sender. The case of a downwardly-biased sender is symmetric and pooling can occur only in high states. In the case where the sender's bias switches sign at some state, extreme conflict can occur at both high and low states. If the receiver is not prone to overreaction, $d'_R(\theta) < 2d'_{FB}(\theta)$, then full separation is optimal and the self-enforcement constraint binds for states with extreme conflict (see Figure 3a). If the receiver is prone to overreaction, $d'_R(\theta) > 2d'_{FB}(\theta)$, then states with extreme conflict are pooled, for the same reasons as in Proposition 4. Intermediate states, where conflict of interest is moderate, may or may not be separated (see Figures 3b and 3c).

5 The Value of Asymmetric Information

Now, we augment the model so that at the start of each period, the receiver observes a public state-dependent signal realization σ. Just as with message rules, we restrict attention (without loss) to direct signal rules such that $\sigma = \mathbb{E}[\theta|\sigma]$ for any realization σ. We assume that the signal rule σ is monotone in the sense that $\sigma(\theta)$ is non-decreasing in θ.

We say that σ is *more informative* than σ' if for each realization of s' of σ' there exists a realization s of σ such that the set $\sigma^{-1}(s)$ is a subset of $\sigma'^{-1}(s')$. For monotone signal rules, this notion coincides with the informativeness criterion of Blackwell (1953). Signal rule σ is *strictly* more informative than σ' if σ is more informative than σ' and the set of states where $\sigma(\theta) \neq \sigma'(\theta)$ has a strictly positive Lebesgue measure.

In this section, we focus on *monotone* equilibria. An equilibrium is monotone if after any history (on and off equilibrium path), the sender's message rule $m(\theta)$ is monotone and the receiver's decision rule $d(m)$ is strictly increasing. We discuss to what extent our results extend to non-monotone equilibria in Appendix C.

The second main result of the paper is that asymmetric information improves the relationship. Specifically, the best joint equilibrium surplus $\overline{v}$ strictly increases as the public signal becomes less informative (see Figure 4).

Proposition 5 *Let σ and σ' be signal rules with corresponding monotone equilibrium payoff sets V and V'. If σ is strictly more informative than σ', then $V \subsetneq V'$. If, in addition, $\delta > 0$ and the first-best outcome is not attainable under σ, then V lies in the interior of V'.*

To build intuition for this result, we start with the myopic benchmark. We will argue that the set of equilibrium payoffs V expands when moving from a fully informative public signal ($\sigma_f(\theta) = \theta$) to a completely uninformative public signal ($\sigma_u(\theta) = \mathbb{E}[\theta]$). Specifically, $\underline{v}_S$ and $\underline{v}_R$ strictly decrease and $\bar{v}$ weakly increases.

The receiver's worst equilibrium payoff $\underline{v}_R$ is lower under σ_u than σ_f. In the receiver's worst equilibrium, the receiver always chooses her preferred decision $d_R(\sigma)$ given the public signal σ 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 σ_u than σ_f. The basic idea is that any monotone equilibrium decision outcome implemented under σ_f (and thus a fully informed receiver) can also be implemented under σ_u by inducing the sender to fully reveal the state to the receiver. The sender's payoff $\underline{v}_S$ is strictly smaller under σ_u because inducing full separation requires the sender to make positive interim transfers to the receiver.

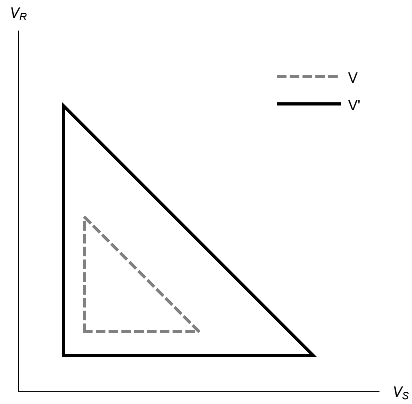

Figure 4: Equilibrium Payoff Sets

Joint surplus $\bar{v}$ is weakly higher under σ_u than σ_f, again because any monotone equilibrium under σ_f can be implemented under σ_u. In fact, joint surplus may be strictly higher under σ_u than σ_f. Under σ_u, the joint surplus is maximized under complete pooling of the states if the receiver's preferred decision is too steep (see Section 3.2). Such pooling, however, is precluded under σ_f (and thus a fully informed receiver).

In the non-myopic case, these effects are preserved, and the self-enforcement constraint produces a new effect. Moving from σ_f to σ_u expands V and thus increases the relational leeway Δ (which increases with $\bar{v}$ and decreases with $\underline{v}_S$ and $\underline{v}_R$). This in turn relaxes constraints on decision-making and expands the set V even further.

6 Discussion of Model Assumptions

We discuss some of our modelling assumptions and highlight the extent to which our results depend on these assumptions.

Non-Quadratic Payoffs

We have assumed that payoffs are quadratic. Instead, we now assume that payoff functions satisfy Crawford and Sobel (1982)'s assumptions.

Assumption 1 *For $i \in \{S, R\}$, $\partial^2 u_i(d,\theta)/\partial d^2 < 0$ and $\partial u_i(d,\theta)/\partial d = 0$ for some $d = d_i(\theta)$, so that u_i is concave and has a unique maximum in d for each θ. Moreover, for $i \in \{S, R\}$, $\partial^2 u_i(d,\theta)/\partial d\,\partial\theta > 0$, so that the sender's and receiver's preferred decisions $d_S(\theta)$ and $d_R(\theta)$ are strictly increasing in θ.*

Lemma 1 continues to hold: a message rule $m(\theta)$ is supported in some static equilibrium if and only if $m(\theta)$ is non-decreasing in θ. Therefore, optimal static equilibria still solve the problem of a social planner who wants to maximize joint surplus and can choose any monotone message rule. In contrast to Proposition 2, the optimal static message rule may involve partial pooling.

Similarly to Proposition 3, an optimal equilibrium achieves the first-best outcome ($m(\theta) = \theta$ and $d_{FB}(m) = \arg\max_d\{u_S(d,m) + u_R(d,m)\}$) if and only if

$$u_R(d_R(\theta),\theta) - u_R(d(\theta),\theta) \leq \frac{\delta}{1-\delta}(\overline{v} - \underline{v}_S - \underline{v}_R) \text{ for all } \theta \in [0,1)\,.$$

Proposition 5 continues to hold as well: public information hurts the relationship.

The second-best message rule is harder to characterize than in Proposition 4, because we have to work with infinite-dimensional messages $m = \mu(.|m)$, where $\mu(.|m)$ denotes the posterior distribution of θ given m, instead of one-dimensional messages $m = \mathbb{E}[\theta|m]$. Nevertheless, similarly to Proposition 4, we may show that the second-best message rule involves pooling of states where preferences are sufficiently divergent.

We may recover much of the structure of the second-best message rule if we further assume that $u_i(d,\theta) = \gamma_i(\theta) + \kappa_i(d) + \lambda_i(d)\,\theta$, so that expected payoffs depend on m only through $\mathbb{E}[\theta|m]$.[19] This assumption ensures that, without loss, we can restrict attention to message rules satisfying $m = \mathbb{E}[\theta|m]$; see footnote 15 and the preceding discussion.

