The Use and Abuse of Coordinated Punishments

Daniel Barron and Yingni Guo*

August 31, 2018

Abstract

A principal plays trust games with a sequence of short-lived agents who must communicate deviations to sustain cooperation. We argue that agents can abuse communication by shirking and threatening to report a deviation unless the principal nevertheless pays them. To make this threat credible, we allow agents to probabilistically commit to messages as a function of the principal's action. In equilibrium, this threat increases the agents' share of surplus but decreases total surplus; no cooperation can be sustained if agents can always commit. We show how evidence or strong bilateral principal-agent relationships mitigate this problem and partially restore cooperation.

*Barron: Northwestern University Kellogg School of Management, Evanston IL 60208; email: dbarron@kellogg.northwestern.edu. Guo: Northwestern University Economics Department, Evanston IL 60208; email: yingni.guo@northwestern.edu. The authors would like to thank Joyee Deb, Marina Halac, Ilan Kremer, Nicolas Lambert, Jin Li, Elliot Lipnowski, Joshua Mollner, Michael Powell, Luis Rayo, and Takuo Sugaya.

1 Introduction

The threat of sanctions motivates individuals and firms to work hard and reward the hard work of their partners (Malcomson [2013], Gibbons and Henderson [2013]). Word-of-mouth can be an essential tool for carrying out these sanctions, as news of misbehavior spreads far beyond those who directly observed that misbehavior. For instance, communities collectively punish members accused of not contributing to public goods (Ostrom [1990], Ali and Miller [2016]). Managers implement promised workplace improvements because employees threaten to otherwise file grievances or coordinate strikes (Freeman and Lazear [1995], Levin [2002]). Business associations keep their partners in check by sharing information about past misbehaviors (Greif et al. [1994], Robinson and Stuart [2006], Bernstein [2015]). Firms provide exemplary service to avoid negative reviews and lost sales (Hörner and Lambert [2017]).

Communication is valuable in these settings because it allows parties who do not observe a transgression to nevertheless punish the transgressor. But this feature means that coordinated punishments can be abused: parties might engage in "hold up" by threatening to falsely report a deviation unless they receive undeserved pay or other perks. So, for example, influential members of close-knit associations can threaten to ostracize anyone who disagrees with their policies.[1] Employees can threaten strikes to demand above-market wages or rigid employment rules.[2] Suppliers can threaten to complain to a business association unless their buyers agree to excessively generous terms. Customers can threaten to leave negative reviews unless a business inefficiently favors them.[3]

This paper studies how participants might misuse communication systems to hold one another up. We argue that the possibility of hold up makes cooperation more challenging

[1]Businesses sometimes "blacklist" employees, which prevents them from seeking employment from other companies in the industry. NPR's Planet Money podcast has reported allegations that Wells Fargo misused a tool designed to share information about unethical workers to blacklist former employees who protested fraudulent sales practices (Arnold and Smith [2016]).

[2]We are far from the first to suggest that unions, and the coordination they facilitate, can be used to pursue both welfare-improving and welfare-destroying goals. See, for instance, Freeman and Medoff [1979] and Freeman [1980] for examples, evidence, and discussion.

[3]For instance, the Canadian Broadcasting Company covered a recent lawsuit against a disgruntled customer who used defamatory online reviews to drive a photography studio out of business (Proctor [2018]).

because it undermines the value of coordinated punishments. In extreme cases, cooperation might collapse entirely, as any attempt at punishing a deviation simply becomes an opportunity for parties to extract resources without exerting effort. We also explore two possible remedies to this problem, each with its own drawbacks: hard evidence, which prevents parties from lying about the payments they receive, and strong bilateral relationships, which allow parties to directly punish hold-up attempts. As we will show, a pessimistic takeaway from our analysis is that coordinated punishments are inherently susceptible to abuse. A more optimistic takeaway is that several features of the environment can mitigate this problem and so encourage cooperation.

To make this point, we consider a model of a long-lived principal who interacts with a sequence of short-lived agents. Each agent exerts costly effort to benefit the principal, who can then pay that agent. Agents observe only their own interaction with the principal but can costlessly communicate. Our crucial modeling assumption is that each agent chooses a **communication protocol,** which associates a message with each possible payment, before the principal pays him. This agent is then committed to that protocol with some probability and otherwise communicates freely.

Agents are willing to exert effort only if they are compensated for doing so, while the principal is willing to compensate an agent only if she is otherwise punished. Communication is therefore essential for inducing effort, since an agent cannot directly punish the principal for reneging. Once armed with messages that can trigger punishment, however, an agent can shirk and then threaten to report a deviation unless he is nevertheless paid. This threat is credible because each agent might be committed to his communication protocol; consequently, the principal prefers to pay at least a small amount rather than face possible future punishment. That is, commitment allows agents to hold up the principal.

We show that the threat of hold-up affects both equilibrium effort and how the resulting surplus is split between the principal and each agent. The higher the probability that an agent is committed to his communication protocol, the more tempting is his deviation to shirk and threaten to report defection unless he is paid. An agent is deterred from deviating in this way only if he is paid a rent. But paying rent to an agent creates a negative intertemporal externality on the principal's relationships with other agents, since it means that the principal earns lower continuation surplus and so is less willing to pay large rewards for effort. In the extreme case where agents always commit to their communication protocols, we show that they never exert any effort in equilibrium. This intuition extends to the entire equilibrium payoff set, which shrinks as the probability of commitment increases and includes only the one-shot equilibrium payoff if agents are always committed.

We then consider two imperfect remedies to this problem, focusing on the case where agents are always committed and so hold-up is particularly severe. Our first remedy introduces hard evidence about the principal's payment. Evidence means that an agent cannot lie about the principal's payments, which limits the kinds of threats he can make and so leads to more cooperation. This mechanism is a little subtle, however, because non-payment generates the same evidence regardless of whether an agent works hard or shirks. In equilibrium, we show that evidence is used to make the principal indifferent between payments, so that she is willing to both pay a hard-working agent and not pay an agent who tries to hold her up. For the principal to be indifferent, the *absence* of evidence must be treated as harshly as evidence of malfeasance. Consequently, if evidence is only probabilistically available, then the principal is periodically punished on the equilibrium path.

Our second remedy introduces a bilateral relationship between the principal and each agent. We model this relationship in an abstract, "reduced form" way by allowing each agent to play a coordination game with the principal *after* he has communicated with the other agents. This coordination game can be used to encourage cooperation by directly punishing the principal or an agent for deviating. In addition, we show that bilateral relationships enable coordinated punishments, since the coordination game can also be used to punish agents for trying to extract rent and the principal for giving in to those threats. However, bilateral relationships do not perfectly solve hold-up, as the strength of the bilateral relationship tightly constrains how severe coordinated punishment can be while remaining effective. Appendix C builds on this intuition by showing that hold-up similarly constrains cooperation if agents are long-lived and interact repeatedly with the principal.

The essential assumption in our model is that agents can commit to their messages, which means that an agent who has shirked can still credibly threaten to spread word that the principal deviated. We show that this comitment assumption selects an equilibrium of the game without commitment, in the sense that our equilibria are outcome-equivalent to equilibria if the agents cannot commit. In other words, commitment is powerful in our setting because it allows agents to specify their messages in advance, not because it results in messages that agents would otherwise be unwilling to send. We also argue that even without commitment, agents act as if they are committed so long as they have a mild intrinsic preference for following through on their communication protocols.

The broader lesson of our analysis is that coordinated punishments are particularly susceptible to abuse by rent-seeking parties. Communication is valuable precisely when parties (i) do not directly observe deviations in other relationships, and (ii) cannot, on their own, severely punish deviations in their own relationships. The lack of direct monitoring means that parties who hear about a defection cannot tell whether it truly occurred or whether that message is instead part of a hold-up attempt. The lack of severe bilateral punishments means that the victims of hold-up attempts have limited recourse to directly punish the perpetrators of those attempts. Both of these features mean that an agent can threaten hold-up at relatively low personal cost. In equilibrium, however, the possibility of these threats limits effort and in some cases can completely undermine cooperation.

Related Literature

Our framework builds on the relational contracting literature. Much of this literature considers interactions between one principal and one agent (Bull [1987]; MacLeod and Malcomson [1989]; Baker et al. [1994]; Levin [2003]), while we focus on the costs and benefits of communication among multiple agents. Recent papers have explored relational contracts with additional frictions in the environment, including limited transfers (Fong and Li [2017]; Barron et al. [2018]), asymmetric information (Halac [2012]; Malcomson [2016]), or both (Li et al. [2017]; Lipnowski and Ramos [2017]; Guo and Hörner [2018]). We focus on a different friction, bilateral monitoring within each principal-agent pair, which implies that communication is essential for cooperation. Other papers that study settings with bilateral monitoring, including Board [2011], Andrews and Barron [2016], and Barron and Powell [2018], do not allow communication among agents. We complement these papers by identifying a reason why communication might be of limited use in sustaining cooperation.

Levin [2002], which studies relational contracts between a principal and a group of agents, is the seminal paper within this literature that touches on both coordinated punishments and bargaining. Proposition 5 of that paper shows how increasing the agents' bargaining power leads to less efficient relational contracts, since the principal has less to lose following a deviation if agents earn more surplus. An analogous negative intertemporal externality arises in our setting. However, we connect an agent's ability to extract rent to how her messages affect continuation play, which leads to new implications about the structure of punishments and new ways to mitigate this externality.

Within the broader literature on cooperation, our approach is related to papers that explore how communication improves outcomes. Ali et al. [2016], which analyzes renegotiation-proof equilibria in settings where players must communicate about deviations, is perhaps the most closely related paper in this literature. That paper shows that renegotiation does not inhibit cooperation if communication is costless. We consider hold-up rather than renegotiation, and we show that, unlike renegotiation, the threat of hold-up can severely limit equilibrium cooperation. Much of this literature focuses on *networks* of players with the goal of identifying network structures or equilibrium strategies that are particularly conducive to cooperation (Lippert and Spagnolo [2011]; Wolitzky [2013]; Ali and Miller [2016]; Ali and Liu [2018]). We do not focus on network structure. Instead we focus on how communication can be abused by rent-seeking agents. We introduce a new ingredient, the communication protocol, to ensure that agents' hold-up threats are credible. More distantly related papers include Awaya and Krishna [2018], which identifies settings in which cooperation requires communication, and Compte [1998] and Kandori and Matsushima [1998], which prove folk theorems in games with private monitoring and communication.

Commitment allows our agents to extract rent by threatening the principal. A relatively small literature analyzes bargaining in repeated games, which is typically modeled as an equilibrium refinement. For instance, several papers impose the constraint that players Nash bargain over either surplus in each period or continuation play (Baker et al. [2002], Halac [2012], Halac [2015], Goldlucke and Kranz [2017]). The closely related papers in this literature are Miller and Watson [2013] and Miller et al. [2018], which define and analyze *contractual equilibria* in relational contracts. Contractual equilibria allow players to propose continuation equilibria to one another and so capture a recursive notion of bargaining in relational contracts. By focusing on how communication among multiple agents in a setting with private monitoring, we focus on a complementary set of issues to these papers, which mostly focus on bargaining in settings with either two players or public monitoring. Our focus on communication means that we highlight a new set of hold-up problems and remedies to those problems.

Our assumption that agents commit to messages is similar to the persuasion literature, as in Rayo and Segal [2010] and Kamenica and Gentzkow [2011]. Particularly related are Min [2017], Lipnowski et al. [2018], and Guo and Shmaya [2018], which similarly assume that players are not perfectly committed to their messages. However, commitment serves a very different role in our setting, since it allows agents to make off-path threats rather than reveal information about payoff-relevant information. Consequently, agents' equilibrium payoffs are non-monotonic in the probability of commitment in our setting.[4]

[4] This non-monotonicity arises because neither effort nor payments are contractible and so is a particular manifestation of the Theory of the Second Best (Lipsey and Lancaster [1956-1957], Bernheim and Whinston [1998]).

2 Model

A long-lived principal ("she") interacts with a sequence of short-lived agents (each "he"). In each period $t \in \{0, 1, 2, ...\}$, the principal and agent t play a favor-trading game: agent t exerts effort and exchanges payments with the principal. Agents do not observe one another's interactions with the principal but can costlessly communicate. The key assumption is that before transfers are paid, each agent chooses a *communication protocol,* which is a mapping from the net transfer to this agent's message. An agent is committed to communicate according to this protocol with probability $\lambda \in [0, 1]$ and is otherwise free to choose any message.

We assume that players have access to a public randomization device (notation for which is suppressed) in every step of the stage game. The stage game in each period t is:

1. Agent t chooses his effort $e_t \in \mathbb{R}_+$ and a communication protocol $\mu_t : \mathbb{R} \to M$, where M is a large, finite message space. Both e_t and μ_t are observed by the principal but not by any other agent.

2. The principal and agent t simultaneously pay non-negative transfers to one another, which are observed by those two players but not by any other agent. Let $s_t \in \mathbb{R}$ be the net transfer to agent t.[5]

3. Agent t sends a message m_t that is publicly observed. With probability $\lambda \in [0, 1]$, agent t is committed to μ_t so $m_t = \mu_t(s_t)$; otherwise, agent t chooses any $m_t \in M$.

The principal's period-t payoff and agent t's utility are $(e_t - s_t)$ and $(s_t - c(e_t))$, respectively, where $c(\cdot)$ is differentiable, strictly increasing, strictly convex, and satisfies $c(0) = c'(0) = 0$. We assume that there exists a unique effort level, denoted e^{FB}, that solves $c'(e^{FB}) = 1$ and so maximizes total surplus. The principal has discount factor $\delta \in [0, 1)$, with corresponding normalized discounted payoffs $\Pi_t = (1 - \delta) \sum_{t'=t}^{\infty} \delta^{t'-t}(e_{t'} - s_{t'})$. We sometimes consider the

[5]Note that the *equilibrium* transfer is negative only in Section 4.2, which considers a setting in which the principal and agent t continue interacting after paying one another.

normalized discounted sum of agents' utilities, $U_t = (1-\delta)\sum_{t'=t}^{\infty}\delta^{t'-t}(s_t - c(e_t))$, which we call the agents' *joint utility*.

Our solution concept is plain Perfect Bayesian Equilibrium (hereafter PBE) as defined in Watson [2016]. Many of our results focus on *principal-optimal equilibria*, which maximize the principal's *ex ante* expected payoff $\mathbb{E}[\Pi_0]$ among PBE.

Our key modeling ingredient is the assumption that agents can commit to a communication protocol. We view this assumption as a parsimonious way to allow an agent to credibly threaten to condition his message on the transfer even he has deviated. Communication can then be abused by agents who shirk and then demand payment from the principal. Section 5 explores this assumption in more detail.

Agents are short-lived in our setting. Communication is therefore necessary to punish the principal for deviating, which lets us cleanly highlight how communication can lead to hold-up. Section 4.2 shows how future bilateral relationships, modeled as a coordination game played between the principal and each agent at the end of the period, can mitigate hold-up. Appendix C shows how the threat of hold-up undermines cooperation in a repeated game with long-lived agents.

Several other assumptions warrant further comment. First, we do not include a round of transfers at the start of each period. Appendix D.1 shows that allowing such transfers does not change the principal-optimal equilibrium in our baseline model, since requiring an agent to make an up-front payment simply gives them a stronger incentive to deviate and hold up the principal. Second, our model allows the agents but not the principal to send public messages. Appendix D.2 considers what happens if the principal can also send messages. Our result extends so long as the principal knows an agent's message before choosing her own. Finally, we do not allow agents to interact directly. Appendix D.3 allows interaction between each agent and his immediate predecessor. These inter-agent interactions do not improve cooperation, since each agent does not observe, and so cannot reward or punish, his predecessor's effort or communication protocol.

3 Multilateral Punishments Enable Hold-Up

This section proves our main result. Each agent is tempted to shirk and then use the communication protocol to induce the principal to nevertheless pay him. The principal deters this deviation by paying each agent rent, which limits the credibility of her promises to other agents, decreases effort, and in extreme cases can lead to the complete unravelling of cooperation. We build intuition using the full- and no-commitment cases before proving a general characterization.

Cooperation requires informative communication among the agents. Without communication, an agent would have no way to punish the principal for reneging on a payment and would therefore expect no payment and have no incentive to exert effort. We show that communication does indeed sustain cooperation in a setting without commitment.

