Relational Communication with Transfers*

Anton Kolotilin and Hongyi Li†

October 13, 2017

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 the receiver's decision-making is sufficiently state-sensitive. In this case, the receiver's decision-making is disciplined by pooling extreme states, where she is most tempted to renege.

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, Marco Ottaviani, Michael Powell, and participants at various seminars and conferences for helpful comments. Kolotilin acknowledges support from the Australian Research Council.

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

1 Introduction

Decision-makers and informed parties often develop relationships in which communication and decision-making are governed by informal agreements. We study how such interactions can be disciplined using relational contracts: discretionary compensation schemes that are self-enforcing, so that the value of the future relationship outweighs the players' temptations to renege. We characterize communication and decision-making patterns in optimal relational contracts.

The interaction between lobbyists and politicians provides an example of such relational communication. Lobbyists seek to influence politicians' policy decisions.[1] They provide politicians with information about the electoral and economic consequences of various policy choices, such as focus group attitudes towards potential tobacco regulations, or the impact of cigarette smoking on health outcomes. Lobbyists also make transfers to politicians, in the form of political contributions. Such transfers serve as contingent contributions for favorable policy decisions (Grossman and Helpman 1994, 1996) and credible signals of lobbyists' information (Austen-Smith 1995 and Lohmann 1995).

While political contributions are legal in many countries, explicit payments for policy decisions would usually constitute illegal bribery and political corruption. Consequently, agreements between politicians and lobbyists are largely implicit and supported by trust and reputation. Indeed, lobbyists often maintain long-standing relationships with politicians.

Another example of relational communication is organizational decision-making, which is often governed by informal agreements — "firms are riddled with relational contracts" (Baker, Gibbons and Murphy 2002). Consider a subordinate who implements a project, and a superior who has relevant information. The superior may advise or even instruct the subordinate, but it is the subordinate who decides how to implement the project.[2] Besides giving advice, the superior often pays the subordinate to influence implementation. Payments may take the form of wages, bonuses, raises, and gifts. Payments may directly reward the subordinate for compliant implementation. Payments may also give credibility to the superior's advice — "the leader offers gifts to the followers ... because the leader's sacrifice convinces them that she must truly consider this to be a worthwhile activity" (Hermalin 1998).

[1] See Grossman and Helpman (2001) and Persson and Tabellini (2002) for reviews.

[2] Similar to us, Landier, Sraer and Thesmar (2009) and Van den Steen (2010) consider situations with a subordinate as a decision-maker implementing a project and a superior as an informed party giving advice. See Section 3.4 of Gibbons, Matouschek and Roberts (2013) for a review.

Our analysis of relational communication is based on an infinitely-repeated cheap-talk game, played by a sender and receiver who can make voluntary transfers to each other at any point in the game. 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. The players' preferred decisions are increasing in the state, but the magnitude and sign of the sender's bias may depend on the state.[3] 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.[4]

In our setting, transfers play two key roles in the relational contract. First, discretionary transfers motivate the receiver to make surplus maximizing decisions. Second, transfers allow the sender to credibly signal his private information at no welfare cost. 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.

One implication 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.[5] 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 costlessly achieved.

Because utility is transferable in our setting, optimality 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

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

[4]Specifically, in Section 5, we consider the following extensions: non-quadratic payoff functions; imperfect monitoring of the state and of the receiver's decision; and correlation of states across time.

[5]In Section 6, we discuss the implications of this point for the separation of information and control in organizations. We show that increasing organizational transparency and delegating the decision right to an informed player generally decreases the efficiency of informal relationships.

outcomes, and differ only in transfers (which determine the division of surplus).

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 main contribution is the complete characterization of the optimal relational contract when the first best is unattainable, because players are insufficiently patient. Our key insight is that the optimal relational contract may involve (partial or complete) pooling of information. Pooling is optimal only at states of *extreme* conflict (where the sender's bias is sufficiently large) and only if the receiver is very *sensitive* (the derivative of the receiver's preferred decision with respect to the state is sufficiently large relative to that of the sender).

The receiver's self-enforcement constraint – representing her temptation to deviate to her preferred decisions – requires that any feasible decision be close to the receiver's preferred decision. For non-extreme states, first-best decision-making is feasible. But for extreme states, the receiver's decision-making can only be partially disciplined. In particular, if the receiver is very sensitive, then second-best decision-making is too sensitive to information about extreme states relative to the first-best; so extreme states are optimally pooled.