A (weaker) version of Proposition 4 holds under this additional assumption. Suppose that the sender is upwardly-biased. Denote the relational leeway as the maximum value $\Delta(\theta)$ that satisfies

$$u_R(d_R(\theta),\theta) - u_R(d_R(\theta) + \Delta(\theta),\theta) \leq \frac{\delta}{1-\delta}(\overline{v} - \underline{v}_S - \underline{v}_R).$$

Parallel to Section 4.2, assume that there exists a threshold state $\widehat{\theta}$ such that $d_{FB}(\theta) \leq d_R(\theta)+\Delta(\theta)$ if and only if $\theta \geq \widehat{\theta}$, and that $G(m) = \sum_i \kappa_i(d(m))+\lambda_i(d(m))\,m$ is either convex or concave on $[0,\widehat{\theta}]$. Then the second-best equilibrium is as follows. First, $d(m) = d_{FB}(m)$ if $m \geq \widehat{\theta}$ and $d(m) = d_R(m) + \Delta(m)$ otherwise. The second-best message rule $m(\theta)$ maximizes $\mathbb{E}_\theta[G(m(\theta))]$. G is convex on $[\widehat{\theta},1]$, being an upper envelope of linear functions. Consequently, if G is also convex on $[0,\widehat{\theta}]$, then $m(\theta) = \theta$, but if G is concave on $[0,\widehat{\theta}]$, then $m(\theta)$ is given by (5).[20]

[19] Quadratic payoffs satisfy this requirement. This requirement is also satisfied if the state is binary.

[20] However, if G is neither convex nor concave on $[0,\widehat{\theta}]$, the second-best message rule may be non-monotone, in which case the self-enforcement constraint has to also account for the sender's temptation to deviate from the optimal message rule, so that $d(m)$ may differ from $d_R(m) + \Delta(m)$ on $[0,\widehat{\theta}]$.

Allocation of Authority

We have assumed that decision-making authority always resides with the receiver and is not transferable (*receiver-authority*). Some papers explore the tradeoff between allocating authority to an uninformed receiver versus an informed but biased sender in organizational settings (e.g., Dessein 2002, Alonso, Dessein and Matouschek 2008, and Rantakari 2008).[21]

In this vein, consider a variation where the sender chooses the decision instead of the receiver; call this variation *sender-authority*. Focus on the case $\alpha_S = \alpha_R$, so that the sender has the same temptation to renege on the first-best decision under sender-authority as the receiver has under receiver-authority. In this case, full separation is always optimal under receiver-authority. It turns out that allocating decision authority to the sender strictly decreases the best joint equilibrium surplus. This is because the worst joint equilibrium surplus is strictly higher,[22] and thus the self-enforcement constraint is strictly tighter, under sender-authority. Importantly, the standard rationale for delegation – that the better-informed player can more effectively adapt the decision – no longer applies in our setting because interim transfers can credibly achieve arbitrary communication outcomes at no welfare cost. This implies that all our results continue to hold even if decision-making authority could be allocated to either player at the beginning of the game, because the players would always choose receiver-authority over sender-authority.

When $\alpha_S \neq \alpha_R$, the comparison between sender- and receiver-authority is more nuanced. For example, if $\alpha_S > \alpha_R$, then in addition to the effect identified above, which favours receiver-authority, two additional effects emerge, the first (second) of which favours sender-authority (receiver-authority). First, the decision-maker's temptation to renege on the first-best decision is larger under receiver-authority than under sender-authority. Second, under receiver-authority, the optimal equilibrium may involve pooling to discipline decision-making; this tool is unavailable under sender-authority.

[21] Relatedly, Holmstrom (1984), Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), Mylovanov (2008), Goltsman et al. (2009), Kováč and Mylovanov (2009), and Amador and Bagwell (2013) study the optimal delegation problem. Lim (2012) allows authority to be allocated after the sender observes the state.

[22] Under sender-authority, $\underline{v}_S = 0$ because the sender can always choose his preferred decision in each state, whereas $\underline{v}_R = -\alpha_R \mathbb{E}\left[(a\theta + b + \Delta - \theta)^2\right]$ because $a\theta + b + \Delta$ is the worst possible decision for the receiver that is self-enforcing for the (upwardly-biased) sender. On the other hand, under receiver-authority, $\underline{v}_R = -\alpha_R \mathrm{Var}[\theta]$ because the receiver can always choose the uninformed decision $d = \mathbb{E}[\theta]$, whereas $\underline{v}_S \leq -\alpha_S \mathbb{E}\left[(\theta - \Delta - (a\theta + b))^2\right] = -\alpha_R \mathbb{E}\left[(a\theta + b + \Delta - \theta)^2\right]$ because full separation, non-negative interim transfers by the sender, and decision $\theta - \Delta$ can always be achieved in equilibrium. Therefore, the worst joint equilibrium surplus is strictly higher under sender-authority.

Consider another variation where decision-making authority is allocated at the beginning of each period (*short-term-authority*). Specifically, following Baker, Gibbons and Murphy (2011), suppose that at the beginning of each period, the receiver has decision-making authority by default, and can make a take-it-or-leave-it offer to transfer authority to the sender for that period in exchange for a transfer payment. As above, focus on the case $\alpha_S = \alpha_R$. We know from above that the best (worst) joint equilibrium surplus is higher (lower) under receiver-authority than under sender-authority. This implies that relative to receiver-authority, short-term-authority does not improve on the best joint equilibrium surplus (because the players cannot do better than to allocate authority to the receiver in each period), but increases the worst joint equilibrium surplus (because the players always have the option to allocate authority to the sender in each period). This then implies that the relational leeway, and thus the best joint equilibrium surplus, is strictly lower under short-term-authority than under receiver-authority.

Exogenous Outside Options

We follow Abreu (1988) and Abreu, Pearce and Stacchetti (1990) in characterizing the entire set of equilibrium payoffs. In particular, optimal equilibria utilize the worst possible equilibria as off-path punishments. One alternative modeling approach taken by Levin (2003) is to specify exogenous outside option payoffs $\underline{u}_S$ and $\underline{u}_R$ for both players; so that players are punished for deviations by receiving their outside option payoffs thereafter. In this approach, at the beginning of each period, the receiver makes an offer to the sender consisting of a contractible commitment to an ex-ante transfer payment. If the sender rejects this offer, the players receive their outside option payoffs, and time moves on to the next period. Another alternative modeling approach taken by Baker, Gibbons and Murphy (1994, 2011) is to restrict attention to equilibria where off-path punishments correspond to some static equilibria.

Propositions 1 – 5 continue to hold in these settings, with the worst equilibrium payoffs equal to either the outside option payoffs or the static equilibrium payoffs.[23] Moreover, in Proposition 5, we no longer have to restrict attention to monotone equilibria, because (i) static equilibria (which are used as punishments) are monotone by Lemma 2 and (ii) optimal equilibrium play is monotone by Proposition 4.