Proposition 1. *Suppose that $\lambda = 0$. Let e^* equal the minimum of e^{FB} and the positive root of $c(e) = \delta e$. In any principal-optimal equilibrium, $e_t = e^*$ and $s_t = c(e^*)$ in every $t \geqslant 0$.*

Proof: Let $\overline{\Pi}$ be the principal's maximum equilibrium payoff, and let e^* be the maximum equilibrium effort in an principal-optimal equilibrium. The principal pays s_t only if $(1-\delta)s_t \leqslant \delta\overline{\Pi}$, while agent t chooses $e_t = e^*$ only if $s_t - c(e^*) \geqslant 0$. Effort e^* must therefore satisfy $(1-\delta)c(e^*) \leqslant \delta\overline{\Pi}$. In particular, $\overline{\Pi} = e^{FB} - c(e^{FB})$ whenever $c(e^{FB}) \leqslant \delta e^{FB}$. Otherwise, e^* is bounded above by $(1-\delta)c(e^*) = \delta\overline{\Pi}$, and the expression $c(e) = \delta e$ follows from setting $\overline{\Pi} = e^* - c(e^*)$.

Consider the following strategy: in each period $t \geqslant 0$ on the equilibrium path, $e_t = e^*$, $s_t = c(e^*)$, and $m_t = m^*$. If agent t deviates in e_t, then $s_t = 0$ and $m_t = m^*$. If the principal deviates, then $m_t \neq m^*$. Once a message $m_t \neq m^*$ is observed, all future periods entail repetition of the stage-game equilibrium $e_t = s_t = 0$. Agents are indifferent over messages and are willing to send the specified messages. The principal is willing to pay $s_t = c(e^*)$ because $(1-\delta)c(e^*) \leqslant \delta(e^* - c(e^*))$ by definition of e^*. Each agent is indifferent between choose $e_t = e^*$ and $e_t = 0$. So this strategy profile is an equilibrium. It is principal-optimal because it maximizes total surplus among all equilibria and gives each agent a payoff of 0. ■

An agent exerts effort in equilibrium only if he expects a high payment if he works and a low payment if he shirks. Proposition 1 shows that communication can induce the principal to pay agents who work hard, which is enough to sustain effort in a setting without commitment. The principal-optimal equilibrium attains first-best effort if the principal is sufficiently patient. Otherwise, agents' equilibrium efforts are limited by a dynamic enforcement constraint that bounds the equilibrium payment from above. Note that the non-commitment case admits many other equilibria; for example, if $\delta > c(e^{FB})/e^{FB}$, then there exist equilibria that induce first-best effort and give the agents strictly positive utility.

Communication can similarly induce the principal to make payments if $\lambda > 0$. However, the possibility of commitment leads to a new problem: the principal might also reward an agent who shirks but threatens to send a message that triggers punishment unless he is paid. In the extreme case where agents always commit to their communication protocols, our next result shows that equilibrium payments are essentially independent of effort. Consequently, agents never exert effort in equilibrium.

Proposition 2. *If $\lambda = 1$, then every equilibrium entails $e_t = s_t = 0$ in each $t \geqslant 0$.*

Proof: Suppose that $e_t > 0$ in equilibrium. Define $\overline{\Pi}$ and $\underline{\Pi}$ as the principal's maximum and minimum continuation payoffs following period t, with corresponding messages $\overline{m}_t$ and $\underline{m}_t$. Since $e_t > 0$, $s_t \geqslant c(e_t) > 0$ and hence $\overline{\Pi} > \underline{\Pi}$.

For small $\varepsilon > 0$, consider the following deviation by agent t: choose $e_t = 0$ and

$$\mu_t(s) = \begin{cases} \overline{m}_t & s \geqslant s_t - \varepsilon \\ \underline{m}_t & \text{otherwise.} \end{cases}$$

The principal's unique best response to this deviation is to pay $s_t - \varepsilon$, since

$$-(s_t - \varepsilon) + \frac{\delta}{1-\delta}\overline{\Pi} > \frac{\delta}{1-\delta}\underline{\Pi}$$

is implied by the fact that the principal is willing to pay s_t on the equilibrium path. For any $\varepsilon \in [0, c(e_t))$, $s_t - \varepsilon > s_t - c(e_t)$ and so this deviation is profitable for agent t. We conclude that no equilibrium can entail $e_t > 0$. But if $e_t = 0$ for all $t \geqslant 0$, then $s_t = 0$ as well. ■

If $s_t > 0$ on the equilibrium path, then agent t can choose zero effort and threaten to send a message that punishes the principal unless she pays him *slightly less* than s_t. Since the principal was weakly willing to pay s_t when faced with the same punishment, she is strictly willing to pay a smaller amount. The agent can therefore shirk and still guarantee nearly the same transfer as if he exerted effort. Commitment essentially severs the link between payment and effort and hence eliminates the agents' incentives to work.

Propositions 1 and 2 illustrate how messages can be used to both deter the principal from reneging on a payment and to hold her up. In equilibrium, the principal must have the incentive to both pay an agent who has worked hard and refrain from paying an agent who has shirked. As in Proposition 1, equilibrium communication enables punishments that induce the principal to pay a working agent. As in Proposition 2, however, an agent might be able to use those same messages to force the principal to pay him even after he exerts no effort. This latter effect increases an agent's temptation to shirk such that no effort can be sustained in equilibrium if $\lambda = 1$.

Our main result in this section shows how principal-optimal equilibria balance these two contrasting roles for $\lambda \in (0,1)$. The former role dominates for $\lambda \leqslant 1/2$, in which case equilibrium effort is the same as in the case without commitment. For higher λ, the latter role becomes increasingly important as λ increases, resulting in agents who can make more credible threats and so exert less effort relative to Proposition 1.

Proposition 3. *Let $e^*(\lambda)$ equal the minimum of the effort level that solves $c'(e) = 2(1-\lambda)$*

and the positive root of $c(e) = 2\delta(1-\lambda)e$. In any principal-optimal equilibrium:

1. *If $\lambda \leqslant \frac{1}{2}$, then $e_t = e^*\left(\frac{1}{2}\right)$ and $s_t = c(e_t)$ in each period on the equilibrium path. The principal's payoff is constant in λ.*

2. *If $\lambda > \frac{1}{2}$, $e_t = e^*(\lambda)$ and $s_t = \frac{c(e_t)}{2(1-\lambda)} > c(e_t)$ in each period on the equilibrium path. The principal's payoff is strictly decreasing in λ.*

Proof: See Appendix A.

To prove Proposition 3 for $\lambda \leqslant \frac{1}{2}$, it suffices to show that the principal is willing to pay nothing to an agent who deviates, because in that case the agent's incentives are identical to the case with $\lambda = 0$. Consider play in period t, and let $\overline{\Pi}$ and $\underline{\Pi}$ be the principal's highest and lowest equilibrium continuation payoffs from period $t+1$ onwards. If agent t is not committed to his communication protocol, he is willing to send any message in equilibrium. Suppose that, if he shirks, then agent t's message results in continuation payoff $\overline{\Pi}$ if $s_t = 0$ and continuation payoff $\underline{\Pi}$ for any $s_t > 0$. Given these messages, the principal's continuation surplus can be no worse than $\lambda\underline{\Pi} + (1-\lambda)\overline{\Pi}$ if $s_t = 0$ and no better than $\lambda\overline{\Pi} + (1-\lambda)\underline{\Pi}$ if $s_t > 0$. If $\lambda \leqslant \frac{1}{2}$, the principal's best response is $s_t = 0$, resulting in a principal-optimal equilibrium that is outcome-equivalent to Proposition 1.[6]

By a similar logic, an agent can deviate and induce a strictly positive transfer whenever $\lambda > \frac{1}{2}$. Larger λ or $\overline{\Pi} - \underline{\Pi}$ means that each agent can extract a larger share of the total surplus. Consequently, increasing $\overline{\Pi} - \underline{\Pi}$ might increase effort but definitely increases the share of the output earned by agents. In the principal-optimal equilibrium, punishments are chosen to balance the effect on effort, which increases total surplus, against the fact that agents earn a larger share of that surplus. The resulting equilibrium is either a corner

[6]One seemingly unusual feature of this construction is that an agent who deviates sometimes punishes the principal for paying him. This construction minimizes the amount the principal is willing to pay following a deviation and so minimizes an agent's incentive to deviate. However, we can sustain positive, albeit lower, effort in equilibria that do not share this feature, provided that the principal is punished less harshly for non-payment if the agent deviates than if he does not. For example, there exists an equilibrium with positive effort in which an agent who deviates and is not committed chooses *the same* message regardless of the transfer.

solution that maximizes $\overline{\Pi} - \underline{\Pi}$, or it sets this difference so that equilibrium effort maximizes $e - \frac{c(e)}{2(1-\lambda)}$, which is the principal's payoff after we subtract the rents she must pay to the agents.

Increasing λ decreases total surplus but increases the agents' share of that surplus. Indeed, agents earn no rent if either $\lambda = 0$, in which case they cannot hold up the principal, or $\lambda = 1$, in which case no effort is possible in equilibrium. This intuition suggests that agents' rent is maximized at an interior probability of commitment, which is confirmed in the following example.

Example 1. *Suppose that $c(e) = e^2$. Then effort $\hat{e} = 1 - \lambda$ solves $c'(e) = 2(1 - \lambda)$, so the condition $\delta \geqslant \frac{c(\hat{e})}{2(1-\lambda)\hat{e}}$ is equivalent to $\delta \geqslant \frac{1}{2}$. If $\delta \geqslant \frac{1}{2}$, the agents' joint utility in the principal-optimal equilibrium equals $(1 - \lambda)(2\lambda - 1)/2$, which is maximized at $\lambda = \frac{3}{4}$.*

Thus far, we have focused on the principal's maximum equilibrium payoff. Our next result shows the intuition extends to the entire equilibrium set, in the sense that increasing the probability of commitment decreases the set of equilibrium outcomes. Note that Proposition 2 has already shown that this set contains a single outcome if $\lambda = 1$: the agents shirk and the principal pays nothing.

Proposition 4. *For any $\delta \in [0, 1)$ and $\lambda' > \lambda$, for any equilibrium of the game with commitment probability λ', there exists an equilibrium of the game with commitment probability λ which induces the same distribution over $\{e_t, s_t\}_{t=0}^{\infty}$.*

Proof: See Appendix A.

Given an equilibrium in the game with a high commitment probability, we can construct an outcome-equivalent equilibrium in the game with a low commitment probability. This construction relies on the fact that an agent who is not committed is willing to choose any message, including the message specified by the communication protocol. Therefore, agents are always willing to act as if they are committed to their protocols, so we can replicate the

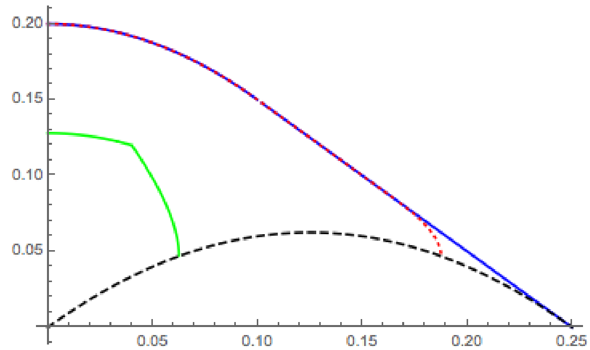

Figure 1: Payoff frontier if $c(e) = e^2$ and $\delta = 8/10$. The x-axis is the principal's payoff and the y-axis is the agents' joint utility. The blue curve corresponds to $\lambda \leqslant 1/2$; the red-dotted curve to $\lambda = 5/8$; and the green curve to $\lambda = 7/8$. The black-dashed curve traces the principal's payoff and agents' joint utility in the principal-optimal equilibrium as λ increases from $1/2$ to 1, where higher λ correspond to points closer to the origin.

equilibrium messages from the game with high commitment probability in the game with low commitment probability. An identical mapping from transfers to messages induces identical transfers and hence identical efforts, which proves the result.

Appendix B expands Proposition 4 by fully characterizing the equilibrium payoff frontier for any $\lambda \in [0, 1]$. Figure 3 illustrates how the payoff frontier changes as λ varies. In general, the principal-optimal equilibrium does not maximize total welfare, since making coordinated punishments more severe can increase the agents' joint utility faster than it decreases the principal's payoff. Pareto efficient equilibria are not necessarily stationary, either, though any non-stationarity occurs only in the first period of play.

Proposition 4 illustrates a negative intertemporal externality that arises when an agent extracts surplus. If agent t is paid a rent, then the principal has less to lose from reneging on her promises in all periods $t' < t$. She is therefore less willing to reward effort. An increase in λ increases each agent's ability to extract rent, which exacerbates this negative externality and so decreases both effort and total surplus.

4 Two Remedies to the Hold-Up Problem

This section studies two imperfect solutions to the hold-up problem identified in Section 3. In Section 4.1, we introduce the possibility of hard evidence about the transfer, which limits the messages that an agent who deviates can use to threaten the principal. Section 4.2 allows each agent to play a coordination game with the principal after sending a message, which can be used to punish that agent for threatening hold-up or the principal for giving in to those threats. Both of these sections restrict attention to $\lambda = 1$ so that the hold-up problem is particularly severe.

4.1 Remedy 1: Evidence

The *game with evidence* modifies the baseline game by introducing evidence about the transfer that can be concealed but cannot be distorted. After the principal pays s_t, agent t gets hard evidence of this payment with probability $p \in [0, 1]$. Let $x_t \in \{\emptyset, s_t\}$ denote this evidence, where $x_t = \emptyset$ if no evidence is available. When agent t chooses a communication protocol, he also chooses a *disclosure protocol* $\kappa_t : \mathbb{R} \times \{\emptyset\} \to \mathbb{R} \cup \{\emptyset\}$ subject to the restriction that $\kappa_t(x) \in \{\emptyset, x\}$ for any $x \in \mathbb{R} \cup \{\emptyset\}$. The disclosure protocol is observed by the principal but not by any other agents. The resulting disclosure $k_t = \kappa_t(x_t)$ is then publicly observed by all agents.

Our main result in this section characterizes the principal-optimal equilibrium in the game with evidence. The good new is that, unlike Proposition 2, positive effort is possible in equilbrium. However, the principal-optimal equilibrium punishes the principal whenever an agent does not reveal evidence, which limits the gains from cooperation for any $p < 1$.

Proposition 5. *Consider the game with evidence and let $\lambda = 1$. In any principal-optimal equilibrium, on-path play is non-stationary and the principal's expected payoff is*

$$e^* - \frac{c(e^*)}{p},$$

where e^* is the minimum of the effort that maximizes $e - \frac{c(e)}{p}$ and the positive root of $\delta pe = c(e)$.

There exists a principal-optimal equilibrium with the following features. Play starts in the "production" phase. In this phase, $e_t = e^*$ and $s_t = c(e^*)$. If $x_t = c(e^*)$, play stays in the production phase. Otherwise, play moves to a "frozen" phase with strictly positive probability and otherwise stays in the production phase. The frozen phase is absorbing and features $e_t = s_t = 0$.

Proof: See Appendix A.

Proposition 5 says that evidence solves the hold-up problem if $p = 1$, in the sense that the principal-optimal equilibrium induces the same effort and payoffs as in the setting without commitment. The logic for why it does so, however, is a little subtle. To see why, note that we would like the principal to not pay an agent who has shirked. Evidence, however, does not reveal anything about an agent's effort, and in particular it does not distinguish a principal who deviates from one who refuses to pay an agent who has deviated. However, evidence does allow the construction of an equilibrium in which the principal is *indifferent* between paying an agent and suffering punishment.

Suppose the principal's continuation payoff equals Π_{t+1} if the agent reveals that $s_t = c(e^*)$ and otherwise equals $\Pi_{t+1} - \frac{1-\delta}{\delta} c(e^*)$. The principal is indifferent between paying $c(e^*)$ and 0 and so is willing to pay $c(e^*)$ on the equilibrium path. If the agent deviates, however, the principal is equally willing to pay 0 in order to punish agent t. This indifference condition essentially uniquely pins down the possible continuation payoffs. Evidence helps cooperation by ensuring that, unlike Proposition 2, an agent cannot shirk and then ask the principal to pay an amount that is *less* than the on-path transfer in exchange for a high continuation payoff. For evidence to serve this role, however, the principal must be punished whenever an agent does not reveal evidence, since otherwise a shirking agent could use non-disclosure to reward a principal who pays him.

A similar logic implies that the equilibrium must punish the principal following non-disclosure if $p < 1$. Principal-optimal equilibria are therefore non-stationary, since a punishment phase is triggered whenever evidence is not available. This punishment phase also lowers the principal's payoff and implies that principal-optimal equilibrium effort will be lower than first-best even if players are very patient. The equilibrium described in Proposition 5 illustrates a particularly simple form of this punishment. Decreasing p makes the principal worse off for two reasons in this equilibrium. First, decreasing p makes it more likely that the game transitions to the frozen phase, which entails permanent low output. Moreover, decreasing p decreases equilibrium effort even in the production phase.