The result that the sender reveals (hides) information when conflict of interest is moderate (extreme) seems to be a natural pattern of communication in relationships. Lobbyists often discuss in detail the costs and benefits of potential legislation with politicians, but may hide their private information in cases that are particularly controversial or consequential. For example, the tobacco lobby concealed and distorted evidence from internal studies that cigarettes were addictive and caused lung cancer (Hilts 1994 and Harris 2008), to soften regulation of tobacco products by Congress. In organizations, superiors provide honest advice and subordinates comply when their preferences are largely aligned, but superiors may hide information when subordinates are most tempted to dissent or disobey.

1.1 Related Literature

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.[6]

Similar to us, Alonso and Matouschek (2007) consider repeated communication.[7] 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.[8] Austen-Smith and Banks (2000) and Kartik (2007) consider a related (albeit static) setting where the sender burns money to signal information.[9] They show that virtually any message rule can be implemented by appropriately burning money (as can be achieved by transferring money in our setting). However, unlike burning money, signaling information with transfers incurs no welfare cost. This leads to a clean characterization of the set of optimal equilibria; in particular, all optimal equilibria in our model produce identical communication outcomes.[10]

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

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

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

[8]Ottaviani (2000) and Krishna and Morgan (2008) consider communication games with contractible transfers (in contrast to the voluntary transfers in our setting). In their settings, transfers are from receiver to sender, and thus cannot be used to signal information. Bester and Krähmer (2017) consider a related setting with contractible transfer schemes and study the optimal allocation of authority, similar to Dessein (2002).

[9]Kartik, Ottaviani and Squintani (2007) and Kartik (2009) consider related models with lying costs instead of money burning.

[10]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 (2017) characterize the sender's optimal equilibrium.

Persuasion (Rayo and Segal 2010 and Kamenica and Gentzkow 2011).[11] 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 (2017).

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 Morgan (2008), discussed in Footnote 8, 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), 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.

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

[11] 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 2017, Ely 2017, Hörner and Skrzypacz 2016, and Orlov, Skrzypacz and Zryumov 2016).

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 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.

[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.

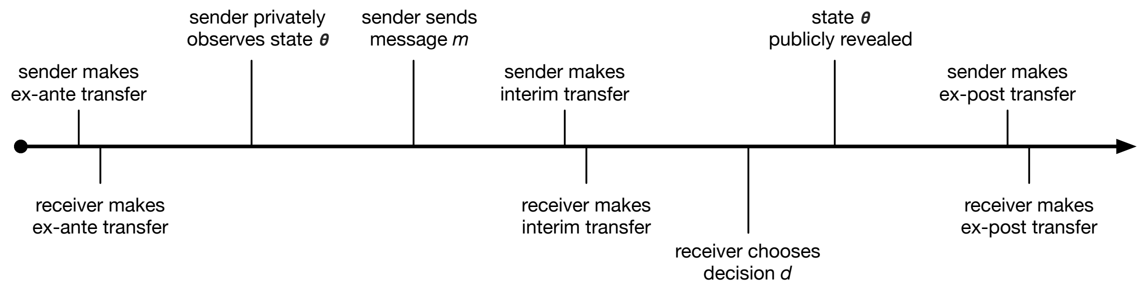

Figure 1: Timing of stage game

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 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*

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

[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.

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.

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 $\underline{v}_R$ 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

[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. See Appendix E for discussion.

in equilibrium by using interim transfers to signal information.[17] To illustrate this point, 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 $t_S(m) = 2\alpha_S bm = -t_R(m)$ 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]

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

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

[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.

[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 $t(0) = 0$.

3.2 Optimal Communication

The optimal myopic equilibrium involves either complete pooling or full separation.

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

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

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*

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

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*

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

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

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

for some $\theta^* \geq \widehat{\theta} = \max\left\{\theta \in [0,1] : d_{FB}(\theta) \geq \theta + \Delta\right\}$; 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

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

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.

