

Renegotiation-Proof Multilateral Enforcement

S. Nageeb Ali, David A. Miller, and David Yilin Yang*

July 31, 2016

Abstract

In multilateral enforcement, a player who cheats on one partner is punished by many partners. But renegotiation might subvert the threat of multilateral punishment. We consider renegotiation proofness in multilateral enforcement games with public monitoring, and also introduce the notion of “bilateral renegotiation proofness” for games with private monitoring. With public monitoring, renegotiation proofness does not impede multilateral enforcement at all; even with private monitoring, bilateral renegotiation imposes no cost when a principal interacts with many agents who can communicate with each other. For community enforcement games with private monitoring, players’ ability to renegotiate bilaterally has some cost, but this cost is relatively small in large communities.

*Ali: Pennsylvania State University. Miller: University of Michigan. Yang: Stanford GSB. This research is financially supported by NSF grant SES–1127643.

Contents

1	Introduction	1
1.1	Our Approach	1
1.2	Relationship to Prior Work	2
2	The Community Enforcement Game	4
2.1	Environment	4
2.2	Histories and strategies	5
2.3	Equilibrium concept	6
2.4	Renegotiation proofness	6
2.5	Bilateral renegotiation proofness	7
3	Renegotiation in the Community Enforcement Game	8
3.1	Outline	8
3.2	Two-Player Games: A Building Block	8
3.3	Many-Player Games with Perfect Public Monitoring	11
3.4	Many-Player Games with Perfect Private Monitoring	13
4	Renegotiation in Multilateral Relational Contracts	15
5	Conclusion	17
	References	17
A	Main Appendix	19
A.1	Proof of Lemma 3 on p. 11	19
A.2	Proof of Proposition 2 on p. 12	21
A.3	Proof of Proposition 3 on p. 14	22
A.4	Proof of Proposition 4 on p. 15	27
A.5	Proof of Proposition 5 on p. 17	30
B	Supplementary Appendix: Definition of Plain PBE	34

1 Introduction

A firm hires two workers but cannot commit to paying their wages. A typical enforcement scheme involves both workers quitting if the firm reneges on a wage payment to either worker. Because being punished by both workers is worse than being punished by only the mistreated worker, such multilateral enforcement is potentially more powerful than bilateral enforcement. But is the threat credible? When the firm is supposed to be punished by both workers, it has a motive to secretly renegotiate with each worker to recover some production. Of course, if the firm anticipates that it can renegotiate its punishments, it may not be so deterred by multilateral enforcement. Thus, the question that motivates our study is:

If parties may renegotiate, how effective is multilateral enforcement?

That multilateral enforcement might be weakened by renegotiation is a concern expressed by previous work. [Greif, Milgrom, and Weingast \(1994\)](#) discuss how bilateral renegotiation may impede a merchant guild from credibly punishing a ruler that does not protect their property rights. [Karlán, Möbius, Rosenblat, and Szeidl \(2009\)](#) study how the measure of “trust” on a network should anticipate coalitions of players may deviate from punishing a defector; [Ambrus, Möbius, and Szeidl \(2014\)](#) use this approach to understand risk-sharing networks. [Jackson, Rodriguez-Barraquer, and Tan \(2012\)](#) highlight how particular network structures may fail to sustain favor exchanges in a way that is renegotiation-proof.

We revisit this issue and emerge with a contrasting message: renegotiation proofness need not conflict with multilateral enforcement. In two commonly studied settings, we show that the degree of cooperation supported by the best equilibrium can also be supported by one that is renegotiation-proof. Renegotiation costs emerge only when players hold private information about their past actions, but even then the cost of renegotiation is proportionally small and vanishes in large groups.

Apart from offering a contrasting message for renegotiation-proof multilateral enforcement, our formalism emphasizes the importance of “penance”, which is familiar from the study of renegotiation in two player games (e.g., [van Damme 1989](#)). A credible social norm forces defectors to perform favors in penance for others while receiving nothing in return.

1.1 Our Approach

We study renegotiation proofness in two familiar multilateral enforcement games, in each of which the players’ levels of effort are variable. One game, corresponding to that described above, is the “agents-and-principal game” in which a single firm interacts with multiple workers who can communicate about the firm’s past history (e.g., [Levin 2002](#); [Greif, Milgrom, and Weingast 1994](#)). The other game that we study is the “community enforcement game,” wherein a community is made up of bilateral partnerships, each modeled as a repeated two-sided Prisoner’s Dilemma (e.g., [Kandori 1992](#);

Jackson, Rodriguez-Barraquer, and Tan 2012). For both games, we consider perfect *public* monitoring, as well as perfect *private* monitoring—i.e., in which each player monitors her own partnerships perfectly, but observes nothing that happens in other partnerships.

Our results emerge most crisply in the agents-and-principal game (for both public and private monitoring) and with perfect public monitoring in the community enforcement game. In these settings, the highest level of cooperation that can be achieved by *any* equilibrium can be achieved by a renegotiation-proof equilibrium. Our result is partially constructive: we fully characterize a weakly renegotiation-proof equilibrium that can support the best equilibrium outcome, and use it to partially characterize an equilibrium that itself does not invite renegotiation to another weakly renegotiation-proof equilibrium at any history.

For the community enforcement game with perfect private monitoring, we have to enhance standard renegotiation concepts to accommodate private information. Although renegotiation is a well-studied problem, there is no established notion of how partners may secretly renegotiate in private monitoring environments. Even the act of initiating renegotiation could reveal information, and would have to be treated as an action undertaken in equilibrium. A formal treatment would require a general solution to the problem of bargaining with incomplete information, which has not been undertaken for even two players in a dynamic private monitoring environment.

Instead, our approach imposes the restriction that players can renegotiate using only information that is commonly known to them. When monitoring is public, our notion coincides with standard renegotiation proofness. However, under private monitoring our notion is best described as *bilateral renegotiation proofness*, because it rules out equilibria in which any pair of partners meeting within the game can bilaterally renegotiate, in secret, to gain a strict improvement for both of them. All such renegotiations take place only on the basis of what is commonly known between them, and when it is commonly known that such improvements are possible.

Using this solution concept, we compare the payoffs that emerge from an optimal equilibrium without renegotiation proofness—namely contagion (Kandori 1992; Ellison 1994)—with the payoffs that emerge from a variant of contagion that is bilateral renegotiation-proof. In the community enforcement game with private monitoring, bilateral renegotiation does reduce the level of cooperation that can be sustained. Nonetheless, the fraction of surplus that is lost from having to assure renegotiation proofness converges to zero in large communities.

1.2 Relationship to Prior Work

We build on research in both the renegotiation proofness and the multilateral enforcement literatures. Our motivation comes from contagion equilibria—introduced by Kandori (1992), Ellison (1994), and Harrington (1995)—in which a player shirks on all other players as soon as she observes a defection. While such a social norm provides incentives to exert effort, its punishments are peculiar in two distinct respects. First, cooperation between “innocent” players dissolves once they

become aware that a guilty party has shirked somewhere in the population, whereas it is more natural for innocent players to continue to cooperate among themselves while ostracizing a defector.¹ But a second issue, perhaps equally troubling, is that even between an innocent and a guilty player, it may not be credible that they should suspend cooperation and destroy surplus. The partners are jointly tempted to resume cooperation in a self-enforcing way—to “let bygones be bygones”.

Prior work has examined renegotiation in multilateral enforcement, limiting attention to perfect public monitoring environments.² The broad theme in this prior work is that renegotiation considerations limit the degree of cooperation. Greif, Milgrom, and Weingast (1994) in their study of medieval merchant guilds, allow individual merchants to trade at a bilateral level even with a city being embargoed by their guild. Since their notion of bilateral trade does not allow for asymmetric punishments, such embargo-breaking limits the power of multilateral enforcement. Most closely related, Jackson, Rodriguez-Barraquer, and Tan (2012) investigate the impact of renegotiation on favor-trading networks with fixed favor sizes. They argue that the combination of renegotiation proofness and “robustness to social contagion” selects particular network structures (“social quilts”) that support favor exchanges.

Our paper offers a contrasting message to these prior results: we show that if every partnership can renegotiate and adjust its level of cooperation, then high cooperation can be sustained in a way that is both renegotiation-proof and robust to social contagion.³ An important difference is that we use asymmetric behavior off the equilibrium path in a public monitoring environment to ensure that deviators are punished in a way that is renegotiation-proof.

One direction in which our work departs from this vein is that we also consider settings in which each partnership privately observes its history of play. These settings take us outside the realm of prior work on renegotiation proofness more generally, which restricted attention public (if sometimes imperfect) monitoring.⁴ The main differences among the various notions of renegotiation proofness in the literature regard which alternative equilibria are considered valid renegotiation targets. Our notion of bilateral renegotiation proofness builds on this tradition. When a pair of players renegotiates bilaterally, any bilateral continuation equilibrium satisfying certain conditions is a valid

¹This issue motivates the focus on ostracism and communication incentives in Ali and Miller (2016).

²The issue of renegotiation is also discussed in Appendix C of Greif (2006), which explains why renegotiation isn't a problem when workers can be replaced via thick labor markets, or when merchant guilds could punish merchants for renegotiating with a ruler. Kletzer and Wright (2000) analyze renegotiation proofness between a sovereign borrower and lender, and construct a renegotiation-proof equilibrium that is invulnerable to side deals between the borrower and other potential lenders. A separate line of work examines deviations by coalitions of players smaller than the grand coalition. Ray and Vohra (2001) highlight how the presence of inefficient equilibria may prevent coalitions of players from credibly punishing free-riders in a public good environment with binding contracts. Analogously, Genicot and Ray (2003) describe risk-sharing arrangements that are immune to deviations by sub-groups that can sustain risk-sharing arrangements on their own. Karlan, Möbius, Rosenblat, and Szeidl (2009) and Ambrus, Möbius, and Szeidl (2014) study borrowing and risk-sharing practices when individuals may default and coalitions of players may deviate in their obligations.

³For simplicity, we limit attention to symmetric graphs, but extending the result to asymmetric graphs would not pose a challenge.

⁴See, for example, Rubinstein (1980), Bernheim and Ray (1989), Farrell and Maskin (1989), van Damme (1989), Asheim (1991), Ray (1994), Goldlücke and Kranz (2013), Miller and Watson (2013), and Safronov and Strulovici (2014).

renegotiation target if it is common knowledge between them that it would deliver a strict Pareto improvement. The certain conditions that must be satisfied are essentially the same as those that must be satisfied in a “stationary Pareto perfect” equilibrium (Asheim 1991) of a two-player game.

The perfect private monitoring structure of our environment restricts the scope for all renegotiations to be bilateral. Several works have featured bilateral renegotiation in related contexts. Ghosh and Ray (1996) study a model with myopic and non-myopic types and introduce “bilateral rationality” as a restriction on play once each player in a partnership recognizes that the other isn’t myopic. Bilateral rationality requires that the pair then renegotiates to the best equilibrium of a model in which it is common knowledge that both players are non-myopic. Choy (2015) also assumes bilateral rationality in studying how segregation can be maintained through community enforcement. Less closely related, Serrano and Zapater (1998) impose bilateral renegotiation proofness on a finite-horizon environment in which partnerships impose externalities on each other.

2 The Community Enforcement Game

2.1 Environment

Consider a society of players $N = \{1, \dots, n\}$ in which each pair of players i and j is engaged in a partnership ij that meets at random times generated by a Poisson process of rate $\lambda > 0$. Meetings are i.i.d. across relationships and time. Whenever partnership ij meets, they first have the opportunity to renegotiate, and then they play a stage game in which they simultaneously choose effort levels ϕ_i, ϕ_j in $[0, \infty)$. Each pair has access to randomization devices whose realizations are observed only by that pair; the players can also access randomization devices that are public to all. Players discount their payoffs over time using a common discount rate $r > 0$.

Our variable-stakes formulation follows Ghosh and Ray (1996): player i ’s stage game payoff function when partnership ij meets is $b(\phi_j) - c(\phi_i)$, where $b(\phi_j)$ is the benefit from her partner j ’s effort and $c(\phi_i)$ is the cost she incurs from her own. The benefit function b and the cost function c are smooth functions satisfying $b(0) = c(0) = 0$. The social value of effort is $\pi(\phi) \equiv b(\phi) - c(\phi)$. We make two assumptions below. Assumption 1 implies that higher effort is always socially beneficial, and that holding average effort fixed, it’s always better (from a utilitarian standpoint) for partners to exert a symmetric level of effort. Assumption 2 guarantees that equilibrium effort is bounded, and that a partnership is willing to exert higher effort only if there are stronger incentives to do so.

Assumption 1. π is weakly concave and there exists $\psi > 0$ such that $\pi'(\phi) > \psi$ for every ϕ .

Assumption 2. c is strictly increasing and strictly convex, with $c'(0) = 0$ and $\lim_{\phi \rightarrow \infty} c'(\phi) = \infty$. The “relative cost” $c(\phi)/\pi(\phi)$ is strictly increasing.