We can re-interpret the disclosure protocol as a restriction on the communication protocol, in the sense that it constrains the kinds of messages an agent can choose for each transfer. Since an agent can only reveal or conceal the actual payment, he cannot promise the principal a high continuation payoff unless she actually pays the amount she would on the equilibrium path, which means the principal has to incur a larger cost in order to give in to a hold-up attempt. Evidence limits the deviations available to agents and so decreases their power to extract surplus relative to the baseline model.

The inefficiency illustrated in Proposition 5 is related to the moral hazard in teams problem in Holmstrom [1982]. Both settings suffer from a lack of what Fudenberg et al. [1994] calls "pairwise identifiability:" given an off-path public signal, players cannot tell whether that signal is the result of a deviation by the principal or a deviation by an agent. It is therefore not surprising that the principal's maximum payoff is constrained even if players are patient. Proposition 5 shows that evidence helps because it means that the principal cannot pay a smaller-than-equilibrium payment without triggering punishment. For evidence to play this role, however, coordinated punishments rely *entirely* on hard evidence, which means that evidence is of limited value to the principal unless it is always available.

4.2 Remedy 2: Bilateral Relationships

In the baseline game, the principal can refuse to pay an agent but cannot otherwise punish him. This section explores how future *bilateral* interactions between the principal and each agent can deter hold-up. We argue that these interactions do indeed mitigate the hold-up threat and so facilitate coordinated punishments, but in a way that is limited by the strength of the bilateral relationship.

The **game with bilateral relationships** allows the principal to play a coordination game with each agent *after* that agent sends a message. At the end of each period in the baseline game, suppose that the principal and agent t simultaneously choose actions $a_t^P \in \mathcal{A}^P$ and $a_t^A \in \mathcal{A}^A$, respectively, with $a_t \equiv (a_t^P, a_t^P) \in \mathcal{A}$. These actions are observed by the principal and agent t but not by any other agent. The principal's and agent t's payoffs are $\pi_t = e_t - s_t + v(a_t)$ and $u_t = s_t - c(e_t) + w(a_t)$, respectively.

The action space $\mathcal{A}$ and payoffs $(v(\cdot), w(\cdot))$ define a simultaneous move game, which we call the **bilateral relationship**. In equilibrium, a_t must correspond to a Nash equilibrium of the bilateral relationship. Let $\mathcal{V}^*$ be the set of Nash equilibrium payoffs in the bilateral relationship. An equilibrium of the game with bilateral relationships essentially selects a payoff $(v, w) \in \mathcal{V}^*$ at the end of each period t. This payoff can be used to reward or punish the principal or agent t for their earlier actions in the period. Note that $\mathcal{V}^*$ is convex because the players have access to a public randomization device.

Our next result characterizes the principal-optimal equilibrium for $\lambda = 1$ in the game with bilateral relationships.

Proposition 6. *Consider the game with bilateral relationships and suppose $\lambda = 1$. Define*

$$\begin{aligned}\overline{v} &= \max_{(v,w)\in\mathcal{V}^*}\{v\},\\ \underline{v} &= \min_{(v,w)\in\mathcal{V}^*}\{v\},\\ \underline{w} &= \min_{(v,w)\in\mathcal{V}^*}\{w\},\\ \overline{L} &= \max_{(v,w)\in\mathcal{V}^*}\{v+w\},\\ \underline{L} &= \underline{v}+\underline{w}.\end{aligned}$$

Suppose $\overline{L} > \underline{L}$. There exists a $\overline{\delta} < 1$ such that if $\delta \geqslant \overline{\delta}$, then $e_t = e^*$ for each $t \geqslant 0$ on the equilibrium path of any principal-optimal equilibrium, where

$$c(e^*) = \min\left\{c(e^{FB}), (\overline{L} - \underline{L}) + (\overline{v} - \underline{v})\right\}. \tag{1}$$

The principal's payoff is $\Pi^* = e^* - c(e^*) + \overline{L} - \underline{w}$ and each agent earns $\underline{w}$.

Proof: See Appendix A.

As in the baseline game, the proof of Proposition 6 identifies two incentive constraints that must be satisfied in any equilibrium. First, the principal must be willing to pay the equilibrium transfer rather than renege, where the punishment for reneging potentially includes lower payoffs in both the bilateral relationship and future interactions with agents who learn of the deviation. Second, each agent must be willing to exert the equilibrium effort rather than shirking and holding up the principal. The bilateral relationship makes shirking less attractive by both punishing the agent for shirking and by punishing the principal for paying an agent who has shirked. We then construct an equilibrium that maximizes total surplus subject to these necessary conditions and then gives all of that surplus to the principal.

The expression (1) highlights two ways in which bilateral relationships encourage effort in equilibrium. First, as represented by $(\overline{L} - \underline{L})$ in (1), a_t can directly reward or punish agent t's effort choice or the principal's payment. In addition to this direct effect, which is familiar from the literature on relational contracting, bilateral relationships can also render coordinated punishments more effective by facilitating truthful communication among agents. This second effect is represented by $(\overline{v} - \underline{v})$ and arises because a_t can be used to reward the principal to refusing to pay, or punish her for paying, an agent who has deviated. An agent therefore cannot demand the entire difference between the highest and lowest continuation payoff after he deviates, which means that bilateral relationships facilitate truthful communication and hence coordinated punishments. If these coordinated punishments are too severe, however, an agent can still hold up the principal. Equilibrium effort is therefore limited by the strength of the bilateral relationship.

Rather than modeling the bilateral relationship as an abstract simultaneous-move game, we could have alternatively considered a setting with long-lived agents who interact repeatedly with the principal. Appendix C takes this approach. We develop a repeated game with private monitoring with N agents, one of whom is randomly selected in each period to play a favor-trading game with the principal. Agents do not observe one another's relationships but can commit to a communication protocol whenever they are active. Long-lived agents substantially complicate the analysis. As in Proposition 6, however, each agent's bilateral relationship with the principal facilitates coordinated punishments, but the strength of that bilateral relationship tightly constrains the severity of coordinated punishments in equilibrium. Thus, while some of the details change, our basic intuition extends to a setting with long-lived agents.

5 Interpreting the Communication Protocol

This section interprets the assumption that agents can commit to a communication protocol.

Commitment is powerful in our setting because it allows agents to credibly signal which messages they will send following each transfer. In our baseline model, for instance, agents are indifferent among all possible messages. Without commitment, we can always construct an equilibrium in which an agent who shirk never follows through on his threats, in which case he cannot hold up the principal. Even in the game without commitment, however, there exist equilibria in which agents' hold-up threats are credible. As our next result shows, we can therefore interpret our commitment assumption as an equilibrium refinement in the game without commitment.

Proposition 7. *In the baseline game, the game with evidence, or the game with bilateral punishments, for any equilibrium of the game with* $\lambda \in [0, 1]$*, there exists an equilibrium of the game with* $\lambda = 0$ *which induces the same distribution over* $\{e_t, s_t\}_{t=0}^{\infty}$.

Proof: See Appendix A.

Proposition 7 immediately follows from Proposition 4 for the baseline game, while a nearly identical argument proves the result for the game with evidence. To prove the result for the game with bilateral relationships, we construct an equilibrium in which agents are punished for threatening hold-up regardless of whether or not they follow through on that threat. Agents are then willing to follow their communication protocols both on- and off-path.

It is useful to compare our commitment assumption to other refinements that capture bargaining in repeated games. For instance, our analysis provides a complementary view to Miller and Watson [2013]. We model bargaining in terms of a commitment assumption, an approach that disciplines our study of solutions to the hold-up problem in Section 4. Moreover, while Miller and Watson [2013] embed a notion of renegotiation-proofness in their model of bargaining, our analysis separates hold-up threats from renegotiation and so our equilibria may entail inefficient continuation play.

Finally, our analysis highlights two reasons why punishments that rely on word-of-mouth are particularly susceptible to abuse. First, an agent who learns of a deviation through a message cannot disentangle two possibilities: either the principal truly deviated, or the communicator deviated by trying to hold up the principal. Therefore, that agent cannot tell whether the principal should actually be punished or not. Compounding this problem is the fact that coordinated punishments are valuable only if an agent's bilateral relationship with the principal is relatively weak. But then the principal has little recourse to directly punish an agent who engages in hold-up.

To illustrate this second point, recall that each agent is indifferent among messages in the baseline game. Consequently, if agents have even a mild preference for following through on their threats, they will act as if they are committed to those threats. To make this point, we consider a game with **ε-compliance preferences**, which is identical to the baseline game except that an agent does not commit to his communication protocol. Instead, immediately before each agent chooses his message, he observes his own preference for sending the message specified by his protocol. With probability $\lambda \in [0,1]$, he earns an extra $\varepsilon > 0$ if he chooses $m_t = \mu_t(s_t)$; otherwise, his utility is independent of his message.

We show that the set of equilibria in the game with ε-compliance preferences is exactly the same as in the baseline game with commitment probability λ.

Proposition 8. *For any $\varepsilon > 0$ and $\lambda \in [0,1]$, there exists an equilibrium of the baseline game that induces a distribution over actions if and only if there exists an equilibrium of the corresponding game with ε-compliance preferences with that distribution over actions.*

Proof: See Appendix A.

Proposition 8 follows from the fact that even a mild preference for choosing $m_t = \mu_t(s_t)$ makes this message the agent's unique best response. Consequently, a given communication protocol generates an identical mapping from payments to messages regardless of whether an agent is committed or simply has ε-compliance preferences. This result extends to the game with evidence so long as agents also have a mild preference for following their disclosure protocols. It does not extend to the game with bilateral relationships, since a punishment in the coordination game can overcome a mild intrinsic preference for following through on a threat. In that setting, Proposition 8 would require that agents' preferences for following the communication protocol are stronger than the *most severe* punishment in the bilateral relationship.

6 Discussion and Conclusion

In many settings, businesses and individuals cooperate with one another only because they fear that partners will spread word of any misbehavior. This paper proposes an under-explored downside to communication as a way to coordinate punishments. We argue that agents may abuse messages intended to report deviations to instead threaten the principal, increasing their payoffs from shirking and so limiting equilibrium cooperation. We have also explored evidence and bilateral relationships as two features of the environment that can mitigate this problem.

In our setting, the principal and agents play highly asymmetric roles and so rely on communication to very different extents. A natural next step would be to consider settings with more symmetric interactions, as in, for example, a network of relationships (e.g., Ali and Miller [2016]). In contrast to our setting, in a more symmetric transaction *both* sides have the opportunity to hold one another up. How do players cooperate in the presence of two-sided hold-up? What networks best facilitate cooperation, and how are rents shared within those networks? How should business associations, communities, and firms structure their communication channels to support strong relational contracts? We hope that our analysis provides both a useful building block and an informative first step towards analyzing these, and other, questions.

References

Nageeb Ali and Ce Liu. Conventions and coalitions in repeated games. Working Paper, 2018.

Nageeb Ali and David Miller. Ostracism and forgiveness. *American Economic Review*, 106 (8):2329–2348, 2016.

Nageeb Ali, David Miller, and David Yang. Renegotiation-proof multilateral enforcement. 2016.

Isaiah Andrews and Daniel Barron. The allocation of future business: Dynamic relational contracts with multiple agents. *American Economic Review*, 106(9):2742–2759, 2016.

Chris Arnold and Robert Smith. Bad form, wells fargo. NPR, 10 2016. URL `https://www.npr.org/templates/transcript/transcript.php?storyId=499805238`.

Yu Awaya and Vijay Krishna. Communication and cooperation in repeated games. Working Paper, 2018.

George Baker, Robert Gibbons, and Kevin Murphy. Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*, 109(4):1125–1156, 1994.

George Baker, Robert Gibbons, and Kevin J. Murphy. Relational contracts and the theory of the firm. *The Quarterly Journal of Economics*, 117(1):39–84, 2002.

Daniel Barron and Michael Powell. Policies in relational contracts. Working Paper, 2018.

Daniel Barron, Jin Li, and Michal Zator. Productivity and debt in relational contracts. Working Paper, 2018.

B. Douglas Bernheim and Michael Whinston. Incomplete contracts and strategic ambiguity. *The American Economic Review*, 88(4):902–932, 1998.

Lisa Bernstein. Beyond relational contracts: Social capital and network governance in procurement contracts. *Journal of Legal Analysis*, 7(2):561–621, 2015.

Simon Board. Relational contracts and the value of loyalty. *American Economic Review*, 101(7):3349–3367, 2011.

Clive Bull. The existence of self-enforcing implicit contracts. *The Quarterly Journal of Economics*, 102(1):147–159, 1987.

Olivier Compte. Communication in repeated games with imperfect private monitoring. *Econometrica*, 66(3):597–626, 1998.

Yuk-Fai Fong and Jin Li. Relational contracts, limited liability, and employment dynamics. 2017.

Richard Freeman. The exit-voice tradeoff in the labor market: Unionism, job tenure, quits, and separations. *The Quarterly Journal of Economics*, 94(4):643–673, 1980.

Richard Freeman and Edward Lazear. An economic analysis of works councils. In Joel Rogers and Wolfgang Streeck, editors, *Works Councils: Consultation, Representation, and Cooperation in Industrial Relations*, pages 27–52. 1995.

Richard Freeman and James Medoff. The two faces of unionism. *The Public Interest*, 57: 69–93, 1979.

Drew Fudenberg, David Levine, and Eric Maskin. The folk theorem with imperfect public monitoring. *Econometrica*, 62(5):997–1039, 1994.

Robert Gibbons and Rebecca Henderson. What do managers do? exploring persistent performance differences among seemingly similar enterprises. In Robert Gibbons and John Roberts, editors, *Handbook of Organizational Economics*, pages 680–731. 2013.

Susanne Goldlucke and Sebastian Kranz. Reconciliating relational contracting and hold-up: A model of repeated negotiations. Working Paper, 2017.

Avner Greif, Paul Milgrom, and Barry Weingast. Coordination, commitment, and enforcement: The case of the merchant guild. *Journal of Political Economy*, 102(4):745–776, 1994.

Yingni Guo and Johannes Hörner. Dynamic allocation without money. Working Paper, 2018.

Yingni Guo and Eran Shmaya. Costly miscalibration. Working Paper, 2018.

Marina Halac. Relational contracts and the value of relationships. *American Economic Review*, 102(2):750–779, 2012.

Marina Halac. Investing in a relationship. *RAND Journal of Economics*, 46(1):165–186, 2015.

Bengt Holmstrom. Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340, 1982.

Johannes Hörner and Nicolas Lambert. Motivational ratings. Working Paper, 2017.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

Michihiro Kandori and Hitoshi Matsushima. Private observation, communication, and collusion. *Econometrica*, 66(3):627–652, 1998.

Jonathan Levin. Multilateral contracting and the employment relationship. *The Quarterly Journal of Economics*, 117(3):1075–1103, 2002.

Jonathan Levin. Relational incentive contracts. *The American Economic Review*, 93(3):835–857, 2003.

Jin Li, Niko Matouschek, and Michael Powell. Power dynamics in organizations. *American Economic Journal: Microeconomics*, 2017. Forthcoming.

Elliot Lipnowski and Joao Ramos. Repeated delegation. 2017.

Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Persuasion via weak institutions. Working Paper, 2018.

Steffen Lippert and Giancarlo Spagnolo. Networks of relations and word-of-mouth communication. *Games and Economic Behavior*, 72(1):202–217, 2011.

R.G. Lipsey and Kelvin Lancaster. The general theory of second best. *The Review of Economic Studies*, 24(1):11–32, 1956-1957.

Bentley MacLeod and James Malcomson. Implicit contracts, incentive compatibility, and involuntary unemployement. *Econometrica*, 57(2):447–480, 1989.

James Malcomson. Relational incentive contracts. In Robert Gibbons and John Roberts, editors, *Handbook of Organizational Economics*, pages 1014–1065. 2013.

James Malcomson. Relational contracts with private information. *Econometrica*, 84(1): 317–346, 2016.

David Miller and Joel Watson. A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350, 2013.

David Miller, Trond Olsen, and Joel Watson. Relational contracting, negotiation, and external enforcement. Working Paper, 2018.

Daehong Min. Bayesian persuasion under partial commitment. Working Paper, 2017.

Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press, Cambridge, UK, 1990.

Jason Proctor. Disgruntled bride ordered to pay 115k after defamatory posts ruin chinese wedding-photo business, 2018. URL `http://www.cbc.ca/news/canada/british-columbia/chinese-wedding-weibo-defamation-1.4556433`.

Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010.

David Robinson and Toby Stuart. Network effects in the governance of strategic alliances. *Journal of Law, Economics, and Organizations*, 23(1):242–273, 2006.

Joel Watson. Perfect bayesian equilibrium: General definitions and illustrations. 2016.

Alexander Wolitzky. Cooperation with network monitoring. *The Review of Economic Studies*, 80(1):395–427, 2013.

A Proofs

A.1 Proof of Proposition 3

Each player's equilibrium payoff is at least 0: the principal can choose not to pay the agents, and agent t can choose not to exert effort or pay the principal. Let $\overline{\Pi}$ be the highest equilibrium payoff for the principal. If $\overline{\Pi} = 0$ then we are done, so suppose $\overline{\Pi} > 0$ and consider the equilibrium profile that gives the principal $\overline{\Pi}$.

In $t = 0$, $e_0 > 0$, since otherwise $\overline{\Pi} \leqslant \delta\overline{\Pi}$ which leads to a contradiction. For each $m \in M$, let $\Pi_1(m)$ be the principal's continuation payoff if $m_0 = m$. Let $\overline{\Pi}_1 = \sup_m \Pi_1(m)$ and $\underline{\Pi}_1 = \inf_m \Pi_1(m)$, with corresponding messages $\overline{m}$ and $\underline{m}$. Agent 0 is willing to choose $e_0 > 0$ only if $s_0 - c(e_0) \geqslant 0$, so $s_0 > 0$. The principal is willing to pay s_0 only if

$$s_0 \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right), \tag{2}$$

so $\overline{\Pi}_1 > \underline{\Pi}_1$.

Consider the following deviation strategy by agent 0: $e_0 = 0$ and

$$\mu_0(s) = \begin{cases} \underline{m} & s_0 < s^*, \\ \overline{m} & s_0 \geqslant s^*, \end{cases}$$

for some $s^* \geqslant 0$. Following this deviation, the principal's unique best response is to pay s^* if

$$(1-\delta)(-s^*) + \delta\left(\lambda\overline{\Pi}_1 + (1-\lambda)\underline{\Pi}_1\right) > \delta\left(\lambda\underline{\Pi}_1 + (1-\lambda)\overline{\Pi}_1\right),$$

where the left-hand side of this inequality is the *smallest* principal payoff if she pays s^* and the right-hand side is the *largest* payoff if she does not. Therefore, the principal's unique best response to this deviation is to play s^* so long as

$$s^* < (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right).$$

Consequently,

$$u_0 \equiv s_0 - c(e_0) \geqslant \max\left\{0, (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right)\right\} \tag{3}$$

in any equilibrium.

Combining (2) and (3) implies that a necessary condition for e_0 and s_0 is that

$$c(e_0) + \max\left\{0, (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right)\right\} \leqslant s_0 \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right). \tag{4}$$

We claim that (4) is sufficient for e_0 and s_0 to be equilibrium choices in $t = 0$, given $\Pi_1(\cdot)$. Indeed, consider the following strategy profile:

1. Agent 0 chooses e_0 and $\mu_0(s) = \begin{cases} \underline{m}, & \text{if } s_0 < c(e_0) + \max\left\{0, (2\lambda-1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right)\right\}, \\ \overline{m}, & \text{otherwise.} \end{cases}$
2. The principal pays s_0 satisfying (4) if agent 0 has not deviated, and some $\tilde{s}_0 \leqslant \max\left\{0, (2\lambda-1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right)\right\}$ that maximizes her continuation payoff otherwise.
3. If agent 0 can choose m_0 (i.e., he is not committed), he chooses $m_0 = \overline{m}$ if no deviation has occurred or if only agent 0 himself deviated, and $m_0 = \underline{m}$ if the principal has deviated.

If agent 0 follows his equilibrium strategy, the principal has no profitable deviation because $\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right) \geqslant s_0$. If agent 0 deviates from (e_0, μ_0), then we need to show that the principal can best respond with some $\tilde{s}_0 \leqslant \max\left\{0, (2\lambda-1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right)\right\}$. Indeed, the principal is willing to pay some $\tilde{s}_0 > 0$ if

$$-\tilde{s}_0 + \frac{\delta}{1-\delta}\left(\lambda\overline{\Pi}_1 + (1-\lambda)\underline{\Pi}_1\right) \geqslant \frac{\delta}{1-\delta}\left(\lambda\underline{\Pi}_1 + (1-\lambda)\overline{\Pi}_1\right),$$

or equivalently,

$$\frac{\delta}{1-\delta}(2\lambda-1)\left(\overline{\Pi}_1-\underline{\Pi}_1\right)\geqslant\tilde{s}_0.$$

The left-hand side is weakly less than agent 0's on-path payoff. So agent 0 has no profitable deviation from (e_0,μ_0). Agent 0 has no profitable deviation from m_0 because he is indifferent among all messages. Continuation play is an equilibrium by assumption, so this strategy profile is an equilibrium, which proves that (4) is sufficient.

The principal-optimal equilibrium therefore solves

$$\begin{array}{rcl}\overline{\Pi}=\max_{e_0,s_0,\overline{\Pi}_1,\underline{\Pi}_1} & & e_0-s_0+\frac{\delta}{1-\delta}\overline{\Pi}_1\\ \text{s.t.} & & (4)\quad\text{and}\quad 0\leqslant\underline{\Pi}_1\leqslant\overline{\Pi}_1\leqslant\overline{\Pi}.\end{array}\tag{5}$$

If $\lambda\leqslant\frac{1}{2}$, then (4) simplifies to

$$c(e_0)\leqslant s_0\leqslant\frac{\delta}{1-\delta}\left(\overline{\Pi}_1-\underline{\Pi}_1\right).$$

Clearly, $s_0=c(e_0)$ maximizes the principal's payoff, in which case either $e_0=e^{FB}$ or $c(e_0)=\frac{\delta}{1-\delta}\left(\overline{\Pi}_1-\underline{\Pi}_1\right)$. It is therefore optimal to set $\overline{\Pi}_1=\overline{\Pi}$, which implies that play is stationary on-path, and (without loss) $\underline{\Pi}_1=0$. Consequently, $\overline{\Pi}=e^{FB}-c(e^{FB})$ if $c(e^{FB})\geqslant\delta e^{FB}$, and otherwise

$$c(e_0)=\frac{\delta}{1-\delta}\overline{\Pi}=\frac{\delta}{1-\delta}(e_0-c(e_0))$$

or equivalently, e_0 is the positive root of $c(e_0)=\delta e_0$, as desired.

Consider (5) for $\lambda>\frac{1}{2}$. The transfer s_0 is optimally as small as possible, which by (4) implies that $s_0=c(e_0)+(2\lambda-1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1-\underline{\Pi}_1\right)$. It is optimal to set $\overline{\Pi}_1=\overline{\Pi}$ because we can otherwise increase $\overline{\Pi}_1,\underline{\Pi}_1$ holding $\overline{\Pi}_1-\underline{\Pi}_1$ constant. So $e_t=e_0$ for all $t\geqslant 0$ in any principal-optimal equilibrium.

Since s_0 is increasing in $\frac{\delta}{1-\delta}\left(\overline{\Pi}_1-\underline{\Pi}_1\right)$, it is optimal to set $\underline{\Pi}_1$ so that both inequalities in (4) hold with equality, which implies that $s_0 = \frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right)$. Therefore, we must have

$$\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}_1\right) = c(e_0) + (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}_1\right)$$

or

$$2(1-\lambda)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}_1\right) = c(e_0).$$

We can plug expressions for s_0 and $c(e_0)$ into (5) to rewrite this as an optimization with a single choice variable, $\underline{\Pi}_1$. That is,

$$\begin{array}{ll} \overline{\Pi} = \max_{\underline{\Pi}_1} & c^{-1}\left(2(1-\lambda)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}_1\right)\right) + \frac{\delta}{1-\delta}\underline{\Pi}_1 \\ \text{s.t.} & 0 \leqslant \underline{\Pi}_1 \leqslant \overline{\Pi} \end{array}.$$

Since c^{-1} is concave and its derivative at 0 is $+\infty$, the optimal $\underline{\Pi}_1$ is either interior or equal to 0. If the optimal $\underline{\Pi}_1$ is interior, it is given by the first-order condition

$$(c^{-1})'\left(2(1-\lambda)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}_1\right)\right) = \frac{1}{2(1-\lambda)}.$$

In that case, e_0 solves $c'(e_0) = 2(1-\lambda) < 1$. If $\underline{\Pi}_1 = 0$, then $c(e_0) = (2-2\lambda)\frac{\delta}{1-\delta}\overline{\Pi}$. This gives the equilibrium effort $e^*(\lambda)$.

In either case, we can then solve for $\overline{\Pi}$:

$$\overline{\Pi} = e_0 - \frac{c(e_0)}{2-2\lambda},$$

as desired. Since $e - \frac{c(e)}{2-2\lambda}$ strictly increases in e for any e such that $c'(e) \leqslant 2(1-\lambda)$, and $e^*(\lambda)$ strictly decreases in λ, the principal's payoff also strictly decreases in λ. ■

A.2 Proof of Proposition 4

Let $\lambda' > \lambda$, and consider an equilibrium of the game with commitment probability λ'. We refer to this equilibrium as the *actual equilibrium.* Define the following strategy profile, which we refer to as the *candidate profile,* of the game with commitment probability λ. In each history in each period t:

1. Agent t chooses e_t and μ_t as in the actual equilibrium.

2. The principal chooses s_t as in the actual equilibrium.

3. For any μ_t, agent t plays the following mixed strategy whenever he can freely choose m_t: with probability $\frac{\lambda'-\lambda}{1-\lambda}$, he sends $m_t = \mu_t(s_t)$. Otherwise, he chooses m_t as in the actual equilibrium.

It suffices to show that players cannot profitably deviate in any period $t \geqslant 0$.

By construction, for any realized m_t, continuation payoffs from $t+1$ are identical in the candidate profile and the actual equilibrium. In stage (3), agent t is indifferent among messages and so is willing to follow the specified mixed strategy. Consequently, for any e_t, μ_t, the mapping from s_t to m_t is identical in the candidate profile and the actual equilibrium. The principal is therefore willing to pay the specified s_t. But then any choice of e_t and μ_t induces the same s_t in the candidate profile and the actual equilibrium, so agent t is willing to play the specified e_t and μ_t.

We conclude that the candidate profile is in fact an equilibrium of the game with commitment probability λ for any $\lambda < \lambda'$, as desired. ■

A.3 Proof of Proposition 5

Let $\overline{\Pi}$ be the highest equilibrium payoff for the principal. Let e^* and s^* be the equilibrium effort and transfer, respectively, in period 0 of the principal-optimal equilibrium. Let $\overline{\Pi}_x$ and $\underline{\Pi}_x$ be the highest and lowest equilibrium continuation payoffs, respectively, if the evidence realization is x for $x \in \{\emptyset\} \cup \mathbb{R}_+$.

Define

$$\hat{s} = \min\left\{\arg\max_{s\in[0,s^*]}\left\{-s + \frac{\delta}{1-\delta}\left((1-p)\underline{\Pi}_\emptyset + p\min\left\{\underline{\Pi}_s, \underline{\Pi}_\emptyset\right\}\right)\right\}\right\}.$$

Intuitively, $\hat{s}$ is the *smallest* transfer that maximizes the principal's payoff if her continuation payoff is as low as possible given her transfer.

The principal is willing to pay s^* only if

$$-s^*+\frac{\delta}{1-\delta}\left((1-p)\overline{\Pi}_\emptyset + p\max\left\{\overline{\Pi}_{s^*}, \overline{\Pi}_\emptyset\right\}\right) \geqslant -\hat{s}+\frac{\delta}{1-\delta}\left((1-p)\underline{\Pi}_\emptyset + p\min\left\{\underline{\Pi}_{\hat{s}}, \underline{\Pi}_\emptyset\right\}\right). \quad (6)$$

Agent 0 is willing to choose e^* only if

$$s^* - c(e^*) \geqslant \frac{\delta}{1-\delta}\left(\overline{\Pi}_\emptyset - \left((1-p)\underline{\Pi}_\emptyset + p\min\left\{\underline{\Pi}_{\hat{s}}, \underline{\Pi}_\emptyset\right\}\right)\right). \quad (7)$$

If this inequality does not hold, then agent 0 could deviate by choosing $e_0 = 0$ and choosing the following communication protocol: if $s_0 = s^* - c(e^*) + \varepsilon$ for a small $\varepsilon > 0$, send the message that leads to $\overline{\Pi}_\emptyset$; otherwise, send the message that minimizes the principal's payoff. Faced with this communication protocol, the principal either pays $s^* - c(e^*) + \varepsilon$ or $\hat{s}$. If (7) did not hold, then the principal would strictly prefer to pay $s^* - c(e^*) + \varepsilon$ for sufficiently small $\varepsilon > 0$.

Let us suppose that (6) and (7) are sufficient conditions for equilibrium. In that case, it is without loss to set $\hat{s} = 0$. Indeed, if $\hat{s} > 0$, then

$$-\hat{s} + \frac{\delta}{1-\delta}\left((1-p)\underline{\Pi}_\emptyset + p\min\left\{\underline{\Pi}_{\hat{s}}, \underline{\Pi}_\emptyset\right\}\right) > \frac{\delta}{1-\delta}\left((1-p)\underline{\Pi}_\emptyset + p\min\left\{\underline{\Pi}_0, \underline{\Pi}_\emptyset\right\}\right), \quad (8)$$

which holds only if $\underline{\Pi}_0 < \min\left\{\underline{\Pi}_{\hat{s}}, \underline{\Pi}_\emptyset\right\}$. We can therefore set $\underline{\Pi}_0 \in [0, \min\left\{\underline{\Pi}_{\hat{s}}, \underline{\Pi}_\emptyset\right\}]$ so that (8) holds with equality. This perturbation relaxes (7), does not affect (6), and results in $\hat{s} = 0$. So we can set $\hat{s} = 0$.

Given this perturbation, we can combine (6) and (7) into a single necessary condition:

$$\frac{\delta}{1-\delta}\left(p\left(\max\left\{\overline{\Pi}_{s^*},\overline{\Pi}_{\emptyset}\right\}-\min\left\{\underline{\Pi}_0,\underline{\Pi}_{\emptyset}\right\}\right)+(1-p)\left(\overline{\Pi}_{\emptyset}-\underline{\Pi}_{\emptyset}\right)\right)\geqslant s^*$$
$$\geqslant c(e^*)+\frac{\delta}{1-\delta}\left(p\left(\overline{\Pi}_{\emptyset}-\min\left\{\underline{\Pi}_0,\underline{\Pi}_{\emptyset}\right\}\right)+(1-p)\left(\overline{\Pi}_{\emptyset}-\underline{\Pi}_{\emptyset}\right)\right).$$

In a principal-optimal equilibrium, s^* is as small as possible, so

$$s^*=c(e^*)+\frac{\delta}{1-\delta}\left(p\left(\overline{\Pi}_{\emptyset}-\min\left\{\underline{\Pi}_0,\underline{\Pi}_{\emptyset}\right\}\right)+(1-p)\left(\overline{\Pi}_{\emptyset}-\underline{\Pi}_{\emptyset}\right)\right).$$

Our necessary condition can then be simplified to

$$\frac{c(e^*)}{p}\leqslant\frac{\delta}{1-\delta}\left(\max\left\{\overline{\Pi}_{s^*},\overline{\Pi}_{\emptyset}\right\}-\overline{\Pi}_{\emptyset}\right). \tag{9}$$

Since neither $\underline{\Pi}_0$ nor $\underline{\Pi}_{\emptyset}$ appear in (9), setting $\underline{\Pi}_0\geqslant\underline{\Pi}_{\emptyset}=\overline{\Pi}_{\emptyset}$ minimizes s^*. Similarly, setting $\overline{\Pi}_{s^*}=\overline{\Pi}$ is clearly optimal, since we can adjust $\overline{\Pi}_{\emptyset}$ so that $\overline{\Pi}_{s^*}-\overline{\Pi}_{\emptyset}$ is constant. Consequently, an upper bound on the principal's payoff is given by

$$\begin{aligned}\overline{\Pi} &= \max_{e^*,\overline{\Pi}_{\emptyset}}\left((1-\delta)(e^*-c(e^*))+\delta\left(p\overline{\Pi}+(1-p)\overline{\Pi}_{\emptyset}\right)\right)\\ \text{s.t.} &\qquad (9)\quad\text{and}\quad 0\leqslant\overline{\Pi}_{\emptyset}\leqslant\overline{\Pi}.\end{aligned} \tag{10}$$

If $\overline{\Pi}>0$, then $e^*>0$, in which case $\overline{\Pi}>\overline{\Pi}_{\emptyset}$. In that case, (9) holds with equality, since otherwise we could increase $\overline{\Pi}_{\emptyset}$ and improve the principal's payoff. But then (9) implies that $\frac{\delta}{1-\delta}\overline{\Pi}_{\emptyset}=\frac{\delta}{1-\delta}\overline{\Pi}-\frac{c(e^*)}{p}$, which means that the objective of (10) simplifies to

$$\begin{aligned}\overline{\Pi} &= (1-\delta)(e^*-c(e^*))+\delta\overline{\Pi}-(1-\delta)(1-p)\left(\tfrac{c(e^*)}{p}\right)\Longrightarrow\\ \overline{\Pi} &= e^*-\tfrac{c(e^*)}{p},\end{aligned}$$

as desired.