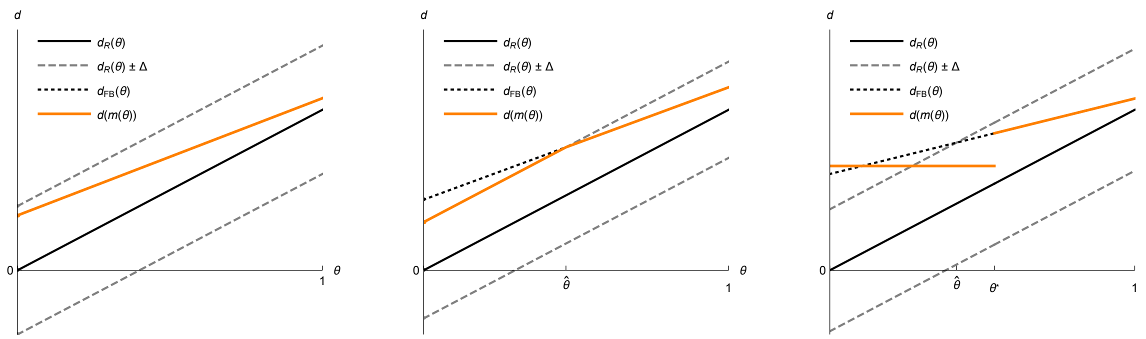

(a) First-best (b) $d'_{FB}(\theta) > 1/2$ (c) $d'_{FB}(\theta) < 1/2$

Figure 2: Second-best outcomes with upwardly biased sender

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 ($|d_{FB}(\theta) - d_R(\theta)| < \Delta$), so the first-best decision is successfully implemented: $d(m(\theta)) = d_{FB}(\theta)$. For low states, the conflict is so extreme 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 $\mathbb{E}[\theta \mid \theta \leq \hat{\theta}] + \Delta$ 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 newly-added states $\left[\hat{\theta}, \hat{\theta}+d\hat{\theta}\right]$ 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 of extreme conflict ($|d_{FB}(\theta) - d_R(\theta)|$ is large) and the receiver's preferred decision is too sensitive to information about the state ($d'_{FB}(\theta) < 1/2$, or equivalently $d'_R(\theta) > 2d'_{FB}(\theta)$). Under the assumption that the sender is upwardly biased, if the receiver is too sensitive, 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 too sensitive, $d'_R(\theta) < 2d'_{FB}(\theta)$, then full separation is optimal and the self-enforcement constraint binds for states of extreme conflict (see Figure 3a). If the receiver is too sensitive, $d'_R(\theta) > 2d'_{FB}(\theta)$, then states of extreme conflict are pooled, for the same reasons as in Proposition 4. Intermediate states, where conflict is moderate, may or may not be separated (see Figures 3b and 3c).

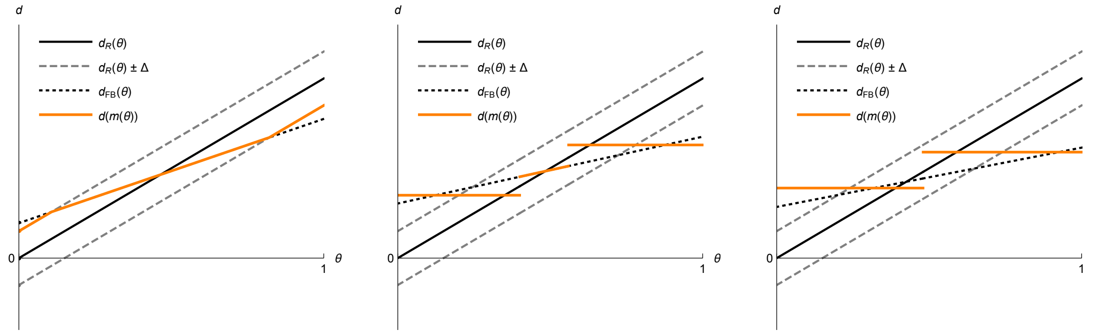

(a) Full separation (b) Pooling at extreme states (c) No separation in middle

Figure 3: Second-best outcomes when bias switches sign

5 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

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

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, the second-best message rule may pool states of extreme conflict.

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]

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

[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 $[0,\widehat{\theta}]$.

Murphy (1994, 2002) is to restrict attention to *trigger strategy* equilibria where off-path punishments correspond to some static equilibria of the stage game. Propositions 1 – 4 continue to hold in these settings, with the worst equilibrium payoffs equal to either the outside option payoffs or the static 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 the players observe neither the state nor their payoffs at the end of each period. 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. Therefore, Propositions 1 – 4 continue to hold verbatim.

The assumption that the players do not observe their payoffs can be interpreted as follows. Payoffs are revealed only at the end of a finitely repeated game with a random number of periods, where δ is the probability of continuation following each stage game rather than the discount factor. This assumption and interpretation, however, are nonstandard.