We consider two monitoring environments. If there is *perfect public* monitoring, then every interaction is immediately observed by all players. If instead there is *perfect private* monitoring, then

players observe only the interactions within their own partnerships; they observe neither when other partnerships are recognized nor what happens when they are recognized. In both cases, monitoring is “perfect” in the sense that whatever players observe they observe without error; we occasionally refer to these as public and private monitoring below.

2.2 Histories and strategies

An *interaction* between players i and j at time t comprises the time t at which their partnership meets, their names, and their effort choices. We focus first on the case of private monitoring, where any interaction of a partnership ij is perfectly observed by partners i and j , but is unobserved by any other player, deferring the simpler case of public monitoring to the end of this subsection. Player i 's *private history*, denoted h_i^t , is the set of all interactions in which she has participated at any time $\tau < t$. We consider only histories that are *regular*—those histories in which no two links have ever been recognized simultaneously. (The set of histories that are not regular is reached with zero probability, regardless of the players' strategies.) A strategy for player i determines the mixed action she should choose when meeting a partner, given her private history. Specifically, if $\sigma_{ij}(h_i^t) \in \Delta[0, \infty)$ is the mixed action she chooses when meeting partner j at private history h_i^t , then $\sigma_i = (\sigma_{ij})_{j \neq i}$ is her strategy and $\sigma = (\sigma_i)_{i \in N}$ is the strategy profile.

Let $U_{ij}(\sigma|h_i^t)$ be the expected continuation payoff that player i obtains within partnership ij from strategy profile σ , starting from meeting partner j at private history h_i^t . We denominate U_{ij} in flow terms; i.e., given the discount rate r , $U_{ij}(\sigma|h_i^t)$ is the constant payoff arriving at Poisson rate λ such that at history h_i^t player i is indifferent between this constant flow and strategy profile σ .

The *ij-partnership history*, denoted h_{ij}^t , is the set of all interactions within partnership ij up to time t . Since players observe only what happens in their own partnerships, when players i and j meet with private histories h_i^t and h_j^t , their partnership history is simply $h_{ij}^t = h_i^t \cap h_j^t$. We say that a strategy σ_i for player i is *bilateral in partnership ij* if $\sigma_{ij}(h_i^t) = \sigma_{ij}(\tilde{h}_i^t)$ for all h_i^t and \tilde{h}_i^t that share the same partnership history h_{ij}^t , and for all t ; i.e., player i 's behavior in partnership ij is measurable with respect to the *ij-partnership history*. A strategy profile σ is *bilateral in partnership ij* if σ_i and σ_j are bilateral in partnership ij ; σ is *bilateral* if it is bilateral in every partnership. Even in a strategy profile σ that is not bilateral in partnership ij , we say that the continuation strategy of player i is *bilateral in partnership ij after partnership history* \hat{h}_{ij}^t if $\sigma_{ij}(h_i^t) = \sigma_{ij}(\tilde{h}_i^t)$ for all h_i^t and \tilde{h}_i^t that share the same partnership history $h_{ij}^t \supset \hat{h}_{ij}^t$, and for all $t \geq \hat{t}$; i.e., if player i 's behavior on link ij is measurable with respect to the *ij-partnership history* for all *ij-partnership histories* that succeed \hat{h}_{ij}^t . If partners i and j are playing bilateral continuation strategies after partnership history \hat{h}_{ij}^t , we write $U_{ij}(\sigma|h_{ij}^t) \equiv U_{ij}(\sigma|h_i^t)$ for any $h_{ij}^t \subset h_i^t$ that succeeds \hat{h}_{ij}^t .

The *ij-joint payoff set* of a strategy profile σ at partnership history h_{ij}^t is the collection of all expected payoff vectors arising from activity in partnership ij given σ at all pairs of private histories that contain h_{ij}^t : $\left\{ \left(U_{ij}(\sigma|\tilde{h}_i^t), U_{ji}(\sigma|\tilde{h}_j^t) \right) : \tilde{h}_i^t \cap \tilde{h}_j^t \supset h_{ij}^t \right\}$. It contains all continuation payoff vectors

the partners could expect within their partnership at pairs of private histories that succeed partnership history h_{ij}^t .

If there is public monitoring, then all players observe all activities. The *public history*, denoted h^t , is simply the union of all private histories: $h_i^t = \cup_{i \in N} h_i^t$. With public monitoring, a strategy for player i determines the mixed action she should choose when meeting a partner, given the public history; i.e., $\sigma_{ij}(h^t) \in \Delta[0, \infty)$ is the mixed action she chooses when meeting partner j at public history h^t . For public monitoring, the *public payoff set* of a strategy profile σ at public history h^t is the expected payoff vector for all players at all public histories that succeed h^t : $\left\{ \left(\sum_{j \neq i} U_{ij}(\sigma|\tilde{h}^t) \right)_{i \in N} : \tilde{h}^t \supset h^t \right\}$.

2.3 Equilibrium concept

We study “plain” *perfect Bayesian equilibrium* (PBE; [Watson 2016](#)), a notion of equilibrium in which each player maintains and updates a belief that is a probability distribution over strategies of other players. This notion imposes sequential rationality at all information sets, Bayesian updating on the equilibrium path, and some constraints on updating off the equilibrium path; [Appendix B](#) details it fully. A plain PBE always satisfies subgame perfection, and is equivalent to subgame perfect equilibrium in games of perfect monitoring. We restrict attention to PBEs in which all effort choices are uniformly bounded across histories,⁵ and effort choices on the equilibrium path are stationary. We refer to these as *equilibria*.⁶

2.4 Renegotiation proofness

A troubling property of many interesting equilibria is that once behavior is off the equilibrium path, the players may commonly understand that the play prescribed by the equilibrium is Pareto dominated by other self-enforcing arrangements, which motivates them to renegotiate their play. This prospect motivates the study of renegotiation-proofness more broadly, and we describe below how we extend prior notions to our settings, beginning with the case of perfect public monitoring.

Following [Farrell and Maskin \(1989\)](#) and [Bernheim and Ray \(1989\)](#), we define “weak renegotiation proofness,” which allows an alternative strategy profile to be a valid renegotiation target if it is available at some public history in their current strategy profile.

Definition 1. An equilibrium σ is *weak renegotiation-proof* (WRP) if for each $t \geq 0$ and each public history h^t , there does not exist any set of alternative public histories $\{\tilde{h}^t\}$ and a randomization $p \in \Delta\{\tilde{h}^t\}$, such that $\mathbb{E}_p \sum_{j \neq i} U_{ij}(\sigma|\tilde{h}^t) > \sum_{j \neq i} U_{ij}(\sigma|h^t)$ for all $i \in N$ and all public histories h^t .

⁵The restriction eliminates unreasonable equilibria in which effort grows with further cooperation, eventually exploding to infinity.

⁶While the restriction to equilibria that are stationary on the path of play is with loss of generality, we know of no approaches to construct optimal history-dependent equilibria for a fixed discount factor in a continuous-action environment. For analogous reasons, [Bernheim and Madsen \(2016\)](#) also restrict attention to equilibria that are stationary on the path of play in a continuous-action collusion environment.

Players should also have the option to renegotiate to a “better” WRP equilibrium, if one exists. We focus on a demanding notion of what it means for one equilibrium to be “better” than another. A set $W \subset \mathbb{R}^n$ *Pareto dominates* another set $W' \neq W$ if

1. For every $a \in W$, there does not exist any $b \in W' \setminus W$ for which $b_i > a_i$ for all $i \in N$;
2. For every $b \in W' \setminus W$, there exists $a \in W$ such that $a_i > b_i$ for all $i \in N$.

Definition 2. An equilibrium σ is *renegotiation proof* (RP) if for each $t \geq 0$ and each public history h^t , there does not exist any alternative WRP equilibrium σ' such that the public payoff set of σ at public history h^t is Pareto dominated by the public payoff set of σ' .

Pareto dominance of the payoff set at a history h^t means the alternative WRP equilibrium must be not merely better for all players at that history, but also not worse for them at any possible future history. In this sense, renegotiation is “forward-looking”: players cannot credibly renegotiate in a way that makes them worse off at some future history because they anticipate that they would renegotiate at that future history.⁷

2.5 Bilateral renegotiation proofness

Players can renegotiate only on the basis of information that is common knowledge. When only members of a partnership (perfectly) observe interactions within that partnership—i.e., *perfect private monitoring*—this restriction implies that all renegotiations must be bilateral.

Definition 3. An equilibrium σ is *weak bilateral renegotiation-proof* (WBRP) if for each partnership ij , each $t \geq 0$, and each partnership history h_{ij}^t , there does not exist any set of alternative partnership histories $\{\tilde{h}_{ij}^t\}$ and a randomization $p \in \Delta\{\tilde{h}_{ij}^t\}$, such that σ is bilateral in partnership ij after each \tilde{h}_{ij}^t , and $\mathbb{E}_p U_{ij}(\sigma|\tilde{h}_{ij}^t) > U_{ij}(\sigma|h_i^t)$ and $\mathbb{E}_p U_{ji}(\sigma|\tilde{h}_{ij}^t) > U_{ji}(\sigma|h_j^t)$ for all h_i^t and h_j^t satisfying $h_i^t \cap h_j^t = h_{ij}^t$.

For purposes of WBRP, for partners at some partnership history, a continuation strategy profile is a valid renegotiation target if it is bilateral, it is available at some partnership history (or randomization over partnership histories) in their current strategy profile, and it would make them both strictly better off no matter what are their private histories. A simpler, but looser phrasing is that the partners should renegotiate if they have common knowledge that they can both be better off switching to some history at which they play bilaterally.

Similarly, partners can renegotiate if they have common knowledge that there exists an alternative WBRP equilibrium that Pareto dominates their joint payoff set at some history.

⁷Because the forward-looking requirement constrains the valid renegotiation targets, this notion is weaker than “strong renegotiation proofness” (Farrell and Maskin 1989) or “consistency” (Bernheim and Ray 1989). Within the renegotiation proofness literature, it is closest to “stationary Pareto perfection” (Asheim 1991).

Definition 4. An equilibrium σ is *bilateral renegotiation proof* (BRP) if for each partnership ij , each $t \geq 0$, and each partnership history h_{ij}^t , there does not exist any alternative WBRP equilibrium σ' such that σ' is bilateral in partnership ij and the ij -joint payoff set of σ at partnership history h_{ij}^t is Pareto dominated by the ij -joint payoff set of σ' .

3 Renegotiation in the Community Enforcement Game

3.1 Outline

This section describes our results in the community enforcement game. We first consider the case of two players and characterize a RP equilibrium that attains the maximum level of cooperation available in any equilibrium (Proposition 1). This equilibrium uses asymmetric punishments to impose minimax payoffs on whichever player deviates.

Next we turn to a larger community of $n > 2$ players. For the case of public monitoring, we show that maximal cooperation is supported by requiring a defector to pay “penance” by working while his next partner shirks, after which he is readmitted to equilibrium-path cooperation.

With private monitoring, such penance is not feasible, since only a deviator’s partner knows he has deviated. Without renegotiation proofness, cooperation is maximized by contagion strategies (Kandori 1992; Ellison 1994). We augment contagion strategies with asymmetric bilateral punishments to characterize a contagion-like BRP equilibrium.

3.2 Two-Player Games: A Building Block

Suppose $n = 2$, so monitoring is public, and no player has private information. Every equilibrium is necessarily bilateral, so the RP and BRP equilibria coincide. We identify a RP equilibrium with several useful properties that we leverage in our later results on many-player games.

Let ϕ^B be the largest solution to a player’s incentive constraint when effort is constant and symmetric on the equilibrium path, but zero off the equilibrium path:

$$c(\phi) \leq \frac{\lambda}{r} \pi(\phi). \quad (1)$$

By Assumption 2, a solution exists and satisfies $\phi^B < \infty$; it is the highest level of symmetric effort that can be supported by any equilibrium. Our result below shows that this is the utilitarian optimal equilibrium effort, and can be supported in a RP equilibrium. Renegotiation proofness, therefore, comes at no cost with two players.

Proposition 1. *The sum of flow payoffs in any equilibrium σ , $U_1(\sigma) + U_2(\sigma)$, is bounded by $2\pi(\phi^B)$. There exists a RP equilibrium in which each player chooses effort ϕ^B at every equilibrium path history.*




FIGURE 1. Bilateral WRP equilibrium σ^* , automaton for player i . The initial phase is outlined in bold, and transitions driven by equilibrium behavior are shown with bold arrows. Not all transitions are shown.

The proof follows three steps (broken into four lemmas below). First, we show that $2\pi(\phi^B)$ is an upper-bound for the sum of flow payoffs in any equilibrium. Then we establish the existence of a WRP equilibrium in which both players choose effort ϕ^B at every equilibrium path history. This equilibrium is depicted in Fig. 1. This equilibrium involves a player being punished maximally in a single period by being forced to exert effort ϕ^B while her opponent shirks, and following that punishment phase the players return to bilateral cooperation. (This WRP equilibrium is a straightforward generalization of van Damme (1989) for fixed-stakes prisoners’ dilemmas.) In principle, this equilibrium could be Pareto dominated by an equilibrium that “smooths” out the punishments in this punishment phase, so we cannot guarantee that this equilibrium is renegotiation-proof. Our final (non-constructive) step establishes that if this equilibrium is not renegotiation-proof, then a renegotiation-proof equilibrium exists that supports the same equilibrium path behavior. Because Proposition 1 is the building block of subsequent results, we include its formal proof below.