[23] In the setting with exogenous outside options, one of the two economic mechanisms driving Proposition 5 vanishes: public information no longer affects the players' worst possible equilibrium payoffs.

Imperfect Monitoring

We have also assumed perfect monitoring in that the sender's private information and the receiver's decision are public information by the end of each period. First, consider a variation where the sender's type is imperfectly monitored. Specifically, suppose that in each period, the sender privately observes his type $\eta \in [0,1)$ and a public signal $\nu \in [0,1)$ is realized at the end of the period. Suppose that η and ν have some joint distribution so that the marginal distributions of η and ν admit strictly positive densities on $[0,1)$. The sender's and receiver's payoffs are exactly the same as before, except that θ is replaced with ν. So, our model from Section 2 simply corresponds to the case $\theta = \eta = \nu$.

In this setting, the sender's message rule conditions on η. Because the players' payoffs are quadratic, normalizations $\mathbb{E}[\,\nu\,|\,\eta\,] = \eta$ and $\mathbb{E}[\,\eta\,|\,m\,] = m$ are without loss and the following equalities hold:

$$\begin{aligned}\mathbb{E}_\nu\left[u_i(d,\nu)|m\right] &= \mathbb{E}_\eta[u_i(d,\eta) + \tilde{u}_i(\eta)|m] \\ &= u_i(d,m) + \hat{u}_i(m),\end{aligned}$$

where $\tilde{u}_i(\cdot)$ and $\hat{u}_i(\cdot)$ are terms that do not depend on d. This implies that the players' optimization problems remain unchanged. As we discussed in Section 4, the sender's temptation to deviate from the equilibrium message rule, and thus the ex-post observability of the sender's private information, does not contribute to the self-enforcement constraint. Combining these observations, we conclude that Propositions 1 – 5 continue to hold in this variation.[24]

Second, consider a variation in which the receiver's decision is imperfectly monitored. Specifically, suppose that the receiver's (private) decision d stochastically determines an *output* $y = d + \varepsilon$ which is publicly observed and replaces d as an argument in the players' payoff functions: $u_R(y,\theta) = -\alpha_R(y-\theta)^2$ and $u_S(y,\theta) = -\alpha_S(y - a\theta - b)^2$. Assume that $\mathbb{E}[\varepsilon] = 0$ and that the density g of ε satisfies the appropriate Mirrlees-Rogerson conditions (Rogerson, 1985),[25] ensuring that the receiver's decision choice can be represented by a first-order condition.

Focusing on the case where the sender is upwardly-biased, consider the highest decision $\overline{d}(m)$ that can be supported in equilibrium. Parallel to Theorem 6 of Levin (2003), $\overline{d}(m)$

[24] We also have to make a few adjustments in calculating the values of Δ, V, θ^*, and $\hat{\theta}$. In particular, these values depend on the distribution of (ν,η).

[25] For example, these conditions are satisfied if $g(x)$ is log-concave in x (monotone likelihood ratio property) and $g(q^{-1}(x))/q'(q^{-1}(x))$ is increasing in x, where $q(d) = \mathbb{E}\left[-\alpha_R(d+\varepsilon-m)^2\right]$.

can be implemented by the strongest 'one-step' incentive scheme that satisfies the self-enforcement constraint: this scheme may take the form of stationary continuation values $v_S(\theta) = \underline{v}_S$, $v_R(\theta) = \overline{v} - \underline{v}_S$ and the ex-post transfer rule

$$T_R(y) = -T_S(y) = \begin{cases} 0 & \text{if } y \geq \overline{d}(m) + \varepsilon^*, \\ \frac{\delta}{1-\delta}(\overline{v} - \underline{v}_S - \underline{v}_R) & \text{if } y < \overline{d}(m) + \varepsilon^* \end{cases}$$

where ε^* is the point where g' switches from negative to positive. The receiver's (unobserved) decision then satisfies the first-order condition

$$\frac{\partial}{\partial d} \mathbb{E}\left[-\alpha_R(d + \varepsilon - m)^2 - T_R(d + \varepsilon))\right]\Big|_{d=\overline{d}(m)} = 0,$$

which simplifies to $\overline{d}(m) = m + \Delta$ for some $\Delta > 0$. In other words, the self-enforcement constraint effectively specifies that the equilibrium decision cannot exceed the receiver's preferred decision by more than the leeway Δ. (As before, the sender's incentive problem does not contribute to the self-enforcement constraint.) Consequently, retracing the steps of our analysis, Propositions 1 – 5 continue to hold in this variation.

Correlated States

We have assumed that the state θ_t is i.i.d. Consider a variation where θ_t is correlated across periods. Specifically, introduce a finite-valued random variable ω_t that is publicly observed at the beginning of each period t (before ex-ante transfers are made), where ω_t is a Markov chain. The realization of ω_t fully determines the (time-independent) distribution $F(\cdot|\omega_t)$ of the state θ_t. Crucially, given ω_t, θ_t contains no further information about ω_{t+1} (and thus about θ_{t+1}); so the sender and receiver are always symmetrically informed about the future distribution of states. Note that this property would no longer hold if we did not introduce ω_t, but simply assumed that θ_t was a Markov chain.

With this modification, without loss of generality, we can restrict attention to equilibria that are stationary conditional on ω_t (Kwon 2016). Consequently, Propositions 1 – 5 continue to hold verbatim, except that the key objects such as V, Δ, and θ^* are now functions of ω_t.

7 Conclusion

In our model, incomplete information transmission does not reflect communication failure, but instead is an instrument for managing decision-making. This insight 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 non-monetary transfers (such as wages or favours) 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 an ex-ante optimal message rule. This is precisely 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-bias states, where conflict of interest is high and self-enforcement constraints are binding. One implication is that with constant bias, pooling does not occur. In contrast, in the standard constant-bias Crawford and Sobel (1982) framework, information transmission is always incomplete, and this is exacerbated in high (low) states if the sender is upwardly (downwardly) biased.

In our model, an 'arms-length' approach with separation of information and control benefits the relationship. This finding suggests that organizations may seek to discourage information acquisition by decision-makers, and endow other interested parties with relevant information. A related implication is that mediators who control the flow of information from sender to 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 by the sender and receiver.