A more standard way to model imperfect monitoring of the state would be to 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. Again, the sender's temptation to renege does not contribute to the self-enforcement constraint, and thus Propositions 1 – 4 continue to hold.[21]

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),[22] 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), $\overline{d}(m)$ 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 – 4 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

[21] We 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).

[22] 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 $q(d) = \mathbb{E}\left[-\alpha_R(d+\varepsilon-m)^2\right]$.

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 – 4 continue to hold verbatim, except that the key objects such as V, Δ, and θ^* are now functions of ω_t.

6 Separation of Information and Control

In this section, we show how 'arms-length' organizational forms that separate information and control enable effective informal communication and decision-making. We consider two changes to the model that reduce the separation of information and control. In Section 6.1, we introduce formal communication processes that mechanically increase transparency. Specifically, we introduce a public signal about the state. In Section 6.2, we allow for delegation of decision rights to informed players.

One might naively expect that improving public information or delegating decision rights to informed players will enable better informed decision making. This turns out not to be the case: the availability of transfers as a signaling device implies that better informed decision making can always be achieved without imposing additional shadow costs on the relational contract, so an organizational form that brings information and control together adds no informational benefits for the relationship. On the flip side, such an organizational form makes a self-enforcing relational contract more difficult to sustain for two reasons. First, it improves both players' worst possible equilibrium payoffs, and thus limits the severity of off-path punishments. Second, it prevents information pooling, and thus limits the ability to discipline decision-making in states of extreme conflict.

6.1 Transparency

We augment the model of Section 2 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 strictly positive probability.

In this section, we focus on trigger strategy equilibria. Any deviation from equilibrium play is punished by permanent reversion to some static equilibrium of the stage game. With some abuse of notation, let $\underline{v}_S$ and $\underline{v}_R$ denote the sender's and receiver's worst static equilibrium payoffs, and let $\overline{v}$ denote the best joint (trigger strategy) equilibrium surplus.[23]

Assuming that the sender is upwardly biased as in Section 4.2, we show 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 *Suppose that $a\theta + b \geq \theta$ for all θ. Suppose also that σ and σ' are signal rules with corresponding trigger strategy equilibrium payoff sets V and V' and best joint equilibrium surpluses $\overline{v}$ and $\overline{v}'$. 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 $\overline{v} < \overline{v}'$.*

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.[24]

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 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

[23]Trigger strategy equilibria achieve the receiver's worst (perfect Bayesian) equilibrium payoff but not the sender's worst equilibrium payoff (see Appendix E for discussion).

[24]In the proof of Proposition 5, we show that $\underline{v}_R$ strictly decreases and $\underline{v}_S$ weakly decreases when moving from any signal σ to any strictly less informative signal σ'. We also note in the proof that $\underline{v}_S$ strictly decreases for generic σ and any strictly less informative signal σ'.

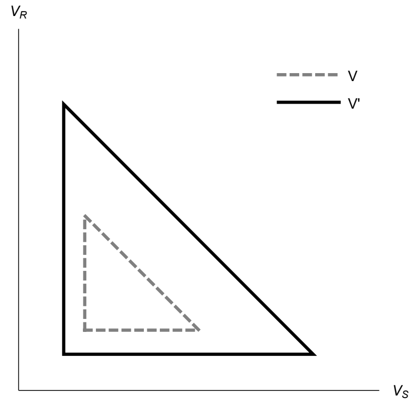

Figure 4: Equilibrium Payoff Sets

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

Joint surplus $\bar{v}$ is weakly higher under σ_u than σ_f, again because any 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 with $\underline{v}_S$ and $\underline{v}_R$). This in turn relaxes constraints on decision-making and expands the set V even further.

The result that public information hurts the relationship relates to various papers that study the social value of public information. Hirshleifer (1971) argues that welfare may be decreasing in the amount of public information available to agents. Bergemann and Morris (2016) 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.[25] 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.[26]

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

[26] This point relates to an insight from Baker, Gibbons and Murphy (1994). There, objective performance