Proof of Proposition 1: We remind the reader that we restrict attention to equilibria that are stationary on the path of play. Lemma 1 bounds the payoffs attained by an equilibrium in this class.

Lemma 1. *For every equilibrium, the sum of equilibrium path payoffs is bounded by $2\pi(\phi^B)$.*

Proof. Let μ be the distribution of equilibrium path efforts: i.e., on the equilibrium path (ϕ_1, ϕ_2) are derived from distribution μ for every history on the equilibrium path. Let Γ_i^μ be the support of ϕ_i in the marginal distribution on ϕ_i induced by μ . Observe that a necessary condition for an equilibrium is that for every ϕ'_1 in Γ_1^μ and ϕ'_2 in Γ_2^μ ,

$$c(\phi'_1) \leq \frac{\lambda}{r} (\mathbb{E}_\mu[b(\phi_2) - c(\phi_1)]), \quad (\text{Player 1's IC})$$

$$c(\phi'_2) \leq \frac{\lambda}{r} (\mathbb{E}_\mu[b(\phi_1) - c(\phi_2)]). \quad (\text{Player 2's IC})$$

In each case, the left-hand side (LHS) includes the one period gain from deviating by choosing $\phi = 0$ rather than ϕ'_i and obtaining the minimax payoff of 0 thereafter. The right-hand side (RHS) for each has the loss from foregoing equilibrium path payoffs. Adding these two together yields that for every

ϕ'_1 in Γ_1^μ and ϕ'_2 in Γ_2^μ , an aggregate constraint must be satisfied:

$$c(\phi'_1) + c(\phi'_2) \leq \frac{\lambda}{r} (E_\mu[\pi(\phi_1) + \pi(\phi_2)]), \quad (\text{Aggregate IC})$$

where we substitute $\pi(\phi) = b(\phi) - c(\phi)$. By [Assumptions 1 and 2](#), the supports Γ_1^μ and Γ_2^μ must have a finite upper bound. Let $\bar{\phi}_i^\mu$ be the highest ϕ_i in Γ_i^μ .

A utilitarian optimal equilibrium maximizes $\mathbb{E}_\mu[\pi(\phi_1) + \pi(\phi_2)]$ subject to all incentive constraints, on and off the path of play, for each player. Imposing only [Player 1's IC](#) and [Player 2's IC](#) is a relaxed problem. Because these constraints together imply [Aggregate IC](#), the following is a further relaxed problem:

$$\max_{\mu} \mathbb{E}_\mu[\pi(\phi_1) + \pi(\phi_2)] \text{ subject to } \text{Aggregate IC.}$$

A solution to the above problem involves $\Gamma_i^\mu = \{\bar{\phi}_i^\mu\}$ (i.e., a degenerate distribution), because any non-degenerate distribution that satisfies [Aggregate IC](#) can be improved by putting all of its mass on its highest realization. Therefore, the relaxed problem can be re-written as

$$\max_{\phi_1, \phi_2} \pi(\phi_1) + \pi(\phi_2) \text{ subject to } c(\phi_1) + c(\phi_2) \leq \frac{\lambda}{r} (\pi(\phi_1) + \pi(\phi_2)).$$

We argue, using a proof by contradiction, that this relaxed problem has a symmetric solution. Suppose that (ϕ'_1, ϕ'_2) satisfies the above constraint and $\phi'_1 \neq \phi'_2$. Consider the vector $(\bar{\phi}, \bar{\phi})$ where $\bar{\phi} = \frac{\phi_1 + \phi_2}{2}$. Because c is strictly convex and π is weakly concave, Jensen's Inequality implies that

$$2c(\bar{\phi}) < c(\phi'_1) + c(\phi'_2) \leq \frac{\lambda}{r} (\pi(\phi'_1) + \pi(\phi'_2)) \leq \frac{2\lambda}{r} \pi(\bar{\phi}).$$

Therefore, shifting to a symmetric solution generates slack in the constraint and (weakly) improves the objective. Accordingly, the optimal solution is symmetric, and binds the constraint. Notice that the symmetric solution to the binding constraint is ϕ^B , which establishes our claim. \square

Lemma 2. *There exists a WBRP equilibrium that supports ϕ^B on the equilibrium path.*

Proof. Consider a strategy profile σ^* in which player i 's strategy follows the three-phase automaton:

1. Bilateral Cooperation: Choose stakes ϕ^B .
2. Punishment: Choose stakes ϕ^B .
3. Reward: Choose stakes 0.

In any phase, if neither player deviates from expected play or both players deviate simultaneously, transition to the Bilateral Cooperation phase. In any phase, if player i deviates, transition to the Punishment phase, while if player $-i$ deviates then transition to the Reward phase. This automaton is illustrated in [Fig. 1](#) on p. 9.

We verify that σ^* is an equilibrium by using the one-shot deviation principle. Let us first describe expected (flow) payoffs in each phase. In the cooperation phase, the constant flow payoff is $\pi(\phi^B)$. In the punishment phase, because $c(\phi^B) = \frac{\lambda}{r}\pi(\phi^B)$, the average flow payoff is 0. In the reward phase, the average flow payoff is

$$z^* = \frac{r}{\lambda} \left[\frac{\lambda}{r+\lambda} (b(\phi^B) + \frac{\lambda}{r}\pi(\phi^B)) \right] = \pi(\phi^B) + \frac{r}{r+\lambda} b(\phi^B).$$

By construction, no player has an incentive to deviate in the Bilateral Cooperation phase: because each player obtains her minimax payoff following a deviation, **Player 1's IC** and **Player 2's IC** characterize the relevant incentive constraints and are satisfied. In the punishment phase, the player has an incentive to choose ϕ^B and obtain a payoff of 0 because no choice of stakes offers her a strictly positive payoff. In the reward phase, the player is supposed to choose stakes of 0, which maximizes both her stage game payoff and her continuation payoff.⁸

Finally, we verify that the equilibrium satisfies WBRP. Let W^* be the convex hull of $(z^*, 0)$, $(0, z^*)$, and $(\pi(\phi^B), \pi(\phi^B))$. Since $z^* > \pi(\phi^B) > 0$, the equilibrium chooses points in W^* that are on the Pareto frontier of W^* . \square

Lemma 3. *Consider any sequence of 1, 2-joint payoff sets $\{W_k\}_{k=1}^\infty$ of WBRP equilibria, starting with W_1 being the Pareto frontier of W^* , such that for each k , W_{k+1} Pareto dominates W_k . There exists a limit set $W = \lim_{k \rightarrow \infty} W_k$ that is the 1, 2-joint payoff set of a WBRP equilibrium such that both players choose ϕ^B along the equilibrium path, and such that partner i earns a continuation payoff of zero after deviating unilaterally.*

Because this proof is involved and technical, we relegate it to the Appendix.

Lemma 4. *There exists a BRP equilibrium that supports ϕ^B on the equilibrium path.*

Proof. Suppose not; then there must exist an infinite sequence of distinct WBRP equilibria, starting with σ^* , such that the joint payoff set of each Pareto dominates its predecessors, but such that there does not exist any WBRP equilibrium whose joint payoff set Pareto dominates every member of the sequence. By **Lemma 3**, however, the limit W of this sequence itself is the 1, 2-joint payoff set of a WBRP equilibrium that Pareto dominates every member of the sequence, a contradiction. \square

3.3 Many-Player Games with Perfect Public Monitoring

Next we address a multiplayer setting with public monitoring. A natural equilibrium is one in which all players exert ϕ so long as all players have done so in the past, and otherwise all players exert zero

⁸Since monitoring is perfect, the constructed subgame perfect equilibrium necessarily corresponds to a plain PBE where each player's beliefs about her opponent's play is correct in every subgame.

effort. Such an equilibrium features the following incentive constraint on the equilibrium path:

$$(n-1)\frac{\lambda}{r}\pi(\phi) \geq c(\phi). \quad (\text{IC}_{ij}^{\text{Perfect}})$$

The left-hand side captures the future gains from cooperating with $n-1$ neighbors while the right-hand side captures the one-time gain from shirking. Maximizing ϕ subject to $\text{IC}_{ij}^{\text{Perfect}}$ yields an equilibrium that is (utilitarian) optimal; we denote this level of effort by ϕ_n^{PM} .

This equilibrium invites renegotiation: why should all partnerships cease cooperation as soon as a single player has shirked? Even further, why should players cease cooperation even with the first defector when all arrangements are open to renegotiation? We establish that other punishments—which are renegotiation-proof—can also support this level of effort.

Proposition 2. *With perfect public monitoring, the highest equilibrium path effort profile can be supported by an RP equilibrium.*

The proof, in [Appendix A.2](#), adapts the equilibrium of [Section 3.2](#). A player who deviates is, upon meeting her next partner, required to pay “penance” by exerting effort ϕ_n^{PM} while her partner shirks. Thereafter all players return to equilibrium-path behavior.

Comparison with prior results: Insofar as the prior literature has articulated how renegotiation considerations can impede multilateral enforcement with public monitoring, we describe here how we reach a different conclusion. The immediate comparison is to [Jackson, Rodriguez-Barraquer, and Tan \(2012\)](#) who also study renegotiation in the context of community enforcement with perfect monitoring, and argue that certain networks feature renegotiation-proof favor exchanges but others do not. To the contrary, we show in [Appendix A.2](#) that our result applies to any regular graph, and it is straightforward to extend this logic to any graph.

Two conceptual differences are responsible for generating our contrasting results. First, we study different variants of renegotiation proofness. They propose a variant in which some punishments are not open to renegotiation: if player i shirks on player j , the relationship ij is mechanically severed, so i and j are forced to punish each other. All other relationships, however, are open to renegotiation. By contrast, our notion of renegotiation proofness allows a defector to renegotiate with her victims and third-parties, adhering to conventional approaches to renegotiation. Second, they study a fixed-stakes environment, where the level of cooperation cannot respond to past history. By contrast, our model allows partners to adjust their level of cooperation to suit the situation, enabling them to discourage renegotiation through asymmetric play.⁹ Our analysis illustrates that giving players more flexibility—to renegotiate all punishments and to adjust their cooperation—enables multilateral enforcement to sidestep issues of renegotiation.¹⁰

⁹Lippert and Spagnolo (2011) construct related, but more complex, multilateral repentance equilibria for networks of fixed-stakes, asymmetric relationships.

¹⁰Jackson, Rodriguez-Barraquer, and Tan (2012) also impose a concept of “robustness to social contagion” whereby

3.4 Many-Player Games with Perfect Private Monitoring

Private monitoring impedes all the neighbors of a deviator from coordinating their punishment, which was necessary for [Proposition 2](#). Contagion sidesteps the issue of coordination by destroying all cooperation in the community as quickly as possible. Of course, such destruction is not (bilateral) renegotiation-proof. In this section we adapt contagion strategies to employ the two-player RP equilibrium from [Section 3.2](#) within each partnership once both partners are off the equilibrium path. Each player cooperates at constant effort along the equilibrium path, but becomes contagious and starts shirking as soon as he observes a deviation. But unlike contagion, a contagious player shirks on each partner *just once*. Upon the first time either partner shirks in a relationship, it becomes common knowledge between them that they are off the equilibrium path. Then the partners play the two-player RP equilibrium from [Proposition 1](#) within their relationship, without affecting what happens in other relationships. Such an equilibrium is immune to bilateral renegotiation, although it does not attain the same level of cooperation as ordinary contagion.

To maximize cooperation under bilateral renegotiation proofness, we tailor the off-path behavior in each partnership depending on who shirked first. Off the equilibrium path, if player i shirked first on player j , then in their subsequent play, i is punished while j is rewarded. However, if both partners simultaneously first shirked on each other, then they cooperate symmetrically.

To state our result, we develop notation for the “rate” at which contagion spreads: namely, once player i shirks on player j at time 0, what is the discounted probability that player k is still in the cooperation phase when players i and k next meet? Any payoff that player i may gain from shirking on player k at that time has to be normalized by the discount rate, and so our interest is in

$$X_n \equiv \int_0^\infty e^{-rt} \lambda x_n(t) dt,$$

where $x_n(t)$ is the probability that player $k \notin \{i, j\}$ is not contagious at time t , after player i shirks on player j at time 0.¹¹

Consider what would happen in a conventional contagion equilibrium that is symmetric on the path of play. Suppose that on the equilibrium path each player always chooses effort ϕ , but any player who has observed any deviation shirks on all her neighbors. Without the possibility of renegotiation, the equilibrium path incentive constraint for player i when meeting player j would be

$$\frac{(n-1)\lambda}{r} \pi(\phi) \geq c(\phi) + (n-2)X_n b(\phi). \quad (\text{IC}^{\text{NoRNP}})$$

The left-hand side is the total continuation payoff from following equilibrium. The first term on the

punishments are contained to the local neighborhood of the link in which a defection took place. The equilibrium that we describe in the proof of [Proposition 2](#) is robust to social contagion.