Appendix A Optimality and Stationarity

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

Notice that the set V of equilibrium payoffs is compact by continuity and Tychonoff's theorem (see Abreu 1988). Thus, there exist worst and best equilibrium payoffs $\underline{v}_i$ and $\bar{v}_i$ for each player $i \in \{S, R\}$ and a best joint equilibrium surplus $\overline{v}$ which maximizes $v_S + v_R$ over $(v_S, v_R) \in V$. By Abreu (1988), $(v_S, v_R) \in V$ if and only if there exist (admissible) functions τ_i, $m(\theta)$, $t_i(m)$, $d(m)$, $T_i(\theta)$, $v_S(\theta)$, $v_R(\theta)$, and (punishment) variables $\theta^p \in [0, 1]$, $d^p \in \mathbb{R}$, such that the following seven conditions hold:

C1. Both players are willing to make the ex-ante transfer payment τ:

$$v_S \equiv (1-\delta)[-\tau_S + \mathbb{E}[u_S(d(m(\theta)), \theta) - t_S(m(\theta)) - T_S(\theta)] + \delta\mathbb{E}[v_S(\theta)] \geq \underline{v}_S$$

$$v_R \equiv (1-\delta)[-\tau_R + \mathbb{E}[u_R(d(m(\theta)), \theta) - t_R(m(\theta)) - T_R(\theta)] + \delta\mathbb{E}[v_R(\theta)] \geq \underline{v}_R.$$

C2. For every state θ, the sender is willing to send message $m(\theta)$ and to make interim transfer payment $t_S(m(\theta))$. Specifically,

(a) There is no profitable deviation to another message – interim-transfer pair $(m(\theta'), t_S(m(\theta')))$ that is observed on the equilibrium path:

$$\begin{aligned} & (1-\delta)[u_S(d(m(\theta)), \theta) - t_S(m(\theta)) - T_S(\theta)] + \delta v_S(\theta) \\ \geq\ & (1-\delta)[u_S(d(m(\theta')), \theta) - t_S(m(\theta'))] + \delta\underline{v}_S \text{ for all } \theta, \theta'. \end{aligned}$$

(It is without loss of generality to let t_S depend on $m(\theta)$ but not directly on θ; since the sender makes his interim transfer choice before the receiver, we can always modify $m(\theta)$ to incorporate any additional information contained in t_S without changing the receiver's information set.)

(b) There is no profitable deviation to some pair (m', t'_S) that is never observed on the equilibrium path:

$$(1-\delta)[u_S(d(m(\theta)), \theta) - t_S(m(\theta)) - T_S(\theta)] + \delta v_S(\theta) \geq (1-\delta)[u_S(d^p, \theta)] + \delta\underline{v}_S \text{ for all } \theta.$$

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

C3. For any message m, the receiver is willing to make interim transfer payment $t_R(m)$:

$$(1-\delta)[-t_R(m) + \mathbb{E}[u_R(d(m), \theta) - T_R(\theta)|m]] + \delta\mathbb{E}[v_R(\theta)|m] \geq (1-\delta)\mathbb{E}[u_R(d', \theta)|m] + \delta\underline{v}_R \text{ for all } m, d'.$$

C4. The receiver is willing to choose decisions $d(m)$ (on-path) and d^p (off-path):

$$\begin{aligned}(1-\delta)\mathbb{E}[u_R(d(m),\theta)-T_R(\theta)|m]+\delta\mathbb{E}[v_R(\theta)|m] &\geq (1-\delta)\mathbb{E}[u_R(d',\theta)|m]+\delta\underline{v}_R \text{ for all } m,d'; \\ (1-\delta)u_R(d^p,\theta^p)+\delta\bar{v}_R &\geq (1-\delta)u_R(d',\theta^p)+\delta\underline{v}_R \text{ for all } d'.\end{aligned}$$

Here, we specify that following any deviation by the sender, the receiver believes that $\theta=\theta^p$.

C5. The players are willing to make ex-post transfer payments $T_i(\theta)$:

$$-(1-\delta)T_S(\theta)+\delta v_S(\theta)\geq\delta\underline{v}_S;$$

$$-(1-\delta)T_R(\theta)+\delta v_R(\theta)\geq\delta\underline{v}_R.$$

C6. The continuation payoffs are admissible:

$$(v_S(\theta),v_R(\theta))\in V \text{ for all } \theta.$$

C7. There is no creation of money:

$$\begin{aligned}\tau_S+\tau_R &\geq 0; \\ t_S(m)+t_R(m) &\geq 0 \text{ for all } m; \\ T_S(\theta)+T_R(\theta) &\geq 0 \text{ for all } \theta.\end{aligned}$$

Proof of Proposition 1. Consider an optimal equilibrium payoff vector (v_S^*,v_R^*) with $v_S^*+v_R^*=\bar{v}$, and let σ^* be an equilibrium supporting (v_S^*,v_R^*). Let (v_S,v_R) be any point in the simplex V defined by (1). Notice that we can modify σ^* to produce (v_S,v_R) by changing the ex-ante transfer payments from τ_i to $\tau_i+v_i^*-v_i$ for each $i\in\{S,R\}$. This modification affects only Conditions C1 and C7. Condition C1 still holds because $v_S\geq\underline{v}_S$ and $v_R\geq\underline{v}_R$ by definition of V. Condition C7 still holds because $v_S+v_R\leq v_S^*+v_R^*$, again by definition of V. Thus, the modified strategy profile is an equilibrium. Conversely, it is easy to see that any (v_S,v_R) not in V cannot be supported in equilibrium. We conclude that V is the set of equilibrium payoffs.

In any optimal equilibrium, continuation is optimal: (i) $v_S(\theta)+v_R(\theta)=\bar{v}$ for all θ, and (ii) money is not burned, i.e., the constraints of Condition C7 hold with equality. Otherwise, one could (i) increase $v_i(\theta)$ and (ii) decrease transfers τ_i, $t_R(m)$, and $T_i(\theta)$, thereby relaxing the constraints of Conditions C1–C5 and increasing joint surplus v_S+v_R.

An optimal equilibrium σ with zero first-period ex-ante transfers clearly exists. Let (v_S, v_R) be the payoff profile under σ. We will modify σ to construct an optimal stationary equilibrium with the same payoff profile. For each player $i \in \{S, R\}$, let $m(\theta)$, $t_i(m)$, $d(m)$, $T_i(\theta)$ and $v_i(\theta)$ be the message rule, interim transfer rule, decision rule, ex-post transfer rule and continuation payoff function in the first period on the equilibrium path of σ. Define $T_i^*(\theta)$ and T_i^p by

$$-(1-\delta)\, T_i^*(\theta) + \delta v_i = -(1-\delta) T_i(\theta) + \delta v_i(\theta),$$
$$-(1-\delta)\, T_i^p + \delta v_i = \delta \underline{v}_i.$$

Consider a stationary strategy profile σ^* specified as follows. First, first-period ex-ante transfers are zero. Second, on the equilibrium path, $\tau_i = 0$, $m(\theta)$, $t_i(m)$, $d(m)$, and $T_i^*(\theta)$ are played in each period. Third, following a deviation, punishment is implemented by specifying ex-post transfer T_i^p for deviating player i and $-T_i^p$ for the non-deviating player, then reverting to the equilibrium path in subsequent periods. By construction, the payoff profile under σ^* is 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. To see this, notice that $-(1-\delta)T_i^*(\theta)+\delta v_i$ replaces $-(1-\delta)T_i(\theta)+\delta v_i(\theta)$ and $-(1-\delta)\, T_i^p + \delta v_i$ replaces $\delta \underline{v}_i$ in the constraints of Conditions C1 – C5. Condition C6 holds because (v_S, v_R) belongs to V by supposition. Further, since $v_S + v_R = v_S(\theta) + v_R(\theta) = \overline{v}$ and $T_S(\theta) + T_R(\theta) = 0$ by optimality of σ, we have $T_S^*(\theta) + T_R^*(\theta) = 0$, so Condition C7 holds on the equilibrium path. Similarly, since the sum of ex-post transfers following a deviation by player i is $T_i^p + (-T_i^p) = 0$, Condition C7 holds in the continuation path following a deviation as well.