6.2 Allocation of Authority

In our model of Section 2, decision-making authority always resides with the receiver and is not transferable (*receiver-authority*). Consider a variation of the model 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,[27] and thus the self-enforcement constraint is strictly tighter, under sender-authority. 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 $\alpha_S \neq \alpha_R$, two additional effects make the comparison between sender- and receiver-authority more nuanced. The first effect favours giving authority to the player who cares more about the decision. Under i-authority where $i \in \{S, R\}$, the temptation to renege on the first-best decision is $\alpha_i \left(d_i(\theta) - d_{FB}(\theta)\right)^2 = \alpha_i \alpha_{-i}^2 \left(a\theta + b - \theta\right)^2$, which is higher than the corresponding temptation under $(-i)$-authority when $\alpha_i < \alpha_{-i}$. The second effect weakly favours receiver-authority. Under receiver-authority, when $\alpha_S > \alpha_R$, the optimal equilibrium may involve pooling to discipline decision-making; this tool is unavailable 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

measures, rather than transparency, improve the players' outside options and make cooperation within the relationship more difficult to sustain.

[27] 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.

$\alpha_S = \alpha_R$. 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.

The standard rationale for delegation is that the better-informed player can more effectively adapt the decision. For example, Dessein (2002), Alonso, Dessein and Matouschek (2008), and Rantakari (2008) explore the tradeoff between allocating authority to an uninformed receiver versus an informed but biased sender.[28] The standard rationale for delegation no longer applies in our setting because interim transfers can credibly achieve arbitrary communication outcomes at no welfare cost.

7 Conclusion

In our model, incomplete information transmission does not reflect communication failure, but instead is an instrument for managing decision-making. This finding relies on the capacity of voluntary transfers to credibly support any monotone message rule at no welfare cost. It suggests that when modeling strategic communication in applied settings, it is crucial to understand whether monetary or 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

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

interactions. This analysis produces a rich and intuitive set of results. In particular, incomplete information transmission is implemented only for states of extreme conflict, and only if the receiver's decision-making is relatively state-sensitive. 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 provides a rationale for opaque organizations which put information in the hands of superiors and prevent subordinates from acquiring information elsewhere. 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 $v_S^* + v_R^* = \overline{v}$, 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) $v_S(\theta) + v_R(\theta) = \overline{v}$ 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

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

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 $\zeta_i < \xi_i \leq \zeta_{i+1}$ for all $i \geq 1$. Moreover, let $v_S(\theta) \geq \underline{v}_S$, $v_R(\theta) \geq \underline{v}_R$, and $T_S(\theta) \leq 0$ for all θ. Define, iteratively over $i \geq 1$,*

$$h_i^{pool} = u_S(d(m_i), \zeta_i) - u_S(d(\zeta_i), \zeta_i) + h_{i-1}^{sep}(\zeta_i), \tag{7}$$

$$h_i^{sep}(m) = \int_{\xi_i}^{m} \frac{\partial u_S(d(\vartheta), \vartheta)}{\partial d} d'(\vartheta)\, d\vartheta + u_S(d(\xi_i), \xi_i) - u_S(d(m_i), \xi_i) + h_i^{pool} \quad \textit{for } m \in [\xi_i, \zeta_{i+1}], \tag{8}$$

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

$$h(m) = \begin{cases} h_{i-1}^{sep}(m) \ \textit{if } m \in [\xi_{i-1}, \zeta_i), \\ h_i^{pool} \ \textit{if } m = m_i, \end{cases} \tag{9}$$

$$t_S(m) = h(m) - \min_{\boldsymbol{m}} h(\boldsymbol{m}), \tag{10}$$

$$\theta^p \in \arg\min_{\boldsymbol{m}} t_S(\boldsymbol{m}). \tag{11}$$

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

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

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 (2017)'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 2017).

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

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

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

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

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 $m\left(\theta\right) = \mathbb{E}\left[\theta\right]$ (full separation $m\left(\theta\right) = \theta$) 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 rule can be supported in equilibrium. Start with the relaxed problem. Given $m(\theta)$, rewrite expected joint surplus as

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

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 $d_{FB}(\widehat{\theta}) = d(\widehat{\theta})$, we may check that $G'_-(\widehat{\theta}) = G'_+(\widehat{\theta}) = 2d_{FB}(\widehat{\theta})d'_{FB}(\widehat{\theta})$; so G is continuously differentiable. Morever, if $\alpha_R + \alpha_S a < 1/2$, then G is concave on $[0, \widehat{\theta}]$ and convex on $[\widehat{\theta}, 1]$, where $\widehat{\theta} > 0$ because $\delta \in (0, \delta_{FB})$. Then by Proposition 3 of Kolotilin (2017), the message rule $m(\theta)$ specified by (5) is optimal; further, either θ^* solves