¹¹[Ali and Miller \(2013\)](#) define X_n as the “viscosity factor” and offer a closed-form formula for the special case of a complete graph. Since we do not use that formula in our derivations in the main text, we do not reproduce it here.

right-hand side is the instantaneous gain from shirking on player j (because choosing an effort of 0 saves cost $c(\phi)$), and the second term accounts for the possibility that player i may meet his $n - 2$ other partners while they still think they are on the equilibrium path. Let ϕ_n^C be the effort level that binds this constraint; Ali and Miller (2013) show that if ϕ_n^C is the equilibrium path effort then all off-path incentive constraints are satisfied, and this ordinary contagion equilibrium is optimal.

However, ordinary contagion is susceptible to bilateral renegotiation because players perpetually shirk on each other off the equilibrium path. Each pair of players can privately renegotiate, and sustain—between themselves—cooperation at level of ϕ^B for a flow payoff of $\pi(\phi^B)$. In an equilibrium that is immune to such bilateral renegotiations, punishments cannot be so severe. Instead, we construct a WBRP contagion equilibrium as follows. Players always choose effort ϕ on the equilibrium path. When a player observes his first deviation, he becomes contagious and begins shirking on all his partners. But in each partnership, once both partners know they are off the equilibrium path, they renegotiate to the RP equilibrium from Section 3.2, such that if one of them shirked first then that partner is punished, but if both simultaneously shirked first then they cooperate symmetrically. The equilibrium path incentive constraint is:

$$\begin{aligned} (n-1)\frac{\lambda}{r}\pi(\phi) &\geq c(\phi) + (n-2) \int_0^\infty e^{-rt} \lambda \left(x_n(t)b(\phi) + (e^{-\lambda t} - x_n(t))\frac{\lambda}{r}\pi(\phi^B) \right) dt \\ &= c(\phi) + (n-2) \left(b(\phi)X_n + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r}\pi(\phi^B) \right). \end{aligned} \quad (\text{IC}_{ij}^{\text{RNP}})$$

The left-hand side is the same as before. The right-hand side comprises the instantaneous gain from shirking on player j , and the future prospects. Being off the equilibrium path leads to the following consequences when player i next meets player k at time t : if player k thinks she is still on the equilibrium path (which occurs with probability $x_n(t)$), then she chooses ϕ while player i shirks, which leads player i to earn a payoff of $b(\phi)$ at time t and then zero thereafter; otherwise, if player k knows she is off the equilibrium path (which occurs with probability $e^{-\lambda t} - x_n(t)$) then she shirks simultaneously with player i , and from that point on both cooperate at effort ϕ^B .¹² We find that the best symmetric BRP equilibrium binds this incentive constraint.

Proposition 3. *With private monitoring and $n > 2$, there exists a BRP equilibrium in which players always exert effort ϕ_n^R that binds Eq. (IC_{ij}^{RNP}). This BRP equilibrium outperforms bilateral enforcement, but underperforms the best ordinary contagion equilibrium—i.e., $\phi^B < \phi_n^R < \phi_n^C$. No other symmetric BRP equilibrium attains higher total payoffs on the equilibrium path.*

The proof is in Appendix A.3. As in the case of Proposition 1, although the WBRP contagion equilibrium we construct may not satisfy BRP, it implies the existence of an BRP equilibrium that

¹²This expression bears further explanation. Suppose, as we have, that player i deviates from the equilibrium path by shirking on link ij at time 0. Let event P_t be that at time t , players i and k have not met since time 0, and let event Q_t be that at time t , player k has seen any partner other than player i shirk. Then we claim that $e^{-\lambda t} - x_n(t)$ is the probability of event $P_t \cap Q_t$: The viscosity factor can be rewritten as $X_n = \int_0^\infty e^{-rt} \lambda e^{-\lambda t} \tilde{x}_n(t) dt$, where $e^{-\lambda t} = \Pr(P_t)$ and $\tilde{x}_n(t) = 1 - \Pr(Q_t|P_t)$, and therefore $\Pr(P_t \cap Q_t) = e^{-\lambda t}(1 - \tilde{x}_n(t)) = e^{-\lambda t} - x_n(t)$.

attains the same payoffs on the equilibrium path. While for expositional simplicity we have described behavior as operating on a complete network, our conclusions extend to any symmetric network.

Cost of Renegotiation: The *proportional cost of renegotiation* is the fraction by which the players' payoffs decrease when moving from an ordinary contagion equilibrium (without renegotiation) to the BRP equilibrium described in [Proposition 3](#). We prove that the proportional cost of renegotiation is small in large communities.

Proposition 4. *Suppose there exists $\epsilon > 0$ such that $\frac{1}{\epsilon} > b''(\phi) > \epsilon$ for all $\phi \geq 0$. Then*

$$\lim_{n \rightarrow \infty} \frac{\phi_n^R}{\phi_n^C} = 1, \quad (2)$$

where ϕ_n^C is the level of cooperation in the best contagion equilibrium (without renegotiation).

The proof, in [Appendix A.4](#), uses the condition on b'' and the fact that the viscosity factor X_n vanishes as n diverges to infinity to bound the differences between the incentive constraints [Eqs. \(IC^{NoRNP}\)](#) and [\(IC^{RNP}\)](#). The cost of renegotiation arises entirely from the off-path events in which both partners first shirk on each other at the same time, in which case they both receive strictly positive continuation rewards. This means if one of the partners happens to be the original deviator, in this relationship he is not being punished as harshly as he would be under ordinary contagion.

Contagion spreads more quickly in larger communities, increasing the chance that after a deviation each player $k \notin \{i, j\}$ will be contagious when player i first meets k . Since this is the case in which they both shirk on each other at the same time, it is this case in which the punishment in a contagion-like BRP equilibrium is less powerful than under contagion. However, as the community grows, the number of partners increases faster than the reduction in punishment power per partner. In the limit, the reduction becomes proportionally negligible.¹³

The conclusion of [Proposition 4](#) also holds for any class of symmetric networks in which the viscosity factor vanishes in the limit as $n \rightarrow \infty$. This condition is easily satisfied; indeed it can only be violated if average path lengths in the network grow very quickly as the number of players grows, which is contrary to the nature of large social networks that are typically observed (see, e.g., [Jackson 2008](#); [Goyal 2007](#)).

4 Renegotiation in Multilateral Relational Contracts

Here, we study a single principal interacting with multiple agents, as in [Greif, Milgrom, and Weingast \(1994\)](#), [Levin \(2002\)](#), [Rayo \(2007\)](#), and [Andrews and Barron \(2016\)](#). Multilateral enforcement, in

¹³Indeed, the conclusion of [Proposition 4](#) would still hold even if players simply renegotiated directly to the bilateral cooperation phase, though such an equilibrium would not be optimal among BRP equilibria on any given network.

this context, is used to punish the Principal in all of her relationships if she deviates on a single agent. The concern for renegotiation is that she may bilaterally renegotiate with each agent after such a defection so that the multilateral punishment is no longer credible. Below we describe a model and set-up in which punishments can be used to enforce the optimal arrangement without inviting renegotiation.

The principal (“she”) interacts with n agents (each a generic “he”), indexed by $1, \dots, n$. Each agent meets the principal at Poisson rate λ to play the following stage game: agent i chooses an effort level $\phi_i \geq 0$ and the principal simultaneously choose a wage payment $w_i \geq 0$. Their stage game payoffs are $w_i - c(\phi_i)$ for the agent and $b(\phi_i) - w_i$ for the principal. An agent’s effort is beneficial to the principal but costly to him; wage payments are pure (linear) transfers from the principal to the agent. We maintain [Assumptions 1](#) and [2](#).

The agents have no payoff-relevant interactions with each other, but can communicate about the principal’s behavior. This word-of-mouth communication across buyers may be used to discipline the principal, as is common in the literature on multilateral enforcement ([Klein 1992](#); [Dixit 2003](#); [Lippert and Spagnolo 2011](#); [Ali and Miller 2016](#)). To facilitate comparison with [Section 2](#), we maintain the same speed of information diffusion: all communication is bilateral and each pair of agents meets at Poisson rate λ . We assume that all communication between agents is cheap talk (and as we highlight, there are no communication incentive problems here).

We now describe the structure of equilibria. Because each agent has payoff interactions with only one party, namely the principal, no agent can be motivated to exert effort beyond “bilateral enforcement.” So if agent i is expected to exert ϕ_i on the equilibrium path and is paid w_i every time that she does so, the following incentive constraint must be satisfied:

$$w_i - c(\phi_i) + \frac{\lambda}{r} (w_i - c(\phi_i)) \geq w_i. \quad (\text{Agent's IC})$$

The left-hand side above captures the worker’s payoff from working today and in the future whereas the right-hand side captures the one-time gain from shirking. Agent i ’s incentive to exert effort ϕ_i for wage w_i is uninfluenced by the presence of other agents.

The principal, by contrast, is motivated by multilateral enforcement. In an optimal equilibrium, absent renegotiation concerns, each agent stops working for the principal once he learns that she has reneged on any payment to any agent; correspondingly, the principal never pays an agent who has ever shirked. Since news of the principal’s deviation spreads at the same rate as in the community enforcement game with $n + 1$ players, such an equilibrium generates the incentive constraint

$$b(\phi_i) - w_i + \frac{\lambda}{r} \sum_{j=1}^n (b(\phi_j) - w_j) \geq b(\phi_i) + \sum_{j \neq i} b(\phi_j) X_{n+1} \quad (\text{Principal's IC})$$

However, both [Agent’s IC](#) and [Principal’s IC](#) ignore renegotiation: once an agent or a principal has deviated, there is the temptation to bilaterally renegotiate. We prove in this setting that—unlike

the community enforcement game with private monitoring studied in [Section 3.4](#)—imposing bilateral renegotiation proofness does not result in any loss of equilibrium welfare.

Proposition 5. *The highest equilibrium path effort profile can be supported by a BRP equilibrium.*

The building block for this result is a RP equilibrium for two players, much like the one described in [Section 3.2](#); the details are given in [Lemma 10](#) in the [Appendix A.5](#). In this equilibrium, asymmetric play off the equilibrium path (either working for free, or paying wages for no effort) is used to achieve minimax payoffs within a single principal-agent partnership. The second step is to incorporate this equilibrium as the off-path behavior in a multilateral enforcement scheme. In this scheme, if an agent shirks first on the principal, then she works for free at high stakes before she and the principal return to the equilibrium path level of cooperation. By contrast, if the principal deviates, then the agents communicate among themselves about the principal’s deviation, to impose multilateral punishment on the principal. Once the principal has reneged on an agent, she and that agent switch to the bilateral enforcement equilibrium within their relationship, starting in the phase that punishes the principal with a continuation payoff of 0. These punishments are just as severe as if renegotiation were not possible.

5 Conclusion

Renegotiation is central to multilateral enforcement. Multilateral enforcement presumes coordination on an equilibrium in which a deviator is punished by multiple players. But players who find themselves at a “bad history” in which surplus is being destroyed may be able to renegotiate to a better outcome. Anticipating such renegotiation may temper or negate the threat of community enforcement, a concern that prior work has explicated.

Our main conclusion is that in multilateral enforcement against a single principal, or in community enforcement with public monitoring, renegotiation does not impinge on the level of cooperation that can be supported in equilibria. It is only when monitoring is private that renegotiation imposes some cost on community enforcement, but this cost proportionally vanishes in large societies.

References

- S. Nageeb Ali and David A. Miller. Enforcing cooperation in networked societies. Working paper, 2013.
- S. Nageeb Ali and David A. Miller. Ostracism and forgiveness. *American Economic Review*, 2016.
- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer, Berlin, 2nd edition, 1999.
- Attila Ambrus, Markus Möbius, and Adam Szeidl. Consumption risk-sharing in social networks. *American Economic Review*, 104(1):149–182, 2014.