Finally, by appropriately modifying the first-period ex-ante transfer in σ^*, we can support any equilibrium payoffs in V. ■

Appendix B Myopic Benchmark

Before proving Lemma 1, we present the following result.

Lemma 2 *Let decision rule $d(m)$ be continuous and strictly increasing in m. Let message rule $m(\theta)$ be such that states in $[\zeta_i, \xi_i)$ are pooled in message m_i and states in $[\xi_i, \zeta_{i+1})$ are separated so that $m(\theta) = \theta$, where $\zeta_i < \xi_i \le \zeta_{i+1}$ for all $i \ge 1$. Moreover, let $v_S(\theta) \ge \underline{v}_S$,*

$v_R(\theta) \geq \underline{v}_R$, and $T_S(\theta) \leq 0$ for all θ. Define, iteratively over $i \geq 1$,

$$\begin{aligned} h_i^{pool} &= u_S\left(d\left(m_i\right), \zeta_i\right) - u_S\left(d\left(\zeta_i\right), \zeta_i\right) + h_{i-1}^{sep}(\zeta_i), &(7)\\ h_i^{sep}(m) &= \int_{\xi_i}^{m} \frac{\partial u_S\left(d\left(\vartheta\right), \vartheta\right)}{\partial d} d'\left(\vartheta\right) d\vartheta + u_S\left(d\left(\xi_i\right), \xi_i\right) - u_S\left(d\left(m_{i-1}\right), \xi_i\right) + h_i^{pool} &(8)\\ &\quad \text{for } m \in [\xi_i, \zeta_{i+1}], \end{aligned}$$

with initialization $h_0^{sep}(m) = \int_0^m \frac{\partial u_S(d(\vartheta),\vartheta)}{\partial d} d'\left(\vartheta\right) d\vartheta$. Then Conditions C2 and C3 are satisfied by the following interim transfer rule and punishment variables: $t_R\left(m\right) = -t_S\left(m\right)$ and $d^p = d\left(\theta^p\right)$, where (defining $\xi_0 = 0$)

$$\begin{aligned} h(m) &= \begin{cases} h_{i-1}^{sep}(m) \text{ if } m \in [\xi_{i-1}, \zeta_i), \\ h_i^{pool} \text{ if } m = m_i, \end{cases} &(9)\\ t_S\left(m\right) &= h(m) - \min_{\boldsymbol{m}} h(\boldsymbol{m}), &(10)\\ \theta^p &\in \arg\min_{\boldsymbol{m}} t_S\left(\boldsymbol{m}\right). &(11) \end{aligned}$$

Proof. Notice that our construction of $t_S(m)$ satisfies local incentive compatibility:

$$\begin{aligned} \frac{\partial}{\partial\vartheta}\left(u_S(d(\vartheta), \theta) - t_S(\vartheta)\right) &= 0 \text{ for } \vartheta \in (\xi_i, \zeta_{i+1}), \\ u_S(d(m_i), \theta) - t_S(m_i) &= u_S(d(\theta), \theta) - t_S(\theta) \text{ for } \theta \in \{\zeta_i, \xi_i\}. \end{aligned}$$

Consider the on-path constraint of Condition C2. To verify this condition, note that $T_S(\theta) \leq 0$, so $-T_S(\theta) + \delta v_S(\theta) \geq \delta \underline{v}_S$ for all θ; thus it is sufficient to show that for each $\theta \in [0, 1]$, within the message space defined by $m(\theta)$, $u_S(d(m(\vartheta)), \theta) - t_S(m(\vartheta))$ is maximized at $\vartheta = \theta$. This claim follows from the proof of Karamychev and Visser (2016)'s Proposition 1. (Although they assume that the sender is upwardly biased and $d(m) = m$, their proof remains valid without these assumptions.)

Further, consider the off-path constraint of Condition C2. The sender does not want to deviate to a message-transfer pair (m', t'_S) that is not observed on the equilibrium path; by doing so, he would induce $d^p = d\left(m\left(\theta^p\right)\right)$, which he could induce more cheaply on the equilibrium path with message $m(\theta^p) = \theta^p$ and zero interim transfer $t_S\left(m(\theta^p)\right) = 0$. (This point relies again on the fact that $-T_S(\theta) + \delta v_S(\theta) \geq \delta \underline{v}_S$.)

Therefore, Condition C2 holds. Finally, Condition C3 holds because $t_R(m) = -t_S(m) \leq 0$ by construction. ∎

Proof of Lemma 1. Although Lemma 2 considers monotone message rules with the restriction that all pooling intervals take the form $[\zeta_i, \xi_i)$, it holds for all monotone message rules. Moreover, it holds if the payoff functions satisfy Assumption 1 (as in Crawford and Sobel 1982, Austen-Smith and Banks 2000, Kartik 2007, and Karamychev and Visser 2016).

The 'if' part of Lemma 1 then follows from Lemma 2. The 'only if' part of Lemma 1 follows from Austen-Smith and Banks (2000, p. 7). ■

Proof of Proposition 2. Given that $\delta = 0$, conditional on message m, the receiver chooses decision

$$d(m) = \arg\max_d \mathbb{E}_\theta \left[u_R(d, \theta) | m \right] = \mathbb{E}_\theta [\theta | m] = m.$$

Consider the relaxed problem of finding the message rule that maximizes joint surplus

$$\begin{aligned} & \mathbb{E}\left[\sum\nolimits_{i \in \{S,R\}} u_i \left(d\left(m(\theta)\right), \theta \right)\right] \\ = \; & \mathbb{E}_m \left[d(m) \left(2 d_{FB}(m) - d(m) \right) \right] - \mathbb{E}\left[\alpha_R \theta^2 + \alpha_S (a\theta + b)^2 \right] \\ = \; & \mathbb{E}_m [\underbrace{2\left(\alpha_R + \alpha_S a - 1/2\right) m^2}_{A}] - \mathbb{E}[\underbrace{\alpha_R \theta^2 + \alpha_S (a\theta + b)^2 - 2\alpha_S b \mathbb{E}[\theta]}_{B}] \end{aligned}$$

The term B does not depend on m, whereas the term A is concave (convex) in m if $\alpha_R + \alpha_S a \leq 1/2$ ($\alpha_R + \alpha_S a \geq 1/2$). Then from Section V.A of Kamenica and Gentzkow (2011), it follows that complete pooling $m(\theta) = \mathbb{E}[\theta]$ (full separation $m(\theta) = \theta$) maximizes joint surplus. Both complete pooling and full separation correspond to monotone message rules, and thus, by Lemma 1, can be supported in equilibrium. ■

Appendix C Relational Communication

Lemma 3 *Fix decision rule $d(m)$, message rule $m(\theta)$, and punishment variables θ^p and $d^p = d(\theta^p)$. Suppose that*

$$|d(m) - m| \leq \Delta \text{ for all } m \in [0,1], \tag{12}$$

that the ex-ante transfer rule is $\tau_R = \tau_S = 0$, that the ex-post transfer rule is $T_R(\theta) = T_S(\theta) = 0$, and that continuation payoffs are $v_S(\theta) = \underline{v}_S$ and $v_R(\theta) = \overline{v} - \underline{v}_S$. Then Conditions C1, and C4 – C7 are satisfied.