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

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 $\alpha_R + \alpha_S a \geq 1/2$, 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 $\alpha_R + \alpha_S a \geq 1/2$ 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. ■

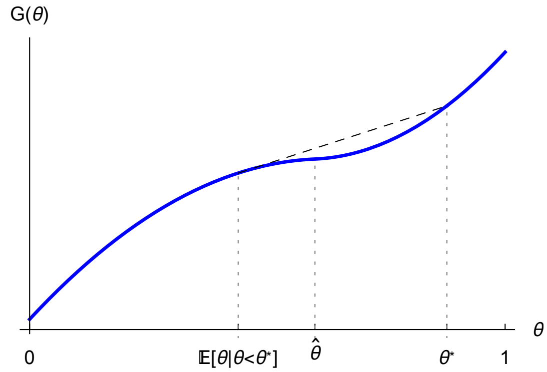

Figure 5: $G(\theta)$ with upwardly biased sender and $\alpha_R + \alpha_S a < 1/2$.

Appendix D Transparency

Proof of Proposition 5. Like a monotone message rule, a monotone signal σ generates a monotone partition of the set of states. We restrict attention to signals 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 signals.

Since signal σ is publicly observed, equilibrium play in each period is conditioned on each realization of the signal. Therefore, the sender's and receiver's worst equilibria are simply the sender's and receiver's worst equilibria conditional on each realization of the signal. Similarly, optimal equilibria are optimal conditional on each realization of the signal.

Lemma 1 implies that any static equilibrium message rule m is monotone conditional on each realization s of signal σ; that is, $m(\theta)$ is nondecreasing in θ on each set $\sigma^{-1}(s)$. Thus, given our restriction to monotone signals and direct message rules, $m(\theta)$ is nondecreasing in θ on the entire set of states $[0, 1)$. Therefore, the sender's and receiver's worst static equilibrium message rules $\underline{m}_S(\theta)$ and $\underline{m}_R(\theta)$ are nondecreasing in θ. Moreover, the static equilibrium decision rule $d(m) = m$ is strictly increasing in m.

Similarly, Proposition 4 implies that the optimal equilibrium message rule is monotone conditional on each realization of signal σ, and, thus is monotone on the entire set of states. Moreover, since the optimal decision rule $d(m)$ minimizes the distance $|d(m) - d_{FB}(m)|$ subject to the self-enforcement constraint $|d(m) - m| \leq \Delta$, it is strictly increasing in m.

In the receiver's worst static equilibrium, the sender sends the uninformative message $m\,(\theta) = \sigma(\theta)$, and the receiver chooses the decision $d(\sigma) = \mathbb{E}[\theta|\sigma]$ and does not pay or receive any transfers, $\tau_R = 0$ and $T_R(\theta) = 0$. 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'.

In the sender's worst static equilibrium, $\tau_S = 0$, $T_S(\theta) = 0$, and $d(m) = m$. For separating intervals $[\varrho_{i-1}, \varsigma_i)$ of σ, the interim transfer rule is $t_S(m) = 0$ and, for pooling intervals $[\varsigma_i, \varrho_i)$ of σ, $t_S(m)$ is given by Lemma 2 after replacing the set of states $[0, 1)$ with $[\varsigma_i, \varrho_i)$. Therefore, $\underline{v}_S = \mathbb{E}[u_S(m(\theta), \theta) - t_S(m(\theta))]$, where $m(\theta)$ is the sender's worst static equilibrium message rule $\underline{m}_S(\theta)$ under σ. Since $m(\theta)$ is nondecreasing in θ, Lemma 1 implies that $m(\theta)$ can be supported in a static equilibrium under the less informative signal σ'. Moreover, the interim transfer rule $t'_S(m)$ that supports $m(\theta)$ under σ' is as follows. For separating intervals $[\varrho'_{i'-1}, \varsigma'_{i'})$ of σ', $t'_S(m) = 0$, and, for pooling intervals $[\varsigma'_{i'}, \varrho'_{i'})$ of σ', $t'_S(m)$ is given by Lemma 2 after replacing the set of states $[0, 1)$ with $[\varsigma'_{i'}, \varrho'_{i'})$. Since σ is more informative than σ', each separating interval $[\varrho'_{i'-1}, \varsigma'_{i'})$ of σ' belongs to some separating interval of σ; so $t'_S(m) = 0 = t_S(m)$ for $\theta \in [\varrho'_{i'-1}, \varsigma'_{i'})$. Moreover, since each pooling interval $[\varsigma'_{i'}, \varrho'_{i'})$ of σ' contains some pooling interval $[\varsigma_i, \varrho_i)$ of σ, $t'_S(m(\theta)) \leq t_S(m(\theta))$ for $[\varsigma'_{i'}, \varrho'_{i'})$, because the minimum of $h(m)$ in Lemma 2 is smaller when taken over $[\varsigma'_{i'}, \varrho'_{i'})$ than over $[\varsigma_i, \varrho_i) \subset [\varsigma'_{i'}, \varrho'_{i'})$. Therefore,