Let $e^{max}(p)=\arg\max_e\left\{e-\frac{c(e)}{p}\right\}$. If $\delta\geqslant\frac{c(e^{max}(p))}{pe^{max}(p)}$, then $e^*=e^{max}(p)$ and $\overline{\Pi}_{\emptyset}$ solves (9) with equality. Otherwise, $\overline{\Pi}_{\emptyset} = 0$, in which case e^* is the positive root of $\delta = \frac{c(e)}{pe}$.

It remains to show that the solution to this relaxed problem can be attained in equilibrium. We do so by showing that an equilibrium of the type described in the proposition attains this payoff. In particular, consider an equilibrium that starts in the "production phase." In this phase, play in each period is:

1. Agent t chooses $e_t = e^*$, $\mu_t(s_t) = m$, and $\kappa_t(x_t) = x_t$.

2. On-path, the principal pays $s_t = c(e^*)$. Following a deviation, the principal pays $s_t = 0$.

3. If $\kappa_t(x_t) = c(e^*)$, then stay in the production phase. Otherwise, stay in the production phase with probability $\gamma \in [0, 1]$ and otherwise transition to the "frozen phase," where γ solves $\overline{\Pi}_{\emptyset} = \gamma\overline{\Pi}$ for the optimal $\overline{\Pi}_{\emptyset}$ from (10).

In the frozen phase, which is absorbing, players play the one-shot equilibrium in each period and the principal earns 0.

Play in the frozen phase is clearly an equilibrium. In the production phase, the principal is willing to pay $s_t = c(e^*)$ on the equilibrium path because

$$-c(e^*) + \frac{\delta}{1-\delta}\left(p\overline{\Pi} + (1-p)\gamma\overline{\Pi}\right) \geqslant \frac{\delta}{1-\delta}\gamma\overline{\Pi}$$

by choice of γ and by (9). Following any deviation by agent t, the principal pays 0. She is willing to pay 0 because (9) holds with equality. Given the principal's actions, the agent can earn no more than 0 and so has no profitable deviation. So this strategy profile is an equilibrium that yields payoff

$$(1-\delta)\left(e^* - c(e^*)\right) + \delta p\overline{\Pi} + \delta(1-p)\overline{\Pi}_{\emptyset} = \overline{\Pi},$$

as desired. ■

A.4 Proof of Proposition 6

Consider period t of an equilibrium in the game with bilateral relationships. Let on-path effort, transfer, and coordination game payoffs equal e^*, s^*, and (v^*, w^*), respectively. Let $\overline{\Pi}$ and $\underline{\Pi}$ be the highest and lowest continuation payoffs for the principal in this equilibrium, and let $\overline{m}$ and $\underline{m}$ be the messages that induce those payoffs.

The payment s^* must satisfy

$$s^* \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right) + v^* - \underline{v}. \tag{11}$$

Agent t can always deviate by choosing $e_t = 0$ and then choosing an optimal communication protocol. By deviating in this way, agent t earns a transfer of no less than $\hat{s}$ in equilibrium, where $\hat{s}$ is the largest transfer that satisfies

$$-s + \overline{v} + \frac{\delta}{1-\delta}\underline{\Pi} < -\hat{s} + \underline{v} + \frac{\delta}{1-\delta}\overline{\Pi}$$

for any $s < \hat{s}$. Therefore, agent t receives a transfer of at least

$$\hat{s} = \max\left\{0, \underline{v} - \overline{v} + \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right)\right\} - \varepsilon$$

for any $\varepsilon > 0$ by deviating.

Given the argument above, a necessary condition for agent t to have no profitable deviation is

$$\max\left\{0, \underline{v} - \overline{v} + \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right)\right\} + \underline{w} \leqslant s^* - c(e^*) + w^*. \tag{12}$$

Combining (11) and (12), a necessary condition for equilibrium is that

$$\max\left\{0, \underline{v} - \overline{v} + \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right)\right\} + \underline{w} + c(e^*) - w^* \leqslant \frac{\delta}{1-\delta}(\overline{\Pi} - \underline{\Pi}) + v^* - \underline{v}. \tag{13}$$

Note that $v^* + w^* \leqslant \overline{L}$. If $\underline{v} - \overline{v} + \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right) < 0$, then increasing $\frac{\delta}{1-\delta}(\overline{\Pi} - \underline{\Pi})$ relaxes (13).

Suppose that we can increase $\frac{\delta}{1-\delta}(\overline{\Pi} - \underline{\Pi})$ without bound. Then a necessary condition for (13) is that

$$\underline{v} - \overline{v} + \underline{w} + c(e^*) \leqslant \overline{L} - \underline{v},$$

or

$$c(e^*) \leqslant \overline{L} - \underline{L} + \overline{v} - \underline{v}. \tag{14}$$

Each agent t can earn no less than $\underline{w}$. The principal's payoff cannot exceed

$$\begin{aligned} \Pi^* \quad &= \quad \max_{e^*} \left\{ e^* - c(e^*) + \overline{L} - \underline{w} \right\} \\ \text{s.t.} \quad & \qquad\qquad (14). \end{aligned}$$

The solution to this problem is clearly the effort level given by (1). To prove the claim, therefore, it suffices to find a discount factor $\overline{\delta} < 1$ such that for any $\delta \geqslant \overline{\delta}$, the principal earns Π^*. This will immediately imply that e_t satisfies (1) in each period on the equilibrium path.

We construct a strategy profile that attains the upper bound and give conditions under which it is an equilibrium. Let e^* satisfy (1) and $(v^*, w^*) \in \arg\max_{(v,w)\in\mathcal{V}^*}\{v + w\}$. Define

$$s^* = \underline{w} - w^* + c(e^*).$$

Let ρ be the solution to

$$\frac{\delta}{1-\delta}\rho\left(\Pi^* - \underline{v}\right) = \overline{v} - \underline{v}. \tag{15}$$

Since $\overline{L} - \underline{L} > 0$ and $\Pi^* \geqslant \overline{L} - \underline{w}$, $\Pi^* - \underline{v} > 0$. Consequently, there exists a $\overline{\delta} < 1$ such that $\rho \in [0, 1]$ if $\delta \geqslant \overline{\delta}$.

For $\delta \geqslant \overline{\delta}$, consider the following strategy. In each t:

1. If $m_{t'} = \overline{m}$ for all $t' < t$:

(a) Agent t chooses $e_t = e^*$ to satisfy (1), and

$$\mu_t(s_t) = \begin{cases} \underline{m} & s_t < s^* \\ \overline{m} & s_t \geqslant s^* \end{cases}.$$

(b) If agent t did not deviate, the principal pays $s_t = s^*$ if $s^* > 0$, or agent t pays $s_t = s^*$ if $s^* < 0$. If agent t has deviated, the principal pays nothing and agent t pays $s_t = -(w^P - \underline{w})$, where $(v^P, w^P) \in \mathcal{V}^*$ satisfies $v^P = \overline{v}$.

(c) If neither the principal nor agent t has deviated in period t, $(v_t, w_t) = (v^*, w^*)$ in the coordination game. If the principal has deviated in s_t (regardless of agent t's actions), (v_t, w_t) satisfies $v_t = \underline{v}$. If agent t has deviated in e_t or μ_t but pays $s_t = -(w^P - \underline{w})$, then $(v_t, w_t) = (v^P, w^P)$; if he has deviated in s_t, then (v_t, w_t) satisfies $w_t = \underline{w}$.

2. If $m_{t-1} \neq \overline{m}$: Using the public randomization device, with probability $(1 - \rho) \in [0, 1]$ continuation play is as if $m_{t-1} = \overline{m}$. Otherwise, play $e_t = 0$ and (v_t, w_t) such that $v_t = \underline{v}$ in this and all future periods.

Under this strategy profile, the principal earns Π^* on the equilibrium path. Therefore, it suffices to show that this strategy profile is an equilibrium.

The players have no profitable deviation from Step 2 of this equilibrium, since there exists an equilibrium of the coordination game with $v_t = \underline{v}$. Therefore, the principal's continuation payoff equals $\underline{v}$ in this punishment phase.

In Step 1(c), all specified payoffs are elements of $\mathcal{V}^*$ and so players cannot profitably deviate.

In Step 1(b), suppose agent t has not yet deviated. If $s^* > 0$, then the principal has no profitable deviation if

$$-s^* + v^* + \frac{\delta}{1-\delta}\Pi^* \geqslant \underline{v} + \frac{\delta}{1-\delta}\left((1-\rho)\Pi^* + \rho\underline{v}\right)$$

or

$$-\left(\underline{w}-w^{*}+c(e^{*})\right)+v^{*}-\underline{v}+\frac{\delta}{1-\delta}\rho\left(\Pi^{*}-\underline{v}\right)\geqslant 0.$$

By the definition of ρ, this inequality is equivalent to

$$-c(e^{*})+\overline{L}-\underline{L}+\overline{v}-\underline{v}\geqslant 0,$$

which holds by (14). If $s^{*}<0$, then agent t has no profitable deviation if

$$-s^{*}\leqslant w^{*}-\underline{w}$$

or

$$-\left(\underline{w}-w^{*}+c(e^{*})\right)\leqslant w^{*}-\underline{w},$$

which holds because $c(e^{*})\geqslant 0$.

Now, consider Step 1(b) and suppose agent t deviated in μ_t or e_t. Agent t has no profitable deviation from $s_t=-(w^{P}-\underline{w})\leqslant 0$ because he earns w^{P} if he does so and $\underline{w}$ otherwise. The principal pays nothing so long as for any $\hat{s}>0$,

$$-\hat{s}+\underline{v}+\frac{\delta}{1-\delta}\Pi^{*}\leqslant\overline{v}+\frac{\delta}{1-\delta}\left((1-\rho)\Pi^{*}+\rho\underline{v}\right).$$

Rearranging and setting $\hat{s}=0$, we require

$$\frac{\delta}{1-\delta}\rho\left(\Pi^{*}-\underline{v}\right)\leqslant\overline{v}-\underline{v},$$

which is implied by (15). Therefore, the players have no profitable deviation in Step 1(b).

If agent t deviates in Step 1(a), then he earns no more than $\underline{w}$ by the argument above. But he also earns $\underline{w}$ on the equilibrium path. So agent t has no profitable deviation from Step 1(a). We conclude that the specified strategy profile is an equilibrium and therefore a principal-optimal equilibrium, and moreover, that in any principal-optimal equilibrium, e_t

satisfies (1) in each period $t \geqslant 0$ on the equilibrium path. ■

A.5 Proof of Proposition 7

This result follows immediately from Proposition 4 in the baseline game. In the game with evidence, the agents are again indifferent over messages and so the result holds by an argument nearly identical to Proposition 4.

Consider the game with bilateral relationships. Let σ^* be an equilibrium for $\lambda \geqslant 0$, and consider the following strategy profile of the game with $\lambda = 0$: in each period $t \geqslant 0$,

1. Agent t chooses e_t, μ_t as in σ^*.

2. The principal chooses s_t as in σ^*.

3. Using the public randomization device: with probability λ, agent t chooses $m_t = \mu_t(s_t)$. Otherwise, agent t chooses m_t as in σ^*.

4. If the agent follows this message strategy, a_t is as in σ^*; otherwise, a_t is chosen from those actions played in σ^* in order to minimize agent t's payoff.

No player has a profitable deviation from a_t because a_t is always an equilibrium of the simultaneous move game at the end of the period. By the choice of a_t following a deviation in m_t, agent t has a weak incentive to follow the specified message strategy m_t. But then the principal and agent t have no profitable deviation from e_t, μ_t, or s_t, since continuation play is exactly as in σ^*. So this strategy profile is an equilibrium of the game without commitment, as desired. ■

A.6 Proof of Proposition 8

Fix $\varepsilon > 0$. Consider an equilibrium of the baseline game with ε-compliance preferences. Agent t is indifferent among messages with probability $1 - \lambda$ and so can send a message exactly as in the baseline game. With probability λ, agent t's unique optimum is to choose

$m_t = \mu_t(s_t)$ and so he will do so in every equilibrium. So the mapping from μ_t and s_t to m_t is identical to that in baseline game, which proves the result. ■

B The Payoff Frontier in the Baseline Game

In this Appendix, we fully characterize the payoff frontier of the baseline game for all $\lambda \in [0,1]$. Define $U(\pi)$ as the *largest* agents' joint utility that can be attained in equilibrium if the principal earns payoff π.

Lemma B.1. *Suppose that $\lambda = 0$ and $\delta \geqslant \frac{c\left(e^{FB}\right)}{e^{FB}}$. Among all efficient equilibria, the one that maximizes agents' joint utility generates $\delta e^{FB} - c\left(e^{FB}\right)$ for the agents, and $(1-\delta)e^{FB}$ for the principal.*

Proof of Lemma B.1. Consider period t after agent t has put effort e^{FB}. The principal's equilibrium payoff should be at least $(1-\delta)e^{FB}$ given that he can pocket the output e^{FB} produced by agent t. We can modify the profile in the proof of Proposition 1, so that the principal's payoff is indeed $(1-\delta)e^{FB}$.

Let $s_t = \delta e^{FB}$ be the payment on path. If the payment is below this level, agent t sends $m_t \neq m^*$. For this to be an equilibrium, it must be true that $\delta e^{FB} \geqslant c\left(e^{FB}\right)$. Under this condition, agent t has no incentive to deviate. ■

We first characterize the payoff frontier for $\lambda \leqslant 1/2$. We show that when π is small, we cannot have an efficient equilibrium even for high δ. Yet, inefficiency occurs only in period 0.

Proposition B.1. *Suppose that $\lambda \leqslant 1/2$. If $\delta \geqslant \frac{c\left(e^{FB}\right)}{e^{FB}}$, the payoff frontier $U(\pi)$ is given by:*

$$\begin{cases} (1-\delta)\left(\frac{\pi}{1-\delta} - c\left(\frac{\pi}{1-\delta}\right)\right) + \delta\left(e^{FB} - c\left(e^{FB}\right)\right) - \pi, & \text{if } \pi \in \left[0, (1-\delta)e^{FB}\right), \\ e^{FB} - c\left(e^{FB}\right) - \pi, & \text{if } \pi \in \left[(1-\delta)e^{FB}, e^{FB} - c\left(e^{FB}\right)\right]. \end{cases}$$

To achieve $(\pi, U(\pi))$, period 0's effort is $\min\left\{\frac{\pi}{1-\delta}, e^{FB}\right\}$, and all future periods' effort is e^{FB}.

If $\delta < \frac{c\left(e^{FB}\right)}{e^{FB}}$, the payoff frontier $U(\pi)$ is given by:

$$(1-\delta)\left(\frac{\pi}{1-\delta} - c\left(\frac{\pi}{1-\delta}\right)\right) + \delta\left(\tilde{e} - c\left(\tilde{e}\right)\right) - \pi, \text{ for } \pi \in \left[0, \tilde{e} - c\left(\tilde{e}\right)\right],$$

where $\tilde{e}$ is the positive root of $c(e) = \delta e$. To achieve $(\pi, U(\pi))$, period 0's effort is $\frac{\pi}{1-\delta}$, and all future periods' effort is $\tilde{e}$. The surplus-maximizing equilibrium is unique and coincides with the principal-optimal equilibrium.

Proof of Proposition B.1. We prove the result for $\lambda = 0$. The argument for $\lambda \in (0, 1/2]$ is similar and hence omitted.