- Isaiah Andrews and Daniel Barron. The allocation of future business: Dynamic relational contracts with multiple agents. *American Economic Review*, 2016.
- Geir G. Asheim. Extending renegotiation-proofness to infinite horizon games. *Games and Economic Behavior*, 3(3):278–294, August 1991.
- B. Douglas Bernheim and Erik Madsen. Price cutting and business stealing in imperfect cartelsbusiness stealing in imperfect cartels. Working paper, 2016.
- B. Douglas Bernheim and Debraj Ray. Collective dynamic consistency in repeated games. *Games and Economic Behavior*, 1(4):295–326, 1989.
- James P. Choy. Constructing social division to support cooperation. Working paper, July 2015.
- Avinash K. Dixit. Trade expansion and contract enforcement. *Journal of Political Economy*, 111(6): 1293–1317, 2003.
- Glenn Ellison. Cooperation in the prisoner’s dilemma with anonymous random matching. *Review of Economic Studies*, 61(3):567–588, July 1994.
- Joseph Farrell and Eric Maskin. Renegotiation in repeated games. *Games and Economic Behavior*, 1(4):327–360, 1989.
- Garance Genicot and Debraj Ray. Group formation in risk-sharing arrangements. *Review of Economic Studies*, 70(1):87–113, January 2003.
- Parikshit Ghosh and Debraj Ray. Cooperation in community interaction without information flows. *Review of Economic Studies*, 63(3):491–519, 1996.
- Susanne Goldlücke and Sebastian Kranz. Renegotiation-proof relational contracts. *Games and Economic Behavior*, 80:157–178, July 2013.
- Sanjeev Goyal. *Connections: An introduction to the economics of networks*. Princeton University Press, 2007.
- Avner Greif. *Institutions and the path to the modern economy: Lessons from medieval trade*. Cambridge Univ Press, New York, N.Y., 2006.
- Avner Greif, Paul Milgrom, and Barry R. Weingast. Coordination, Commitment, and Enforcement: The Case of the Merchant Guild. *Journal of Political Economy*, 102(4):745–776, 1994.
- Joseph E. Harrington, Jr. Cooperation in a one-shot prisoners’ dilemma. *Games and Economic Behavior*, 8:364–377, 1995.
- Matthew O. Jackson. *Social and economic networks*. Princeton University Press, 2008.
- Matthew O. Jackson, Tomas Rodriguez-Barraquer, and Xu Tan. Social capital and social quilts: Network patterns of favor exchange. *American Economic Review*, 102(5):1857–1897, 2012.
- Michihiro Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59(1): 63–80, 1992.
- Dean Karlan, Markus Möbius, Tanya Rosenblat, and Adam Szeidl. Trust and social collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, August 2009.

- Daniel B. Klein. Promise keeping in the Great Society: A model of credit information sharing. *Economics and Politics*, 4(2):117–136, July 1992.
- Kenneth M. Kletzer and Brian D. Wright. Sovereign debt as intertemporal barter. *American Economic Review*, 90(3):621–639, 2000.
- Jonathan Levin. Multilateral contracting and the employment relationship. *Quarterly Journal of Economics*, pages 1075–1103, 2002.
- Steffen Lippert and Giancarlo Spagnolo. Networks of relations and word-of-mouth communication. *Games and Economic Behavior*, 72(1):202–217, May 2011.
- David A. Miller and Joel Watson. A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350, November 2013.
- James R. Munkres. *Topology*. Prentice-Hall, 2nd. edition, 2000.
- Debraj Ray. Internally renegotiation-proof equilibrium sets: Limit behavior with low discounting. *Games and Economic Behavior*, 6(1):162–177, January 1994.
- Debraj Ray and Rajiv Vohra. Coalitional power and public goods. *Journal of Political Economy*, 109(6):1355–1385, 2001.
- Luis Rayo. Relational incentives and moral hazard in teams. *Review of Economic Studies*, 74(3):937–963, July 2007.
- Ariel Rubinstein. Strong perfect equilibrium in supergames. *International Journal of Game Theory*, 9(1):1–12, 1980.
- Mikhail Safronov and Bruno Strulovici. Explicit renegotiation in repeated games. Working paper, November 2014.
- Roberto Serrano and Inigo Zapater. The three-legged race: Cooperating to compete. *Games and Economic Behavior*, 22(2):343–363, 1998.
- Eric van Damme. Renegotiation-proof equilibria in repeated prisoners’ dilemma. *Journal of Economic Theory*, 47(1):206–217, 1989.
- Arthur R. Warburton. Quasiconcave vector maximization: Connectedness of the sets of Pareto-optimal and weak Pareto-optimal alternatives. *Journal of Optimization Theory and Applications*, 40(4):537–557, August 1983.
- Joel Watson. Perfect Bayesian equilibrium: General definitions and illustrations. Working paper, January 2016.

A Main Appendix

A.1 Proof of Lemma 3 on p. 11

Consider a sequence as specified above. First we establish some facts for each W_k .

1. Each W_k contains the point $\pi(\phi^B, \phi^B)$.

Proof: By Step 1, the sum of equilibrium payoffs is less than $2\pi(\phi^B)$ and therefore, W_k cannot contain a point that Pareto dominates (ϕ^B, ϕ^B) . Therefore, $(\phi^B, \phi^B) \in W_k$ for each k .

2. The endpoints of each W_k are $(z_k, 0)$ and $(0, \hat{z}_k)$, for some $z_k, \hat{z}_k \in [z^*, 2\pi(\phi^B)]$.

Proof: $\pi(\phi^B)$ is obtained by both partners cooperating at ϕ^B , which requires that any deviator receive a continuation payoff of no more than zero, which is her stage game minimax. W_k therefore contains some point $(z_k, 0)$. If $z_k < z^*$ then $(z_k, 0)$ would be strictly Pareto dominated by a convex combination of $\pi(\phi^B, \phi^B)$ and $(z^*, 0)$, contrary to the supposition that W_k Pareto dominates W_1 . Finally, by Step 1, we must have $z_k, \hat{z}_k \leq 2\pi(\phi^B)$.

Next we show that the pointwise limit $W = \lim_{k \rightarrow \infty} W_k$ exists and is the joint payoff set of a WBRP equilibrium. Using public randomization, without loss of generality we take each joint payoff set W_k to be the weak-Pareto frontier of the convex hull of the closure of its actual joint payoff set. It then follows from Theorem 4.1 of [Warburton \(1983\)](#) that W_k is connected and closed.

1. The pointwise limit $W = \lim_{k \rightarrow \infty} W_k$ exists, and is connected and internally weak-Pareto incomparable.

Proof: Since each $W_k \subset \mathbb{R}_+^2$ is internally weak-Pareto incomparable, we identify each W_k by a function $F_k(x) = \max\{y : (x, y) \in W_k\} \cup \{0\}$ on the domain $[0, 2\pi(\phi^B)]$. Each F_k is evidently decreasing and integrable, and is concave on the subdomain $[0, z_k]$. Since the sequence is increasing in the Pareto dominance ordering, $F_k(x)$ is increasing in k for all x . By our earlier steps, $F_k(0) \in [z^*, 2\pi(\phi^B)]$, $F_k(2\pi(\phi^B)) = 0$, and $\int_0^{2\pi(\phi^B)} F_k dx \leq 2\pi(\phi^B)$. Now the Monotone Convergence Theorem (e.g., [Aliprantis and Border 1999](#), Theorem 11.17) implies that there exists a decreasing and integrable function F that is the pointwise limit of $\{F_k\}$, with $\int_0^{2\pi(\phi^B)} F dx = \lim_{k \rightarrow \infty} \int_0^{2\pi(\phi^B)} F_k dx$. Since F is the pointwise limit of $\{F_k\}$, there exists $z = \inf\{x : F(x) \leq 0\} \in [z^*, 2\pi(\phi^B)]$ such that F is concave on $[0, z]$.

Define $W \equiv \{(x, F(x)) : x < z\} \cup \{(z, y) : y \in [0, F(z)]\}$. We have already shown that each point $(x, F(x))$ for which $x \leq z$ is a limit point of $\{W_k\}$, as is the point $(z, 0)$. Finally, each point (z, y) such that $y \in (0, F(z)]$ is a limit point of $\{W_k\}$ since each W_k is connected and concave. Therefore W as defined is in fact the pointwise limit of $\{W_k\}$.

2. There exists a WBRP equilibrium whose joint payoff set is W .

Proof: Fix any $v \in W$, and consider a sequence $v_k \rightarrow v$ such that $v_k \in W_k$ for each k . Fix a compact, metrizable state space Ω for the public randomization device. Since each W_k is the joint payoff set of an equilibrium, for each $v_k \in W_k$ there exists a public distribution ζ_{v_k} on Ω , efforts $\{\phi_{i;v_k} : \Omega \rightarrow \mathbb{R}_+\}_{i=1,2}$, and a promised utility function $w_{v_k} : \Omega \rightarrow W_k$ (simplified to deliver only the promised utility expected if neither player deviates) such that

$$v_k = \mathbb{E}_{\zeta_{v_k}} \left(\frac{r}{r+\lambda} \pi(\phi_{1;v_k}, \phi_{2;v_k}) + \frac{\lambda}{r+\lambda} w_{v_k}(\omega) \right) \quad (3)$$

and, for all ω and $i = 1, 2$,

$$\frac{r}{r+\lambda} (b(\phi_{i;v_k}) - c(\phi_{-i;v_k})) + \frac{\lambda}{r+\lambda} w_{i;v_k}(\omega) \geq \max_{\phi_i \in [0, 2\pi(\phi^B)]} \frac{r}{r+\lambda} (b(\phi_i) - c(\phi_{-i;v_k}(\omega))) \quad (4)$$

Now although $\{\zeta_{v_k}, \phi_{v_k}, w_{v_k}\}$ need not converge, since all effort profiles used in equilibrium are bounded, they lie in the product of compact spaces with natural metrics defined for each space.¹⁴ Since a finite product of compact metric spaces is compact and metrizable, the sequence lies in a sequential compact space (see [Munkres 2000](#), Theorems 26.7 and 28.2 and p. 219). Therefore there exists a subsequence $\{\zeta_{v_j}, \phi_{v_j}, w_{v_j}\}$, where $\{j\} \subset \{k\}$, that converges to a limit point (ζ_v, ϕ_v, w_v) that satisfies [Eqs. \(3\) and \(4\)](#). For each $v \in W$, consider the promised utility function that yields $(0, z) \in W$ if player 1 deviates, $(z, 0) \in W$ if player 2 deviates, and $w_v(\omega) \in W$ otherwise. Then we construct an equilibrium using an automaton strategy profile with states W , randomization rules $\{\zeta_v\}_{v \in W}$, effort rules $\{\phi_v\}_{v \in W}$, and transition rules given by the promised utility functions just described. Because W is internally weak-Pareto incomparable, this equilibrium is WBRP.

A.2 Proof of [Proposition 2](#) on p. 12

Instead of proving results for a complete graph with n players, we establish these results for any “regular network” in which each player has d partners. Doing so facilitates comparison with the closest antecedent, [Jackson, Rodriguez-Barraquer, and Tan \(2012\)](#). Of course, a special case of these results is for the complete graph with $n = d + 1$ players; in that case, in what follows let $\phi_n^{\text{PM}} = \tilde{\phi}_{d+1}^{\text{PM}}$.

Lemma 5. *For every equilibrium, the sum of payoffs is bounded above by $\tilde{\phi}_d^{\text{PM}}$, the solution to*

$$d \frac{\lambda}{r} \pi(\phi) = c(\phi). \quad (5)$$

Proof. We denote the (undirected) network by \mathcal{G} , and say that $ij \in \mathcal{G}$ if player i and j are linked in network \mathcal{G} . Generically, we denote by ϕ_{ij} the effort exerted by player i with respect to her neighbor player j , and with a slight abuse of notation, we denote $(\phi_{ij})_{ij \in \mathcal{G}}$ as the vector of stakes across all relationships.

We follow the proof approach of [Lemma 1](#), extending it to n players and for regular networks of degree d . For an equilibrium, let μ be a distribution of equilibrium path stakes and let Γ_{ij}^μ be the support of effort level ϕ_{ij} for the marginal distribution of ϕ_{ij} induced by μ . A necessary condition for an equilibrium is that for every ϕ'_i in Γ_{ij}^μ ,

$$c(\phi'_{ij}) \leq \frac{\lambda}{r} \sum_{k \in N(i)} (\mathbb{E}_\mu [b(\phi_{ki}) - c(\phi_{ik})]), \quad (\text{Player } i\text{'s IC})$$

where $N(i)$ is the set of neighbors of player i . Let $ij \in \mathcal{N}$. Adding these ICs across all relationships generates the inequality

$$\sum_{ij \in \mathcal{G}} (c(\phi'_{ij}) + c(\phi'_{ji})) \leq \frac{\lambda}{r} \sum_{ij \in \mathcal{G}} (\mathbb{E}_\mu [\pi(\phi_{ij}) + \pi(\phi_{ji})]). \quad (\text{Aggregate IC with } n \text{ players})$$

¹⁴ $\{\zeta_{v_k}\}$ is compact and metrizable by Theorem 14.11 of [Aliprantis and Border \(1999\)](#); each $\{\phi_{i,v_k}\}$ lies in the set of probability measures over $[0, 2\phi^B]$; $\{w_{v_k}\}$ may also be identified by 2-dimensional vectors in $[0, 2\phi^B]^2$. Each sequence therefore lies in a compact metrizable space.