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

where the first inequality holds because the sender's worst static equilibrium message rule $m'(\theta)$ under σ' can be worse than $m(\theta)$. As a side note to clarify Footnote 24, the second inequality is strict unless the following (nongeneric) condition holds: there exist contiguous pooling intervals $[\varsigma_i, \varrho_i)$ and $[\varsigma_{i+1}, \varrho_{i+1})$ of σ such that at their common end point state the sender's and receiver's preferred decisions coincide, $\varrho_i = \varsigma_{i+1}$ and $d_S(\varrho_i) = d_R(\varrho_i)$. Indeed, if this condition does not hold, then there exist pooling intervals $[\varsigma_i, \varrho_i)$ of σ and $[\varsigma'_{i'}, \varrho'_{i'})$ of σ' such that $[\varsigma_i, \varrho_i) \subsetneq [\varsigma'_{i'}, \varrho'_{i'})$ and $t'_S(m(\theta)) < t_S(m(\theta))$ for either states $[\varsigma_i, \varrho_i)$ or states $[\varsigma'_{i'}, \varrho'_{i'}) \setminus [\varsigma_i, \varrho_i)$.

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 $\overline{v}' \geq \overline{v}$. Taking into account that $\underline{v}_R > \underline{v}'_R$ and $\underline{v}_S \geq \underline{v}'_S$, Equation (2) implies that the relational leeway under σ' is $\Delta' \geq \Delta$, with strict inequality if $\delta > 0$. The proposition follows easily from this observation.

As we have shown above, in optimal equilibria under σ, the message rule $m(\theta)$ is nondecreasing in θ and the decision rule $d(m)$ is strictly increasing in m. Any separating interval $[\varrho'_{i'-1}, \varsigma'_{i'})$ of the less informative signal 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 $\overline{v}' \geq \overline{v}$. ■

Appendix E Sender's Worst Equilibrium

The sender's worst equilibrium payoff cannot be supported by the repetition of a static equilibrium. In notes available upon request, we consider an example in which $d_S(\theta) = \theta + b$, $f(\theta) = 1$, $b > 3/32$, and $\alpha_S b < \Delta$. We show that the sender's worst static equilibrium payoff is supported by $m(\theta) = \theta$, $d(m) = m$, $\tau = 0$, $t_S(m) = 2\alpha_S bm$, and $T_S(\theta) = 0$; so the sender's worst static equilibrium payoff is $-\alpha_S b^2 - \alpha_S b$. We then construct a first-best equilibrium in which the sender's equilibrium payoff is $-\alpha_S(b+\Delta)^2 - \alpha_S(b+\Delta) < -\alpha_S b^2 - \alpha_S b$. On the equilibrium path, the sender makes a large ex-ante transfer in each period. A sender's deviation is punished by the harshest credible decision $d(m) = m - \Delta$ for one period. In turn, a receiver's deviation is punished by the repetition of a static babbling equilibrium.

But this equilibrium still does not achieve the sender's worst equilibrium payoff. Besides using decisions as punishments, the receiver can make the sender worse off by using credible ex-post transfers to induce higher interim transfers from the sender. To illustrate this point, we construct a (non-static) equilibrium in which the receiver uses the static equilibrium decision rule $d(m) = m$ on and off the equilibrium path, yet the sender is worse off than in any static equilibrium. Let $m(\theta) = \theta$, $\tau = 0$, and $t_S(\theta) = 2\alpha_S b\theta + \epsilon$. Moreover, let $T_S(\theta) = -\epsilon$ if the sender's message coincides with the realized state θ and $\theta \leq \tilde{\theta}$, and $T_S(\theta) = 0$ otherwise, where $\epsilon > 0$ and $\tilde{\theta} \in (0,1)$. Given a sufficiently small ϵ, the sender does not want to deviate from $(m(\theta), t_S(\theta))$ because local deviations are not profitable and the receiver does not want to deviate from $T_R(\theta) = -T_S(\theta)$ because such a deviation is punished by the repetition of a static babbling equilibrium. In this equilibrium, the sender's equilibrium payoff is $-\alpha_S b^2 - \alpha_S b - (1-\tilde{\theta})\epsilon < -\alpha_S b^2 - \alpha_S b$. If the receiver uses ex-post transfers in this way, then the self-enforcement constraint has to account for the receiver's temptation to deviate from both equilibrium decisions and ex-post transfers, and thus no longer takes the form (6).