We choose e, s for period 0, and the continuation principal's payoff and agents' joint utility (Π_p, U_a), to maximize agents' joint utility. The payoff $U(\pi)$ is equal to the maximal:

$$U(\pi) = \max_{e,s,\Pi_p,U_a} (1-\delta)(s - c(e)) + \delta U_a,$$

subject to the constraints:

$$\pi = (1-\delta)(e - s) + \delta \Pi_p, \tag{16}$$

$$s - c(e) \geqslant 0, \tag{17}$$

$$(1-\delta)(-s) + \delta \Pi_p \geqslant 0, \tag{18}$$

$$0 \leqslant U_a \leqslant U(\Pi_p).$$

The first condition says that the principal's ex ante payoff is indeed π. The second says that agent 0 is willing to put effort e. The third says that the principal is willing to pay s. The last says that the continuation payoff (Π_p, U_a) must be generated by an equilibrium.

We begin with the case that $\delta \geqslant \frac{c\left(e^{FB}\right)}{e^{FB}}$. We have argued in Proposition 1 and Lemma B.1 that for $\pi \in \left[(1-\delta)e^{FB}, e^{FB} - c\left(e^{FB}\right)\right]$, there exists an efficient equilibrium that gives the principal π. Therefore, agents' joint utility is $U(\pi) = e^{FB} - c\left(e^{FB}\right) - \pi$ for π in this region. Hence, we now focus on $\pi < (1-\delta)e^{FB}$. Substituting π from (16) into the objective, we have

$$U(\pi) + \pi = (1-\delta)(e - c(e)) + \delta(U_a + \Pi_p).$$

We can treat this as the new objective function. Combining (16) and (18), we have

$$e \leqslant \frac{\pi}{1-\delta}.$$

Therefore, the formula of $U(\pi)$ given in the proposition is indeed an upper bound, since $e \leqslant \frac{\pi}{1-\delta}$ and $U_a + \Pi_p \leqslant e^{FB} - c\left(e^{FB}\right)$. This upper bound can be achieved if we choose

$$e = \frac{\pi}{1-\delta}, \quad s = \frac{\delta}{1-\delta}\Pi_p, \quad \Pi_p = (1-\delta)e^{FB}, \quad U_a = e^{FB} - c\left(e^{FB}\right) - \Pi_p.$$

It is easy to verify that (17) is also satisfied. We have freedom with respect to Π_p. We can choose Π_p to be anywhere between $(1-\delta)e^{FB}$ and $e^{FB} - c\left(e^{FB}\right)$.

We proceed to the case that $\delta < \frac{c\left(e^{FB}\right)}{e^{FB}}$. Given the function of $U(\pi)$, the optimum is achieved by:

$$e = \frac{\pi}{1-\delta}, \quad s = \frac{\delta}{1-\delta}\Pi_p, \quad \Pi_p = \tilde{e} - c\left(\tilde{e}\right), \quad U_a = 0.$$

Both (16) and (18) are satisfied as equalities. Since $\pi \leqslant \tilde{e} - c\left(\tilde{e}\right)$ and $c\left(\tilde{e}\right) = \delta\tilde{e}$, we have that $\frac{\pi}{1-\delta} \leqslant \tilde{e}$. Given that $\frac{\delta}{1-\delta}\left(\tilde{e} - c\left(\tilde{e}\right)\right) \geqslant c\left(\tilde{e}\right)$, (17) is satisfied as well. This is the solution for two reasons: (i) we have argued that e must be smaller than $\frac{\pi}{1-\delta}$; (ii) since $U'(\pi) = -c'\left(\frac{\pi}{1-\delta}\right) > -1$, it is optimal to choose Π_p as high as possible. This completes the proof. ■

This proposition is quite intuitive. Take the high δ case as an example. The principal's equilibrium payoff consists of agent 0's effort and her continuation payoff right before she pays the agent. This continuation payoff right before the principal pays must be weakly positive, since the principal can deviate and pay nothing. Therefore, agent 0's effort is at most $\frac{\pi}{1-\delta}$. This upper bound is lower than e^{FB} when π is sufficiently small. Therefore, the effort level in period 0 is inefficiently low for small π. However, this inefficiency lasts for only one period, since we can ask the principal to pay agent 0 a high amount. In return, we promise the principal a higher continuation payoff from tomorrow on. Thus, efficiency is restored from period 1 on.

For the rest of this section, we focus on the case with $\lambda > 1/2$. For some parameter region with $\lambda > 1/2$, full efficiency can still be achieved in equilibrium.

Lemma B.2. *Suppose that $\lambda > 1/2$. There exists an efficient equilibrium if and only if $c\left(e^{FB}\right) \leqslant (2-2\lambda)\delta e^{FB}$. Among efficient equilibria, the principal-optimal equilibrium gives her $e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}$.*

Proof of Lemma B.2. Consider an arbitrary period t. Let $e > 0$ be agent t's effort on path, and s be the principal's payment. Let $\overline{\Pi}$ and $\underline{\Pi}$ be the highest and lowest continuation payoffs that agent t's messages can induce. We let $\overline{m}$ and $\underline{m}$ be the messages that induce $\overline{\Pi}$ and $\underline{\Pi}$. Consider the following deviation strategy by agent t. He puts effort zero and announces the following communication protocol: if the principal's payment is above s^*, he sends $\overline{m}$; if not, he sends $\underline{m}$. The principal's payoff if he pays s^* is at least:

$$(1-\delta)(-s^*) + \delta\left(\lambda\overline{\Pi} + (1-\lambda)\underline{\Pi}\right).$$

The principal's payoff if he doesn't pay as much as s^* is at most:

$$\delta\left(\lambda\underline{\Pi} + (1-\lambda)\overline{\Pi}\right).$$

Therefore, the principal is willing to pay s^* if:

$$s^* \leqslant (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right).$$

This implies that agent t's payoff is at least $(2\lambda-1)\frac{\delta}{1-\delta}\left(\overline{\Pi}-\underline{\Pi}\right)$. If not, agent t can deviate and get a payoff that is arbitrarily close to $(2\lambda-1)\frac{\delta}{1-\delta}\left(\overline{\Pi}-\underline{\Pi}\right)$. (We can choose the strategy for the non-committed agent so that the principal won't concede to any demand more than this.) Therefore, we have:

$$s - c(e) \geqslant (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right). \tag{19}$$

On the other hand, the principal must be willing to pay s instead of pocketing s:

$$s \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right). \tag{20}$$

In an efficient equilibrium, the effort level in every period must be e^{FB}. Inequality (19) and (20) must hold. This implies that:

$$c\left(e^{FB}\right) + (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right) \leqslant s \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right) \Longrightarrow \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right) \geqslant \frac{c\left(e^{FB}\right)}{2-2\lambda}.$$

The payment s is at least:

$$s \geqslant c\left(e^{FB}\right) + (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right) \geqslant \frac{c\left(e^{FB}\right)}{2-2\lambda}.$$

Therefore, the principal's payoff in an arbitrary period is at most:

$$e^{FB} - s \leqslant e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}.$$

This is also the upper bound on the principal's continuation payoff. There exists $\underline{\Pi}, \overline{\Pi}$ satisfying (19) and (20) if and only if:

$$\frac{\delta}{1-\delta}\left(e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}\right) \geqslant \frac{c\left(e^{FB}\right)}{2-2\lambda}.$$

This is equivalent to the condition in the lemma. ■

The following proposition characterizes the payoff frontier for the parameter region under which an efficient equilibrium exists.

Proposition B.2. *Suppose that* $\lambda > 1/2$ *and that* $\delta \geqslant \frac{c\left(e^{FB}\right)}{(2-2\lambda)e^{FB}}$*, which is higher than* $\frac{c(\hat{e})}{(2-2\lambda)\hat{e}}$ *where* $\hat{e}$ *solves* $c'(e) = 2(1-\lambda)$*. The payoff frontier* $U(\pi)$ *is given by:*

$$\begin{cases} (1-\delta)\left(\frac{\pi}{1-\delta} - c\left(\frac{\pi}{1-\delta}\right)\right) + \delta\left(e^{FB} - c\left(e^{FB}\right)\right) - \pi, & \text{if } \pi \in \left[0, (1-\delta)e^{FB}\right), \\ e^{FB} - c\left(e^{FB}\right) - \pi, & \text{if } \pi \in \left[(1-\delta)e^{FB}, e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}\right], \\ \frac{2\lambda-1}{2-2\lambda}c(e), \quad \text{where } e \text{ solves } \pi = e - \frac{c(e)}{2-2\lambda}, & \text{if } \pi \in \left(e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}, \hat{e} - \frac{c(\hat{e})}{2-2\lambda}\right]. \end{cases}$$

In the third region, $U'(\pi)$ *goes from* -1 *to* $-\infty$*.*

Proof of Proposition B.2. This is obvious for the first two regions (i.e., $\pi \leqslant e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}$) based on Proposition B.1 and Lemma B.2. (For these two regions, we can choose $e = \min\left\{\frac{\pi}{1-\delta}, e^{FB}\right\}$, $s = \frac{\delta}{1-\delta}\Pi_p$, $\Pi_p = e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}$, $\underline{\Pi} = 0$ in the following program.) We also note that $U'(\pi)$ for the first region equals $-c'\left(\frac{\pi}{1-\delta}\right)$, so it is strictly higher than -1 in the first region.

We thus focus on the third region. The payoff $U(\pi)$ is equal to the maximal:

$$U(\pi) = \max_{e,s,\underline{\Pi},\Pi_p,U_a} (1-\delta)(s - c(e)) + \delta U_a,$$

subject to the constraints:

$$\pi = (1-\delta)(e-s) + \delta\Pi_p, \tag{21}$$

$$s - c(e) \geqslant (2\lambda - 1)\frac{\delta}{1-\delta}(\Pi_p - \underline{\Pi}), \tag{22}$$

$$s \leqslant \frac{\delta}{1-\delta}(\Pi_p - \underline{\Pi}), \tag{23}$$

$$0 \leqslant U_a \leqslant U(\Pi_p) \tag{24}$$

$$0 \leqslant \underline{\Pi} \leqslant \Pi_p.$$

First, U_a must equal $U(\Pi_p)$. Second, we argue that $\Pi_p \geqslant (1-\delta)e^{FB}$. If not, we can increase

Π_p by ε and increase s by $\frac{\delta}{1-\delta}\varepsilon$. All the constraints are still satisfied. The change in the objective function is proportional to $(1+U'(\Pi_p))$ which is strictly positive if $\Pi_p < (1-\delta)e^{FB}$, given that $U'(\Pi_p) > -1$ for $\Pi_p < (1-\delta)e^{FB}$.

Third, combining (22) and (23), we obtain that

$$s \geqslant \frac{c(e)}{2-2\lambda}. \tag{25}$$

We now examine a relaxed problem. We maximize the objective with respect to only (21), (25), and that $\Pi_p \geqslant (1-\delta)e^{FB}$. If (25) does not bind, the solution consists of $e = e^{FB}$, $\Pi_p \in \left[(1-\delta)e^{FB}, e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}\right]$, and

$$s = \frac{(1-\delta)e^{FB} + \delta\Pi_p - \pi}{1-\delta} \leqslant \frac{(1-\delta)e^{FB} + \delta\left(e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}\right) - \pi}{1-\delta}.$$

This violates (25) for π in the third region. Therefore, (25) must bind. We thus obtain the following equations:

$$s = \frac{c(e)}{2-2\lambda}, \quad \Pi_p = \frac{\pi - (1-\delta)(e-s)}{\delta}.$$

We substitute these values into the objective. The first-order condition is satisfied when e solves $\pi = e - \frac{c(e)}{2-2\lambda}$. It is easy to show that the objective is concave in e. Therefore, the optimal e is indeed pinned down by the first-order condition. Lastly, we choose $\underline{\Pi} = \frac{\pi-(1-\delta)e}{\delta}$, which is positive for π in the third region. Under these values, the original constraints (21)-(24) are all satisfied. Thus, the solution to the relaxed problem solves the original problem.

Lastly, we show that $U'(\pi)$ goes from -1 to $-\infty$ in the third region. The derivative $U'(\pi)$ equals:

$$\frac{(2\lambda-1)e'(\pi)c'(e(\pi))}{2-2\lambda}.$$

Given that $e(\pi)$ solves $\pi = e - \frac{c(e)}{2-2\lambda}$, we have

$$e'(\pi) = \frac{2(\lambda - 1)}{c'(e(\pi)) + 2\lambda - 2}.$$

Substituting $e'(\pi)$ into $U'(\pi)$, we obtain that

$$U'(\pi) = \frac{(1-2\lambda)c'(e(\pi))}{c'(e(\pi)) + 2\lambda - 2},$$

which equals -1 when $\pi = e^{FB} - \frac{c\left(e^{FB}\right)}{2-2\lambda}$ since $c'\left(e^{FB}\right) = 1$, and equals $-\infty$ when $\pi = \hat{e} - \frac{c(\hat{e})}{2-2\lambda}$ since $c'(\hat{e}) = 2 - 2\lambda$. ■

The following proposition characterizes the payoff frontier for the parameter region under which an efficient equilibrium does not exist but $\delta \geqslant \frac{c(\hat{e})}{(2-2\lambda)\hat{e}}$, where $\hat{e}$ solves $c'(e) = 2(1-\lambda)$. Let $(\underline{\pi}, \underline{e})$ be the positive roots of:

$$\pi = (1-\delta)e, \text{ and } \pi = e - \frac{c(e)}{2-2\lambda}.$$

Proposition B.3. *Suppose that $\lambda > 1/2$ and that $\frac{c\left(e^{FB}\right)}{(2-2\lambda)e^{FB}} > \delta \geqslant \frac{c(\hat{e})}{(2-2\lambda)\hat{e}}$, where $\hat{e}$ solves $c'(e) = 2(1-\lambda)$. The payoff frontier $U(\pi)$ is given by:*

$$\begin{cases} (1-\delta)\left(\frac{\pi}{1-\delta} - c\left(\frac{\pi}{1-\delta}\right)\right) + \delta\left(\underline{e} - c(\underline{e})\right) - \pi, & \text{if } \pi \in [0, \underline{\pi}), \\ \frac{2\lambda-1}{2-2\lambda}c(e), \quad \text{where } e \text{ solves } \pi = e - \frac{c(e)}{2-2\lambda}, & \text{if } \pi \in \left[\underline{\pi}, \hat{e} - \frac{c(\hat{e})}{2-2\lambda}\right]. \end{cases}$$

Proof of Proposition B.3. Note that $U'(\pi)$ for the first region equals $-c'\left(\frac{\pi}{1-\delta}\right)$, so it is strictly higher than -1 in the first region. The proof for the second region uses similar arguments as in the proof of Proposition B.2. We thus focus on the first region. We examine the relaxed problem by maximizing:

$$\max_{e,s,\underline{\Pi},\Pi_p,U_a} (1-\delta)(s - c(e)) + \delta U(\Pi_p) + \pi = (1-\delta)(e - c(e)) + \delta(\Pi_p + U(\Pi_p)),$$

subject to the constraints:

$$\pi = (1-\delta)(e-s) + \delta\Pi_p,$$
$$s \leqslant \frac{\delta}{1-\delta}\Pi_p.$$

Combining these two constraints, we have that $e \leqslant \frac{\pi}{1-\delta}$. The first part of the objective, $e - c(e)$, is maximized by choosing $e = \frac{\pi}{1-\delta}$. The second part is maximized by choosing $\Pi_p = \underline{\pi}$, because $U'(\pi) > -1$ in the first region and $U'(\pi) < -1$ in the second region. By choosing $s = \frac{\delta}{1-\delta}\underline{\pi}$ and $\underline{\Pi} = 0$, we can verify that (22) holds for π in the first region. Hence, the original constraints (21)-(24) are all satisfied. ∎

Lastly, the following proposition characterizes the payoff frontier for the parameter region where $\delta < \frac{c(\hat{e})}{(2-2\lambda)\hat{e}}$, where $\hat{e}$ solves $c'(e) = 2(1-\lambda)$.

Proposition B.4. *Suppose that $\lambda > 1/2$ and that $\delta < \frac{c(\hat{e})}{(2-2\lambda)\hat{e}}$, where $\hat{e}$ solves $c'(e) = 2(1-\lambda)$. Let $\tilde{\pi} = \tilde{e} - \frac{c(\tilde{e})}{2-2\lambda}$ where $\tilde{e}$ is the positive root of $c(e) = 2(1-\lambda)\delta e$. The payoff frontier $U(\pi)$ is given by:*

$$(1-\delta)\left(\frac{\pi}{1-\delta} - c\left(\frac{\pi}{1-\delta}\right)\right) + \delta\left(\tilde{e} - c(\tilde{e})\right) - \pi, \quad \text{for } \pi \in [0, \tilde{\pi}].$$

The surplus-maximizing equilibrium is unique and coincides with the principal-optimal equilibrium.