Instead of solving for the utilitarian optimal equilibrium, let us solve the relaxed problem

$$\max_{\mu} \sum_{ij \in \mathcal{G}} \mathbb{E}_{\mu} [\pi(\phi_{ij}) + \pi(\phi_{ji})] \text{ subject to Aggregate IC with } n \text{ players.}$$

A solution to the above problem must involve a degenerate distribution (because any non-degenerate distribution that is in the constraint set can be improved by putting all of its mass on its highest realizations). Therefore, the relaxed problem can be re-written as

$$\max_{(\phi_{ij})_{ij \in \mathcal{G}}} \sum_{ij \in \mathcal{G}} [\pi(\phi_{ij}) + \pi(\phi_{ji})] \text{ subject to } \sum_{ij \in \mathcal{G}} (c(\phi_{ij}) + c(\phi_{ji})) \leq \frac{\lambda}{r} \sum_{ij \in \mathcal{G}} (\pi(\phi_{ij}) + \pi(\phi_{ji})).$$

We argue that any solution to the above program must be symmetric: $\exists \phi$ such that $\phi_{ij} = \phi$ for every $ij \in \mathcal{G}$. Consider a vector $(\phi_{ij})_{ij \in \mathcal{G}}$ such that $\phi_{ij} \neq \phi_{kl}$ for some ij and kl in \mathcal{G} . Consider a vector of stakes, $\tilde{\phi}$ where each coordinate is the arithmetic average of $(\phi_{ij})_{ij \in \mathcal{G}}$: because π is weakly concave and c is strictly convex, it follows from Jensen's Inequality that $\tilde{\phi}$ improves the objective while strictly relaxing constraints. Therefore, any asymmetric vector of stakes that is in the constraint set can be strictly improved (from the perspective of the objective) without leaving the constraint set. The maximal symmetric solution corresponds to ϕ_d^{PM} . \square

Lemma 6. *There exists a WRP equilibrium that supports $\tilde{\phi}_d^{\text{PM}}$ on the equilibrium path.*

Proof. Consider a strategy profile σ^* that follows the automaton:

1. Cooperation: Exert effort $\tilde{\phi}_d^{\text{PM}}$.
2. Punish i : Player $j \neq i$ chooses effort $\tilde{\phi}_d^{\text{PM}}$ when interacting with player $k \neq i$ and chooses effort 0 when interacting with player i . Player i chooses $\tilde{\phi}_d^{\text{PM}}$ when interacting with all players.

In any phase, if more than one player deviates simultaneously, then remain in the current phase; if player i deviates unilaterally then transition to the Punish i phase. In the Cooperation phase, if no player deviates, then remain in that phase. In the Punish i phase, if no player deviates, transition to the Cooperation phase.

We verify that σ^* is an equilibrium by using the one-shot deviation principle. Player i 's expected payoffs, in flow terms, in each phase are as follows

1. Cooperation Phase: $d\pi(\tilde{\phi}_d^{\text{PM}})$
2. Punish i Phase: A multiple of $-c(\tilde{\phi}_d^{\text{PM}}) + d\frac{\lambda}{r}\pi(\tilde{\phi}_d^{\text{PM}}) = 0$.
3. Punish j Phase (where $j \neq i$): $d\pi(\tilde{\phi}_d^{\text{PM}}) + \frac{r}{r+d\lambda}b(\tilde{\phi}_d^{\text{PM}})$.

By construction, no player has an incentive to deviate in the Cooperation Phase. In the Punish i phase, player i does not gain from deviating because she receives her minimax payoff. In the Punish j phase, player i obtains lower short-term and long-term payoffs from deviating. \square

Lemma 7. *There exists an RP equilibrium that supports $\tilde{\phi}_d^{\text{PM}}$.*

Proof. The argument follows that of Lemma 4, establishing an analogue to Lemma 3, and hence is omitted. \square

A.3 Proof of Proposition 3 on p. 14

We begin by establishing some properties of X_n .

Lemma 8. $X_n < \frac{\lambda}{r+\lambda}$ for all $n \geq 2$, and $\lim_{n \rightarrow \infty} X_n = 0$.

Proof. These follow from Lemma 2 of [Ali and Miller \(2013\)](#), which states that

$$X_n = \frac{1}{n-2} \sum_{m'=2}^{n-1} \left(\frac{1}{m'} \prod_{m=2}^{m'} \frac{\lambda m(n-m)}{r + \lambda m(n-m)} \right). \quad (6)$$

□

To prove [Proposition 3](#), we construct a WBRP contagion equilibrium σ^{**} with the desired properties, and show that it implies the existence of a BRP equilibrium with similar properties. Formally, each player i 's strategy σ_i^{**} is defined by a collection of “partial automata,” one for each of his partnerships. A partial automaton for player i in partnership ij determines player i 's behavior when meeting partner j ; because player i 's various partnerships are strategically interdependent, transitions between phases in partnership ij may be driven by interactions that occur in other partnerships. Let each player i 's behavior and beliefs in partnership ij be governed by the following five-phase partial automaton, illustrated in [Fig. 2](#).

1. Global Cooperation: Exert effort ϕ_n^R . Believe that Player j is in the global cooperation phase.
2. Contagion: Exert effort 0. Believe that Player j is in the global cooperation phase with some probability, and in the contagion phase with complementary probability.
3. Bilateral cooperation: Exert effort ϕ^B . Believe that Player j is in the bilateral cooperation phase.
4. Punishment: Exert effort ϕ^B . Believe that Player j is in the reward phase.
5. Reward: Exert effort 0. Believe that Player j is in the punishment phase.

The bilateral cooperation, punishment, and reward phases in σ^{**} are identical to the corresponding phases described in the proof of [Proposition 1](#). Transitions from the other two phases are a bit more complicated:

- In the global cooperation phase, if both partners exert effort ϕ_n^R , remain in the global cooperation phase. If neither partner exerts effort ϕ_n^R , transition to the bilateral cooperation phase. If player i exerts effort ϕ_n^R but partner j does not, transition to the reward phase; if partner j exerts effort ϕ_n^R but player i does not, transition to the punishment phase. If any player $k \in N \setminus \{i, j\}$ exerts effort other than ϕ_n^R in partnership ik , transition to the contagion phase in partnership ij .
- In the contagion phase, if neither partner exerts effort ϕ_n^R , transition to the bilateral cooperation phase. If player i exerts effort ϕ_n^R but partner j does not, transition to the reward phase; if partner j exerts effort ϕ_n^R but player i does not, transition to the punishment phase. If both partners exert effort ϕ_n^R , transition to the contagion phase.

The global cooperation phase is the initial phase in every partnership.

Proof of Proposition 3. First we confirm that this strategy profile σ^{**} constitutes an equilibrium. By [Proposition 1](#) incentive constraints are satisfied in the bilateral cooperation, punishment, and reward phases. So it suffices to check only the incentive constraints in the global cooperation and contagion phases.

We show that the equilibrium-path incentive constraint [Eq. \(IC_{ij}^{RNP}\)](#), is slack at $\phi = \phi^B$. Since ϕ^B is set to make the bilateral Nash reversion incentive constraint bind, $b(\phi^B) = \frac{r+\lambda}{r} \pi(\phi^B)$. So [Eq. \(IC_{ij}^{RNP}\)](#) is slack at ϕ^B

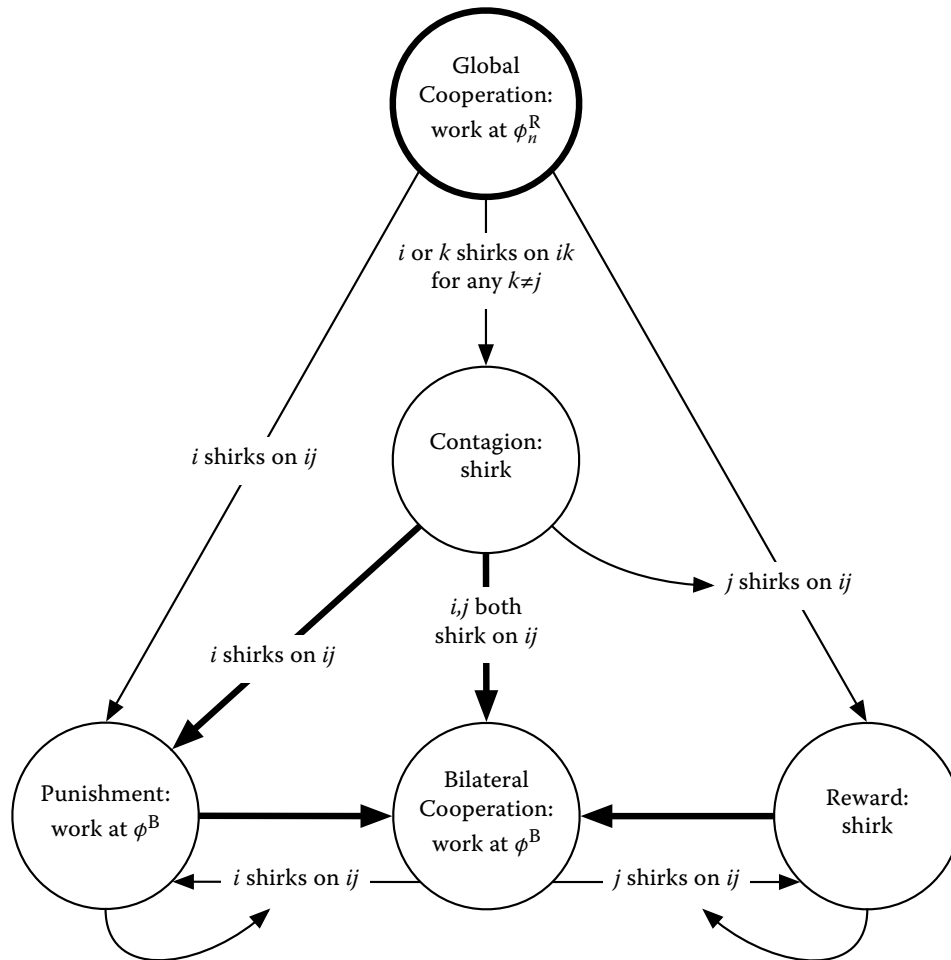


FIGURE 2. Bilateral renegotiation proof contagion partial automaton for player i in partnership ij . The initial phase is outlined in bold, and transitions driven by player i 's equilibrium behavior are shown with bold arrows. Not all transitions are shown.

if and only if

$$\begin{aligned} & \left(\frac{r+\lambda}{r} + (n-2) \left(\frac{r+\lambda}{r} X_n + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r} \right) \right) \pi(\phi^B) < \left(1 + (n-1) \frac{\lambda}{r} \right) \pi(\phi^B) \\ \iff & (n-2) \left(X_n + \frac{\lambda}{r+\lambda} \frac{\lambda}{r} \right) < (n-2) \frac{\lambda}{r} \iff X_n < \frac{\lambda}{r+\lambda}, \end{aligned} \quad (7)$$

which is established by Lemma 8. Because c is continuous and strictly convex (Assumption 2), there exists $\phi_n^R > \phi^B$ that binds Eq. (IC_{ij}^{RNP}), at which each player is indifferent between working at ϕ_n^R and shirking in the global cooperation phase.

To establish that σ^{**} is an equilibrium, it remains to verify the contagion phase incentives. First, player i 's contagion phase incentive constraint when meeting a player j he believes is contagious is

$$-c(\phi) + \int_0^\infty e^{-rt} e^{-\lambda t} \lambda \left(b(\phi^B) + \frac{\lambda}{r} \pi(\phi^B) \right) dt \leq \frac{\lambda}{r} \pi(\phi^B). \quad (8)$$

Since $\phi_n^R > \phi^B$, to show that Eq. (8) is satisfied at $\phi = \phi_n^R$, it suffices to show that Eq. (8) is slack at $\phi = \phi^B$:

$$-c(\phi^B) + \frac{\lambda}{\lambda+r} \left(b(\phi^B) + \frac{\lambda}{r} \pi(\phi^B) \right) < \frac{\lambda}{r} \pi(\phi^B). \quad (9)$$

Recall that $c(\phi^B) = \frac{\lambda}{r} \pi(\phi^B)$ by definition of ϕ^B . Therefore Eq. (8) is slack at $\phi = \phi^B$ if

$$\frac{\lambda}{r} > -\frac{r}{\lambda+r} \frac{\lambda}{r} + \frac{\lambda}{\lambda+r} \left(1 + \frac{\lambda}{r} \right), \quad (10)$$

which is guaranteed.

Next, we address a contagious player i 's incentive constraint when meeting a partner he believes is not contagious. This follows from an identical argument to that of Lemma 5 of Ali and Miller (2013), which proves that player i 's contagion-phase incentive to shirk on a non-contagious partner is strictly larger when he knows contagion has already started than on the equilibrium path. Since he is indifferent on the equilibrium path (Eq. (IC_{ij}^{RNP}) binds), it follows that he strictly prefers to shirk on a non-contagious partner when he is contagious. Intuitively, when he knows more players are already contagious, he has less at stake, and therefore it would be harder to convince him to work, i.e., it is easier to convince him to shirk.

Now we have shown that contagion phase incentives of σ^{**} are satisfied for extremal beliefs. Since expected payoffs are linear in player i 's belief about his current partner, the contagion phase incentive constraints are satisfied for all intermediate beliefs. Given this fact, there is no need to verify plain consistency in the contagion phase. Plain consistency is guaranteed in the bilateral cooperation, reward, and punishment phases because behavior on link ij is independent of behavior on all other links. Evidently σ^{**} is a plain perfect Bayesian equilibrium.