Moreover, the sender's equilibrium payoff may be decreased further by using nonmonotone decision rules. Consider an example in which $d_S(\theta) = 2\theta - 1/2$. Then conditional on the message rule $m(\theta)$, the decision rule that minimizes $\mathbb{E}[u_S(d(m(\theta)), m(\theta))]$ subject to (6) is nonmonotone: $d(m) = m + \Delta$ for $m < 1/2$ and $d(m) = m - \Delta$ for $m > 1/2$. If $d(m(\theta))$ is non-monotone, then the self-enforcement constraint has to account for the sender's temptation to deviate from the equilibrium message rule $m(\theta)$, and thus again no longer takes the form (6).

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 (1995) "Campaign Contributions and Access," *American Political Science Review*, Vol. 89, pp. 566–581.

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

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 (2002) "Relational Contracts and the Theory of the Firm," *Quarterly Journal of Economics*, Vol. 117, pp. 39–84.

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

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 (2017) "The Optimal Allocation of Decision and Exit Rights in Organizations," *RAND Journal of Economics*, Vol. 48, pp. 309–334.

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier (2017) "How to Persuade a Long-Run Decision Maker," 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.

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 (2017) "Beeps," *American Economic Review*, Vol. 107, pp. 31–53.

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

Fong, Yuk-fai and Jin Li (2016) "Information Revelation in Relational Contracts," *Review of Economic Studies*, Vol. 84, pp. 277–299.

Gibbons, Robert, Niko Matouschek, and John Roberts (2013) "Decisions in Organizations," in Robert Gibbons and John Roberts eds. *Handbook of Organizational Economics*: Princeton University Press.

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.

Grossman, Gene M and Elhanan Helpman (1994) "Protection for Sale," *American Economic Review*, Vol. 84, pp. 833–850.

Grossman, Gene M and Elhanan Helpman (1996) "Electoral Competition and Special Interest Politics," *Review of Economic Studies*, Vol. 63, pp. 265–286.

Grossman, Gene M and Elhanan Helpman (2001) *Special Interest Politics*: MIT press.

Harris, Gardiner (2008) "Cigarette Company Paid for Lung Cancer Study," *New York Times*, http://www.nytimes.com/2008/03/26/health/research/26lung.html. Published March 26, 2008.

Hermalin, Benjamin E (1998) "Toward an Economic Theory of Leadership: Leading by Example," *American Economic Review*, Vol. 88, pp. 1188–1206.

Hilts, Philip (1994) "Tobacco Chiefs Say Cigarettes Aren't Addictive," *New York Times*, http://www.nytimes.com/1994/04/15/us/tobacco-chiefs-say-cigarettes-aren-t-addictive.html. Published April 15, 1994.

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 (2017) "Optimal Signaling with Cheap Talk and Money Burning," *International Journal of Game Theory*, Vol. 46, pp. 813–850.

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 (2017) "Optimal Information Disclosure: A Linear Programming Approach," *Theoretical Economics*.

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.

Krähmer, Daniel (2006) "Message-Contingent Delegation," *Journal of Economic Behavior & Organization*, Vol. 60, pp. 490–506.

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.

Landier, Augustin, David Sraer, and David Thesmar (2009) "Optimal Dissent in Organizations," *Review of Economic Studies*, Vol. 76, pp. 761–794.

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

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

Lohmann, Susanne (1995) "Information, Access, and Contributions: A Signaling Model of Lobbying," *Public Choice*, Vol. 85, pp. 267–284.

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.

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

Ottaviani, Marco (2000) "The Economics of Advice," Mimeo, Bocconi 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.

Persson, Torsten and Guido Enrico Tabellini (2002) *Political Economics: Explaining Economic Policy*: MIT press.

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.

Van den Steen, Eric (2010) "Interpersonal Authority in a Theory of the Firm," *American Economic Review*, Vol. 100, pp. 466–490.