Proof of Proposition B.4. Note that $U'(\pi)$ equals $-c'\left(\frac{\pi}{1-\delta}\right)$, which is strictly higher than -1. We examine the relaxed problem by maximizing:

$$\max_{e,s,\underline{\Pi},\Pi_p,U_a} (1-\delta)(s-c(e)) + \delta U(\Pi_p) + \pi = (1-\delta)(e-c(e)) + \delta(\Pi_p + U(\Pi_p)),$$

subject to the constraints:

$$\pi = (1-\delta)(e-s) + \delta\Pi_p,$$
$$s \leqslant \frac{\delta}{1-\delta}\Pi_p.$$

Combining these two constraints, we have that $e \leqslant \frac{\pi}{1-\delta}$. The first part of the objective, $e - c(e)$, is maximized by choosing $e = \frac{\pi}{1-\delta}$. The second part is maximized by choosing $\Pi_p = \tilde{\pi}$, because $U'(\pi) > -1$ for $\pi \in [0, \tilde{\pi}]$. By choosing $s = \frac{\delta}{1-\delta}\tilde{\pi}$ and $\underline{\Pi} = 0$, we can verify that (22) holds for π. Hence, the original constraints (21)-(24) are all satisfied. ■

C Long-Lived Agents

C.1 A Result with Long-Lived Agents

C.1.1 Model, Result, and Discussion

Consider a repeated game with a single principal and N agents with a shared discount factor $\delta \in [0, 1)$. In each period, the following stage game is played:

1. Exactly one agent is publicly selected to be active. For each agent $i \in \{1, ..., N\}$, let $x_{i,t} \in \{0, 1\}$ be the indicator function for agent i being selected. Let $\Pr\{x_{i,t} = 1\} = \rho_i$, where $\sum_i \rho_i = 1$.

2. The active agent chooses $e_t \in \mathbb{R}_+$ and $\mu_t : \mathbb{R} \to M$, which are observed only by the principal and the active agent.

3. The principal and the active agent exchange transfers, with resulting net transfer to the active agent $s_t \in \mathbb{R}$. These transfers are observed only by the principal and the active agent.

4. The message $m_t = \mu_t(s_t)$ is realized and publicly observed.

The principal's and agent i's payoffs in each period t are $\pi_t = e_t - s_t$ and $u_{i,t} = x_{i,t}(s_t - c(e_t))$, respectively, with corresponding expected discounted payoffs $\Pi_t = \sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)(e_t - s_t)$ and $U_{i,t} = \sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)x_{i,t}(s_t - c(e_{i,t}))$. Our solution concept is plain Perfect Bayesian Equilibrium with one additional restriction: at any history h^t such that agent i has observed a deviation, we require that $\mathbb{E}[U_{i,t}|h^t] \geqslant 0$. This restriction rules out pathological off-path behavior that might arise from the fact that an agent's beliefs about the history are essentially arbitrary once he observes a deviation.[7] We also restrict attention to equilibria in pure strategies to simplify agents' beliefs on the equilibrium path.

Proposition 9. *Let e_i^* be the maximum effort attainable in any pure-strategy Perfect Bayesian equilibrium. Letting $s_i^* \equiv \min\{e_i^*, e^{FB}\} - c(\min\{e_i^*, e^{FB}\})$, $e_1^*, e_2^*, ..., e_N^*$ must satisfy the system of inequalities*

$$(1-\delta)c(e_i^*) \leqslant 2\delta\rho_i s_i^* + \frac{2\rho_i\delta}{1-(1-\rho_i)\delta} \sum_{j\neq i} \rho_j s_j^*. \tag{26}$$

The right-hand side of (26) is the analogue to the right-hand side of (1). To translate between settings, define

$$\overline{L} = \overline{v} = \delta\rho_i s_i^*$$

as the maximum payoff the principal can earn from agent i as well as the maximum total surplus produced by agent i. Similarly,

$$\underline{L} = \underline{v} = 0$$

is the minimum total surplus and principal's payoff produced by agent i. Then (26) equals

$$(1-\delta)c(e_i^*) \leqslant (\overline{L} - \underline{L}) + (\overline{v} - \underline{v}) + \frac{2\rho_i\delta}{1-(1-\rho_i)\delta} \sum_{j\neq i} \rho_j s_j^*. \tag{27}$$

[7] This condition is trivially satisfied in any equilibrium that is recursive. It is needed here because this game has private monitoring, which means that equilibria are not necessarily recursive.

The first two terms of the right-hand side of (27) correspond to the right-hand side of (1). The final term represents a new force that arises because long-lived agents have multiple opportunities to communicate with one another. In particular, an agent who has held up the principal in the past might reveal that fact to other agents in subsequent messages, who can then respond by reward the principal for resisting hold-up or punishing her for giving in. Proposition 9 therefore suggests that giving agents multiple rounds to communicate might improve cooperation. The extent to which it does so, however, is still tightly limited by the frequency of those opportunities and hence the probability that agent i is active.

An immediate corollary of Proposition 9 is that, as the probability that an agent interacts with the principal ρ_i approaches zero, that agent's maximum equilibrium effort does too. This implication is similar to our main takeaway from Proposition 6: the strength of each agent's bilateral relationship limits the severity of the coordinated punishments available to him. This result relies on the fact that agents can send messages only when they are active. We can interpret this assumption as the natural extension of our commitment assumption to a setting with long-lived agents; indeed, a result identical to Proposition 9 would hold if agents could communicate in every period but whenever an agent is active, he commits to a *sequence* of messages in each period until he is again active.

C.1.2 Proof of Proposition 9

For each agent $j \in \{1, ..., 2\}$, let e_j^* be the maximum effort that can be attained in any period of any equilibrium. Consider a history h^t right after agent i is chosen to be the active agent in period t. Define four different expectations of Π_{t+1} that follow four different outcomes:

1. $\overline{\Pi}^*$ if no player deviates, with corresponding message $\overline{m}$;

2. $\underline{\Pi}^*$ if the principal deviates but the active agent does not, with corresponding message $\underline{m}$;

3. $\overline{\Pi}^{HU}$ if the active agent deviates and $m_t = \overline{m}$;

4. $\underline{\Pi}^{HU}$ if the active agent deviates and $m_t = \underline{m}$.

We identify necessary conditions for effort e to be attained in equilibrium.

First, the principal must be willing to pay s^* if the active agent does not deviate, which requires

$$s^* \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi}^* - \underline{\Pi}^*\right). \tag{28}$$

Second, the active agent i must be willing to choose effort e and the equilibrium communication protocol μ. Agent i can always deviate by choosing $e_t = 0$ and

$$\mu_t = \begin{cases} \underline{m} & s_t < \hat{s} \\ \overline{m} & \text{otherwise} \end{cases}$$

for some $\hat{s} \geqslant 0$. Following this deviation, the principal's unique best response is to pay $\hat{s}$ so long as

$$-\hat{s} + \frac{\delta}{1-\delta}\overline{\Pi}^{HU} > \frac{\delta}{1-\delta}\underline{\Pi}^{HU},$$

since the principal can earn no less than $\overline{\Pi}^{HU}$ in the continuation game if $m_t = \overline{m}$ and no more than $\underline{\Pi}^{HU}$ if $m_t = \underline{m}$. Therefore, agent i has no profitable deviation of this form only if

$$s^* - c(e) + \frac{\delta}{1-\delta}\overline{U}_i^* \geqslant \max\left\{0, \frac{\delta}{1-\delta}\left(\overline{\Pi}^{HU} - \underline{\Pi}^{HU}\right)\right\}, \tag{29}$$

where $\overline{U}_i^*$ is the agent's expectations about her continuation payoff at the history that yields principal payoff $\overline{\Pi}_i^*$.

Combining (28) and (29) yields the following necessary condition for effort e to be part of equilibrium:

$$c(e) \leqslant \frac{\delta}{1-\delta}\left(\overline{U}_i^* + \overline{\Pi}^* - \underline{\Pi}^*\right) - \max\left\{0, \frac{\delta}{1-\delta}\left(\overline{\Pi}^{HU} - \underline{\Pi}^{HU}\right)\right\} \tag{30}$$

Our next goal is to connect (28) and (29) by studying the relationship between $\overline{U}_i^* + \overline{\Pi}^* -$

$\underline{\Pi}^*$ and $\overline{\Pi}^{HU} - \underline{\Pi}^{HU}$. We do so by bounding $\overline{U}_i^* + \overline{\Pi}^* - \overline{\Pi}^{HU}$ from above and $\underline{\Pi}^* - \underline{\Pi}^{HU}$ from below.

Fix two histories h^{t+1} and $\hat{h}^{t+1}$ at the start of period $t+1$ such that agent i can distinguish h^{t+1} from $\hat{h}^{t+1}$ but no other agents can. For $t' \geqslant t+1$, we will use the notation $h^{t'}$ and $\hat{h}^{t'}$ to represent successor histories to h^{t+1} and $\hat{h}^{t+1}$, respectively. At history $\hat{h}^{t+1}$, the principal can always play the following strategy:

1. At any history $\hat{h}^{t'}$ that the active agent believes is consistent with h^{t+1}, play as in the corresponding successor history to h^{t+1};

2. At any other history, choose $s_t = 0$.

Under this strategy, each agent $j \neq i$ learns that the history is inconsistent with h^{t+1} only when agent i sends a message that is inconsistent with play following h^{t+1}. In a pure-strategy equilibrium, all agents learn this fact at the same time. For each $t' \geqslant t+1$, denote

$$\hat{\mathcal{B}}^{t'} = \left\{ \hat{h}^{t'} | \text{Agents } j \neq i \text{ learn that the history is inconsistent with } h^{t+1} \text{ in period } t'-1, \text{ but not before} \right\}.$$

Where $\hat{\mathcal{B}}^{\infty}$ denotes the event that agents $j \neq i$ never learn that the history is inconsistent with h^{t+1}. Note that these events collectively partition the set of histories following $\hat{h}^{t+1}$. We can define an analogous collection of sets for the event that agents $j \neq i$ learn that the history is inconsistent with $\hat{h}^{t+1}$. We denote this analogous collection $\mathcal{B}^{t'}$.

For each agent $j \in \{1, ..., N\}$, define $\pi_{j,t} = x_{j,t}(e_t - s_t)$ and $\pi_{-j,t} = \sum_{k \neq j} x_{k,t}(e_t - s_t)$ as the principal's payoff from agent j and from all other agents, respectively. Define $\Pi_{j,t} = \sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)\pi_{j,t'}$ and $\Pi_{-j,t} = \sum_{t'=t}^{\infty} \delta^{t'-t}(1-\delta)\pi_{-j,t'}$. Because $\left\{ \hat{\mathcal{B}}^{\tilde{t}} \right\}_{\tilde{t}=t+1}^{t'}$ partitions the histories of length t' following $\hat{h}^{t+1}$,

$$\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right] = \sum_{t'=t+1}^{\infty}(1-\delta)\delta^{t'-t-1}\left(\mathbb{E}\left[\pi_{i,t'}|\hat{h}^{t+1}\right] + \sum_{\tilde{t}=t+1}^{t'}\mathbb{E}\left[\pi_{-i,t'}|\hat{h}^{t+1},\hat{\mathcal{B}}^{\tilde{t}}\right]\Pr\left\{\hat{\mathcal{B}}^{\tilde{t}}\right\}\right). \tag{31}$$

The right-hand side of (31) is absolutely convergent, so we can rearrange the order of summation to yield

$$\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right] = \begin{array}{c} \sum_{t'=t+1}^{\infty}(1-\delta)\delta^{t'-t-1}\mathbb{E}\left[\pi_{i,t'}|\hat{h}^{t+1}\right] + \\ \sum_{\tilde{t}=t+1}^{\infty}\left(\sum_{t'=t+1}^{\tilde{t}-1}(1-\delta)\delta^{t'-t-1}\mathbb{E}\left[\pi_{-i,t'}|\hat{\mathcal{B}}^{\tilde{t}}\right] + \delta^{\tilde{t}-t-1}\mathbb{E}\left[\Pi_{-i,\tilde{t}}|\hat{\mathcal{B}}^{\tilde{t}}\right]\right)\Pr\left\{\hat{\mathcal{B}}^{\tilde{t}}\right\}. \end{array} \tag{32}$$

Under the principal's strategy specified above, the principal and agents $j \neq i$ act identically until those agents learn of a deviation. Therefore, for any $t' < \tilde{t}$,

$$\mathbb{E}\left[\pi_{-i,t'}|\mathcal{B}^{\tilde{t}}\right]\Pr\left\{\mathcal{B}^{\tilde{t}}\right\} = \mathbb{E}\left[\pi_{-i,t'}|\hat{\mathcal{B}}^{\tilde{t}}\right]\Pr\left\{\hat{\mathcal{B}}^{\tilde{t}}\right\}.$$

Moreover, for any $\tilde{t}$, $\Pr\left\{\mathcal{B}^{\tilde{t}}\right\} = \Pr\left\{\hat{\mathcal{B}}^{\tilde{t}}\right\}$, since any message that distinguish h^{t+1} from $\hat{h}^{t+1}$ must also distinguish $\hat{h}^{t+1}$ from h^{t+1}.

Now, $\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right]$ is bounded below by the principal's payoff from the strategy specified above. Therefore, we can use (32) to bound the difference

$$\begin{array}{c} \mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right] - \mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right] \leqslant \\ \sum_{t'=t+1}^{\infty}\delta^{t'-t-1}(1-\delta)\left(\mathbb{E}\left[\pi_{i,t'}|h^{t+1}\right] - \mathbb{E}\left[\pi_{i,t'}|\hat{h}^{t+1}\right]\right) + \\ \sum_{\tilde{t}=t+1}^{\infty}\delta^{\tilde{t}-t-1}\left(\mathbb{E}\left[\Pi_{-i,\tilde{t}}|\mathcal{B}^{\tilde{t}}\right] - \mathbb{E}\left[\Pi_{-i,\tilde{t}}|\hat{\mathcal{B}}^{\tilde{t}}\right]\right)\Pr\left\{\mathcal{B}^{\tilde{t}}\right\} \end{array} \tag{33}$$

Under the specified strategy, $\mathbb{E}\left[\Pi_{-i,\tilde{t}}|\hat{\mathcal{B}}^{\tilde{t}}\right] \geqslant 0$ because the principal pays no transfer to an agent j who knows that the history is inconsistent with h^{t+1}, with $\mathbb{E}\left[\pi_{i,t'}|\hat{h}^{t+1}\right] \geqslant 0$ for a similar reason. A necessary condition for (33) is therefore

$$\begin{aligned}&\mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right]-\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right]\leqslant\\&\textstyle\sum_{t'=t+1}^{\infty}\delta^{t'-t-1}\left((1-\delta)\mathbb{E}\left[\pi_{i,t'}|h^{t+1}\right]+\mathbb{E}\left[\Pi_{-i,t'}|\mathcal{B}^{t'}\right]\Pr\left\{\mathcal{B}^{t'}\right\}\right)\end{aligned}\tag{34}$$

Suppose that h^{t+1} is the on-path history such that $\mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right]=\overline{\Pi}^*$. In a pure-strategy equilibrium, agents correctly infer the true history on the equilibrium path, which means that they must earn non-negative utility. Consequently, the principal earns no more than total continuation surplus, so (34) requires

$$\mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right]-\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right]\leqslant\sum_{t'=t+1}\delta^{t'-t-1}\left((1-\delta)\rho_i s_i^*+\sum_{j\neq i}\rho_j s_j^*\Pr\left\{\mathcal{B}^{t'}\right\}\right).\tag{35}$$

Note than an identical bound holds for the expression $\overline{U}_i^*+\overline{\Pi}^*-\underline{\Pi}^*$ because agents $j\neq i$ earn non-negative continuation utilities on the equilibrium path. If h^{t+1} is instead the history such that $\mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right]=\underline{\Pi}^{HU}$, then agents have observed $\underline{m}$ and so know that play is off-path. Our equilibrium restriction requires their utilities to be non-negative at such a history, so (35) again holds.

Since $s_j^*\geqslant 0$, the right-hand side of (35) is maximized by having the event $\mathcal{B}^{\tilde{t}}$ happen as early as possible. The earliest it can occur is the next time that agent i is the active agent, since agent i can send a message only when he is active. Agent i is active for the first time since period t in period t' with probability $(1-\rho_i)^{t'-t-1}\rho_i$, so (35) requires

$$\begin{aligned}\mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right]-\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right]&\leqslant\textstyle\sum_{t'=t+1}\delta^{t'-t-1}\left((1-\delta)\rho_i s_i^*+(1-\rho_i)^{t'-t-1}\rho_i\sum_{j=1}^{N}\rho_j s_j^*\right)\\&=\rho_i s_i^*+\tfrac{\rho_i}{1-(1-\rho_i)\delta}\textstyle\sum_{j\neq i}\rho_j s_j^*.\end{aligned}\tag{36}$$

As argued above, an identical bound holds for $\mathbb{E}\left[U_{i,t+1}+\Pi_{t+1}|h^{t+1}\right]-\mathbb{E}\left[\Pi_{t+1}|\hat{h}^{t+1}\right]$.