Finally, we show that if σ^{**} is not BRP, then there must exist some BRP equilibrium that attains the same equilibrium path payoffs. Observe that two partners cannot bilaterally renegotiate when one or both of them is in the contagion phase—they cannot arrive in the contagion phase by any transition that is common knowledge between them, and they stay in the contagion phase only for one interaction.¹⁵ Moreover,

¹⁵They could stay in the contagion phase for more than one interaction only if the contagious player deviates (or both contagious players deviate, as the case may be) by working at ϕ_n^R . However, such a deviation does not generate common knowledge of the contagion phase, because it is indistinguishable from the behavior specified for the global cooperation

according to σ^{**} , they bilaterally renegotiate to the bilateral equilibrium σ^* described in [Proposition 1](#) as soon as it becomes common knowledge between them that they are off the equilibrium path. Though σ^* may not be BRP, [Proposition 1](#) establishes that there exists some BRP equilibrium in bilateral strategies that supports cooperation at efforts ϕ^B ; to do so, it must employ minimax punishments. In σ^{**} , the first player to depart from global cooperation behavior along any link is punished with the minimax continuation payoff; the only possible way to relax [Eq. \(IC_{ij}^{RNP}\)](#) would be to punish both partners when they simultaneously depart from global cooperation behavior. However, the reason that each player departed from global cooperation behavior is not common knowledge information and therefore cannot be used to assign blame. In order to be measurable with respect to their common knowledge, their continuation payoff must be symmetric. Due to WBRP, they therefore cannot receive continuation payoffs less than $\frac{\lambda}{r}\pi(\phi^B)$. That is, no BRP equilibrium can impose stronger incentives in the global cooperation phase. Since σ^{**} maximizes equilibrium path payoffs subject to these constraints, there must exist an BRP equilibrium that attains the same payoffs.

Optimality of this equilibrium among all symmetric WBRP equilibria is established by [Lemma 9](#), below. \square

Lemma 9. *No symmetric BRP equilibrium supports cooperation at effort greater than ϕ_n^R .*

Proof. Suppose that ϕ is the effort each player always exerts on the equilibrium path of a symmetric BRP equilibrium. On the equilibrium path, suppose that player i shirks on partner j , and that thereafter player i 's punishment in partnership ij is to receive an average payoff of P . In our BRP contagion equilibrium σ^{**} , $P = 0$; more generally $P \geq 0$. Now consider what happens when player i subsequently meets partner $k \neq j$ for the first time after his initial deviation. If player k still believes that nobody has deviated or is otherwise supposed to cooperate, then in partnership ik player i earns at least $b(\phi) + \frac{\lambda}{r}P$, by shirking immediately and earning P thereafter. If instead player k is off the equilibrium path, then at worst for player i they may both shirk and then continue with each earning a continuation flow payoff of $\pi(\phi^B)$, which we will prove below.

Accordingly, the incentive constraint to cooperate on the equilibrium path is

$$b(\phi) + \frac{\lambda}{r}P + (n-2) \left(X_n \left(b(\phi) + \frac{\lambda}{r}P \right) + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r} \pi(\phi^B) \right) \leq (n-1) \frac{\lambda}{r} \pi(\phi) + \pi(\phi). \quad (11)$$

Notice if $P = 0$ (as in σ^{**}) then [Eq. \(11\)](#) reduces to

$$b(\phi) + (n-2) \left(X_n b(\phi) + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r} \pi(\phi^B) \right) \leq (n-1) \frac{\lambda}{r} \pi(\phi) + \pi(\phi), \quad (12)$$

in which case cooperating at ϕ_n^R is supported on the equilibrium path. If instead $P > 0$, then

$$\begin{aligned} & b(\phi_n^R) + \frac{\lambda}{r}P + (n-2) \left(X_n \left(b(\phi_n^R) + \frac{\lambda}{r}P \right) + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r} \pi(\phi^B) \right) \\ & > b(\phi_n^R) + (n-2) \left(X_n b(\phi_n^R) + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r} \pi(\phi^B) \right) = (n-1) \frac{\lambda}{r} \pi(\phi_n^R) + \pi(\phi_n^R). \end{aligned} \quad (13)$$

Thus the equilibrium effort supported in a symmetric equilibrium in which $P > 0$ must be less than ϕ_n^R .

We still need to show that if player i and k first meet when both are already off the equilibrium path, then letting them first shirk on each other and then renegotiate to the bilateral cooperation phase at efforts ϕ^B is the

phase.

harshest possible punishment that can be imposed on whichever (if either) of them was the original deviator. BRP requires that their renegotiated play must be measurable with respect to their partnership history, which unfortunately does not identify the original deviator. We already showed that no BRP equilibrium payoff is Pareto dominated by the average payoff vector $(\pi(\phi^B), \pi(\phi^B))$ (Proposition 1). Now suppose player i and k negotiate to a bilateral equilibrium strategy with average payoff vector $(U_i, U_k) \neq (\pi(\phi^B), \pi(\phi^B))$. BRP implies that (U_i, U_k) must also be on the Pareto frontier of bilateral WBRP equilibria, so without loss of generality assume that $U_k < \pi(\phi^B) < U_i$. Since the equilibrium must be symmetric, on first deviating player i must expect the same payoff in each of his $(n-2)$ other relationships. Therefore a necessary condition for his incentive constraint on the equilibrium path is

$$b(\phi) + \frac{\lambda}{r}P + (n-2) \left(X_n \left(b(\phi) + \frac{\lambda}{r}P \right) + \left(\frac{\lambda}{r+\lambda} - X_n \right) \frac{\lambda}{r} U_i \right) \leq (n-1) \frac{\lambda}{r} \pi(\phi) + \pi(\phi). \quad (14)$$

Since the LHS of Eq. (14) is greater than LHS of Eq. (11), ϕ_n^R will not satisfy Eq. (14). It follows that cooperating can be supported on the equilibrium path at efforts no greater than ϕ_n^R . \square

A.4 Proof of Proposition 4 on p. 15

We begin with two straightforward claims.

Claim 1. For all $n > 0$, $(n-1)(X_n - \frac{\lambda}{r}) - X_n < 0$.

This is a straightforward consequence of Lemma 8.

Claim 2. $\lim_{n \rightarrow \infty} \frac{(n-1)\frac{\lambda}{r}+1}{1+(n-2)X_n} = \infty$.

This follows from Lemma 8 since $X_n \geq 0$ for all n . Note $\lim_{n \rightarrow \infty} \frac{(n-1)\frac{\lambda}{r}+1}{1+(n-2)X_n} = \infty$ is a necessary and sufficient condition for $\lim_{n \rightarrow \infty} \phi_n^C = \infty$, by Eq. (IC^{NoRNP}).

Manipulating Eq. (IC^{NoRNP}) when it binds reveals that $\frac{b(\phi_n^C)}{\pi(\phi_n^C)} = \frac{(n-1)\frac{\lambda}{r}+1}{1+(n-2)X_n}$, so $\lim_{n \rightarrow \infty} \frac{b(\phi_n^C)}{\pi(\phi_n^C)} = \infty$. Notice $\frac{b(\phi)}{\pi(\phi)}$ is monotone (so its inverse function exists), given Assumption 2. Together with Assumption 1, we have

$$\lim_{\phi \rightarrow \infty} \frac{c(\phi)}{\pi(\phi)} = \infty$$

since both $c(\phi)$ and $\pi(\phi)$ are positive and strictly increasing for $\phi \in (0, \infty)$, $c(\phi)$ is strictly convex, and $\pi(\phi)$ is weakly concave. Thus $\frac{b(\phi)}{\pi(\phi)} = \frac{c(\phi)}{\pi(\phi)} + 1 \rightarrow \infty$ as $\phi \rightarrow \infty$. Thus the inverse function of $\frac{b(\phi)}{\pi(\phi)}$, $(\frac{b}{\pi})^{-1}(t) \rightarrow \infty$ as $t \rightarrow \infty$. It follows that $\lim_{n \rightarrow \infty} \phi_n^C = \infty$.

Letting $A = \frac{\lambda}{r} \pi(\phi^B)$ and $B = \frac{\lambda}{r+\lambda} \frac{\lambda}{r} \pi(\phi^B)$, we find that, by Eqs. (IC^{NoRNP}) and (IC^{RNP}), ϕ_n^R and ϕ_n^C are the largest solutions to the following equations, respectively:

$$b(\phi_n^R) + (n-2)X_n b(\phi_n^R) - \left((n-1) \frac{\lambda}{r} + 1 \right) \pi(\phi_n^R) = (n-2)X_n A - (n-2)B, \quad (15)$$

$$b(\phi_n^C) + (n-2)X_n b(\phi_n^C) - \left((n-1) \frac{\lambda}{r} + 1 \right) \pi(\phi_n^C) = 0. \quad (16)$$

Note that

$$(n-2)X_n A - (n-2)B < 0, \quad (17)$$

where the last inequality is equivalent to $X_n < \frac{\lambda}{r+\lambda}$, implied by Lemma 2 of Ali and Miller (2013).

Define $F_n(\phi) = b(\phi) + (n-2)X_n b(\phi) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi(\phi)$; i.e., the left-hand side of Eqs. (15) and (16). Note that F_n inherits convexity from b and $-\pi$, and that $F_n(\phi_n^R) = (n-2)X_n A - (n-2)B < F_n(\phi^C) = 0$. Now $F_n'(\phi) = b'(\phi) + (n-2)X_n b'(\phi) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi)$ and

$$F_n'(0) = b'(0) + (n-2)X_n b'(0) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi'(0) = \left(d\left(X_n - \frac{\lambda}{r}\right) - X_n\right)b'(0) < 0, \quad (18)$$

where the second equality comes from Assumption 2 and the inequality comes from Claim 1. By the convexity of F_n , it is easy to see that for any fixed $n > 0$, there exists a $\phi_{\text{large}} > 0$ such that $F_n'(\phi_{\text{large}}) > 0$. Also, since $F_n'' > 0$, there exists one unique critical point ϕ_n^{\min} that solves $F_n'(\phi) = 0$, i.e.,

$$b'(\phi_n^{\min}) + (n-2)X_n b'(\phi_n^{\min}) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi_n^{\min}) = 0. \quad (19)$$

Then

$$b'(\phi_n^{\min}) = \frac{\left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi_n^{\min})}{1 + (n-2)X_n} > \frac{\left((n-1)\frac{\lambda}{r} + 1\right)\psi}{1 + (n-2)X_n}, \quad (20)$$

where $\psi > 0$ is the lower bound on π' from Assumption 1. Note that $F_n'(\phi) < 0$ for all $\phi < \phi_n^{\min}$, and $F_n'(\phi) > 0$ for all $\phi > \phi_n^{\min}$.

From Claim 2, the RHS of Eq. (20) goes to infinity as $n \rightarrow \infty$; by monotonicity of $b'(\phi)$, $\lim_{n \rightarrow \infty} \phi_n^{\min} = \infty$. It is easy to check that there exist $\phi_{\text{pos}} > \phi_{\text{neg}} > 0$ such that $F_n(\phi_{\text{neg}}) < 0 < F_n(\phi_{\text{pos}})$; from continuity, there exists ϕ_n^C which solves Eq. (16).

From Proposition 3, Eq. (17), and Eq. (IC^{NoRNP}), $F_n(\phi_n^{\min}) < F_n(\phi_n^{GR}) < F_n(\phi_n^R) < F_n(\phi_n^C)$. Since $F_n'(\phi) > 0$ for all $\phi > \phi_n^{\min}$, by monotonicity it follows that $\phi_n^{\min} < \phi_n^R < \phi_n^C$. Since $\lim_{n \rightarrow \infty} \phi_n^{\min} = \infty$, $\lim_{n \rightarrow \infty} \phi_n^R = \infty$.