Proof. Given the conditions of the current Lemma, the constraints of Condition C4 may be restated as

$$(1-\delta)\mathbb{E}[u_R(d(m), \theta) - u_R(m, \theta) | m] + \delta(\overline{v} - \underline{v}_S) \geq \delta \underline{v}_R,$$

which is simply a restatement of (12) and thus is clearly satisfied on and off the equilibrium path. With zero ex-ante and ex-post transfers, the constraints of Conditions C1 and C5 are trivially satisfied. Finally, Conditions C6 and C7 are immediately satisfied by our construction. ■

Proof of Proposition 3. Start with the observation that the first-best outcome involves $m = \theta$ and $d(m) = d_{FB}(m)$. To prove the "only if" part of the proposition, simply add the constraint of Condition C4 evaluated at $d' = \theta$,

$$(1-\delta)\left(-\alpha_R\left(d_{FB}(\theta)-\theta\right)^2 - T_R(\theta)\right) + \delta v_R(\theta) \geq \delta \underline{v}_R,$$

to the sender's constraint of Condition C5,

$$-(1-\delta)T_S(\theta) + \delta v_S(\theta) \geq \delta \underline{v}_S,$$

and apply the fact that in an optimal equilibrium, $T_S(\theta)+T_R(\theta) = 0$ and $v_S(\theta)+v_R(\theta) = \overline{v}$.

Conversely, to prove the 'if' part of the proposition, suppose that (3) holds; we will show that $m = \theta$ and $d(m) = d_{FB}(m)$ can be supported in equilibrium, i.e., that Conditions C1 – C7 can be satisfied with the appropriate choice of transfer functions, punishment variables and continuation values. Pick the interim transfer function and punishment variables specified in Lemma 2. Pick the ex-post transfer rule and continuation payoffs specified in Lemma 3. And, pick ex-ante transfer function $\tau_S = \tau_R = 0$. Remember that full separation corresponds to a monotone message rule, and note that the ex-post transfer rule and continuation payoffs specified in Lemma 3 satisfy the assumptions of Lemma 2. Then Lemma 2 ensures that Conditions C2 and C3 are satisfied. Also, (3) implies (12), so Lemma 3 ensures that the remaining conditions C1 and C4 – C7 are satisfied. ■

Proof of Proposition 4. Fix a message rule $m(\theta)$. Adding the constraint of Condition C4 evaluated at $d' = m$ and the expectation of the sender's constraint from Condition C5 gives the following necessary condition for equilibrium:

$$|d(m) - m| \leq \Delta. \tag{13}$$

Note that for given $m(\theta)$, decision rule $d(m)$ specified by (4) maximizes joint surplus (over all decision rules that condition only on m) subject to (13).

As in Proposition 2, we first consider a relaxed version of the problem, i.e., we derive the surplus-maximizing message rule given decision rule (4); we then show that this message

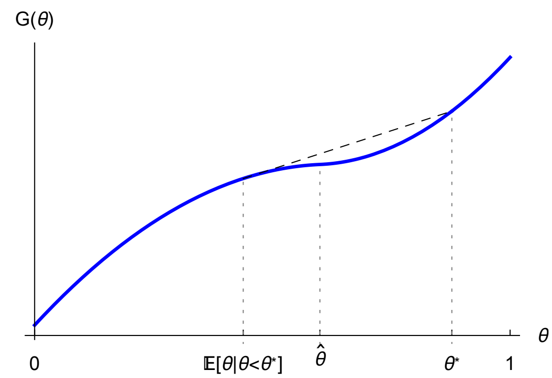

Figure 5: $G(\theta)$ with upwardly-biased sender and $\alpha_R + \alpha_S a < 1/2$.

rule can be supported in equilibrium. Start with the relaxed problem. Given $m(\theta)$, rewrite expected joint surplus as

$$\mathbb{E}\left[G(m(\theta))\right] - \mathbb{E}\left[H(\theta)\right]$$

where $G(m) = d(m)(2d_{FB}(m) - d(m))$ and $H(\theta) = \alpha_R\theta^2 + \alpha_S(a\theta + b)^2$. Now, given the decision rule $d(m)$ specified by (4), consider the corresponding $G(m)$. By using the fact that $d_{FB}\left(\widehat{\theta}\right) = d\left(\widehat{\theta}\right)$, we may check that $G'_-\left(\widehat{\theta}\right) = G'_+\left(\widehat{\theta}\right) = 2d_{FB}\left(\widehat{\theta}\right)d'_{FB}\left(\widehat{\theta}\right)$; so G is continuously differentiable. Morever, if $\alpha_R + \alpha_S a < 1/2$, then G is concave on $\left[0, \widehat{\theta}\right]$ and convex on $\left[\widehat{\theta}, 1\right]$, where $\widehat{\theta} > 0$ because $\delta \in (0, \delta_{FB})$. Then by Proposition 3 of Kolotilin (2016), the message rule $m(\theta)$ specified by (5) is optimal; further, either θ^* solves

$$G'\left(\mathbb{E}\left[\theta|\theta \le \theta^*\right]\right) = \frac{G(\theta^*) - G\left(\mathbb{E}\left[\theta|\theta \le \theta^*\right]\right)}{\theta^* - \mathbb{E}\left[\theta|\theta \le \theta^*\right]} \tag{14}$$

or $\theta^* = 1$ if (14) does not have a root in $(0, 1)$. Figure 5 illustrates this characterization of θ^*. It can be seen that θ^* is uniquely determined and satisfies $\theta^* > \widehat{\theta}$ if $\widehat{\theta} < 1$. On the other hand, if $\alpha_R + \alpha_S a \ge 1/2$, then G is convex on $[0, 1]$, in which case full separation $m(\theta) = \theta$ is optimal as in the proof of Proposition 2.

To show that message rule (5) and decision rule (4) can be supported in equilibrium, we construct transfer rules and punishment variables to satisfy Conditions C1 – C7. Note that in both cases $\alpha_R + \alpha_S a \ge 1/2$ and $\alpha_R + \alpha_S a < 1/2$, the message rule (5) is monotone. Then the construction, and the subsequent verification of Conditions C1 – C7, is replicated from the proof of Proposition 3: pick the interim transfer function and punishment variable specified in Lemma 2, and the ex-ante and ex-post transfer rules and continuation payoffs specified in Lemma 3. ■

Appendix D The Value of Asymmetric Information

Proof of Proposition 5. We prove Proposition 5 for monotone message rules such that states in $[\varsigma_i, \varrho_i)$ are pooled and states in $[\varrho_i, \varsigma_{i+1})$ are separated, where $\varsigma_i < \varrho_i \leq \varsigma_{i+1}$ for all $i \geq 1$. The argument generalizes easily to all monotone message rules.