From (36), we conclude that

$$\overline{U}_i^*+\overline{\Pi}^*-\underline{\Pi}^*\leqslant\overline{\Pi}^{HU}-\underline{\Pi}^{HU}+2\rho_i s_i^*+\frac{2\rho_i}{1-(1-\rho_i)\delta}\sum_{j\neq i}\rho_j s_j^*.$$

A necessary condition for (30) to hold is therefore

$$c(e) \leqslant \left(\begin{array}{c} \frac{\delta}{1-\delta}\left(\overline{\Pi}^{HU} - \underline{\Pi}^{HU} + 2\rho_i s_i^* + \frac{2\rho_i}{1-(1-\rho_i)\delta}\sum_{j\neq i}\rho_j s_j^*\right) - \\ \max\left\{0, \frac{\delta}{1-\delta}\left(\overline{\Pi}^{HU} - \underline{\Pi}^{HU}\right)\right\} \end{array} \right).$$

The right-hand side of this condition is maximized by $\overline{\Pi}^{HU} - \underline{\Pi}^{HU} = 0$, in which case

$$(1-\delta)c(e) \leqslant 2\rho_i s_i^* + \frac{2\rho_i}{1-(1-\rho_i)\delta}\sum_{j\neq i}\rho_j s_j^*,$$

as desired. ■

D Other Extensions

D.1 Up-Front Transfers

The **game with up-front transfers** is identical to the baseline model except that at the start of each period t, the principal and agent t exchange non-negative payments to one another. Denote the resulting net wage to agent t by $w_t \in \mathbb{R}_+$, so that the principal's and agent t's payoffs are $e_t - s_t - w_t$ and $s_t + w_t - c(e_t)$, respectively. Note that agent t chooses a communication mechanism μ_t only *after* these transfers are paid.

Proposition 10. *For any $\lambda \in [0,1]$, the principal-optimal equilibrium of the game with up-front transfers entails the same effort in each period and same principal payoff as Proposition 3.*

Proof of Proposition 10: Setting $w_t = 0$ in each period t recovers the baseline model, which implies that effort and the principal's payoff in the principal-optimal equilibrium of the game with up-front transfers must be weakly higher than in the baseline model.

As in the proof of Proposition 3, let $\overline{\Pi}$ be the highest equilibrium payoff for the principal and suppose $\overline{\Pi} > 0$. Consider an equilibrium that attains $\overline{\Pi}$, and let $\overline{\Pi}_1$ and $\underline{\Pi}_1$, with corresponding messages $\overline{m}$ and $\underline{m}$, be the highest and lowest principal continuation payoffs. The principal is willing to pay s^* only if (2) holds.

The agent can always deviate by paying no upfront transfer, then choosing $e = 0$ and

$$\mu(s) = \begin{cases} \underline{m} & s < \hat{s} \\ \overline{m} & s \geqslant \hat{s} \end{cases}$$

for some $\hat{s} \geqslant 0$. The principal pays $\hat{s}$ so long as

$$\hat{s} < (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right).$$

For this deviation to not be profitable,

$$w - c(e) + s \geqslant \max\{0, w\} + \hat{s},$$

or plugging in the maximum $\hat{s}$,

$$w - c(e) + s \geqslant \max\{0, w\} + (2\lambda - 1)\frac{\delta}{1-\delta}\left(\overline{\Pi}_1 - \underline{\Pi}_1\right) \tag{37}$$

Consequently,

$$\overline{\Pi} \leqslant (1-\delta)\left(e - c(e) + w - \max\{0, w\}\right) - (2\lambda - 1)\delta(\overline{\Pi} - \underline{\Pi}_1) + \delta\overline{\Pi}$$

where w, s, and $c(e)$ must satisfy (2) and (37). But it is clear from this problem that $w = 0$ is optimal, proving the claim. ■

D.2 The Principal Can Send Messages

Let M_p be the set of messages for the principal, and m_p a typical message. We consider two different stage games. First, the principal talks after agent t sends m. Second, the principal talks before agent t. In the second case, we still assume that μ_t is a function of s_t only (so it doesn't depend on m_p); however, if agent t turns out to be noncommitted, his message can depend on m_p.

The principal talks after agent t. Consider some period t. We let $\pi(m, m_p)$ be the principal's continuation payoff if (m, m_p) realizes. Given agent t's message m, the principal always chooses m_p to maximize $\pi(m, m_p)$. We let $\pi(m) := \max_{m_p} \pi(m, m_p)$, so $\pi(m)$ is the principal's continuation payoff after agent t's message m. We let $\overline{\Pi}$ and $\underline{\Pi}$ be the highest and lowest continuation payoffs that agent t's message can induce. Then, the same program as in the baseline model (i.e., the program for Proposition 3) applies. The principal's talk has no impact on the characterization of her optimal equilibrium.

Proposition D.1. *Suppose that in each period t the principal sends $m_p \in M_p$ after agent t sends m. The principal-optimal equilibrium is outcome-equivalent to that in Proposition 3.*

The principal talks before agent t. Consider some period t. We let $\pi(m_p, m)$ be the principal's continuation payoff if (m_p, m) realizes. When agents have full commitment (i.e., $\lambda = 1$), once the principal chooses s_t the principal knows what agent t will say for sure. The principal will choose a message m_p to maximize her continuation payoff given agent t's message.[8] The same argument as in the previous case applies, so the only equilibrium involves no effort and no payment. When agents have weak commitment (i.e., $\lambda \in (0, 1/2)$), the payoff frontier when the principal *cannot* talk is characterized subject to the necessary conditions that are still necessary when the principal *can* talk. Therefore, the ability to talk does not expand the equilibrium payoff frontier or the principal's optimal equilibrium. From now on, we focus on the case that $\lambda \in (1/2, 1)$.

[8] We also allow agents to commit to a mixture over M. In that case, the principal choose m_p to maximize her continuation payoff given the mixture.

We let $\overline{\Pi} := \max_{m_p, m} \pi(m_p, m)$ be the highest payoff for the principal. We let

$$\underline{\Pi} := \min_{\sigma, \tilde{\sigma}(m_p)} \max_{m_p} \lambda \pi(m_p, \sigma) + (1-\lambda)\pi(m_p, \tilde{\sigma}(m_p)),$$

be the lowest payoff for the principal. Here $\sigma \in \Delta(M)$ and $\tilde{\sigma}_1(\cdot) : M_p \to \Delta(M)$. (We allow for mixture by agents since the principal's payoff can be pushed lower if she is uncertain about agents' choice of messages.) Let (e^*, μ^*, s^*) be the equilibrium effort, communication protocol, and payment. It must be true that

$$s^* \leqslant \frac{\delta}{1-\delta}\left(\overline{\Pi} - \underline{\Pi}\right).$$

We next consider how much rent agent t can extract. Agent t can shirk, and ask the principal to pay $\hat{s}$. Agent t will send the mixture of messages, σ_1, if the principal pays at least $\hat{s}$ and the mixture, σ_2, if the principal pays less than $\hat{s}$. The principal believes that agent t, if not committed, will choose $\tilde{\sigma}_1(m_p)$ if he pays at least $\hat{s}$ and $\tilde{\sigma}_2(m_p)$ if he pays less than $\hat{s}$. The principal's continuation payoff if she pays $\hat{s}$ is

$$A := \max_{m_p} \lambda \pi(m_p, \sigma_1) + (1-\lambda)\pi(m_p, \tilde{\sigma}_1(m_p)).$$

Her continuation payoff if she pays less than $\hat{s}$ is

$$B := \max_{m_p} \lambda \pi(m_p, \sigma_2) + (1-\lambda)\pi(m_p, \tilde{\sigma}_2(m_p)).$$

The principal is willing to pay up to $\frac{\delta}{1-\delta}(A-B)$ to a shirking agent. Agent t is choosing σ_1 and σ_2 to make $A-B$ as large as possible. Nature is choosing $\tilde{\sigma}_2(m_p)$ and $\tilde{\sigma}_2(m_p)$ to diminish agent t's ability to extract rent (i.e., to make $A-B$ as small as possible). It is easy to see that when $\lambda = 1$, $A - B = (2\lambda - 1)\left(\overline{\Pi} - \underline{\Pi}\right)$. We have argued in the baseline model that if $\pi(m_p, m)$ doesn't depend on m_p, then for any $\lambda \in (1/2, 1]$, $A - B = (2\lambda - 1)\left(\overline{\Pi} - \underline{\Pi}\right)$.

We want to understand whether it can be the case that $A - B < (2\lambda - 1)\left(\overline{\Pi} - \underline{\Pi}\right)$ for some $\lambda \in (1/2, 1)$. If this is true, then the principal's ability to talk before agent t will mitigate the hold-up problem.

The following example shows that the principal's talking before agent t mitigates the hold-up problem.

Example 2. *Suppose that $M = M_p = \{0, 1\}$, and $\pi(m_p, m) = a$ if $m_p = m$ and $\pi(m_p, m) = b \in [0, a)$ if $m_p \neq m$. On the equilibrium path, agent t commits to send $m = 1$ if the principal pays s^*, and commits to send $m = 1$ and $m = 0$ with equal probability if the principal deviates. Moreover, a noncommitted agent will match the principal's message m_p if the principal pays s^*, and will choose $m \neq m_p$ if the principal deviates. If the principal pays s^*, her optimal message shall be $m_p = 1$. If the principal deviates, both $m_p = 1$ and $m_p = 0$ are her best responses. The principal's continuation payoffs, if he pays s^* and if he deviates, are:*

$$\overline{\Pi} = a, \quad \underline{\Pi} = \lambda \frac{a+b}{2} + (1-\lambda)b.$$

These are also the principal's highest and lowest continuation payoffs.

Agent t can shirk and commit to sending $m = 1$ for sure if the principal pays $\hat{s}$, and to sending $m = 1$ and $m = 0$ with equal probability if the principal pays less than $\hat{s}$. A noncommitted agent will choose a message different than m_p if the principal indeed pays $\hat{s}$ and will match the principal's message if the principal pays less than $\hat{s}$. We choose this strategy for a noncommitted agent in order to deter the principal from giving in to agent t's hold-up. The principal's continuation payoff if she pays $\hat{s}$ is

$$A = \lambda a + (1-\lambda)b.$$

If she pays less than $\hat{s}$ is

$$B = \lambda \frac{a+b}{2} + (1-\lambda)a.$$

The difference $A-B$ is strictly less than $(2\lambda-1)\left(\overline{\Pi}-\underline{\Pi}\right)$ for any $\lambda \in (1/2,1), 0 \leqslant b < a$:

$$A-B=\frac{1}{2}(a-b)(3\lambda-2)<\frac{1}{2}(a-b)(2\lambda-1)(2-\lambda)=(2\lambda-1)\left(\overline{\Pi}-\underline{\Pi}\right).$$

Moreover, $A-B$ is negative when $\lambda \leqslant \frac{2}{3}$. This implies that, if the principal talks before agent t, the principal-optimal equilibrium for $\lambda=0$ can be achieved when $\lambda \leqslant \frac{2}{3}$. This is an improvement compared with the baseline model.

D.3 Agents Interact Directly

The **game with inter-agent rewards** is similar to the baseline game, with the following difference: at the start of each period t, agent t can choose a reward $b_t \in \left[0,\overline{b}\right]$ for agent $t-1$. Payoffs are then $e_t - s_t$ for the principal and $s_t + b_{t+1} - c(e_t)$ for agent t. Note that agent t does not care about his *own* choice of b_t, but does care about the choice of the *next* agent b_{t+1}. This assumption makes these inter-agent rewards particularly powerful, since we do not need to incentivize agent t to reward his predecessor.

Inter-agent rewards can certainly result in higher agent utilities relative to the baseline game. However, we show that they do nothing to improve the principal's payoff. The key problem is that, because each agent can condition only on the previous agents' messages, he cannot tailor the reward to whether or not the previous agent deviated. Consequently, inter-agent rewards do nothing to resolve the hold-up problem.

Proposition 11. *For any $\lambda \in [0,1]$, the principal-optimal equilibrium of the game with inter-agent rewards entails the same effort in each period and same principal payoff as Proposition 3.*

Proof of Proposition 11: Setting $b_t=0$ in each period t recovers the baseline model, which implies that effort and the principal's payoff in the principal-optimal equilibrium of the game with inter-agent rewards must be weakly higher than in the baseline model.

Consider period t of an equilibrium. Define $b(m)$ and $\Pi(m)$ be the equilibrium mappings from m_t to b_{t+1} and $\mathbb{E}\left[\Pi_{t+1}|h^{t+1}\right]$, respectively, and let

$$\bar{b} = \max_m b(m)$$

be the highest b_{t+1} on- or off-path in equilibrium. If agent t is not committed to $\mu_t(\cdot)$, then he must send a message that induces $b_{t+1} = \bar{b}$. Let

$$\begin{aligned}\overline{\Pi}^N &= \max\nolimits_{m|b(m)=\bar{b}} \Pi(m)\\ \underline{\Pi}^N &= \min\nolimits_{m|b(m)=\bar{b}} \Pi(m)\end{aligned}$$

be the maximum and minimum of the set of possible continuation payoffs if agent t is not committed to μ_t. Let $\overline{\Pi}^C$ be the maximum and $\underline{\Pi}^C$ be the minimum continuation payoffs, with corresponding messages $\overline{m}^C$ and $\underline{m}^C$. Note that $\underline{\Pi}^C \leqslant \underline{\Pi}^N \leqslant \overline{\Pi}^N \leqslant \overline{\Pi}^C$.

The principal is willing to pay s_t only if

$$s_t \leqslant \frac{\delta}{1-\delta}\left(\lambda\left(\overline{\Pi}^C - \underline{\Pi}^C\right) + (1-\lambda)\left(\overline{\Pi}^N - \underline{\Pi}^N\right)\right). \tag{38}$$

Agent t can always deviate by choosing $e_t = 0$ and

$$\mu_t(s) = \begin{cases} \overline{m}^C & s \geqslant \hat{s} \\ \underline{m}^C & s < \hat{s} \end{cases}$$

for some $\hat{s} \geqslant 0$. The principal responds to this deviation by paying $\hat{s}$ if

$$-\hat{s} + \frac{\delta}{1-\delta}\left(\lambda\overline{\Pi}^C + (1-\lambda)\underline{\Pi}^N\right) > \frac{\delta}{1-\delta}\left(\lambda\underline{\Pi}^C + (1-\lambda)\overline{\Pi}^N\right).$$

Agent t is therefore willing to choose e_t only if

$$s_t - c(e_t) + \lambda b(\overline{m}^C) + (1-\lambda)\bar{b} \geqslant \hat{s} + \lambda b(\overline{m}^C) + (1-\lambda)\bar{b}$$

for any such $\hat{s}$, which requires

$$s_t - c(e_t) \geqslant \frac{\delta}{1-\delta}\left(\lambda\left(\overline{\Pi}^C - \underline{\Pi}^C\right) + (1-\lambda)\left(\underline{\Pi}^N - \overline{\Pi}^N\right)\right). \tag{39}$$

If $\overline{\Pi}^C > \overline{\Pi}^N$, we can increase $\overline{\Pi}^N$ and decrease $\overline{\Pi}^C$ so that $\lambda\overline{\Pi}^C + (1-\lambda)\overline{\Pi}^N$ is constant. Then (38) is unaffected and (39) is strictly relaxed. Similarly, if $\underline{\Pi}^N > \underline{\Pi}^C$, we can increase $\underline{\Pi}^C$ and decrease $\underline{\Pi}^N$ so that $\lambda\underline{\Pi}^C + (1-\lambda)\underline{\Pi}^N$ is constant, which leaves (38) unaffected and strictly relaxes (39). So it is optimal to set $\overline{\Pi}^C = \overline{\Pi}^N$ and $\underline{\Pi}^C = \underline{\Pi}^N$ in the relaxed problem.

But then (38) and (39) are identical to the necessary conditions for equilibrium identified in the proof of Proposition 3. The principal's payoff is likewise identical. We conclude that effort and the principal's payoff can be no larger in a principal-optimal equilibrium of the game with inter-agent rewards than in a principal-optimal equilibrium of the baseline game.

■