Let the relational leeway under σ be $\Delta \geq 0$. Suppose, for the sake of argument, that the relational leeway under σ' is also Δ. We will show that worst equilibrium payoffs decrease, while best equilibrium payoffs increase: specifically, $\underline{v}'_R < \underline{v}_R$, $\underline{v}'_S \leq \underline{v}_S$, and $\overline{v}' \geq \overline{v}$. Equation (2) then implies that the relational leeway under σ' is $\Delta' \geq \Delta$, with strict inequality if $\delta > 0$. The proposition follows easily from this observation.

The receiver's worst equilibrium payoff is supported by the repetition of a static uninformative equilibrium, whereby the sender sends message $m(\theta) = \sigma(\theta)$ and the receiver chooses the decision $d(\sigma) = \mathbb{E}[\theta|\sigma]$ in each period and does not pay or receive any transfers. Notice that this equilibrium is monotone. Consequently,

$$\underline{v}'_R = -\alpha_R \mathbb{E}[\mathrm{Var}(\theta|\sigma')] < -\alpha_R \mathbb{E}[\mathrm{Var}(\theta|\sigma)] = \underline{v}_R,$$

where the inequality holds because σ is strictly more informative than σ'.

The best joint monotone equilibrium payoff under σ is supported by some monotone message rule $m(\theta)$ and strictly increasing decision rule $d(m)$ that satisfies $|d(m) - m| \leq \Delta$. Consider the less informative signal σ'. Any separating interval $[\varrho'_{i'-1}, \varsigma'_{i'})$ of σ' belongs to some separating interval of σ; so $m(\theta) = \theta$ for all $\theta \in [\varrho'_{i'-1}, \varsigma'_{i'})$. Lemma 3 ensures that $d(m(\theta))$ can be supported on $[\varrho'_{i'-1}\varsigma'_{i'})$ under σ'. Similarly, for any pooling interval $[\varsigma'_{i'}, \varrho'_{i'})$ of σ', Lemmas 2 and 3 ensure that we can support $m(\theta)$ and $d(m(\theta))$ on $[\varsigma'_{i'}, \varrho'_{i'})$ under σ'. Combining these observations, we see that $d(m(\theta))$ can be supported on $[0, 1)$ under σ'. Moreover, the construction in Lemmas 2 and 3 involves no burned money; so $\overline{v}' \geq \overline{v}$.

The sender's worst monotone equilibrium payoff under σ can be supported by $\tau_S = 0$, $T_S(\theta) = 0$ and $v_s(\theta) = \underline{v}_S$; 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 value. Let $m(\theta)$ and $d(m)$ be message and decision rules that support this equilibrium. Then for separating intervals $[\varrho_{i-1}, \varsigma_i)$ of σ, the interim transfer rule is $t_S(\theta) = 0$ and for pooling intervals $[\varsigma_i, \varrho_i)$ of σ, $t_S(m)$ is given by Lemma 2 after replacing the state space $[0, 1)$ with $[\varsigma_i, \varrho_i)$. Therefore, the sender's worst monotone equilibrium payoff can be written as $\underline{v}_S = \mathbb{E}[u_S(d(m(\theta)), \theta) - t_S(m(\theta))]$. Again applying Lemmas 2 and 3, $d(m(\theta))$ can be supported under the less informative signal σ'. Moreover, the interim transfer rule $t'_S(m)$ that supports $d(m(\theta))$ under σ' is zero for separating intervals

$[\varrho'_{i'-1}, \varsigma'_{i'})$ of σ' and given by Lemma 2 for pooling intervals $[\varsigma'_{i'}, \varrho'_{i'})$ of σ'. Since σ is more informative than σ', we have $t'_S(\theta) = 0 = t_S(\theta)$ for $\theta \in [\varrho'_{i'-1}, \varsigma'_{i'})$ and $t'_S(m(\theta)) \leq t_S(m(\theta))$ for $[\varsigma'_{i'}, \varrho'_{i'})$; this follows from Lemma 2 by noting that the minimum of $h(m)$ is smaller when taken over $[\varsigma'_{i'}, \varrho'_{i'})$ than over $[\varsigma_i, \varrho_i) \subset [\varsigma'_{i'}, \varrho'_{i'})$. Therefore,

$$\underline{v}'_S \leq \mathbb{E}[u_S(d(m(\theta)), \theta) - t'_S(m(\theta))] \leq \mathbb{E}[u_S(d(m(\theta)), \theta) - t_S(m(\theta))] = \underline{v}_S.$$

■

Proposition 5 restricts attention to monotone equilibria. This restriction is only needed for one part of the proof: specifically, when comparing the sender's worst equilibrium payoff $\underline{u}_S$ under σ and σ'. Indeed, the other key objects in the proof are always monotone equilibria. The receiver's worst equilibrium is monotone because it is achieved by the repetition of a static uninformative equilibrium. Similarly, the best joint equilibrium is always monotone; it is constructed by applying Proposition 4 and the last paragraph of Section 4.2 to each pooling interval of signal σ.

Without this monotonicity restriction, in the sender's worst equilibrium under σ, $d(m(\theta))$ might be decreasing for some states θ. In this case, we would not be able to apply Lemma 2 and show that $d(m(\theta))$ can be supported under the less informative signal σ'. In fact, if $d(m(\theta))$ is non-monotone, then the self-enforcement constraint may have to account for the sender's temptation to deviate from the equilibrium message rule $m(\theta)$, and thus no longer takes the form (12).

Absent the monotonicity restriction, we have been unable to either prove – or find a counterexample to – Proposition 5. Nonetheless, slightly weaker versions of Proposition 5 hold without the monotonicity condition. If the players are myopic ($\delta = 0$), then in any equilibrium, $d(m(\theta))$ is increasing in θ and Proposition 5 holds without any monotonicity restriction. Similarly, if σ is fully informative and the sender is upwardly biased, then Proposition 5 holds without having to impose monotonicity, because $d(m(\theta)) = \theta - \Delta$ in this case. Finally, as we point out in Section 6, if $\underline{u}_S$ and $\underline{u}_R$ are exogenous outside options, then the restriction to monotone equilibria is without loss of generality.

References

Abreu, Dilip (1988) "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, Vol. 56, pp. 383–396.

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990) "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, Vol. 58, pp. 1041–1063.

Alonso, Ricardo, Wouter Dessein, and Niko Matouschek (2008) "When Does Coordination Require Centralization?," *American Economic Review*, Vol. 98, pp. 145–179.

Alonso, Ricardo and Niko Matouschek (2007) "Relational Delegation," *RAND Journal of Economics*, Vol. 38, pp. 1070–1089.

Alonso, Ricardo and Niko Matouschek (2008) "Optimal Delegation," *Review of Economic Studies*, Vol. 75, pp. 259–293.

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

Au, Pak Hung (2015) "Dynamic Information Disclosure," *RAND Journal of Economics*, Vol. 46, pp. 791–823.

Austen-Smith, David and Jeffrey S. Banks (2000) "Cheap Talk and Burned Money," *Journal of Economic Theory*, Vol. 91, pp. 1–16.

Austen-Smith, David and Jeffrey S Banks (2002) "Costly signaling and cheap talk in models of political influence," *European Journal of Political Economy*, Vol. 18, pp. 263–280.

Baker, George, Robert Gibbons, and Kevin J Murphy (1994) "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics*, Vol. 109, pp. 1125–1156.

Baker, George, Robert Gibbons, and Kevin J Murphy (2011) "Relational Adaptation," Mimeo, MIT.

Benabou, Roland and Jean Tirole (2003) "Intrinsic and Extrinsic Motivation," *Review of Economic Studies*, Vol. 70, pp. 489–520.

Bennedsen, Morten and Sven E Feldmann (2006) "Informational Lobbying and Political Contributions," *Journal of Public Economics*, Vol. 90, pp. 631–656.

Bergemann, Dirk and Stephen Morris (2016) "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, Vol. 11, pp. 487–522.

Bester, Helmut and Daniel Krähmer (2016) "The Optimal Allocation of Decision and Exit Rights in Organizations," *RAND Journal of Economics*.

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier (2016) "Dynamic Bayesian Persuasion with Public News," Mimeo, University of Oxford.

Blackwell, David (1953) "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, Vol. 24, pp. 265–272.

Bull, Clive (1987) "The Existence of Self-Enforcing Implicit Contracts," *Quarterly Journal of Economics*, Vol. 102, pp. 147–159.

Camerer, Colin (1988) "Gifts as Economic Signals and Social Symbols," *American Journal of Sociology*, Vol. 94, pp. S180–S214.

Crawford, Vincent P. and Joel Sobel (1982) "Strategic Information Transmission," *Econometrica*, Vol. 50, pp. 1431–1451.

Crémer, Jacques (1995) "Arm's Length Relationships," *Quarterly Journal of Economics*, Vol. 110, pp. 275–295.

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

Dye, Ronald (1985) "Disclosure of Nonproprietary Information," *Journal of Accounting Research*, Vol. 23, pp. 123–145.

Ely, Jeffrey C (2016) "Beeps," *American Economic Review*.

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica (2015) "Suspense and Surprise," *Journal of Political Economy*, Vol. 123, pp. 215–260.

Esteban, Joan and Debraj Ray (2006) "Inequality, Lobbying, and Resource Allocation," *American Economic Review*, Vol. 96, pp. 257–279.

Fong, Yuk-Kai and Jin Li (2016) "Information Revelation in Relational Contracts," *Review of Economic Studies*.

Goldlücke, Susanne and Sebastian Kranz (2012) "Infinitely Repeated Games with Public Monitoring and Monetary Transfers," *Journal of Economic Theory*, Vol. 147, pp. 1191–1221.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani (2009) "Mediation, Arbitration and Negotiation," *Journal of Economic Theory*, Vol. 144, pp. 1397–1420.

Hirshleifer, Jack (1971) "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, Vol. 61, pp. 561–574.

Holmstrom, Bengt (1984) "On the Theory of Delegation," in M Boyer and R. Kihlstrom eds. *Bayesian Models in Economic Theory*, New York: North-Holland.

Hörner, Johannes and Andrzej Skrzypacz (2016) "Selling Information," *Journal of Political Economy*, Vol. 124, pp. 1515–1562.

Kamenica, Emir and Matthew Gentzkow (2011) "Bayesian Persuasion," *American Economic Review*, Vol. 101, pp. 2590–2615.

Karamychev, Vladimir and Bauke Visser (2016) "Optimal Signaling with Cheap Talk and Money Burning," *International Journal of Game Theory*.

Kartik, Navin (2007) "A Note on Cheap Talk and Burned Money," *Journal of Economic Theory*, Vol. 136, pp. 749–758.

Kartik, Navin (2009) "Strategic Communication with Lying Costs," *Review of Economic Studies*, Vol. 76, pp. 1359–1395.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani (2007) "Credulity, Lies, and Costly Talk," *Journal of Economic Theory*, Vol. 134, pp. 93–116.

Kolotilin, Anton (2015) "Experimental Design to Persuade," *Games and Economic Behavior*, Vol. 90, pp. 215–226.

Kolotilin, Anton (2016) "Optimal Information Disclosure: A Linear Programming Approach," Mimeo, UNSW Australia.

Kováč, Eugen and Tymofiy Mylovanov (2009) "Stochastic Mechanisms in Settings Without Monetary Transfers: The Regular Case," *Journal of Economic Theory*, Vol. 144, pp. 1373–1395.

Krishna, Vijay and John Morgan (2008) "Contracting for Information under Imperfect Commitment," *RAND Journal of Economics*, Vol. 39, pp. 905–925.

Kwon, Suehyun (2016) "Relational Contracts in a Persistent Environment," *Economic Theory*, Vol. 61, pp. 183–205.

Levin, Jonathan (2003) "Relational Incentive Contracts," *American Economic Review*, Vol. 93, pp. 835–857.

Li, Hao and Wei Li (2013) "Misinformation," *International Economic Review*, Vol. 54, pp. 253–277.

Lim, Wooyoung (2012) "Selling Authority," *Journal of Economic Behavior & Organization*, Vol. 84, pp. 393–415.

Macleod, W Bentley and James M Malcomson (1989) "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica*, Vol. 57, pp. 447–480.

Martimort, David and Aggey Semenov (2006) "Continuity in Mechanism Design without Transfers," *Economics Letters*, Vol. 93, pp. 182–189.

Melumad, Nahum D and Toshiyuki Shibano (1991) "Communication in Settings with No Transfers," *RAND Journal of Economics*, pp. 173–198.

Mylovanov, Tymofiy (2008) "Veto-Based Delegation," *Journal of Economic Theory*, Vol. 138, pp. 297–307.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2016) "Persuading the Regulator to Wait," Mimeo, Stanford University.

Ottaviani, Marco and Peter Norman Sørensen (2006a) "Professional Advice," *Journal of Economic Theory*, Vol. 126, pp. 120–142.

Ottaviani, Marco and Peter Norman Sørensen (2006b) "Reputational Cheap Talk," *RAND Journal of Economics*, Vol. 37, pp. 155–175.

Rantakari, Heikki (2008) "Governing Adaptation," *Review of Economic Studies*, Vol. 75, pp. 1257–1285.

Rayo, Luis and Ilya Segal (2010) "Optimal Information Disclosure," *Journal of Political Economy*, Vol. 118, pp. 949–987.

Rogerson, William P (1985) "The First-Order Approach to Principal-Agent Problems," *Econometrica*, Vol. 53, pp. 1357–1367.

Suvorov, Anton and Jeroen van de Ven (2009) "Discretionary Rewards as a Feedback Mechanism," *Games and Economic Behavior*, Vol. 67, pp. 665–681.