Next, we consider linear approximation of F_n near ϕ_n^R and ϕ_n^C : due to the convexity of F_n ,

$$\frac{(n-2)B - (n-2)X_n A}{F_n'(\phi_n^C)} < \phi_n^C - \phi_n^R < \frac{(n-2)B - (n-2)X_n A}{F_n'(\phi_n^R)}. \quad (21)$$

From Eq. (17), $\frac{(n-2)B - (n-2)X_n A}{F_n'(\phi_n^C)} > 0$. Dividing ϕ_n^R and adding 1 on all sides of the inequalities, we get:

$$\frac{(n-2)B - (n-2)X_n A}{\phi_n^R F_n'(\phi_n^C)} + 1 < \frac{\phi_n^C}{\phi_n^R} < \frac{(n-2)B - (n-2)X_n A}{\phi_n^R F_n'(\phi_n^R)} + 1. \quad (22)$$

Taking limits, we get:

$$1 \leq \lim_{n \rightarrow \infty} \frac{\phi_n^C}{\phi_n^R} \leq \lim_{n \rightarrow \infty} \frac{(n-2)B - (n-2)X_n A}{\phi_n^R F_n'(\phi_n^R)} + 1. \quad (23)$$

Since

$$\frac{(n-2)B - (n-2)X_n A}{\phi_n^R F_n'(\phi_n^R)} + 1 < \frac{(n-2)B}{\phi_n^R F_n'(\phi_n^R)} + 1, \quad (24)$$

given Eq. (23) it suffices to show that

$$\lim_{n \rightarrow \infty} \frac{(n-2)B}{\phi_n^R F_n'(\phi_n^R)} = \lim_{n \rightarrow \infty} \frac{\frac{n-2}{n-1}B}{\frac{1}{n-1} \phi_n^R F_n'(\phi_n^R)} = 0 \iff \lim_{n \rightarrow \infty} \frac{1}{n} \phi_n^R F_n'(\phi_n^R) = \infty. \quad (25)$$

Recall $F_n'(\phi) \leq F_n'(\phi_n^R)$ for all $\phi \leq \phi_n^R$. Hence

$$F_n(\phi_n^R) - F_n(\phi_n^{\min}) = \int_{\phi_n^{\min}}^{\phi_n^R} F_n'(\phi) d\phi \leq \int_{\phi_n^{\min}}^{\phi_n^R} F_n'(\phi_n^R) d\phi = F_n'(\phi_n^R)(\phi_n^R - \phi_n^{\min}), \quad (26)$$

where the first equality is the Fundamental Theorem of Calculus, and the inequality is from the monotonicity of F_n' . Transposing $F_n'(\phi_n^R)\phi_n^{\min}$ to the opposite side in Eq. (26), we get

$$\begin{aligned} F_n'(\phi_n^R)\phi_n^R &\geq F_n'(\phi_n^R)\phi_n^{\min} + F_n(\phi_n^R) - F_n(\phi_n^{\min}) \\ &= (1 + (n-2)X_n)b'(\phi_n^R)\phi_n^{\min} - \left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi_n^R)\phi_n^{\min} \\ &\quad + (n-2)X_nA - (n-2)B - \left((1 + (n-2)X_n)b(\phi_n^{\min}) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi(\phi_n^{\min})\right) \\ &= (n-2)X_nA - (n-2)B + (1 + (n-2)X_n)(b'(\phi_n^R)\phi_n^{\min} - b(\phi_n^{\min})) \\ &\quad + \left((n-1)\frac{\lambda}{r} + 1\right)(\pi(\phi_n^{\min}) - \pi'(\phi_n^R)\phi_n^{\min}) \\ &\geq (n-2)X_nA - (n-2)B + (1 + (n-2)X_n)(b'(\phi_n^R)\phi_n^{\min} - b(\phi_n^{\min})), \end{aligned} \quad (27)$$

where the first equality is derived by plugging in

$$F_n'(\phi_n^R) = (1 + (n-2)X_n)b'(\phi_n^R) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi) \quad (28)$$

$$F_n(\phi_n^R) = (n-2)X_nA - (n-2)B \quad (29)$$

$$F_n(\phi_n^{\min}) = (1 + (n-2)X_n)b(\phi_n^{\min}) - \left((n-1)\frac{\lambda}{r} + 1\right)\pi(\phi_n^{\min}); \quad (30)$$

and the last inequality is from the concavity of π :

$$\pi(\phi_n^{\min}) = \int_0^{\phi_n^{\min}} \pi'(\phi) d\phi \geq \int_0^{\phi_n^{\min}} \pi'(\phi_n^R) d\phi = \pi'(\phi_n^R)\phi_n^{\min}. \quad (31)$$

From Eq. (20),

$$b'(\phi_n^{\min}) = \frac{\left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi_n^{\min})}{1 + (n-2)X_n} = \int_0^{\phi_n^{\min}} b''(\phi) d\phi \leq \int_0^{\phi_n^{\min}} \frac{1}{\epsilon} d\phi = \frac{\phi_n^{\min}}{\epsilon}, \quad (32)$$

where the second equality is the Fundamental Theorem of Calculus and the inequality is from the condition $\frac{1}{\epsilon} > b''(\phi) > \epsilon$. Then [Assumption 1](#) implies that

$$\phi_n^{\min} \geq \frac{\left((n-1)\frac{\lambda}{r} + 1\right)\pi'(\phi_n^{\min})\epsilon}{1 + (n-2)X_n} \geq \frac{\left((n-1)\frac{\lambda}{r} + 1\right)\psi\epsilon}{1 + (n-2)X_n}. \quad (33)$$

Notice

$$\begin{aligned}
b'(\phi_n^R)\phi_n^{\min} - b(\phi_n^{\min}) &= \int_0^{\phi_n^{\min}} b'(\phi_n^R) d\phi - \int_0^{\phi_n^{\min}} b'(\phi) d\phi \\
&= \int_0^{\phi_n^{\min}} (b'(\phi_n^R) - b'(\phi)) d\phi \\
&= \int_0^{\phi_n^{\min}} \int_{\phi}^{\phi_n^R} b''(s) ds d\phi \\
&\geq \int_0^{\phi_n^{\min}} \int_{\phi}^{\phi_n^R} \epsilon ds d\phi \\
&= \int_0^{\phi_n^{\min}} \epsilon(\phi_n^R - \phi) d\phi \\
&= \epsilon \left(\phi_n^R \phi_n^{\min} - \frac{1}{2} (\phi_n^{\min})^2 \right) \\
&= \frac{\epsilon}{2} \phi_n^{\min} (2\phi_n^R - \phi_n^{\min}) \\
&\geq \frac{\epsilon}{2} \cdot \frac{((n-1)\frac{\lambda}{r} + 1)\psi\epsilon}{1 + (n-2)X_n} (2\phi_n^R - \phi_n^{\min}),
\end{aligned} \tag{34}$$

where the last inequality is derived from plugging Eq. (33). Using Eq. (34) in Eq. (27), we get

$$F'_n(\phi_n^R)\phi_n^R \geq (n-2)X_n A - (n-2)B + \frac{\epsilon^2 \psi ((n-1)\frac{\lambda}{r} + 1)}{2} (2\phi_n^R - \phi_n^{\min}).$$

Thus, since $\lim_{n \rightarrow \infty} X_n = 0$ (see Lemma 8),

$$\begin{aligned}
\lim_{d \rightarrow \infty} \frac{F'_n(\phi_n^R)\phi_n^R}{n} &\geq \lim_{n \rightarrow \infty} \frac{(n-2)X_n A - (n-2)B}{n} + \frac{\epsilon^2 \psi ((n-1)\frac{\lambda}{r} + 1)}{2n} (2\phi_n^R - \phi_n^{\min}) \\
&= -B + \frac{\epsilon^2 \psi \frac{\lambda}{r}}{2} \lim_{n \rightarrow \infty} (\phi_n^R + \phi_n^R - \phi_n^{\min}) \\
&\geq -B + \frac{\epsilon^2 \psi \frac{\lambda}{r}}{2} \lim_{n \rightarrow \infty} (\phi_n^R).
\end{aligned} \tag{35}$$

Since $\lim_{n \rightarrow \infty} (\phi_n^R) = \infty$, the RHS of Eq. (35) goes to ∞ , so

$$\lim_{d \rightarrow \infty} \frac{F'_n(\phi_n^R)\phi_n^R}{n-1} = \infty. \tag{36}$$

Now Eq. (25) is proved, and plugging it into Eq. (23) gives

$$1 \leq \lim_{n \rightarrow \infty} \frac{\phi_n^C}{\phi_n^R} \leq 1, \tag{37}$$

which concludes our proof.

A.5 Proof of Proposition 5 on p. 17

We begin with a preliminary step for the case of a single agent.

Lemma 10. *Suppose that there is a single agent (i.e. $n = 1$). Then, there exists a BRP equilibrium $\hat{\sigma}$ in which, along the equilibrium path, the agent exerts effort $\phi^B > 0$ that solves*

$$c(\phi) = \left(\frac{\lambda}{r + \lambda}\right)^2 b(\phi), \quad (38)$$

and the principal pays a wage

$$w^B = \frac{r + \lambda}{\lambda} c(\phi^B). \quad (39)$$

Off the equilibrium path, if either party deviates then his or her continuation payoff is zero. No equilibrium sustains higher equilibrium path effort.

Proof. We construct a WBRP equilibrium $\hat{\sigma}$ with the above properties in which a player who deviates is punished for a single period and then cooperation resumes. Specifically, the principal and the agent both play according to the following automaton:

1. Bilateral Cooperation: Exert effort ϕ^B or pay wage w^B .
2. Bilateral Punishment: Exert effort ϕ^B or pay wage w^B .
3. Reward: Exert effort 0 or pay wage 0.

The bilateral cooperation is the initial phase for both players. In any phase, if one player deviates then that player transitions to the punishment phase while the other transitions to the reward phase. A player remains in the reward phase until the other takes the action correspond to the punishment phase. Once that occurs, players transition to the bilateral cooperation phase.

Observe that once the agent is in the punishment phase, the payoff he accrues starting from the next interaction with the principal is $-c(\theta^B) + \frac{\lambda}{r}(w^B - c(\theta^B))$, which given Eq. (39), is 0. Similarly, once the principal is in the punishment phase, the payoff she accrues starting from the next interaction with the agent is also 0, given Eq. (38). Therefore, we may summarize the agent's and principal's ICs in the equilibrium phase as follows:

$$w^B \leq \left(1 + \frac{\lambda}{r}\right)(w^B - c(\phi^B)); \quad (40)$$

$$b(\phi^B) \leq \left(1 + \frac{\lambda}{r}\right)(b(\phi^B) - w^B). \quad (41)$$

We note that Eqs. (40) and (41) ensure that Eqs. (40) and (41) hold with equality. Substituting and simplifying Eq. (40) yields that θ^B is a solution to

$$\frac{c(\phi)}{\pi(\phi)} = \frac{\lambda^2}{r^2 + 2\lambda r}; \quad (42)$$

Assumptions 1 and 2 guarantee that the above equation has a unique solution, and therefore, guarantee existence and uniqueness of ϕ^B and w^B .

Since $w^B > w^B - c(\phi^B) > 0$ and $b(\phi^B) > b(\phi^B) - w^B > 0$, no two points in the 0,1-joint payoff set of this equilibrium are Pareto comparable, so $\hat{\sigma}$ is WBRP. Moreover, since $\hat{\sigma}$ employs minmax punishments, no greater effort can be sustained in any equilibrium. Then, by Lemma 3, there exists a BRP equilibrium that attains the same effort and also employs minimax punishments. \square

We use [Lemma 10](#) to establish the case for n agents and a single Principal.

Proof of Proposition 5. Our proof comprises two steps. The first is deriving the maximal effort profile that can be supported in an equilibrium using only the equilibrium path incentive constraints. The second is constructing a BRP equilibrium that implements that effort profile.

Step 1: Maximal effort in an equilibrium. Recall that we restrict attention to equilibria that are stationary on the path of play. For expositional convenience, we restrict attention here to pure strategy equilibria.¹⁶ In any equilibrium in which on the equilibrium path, agent i exerts effort $\phi_i > 0$ and the principal pays wage $w_i > 0$, both [Agent's IC](#) (on p. 16) and [Principal's IC](#) (on p. 16) must hold. Therefore, optimizing equilibrium path efforts corresponds to the following program

$$\max_{\phi_1, \dots, \phi_n} \sum_{i=1}^n \phi_i \text{ subject to } \forall i, (1) w_i \geq \left(1 + \frac{r}{\lambda}\right) c(\phi_i),$$

(2) [Principal's IC](#),

where (1) is a simplified version of [Agent's IC](#). Because $\frac{\lambda}{r} > X_n$, it follows that (1) must bind for each i ; otherwise, ϕ_i could be increased slightly, which would only relax or keep unchanged other constraints, and increase the objective function. Setting (1) to bind, substituting into [Principal's IC](#), and substituting $\pi(\phi) = b(\phi) - c(\phi)$ yields

$$\max_{\phi_1, \dots, \phi_n} \sum_{i=1}^n \phi_i \text{ subject to } \forall i, \frac{\lambda}{r} \pi(\phi_i) + \left(\frac{\lambda}{r} - X_{n+1}\right) \sum_{j \neq i} \pi(\phi_j) \geq \left(1 + \frac{r}{\lambda}\right) c(\phi_i) + \sum_{j=1}^n c(\phi_j) + X_{n+1} \sum_{j \neq i} c(\phi_j), \quad (43)$$

Adding the d constraints leads to an aggregated constraint

$$\left[\frac{\lambda}{r} + (n-1) \left(\frac{\lambda}{r} - X_{n+1}\right)\right] \sum_{i=1}^n \pi(\phi_i) \geq \left[1 + \frac{r}{\lambda} + n + (n-1)X_{n+1}\right] \sum_{i=1}^n c(\phi_i). \quad (44)$$

We use the above constraints to argue that any solution to (43) involves symmetric efforts. Suppose $(\phi_1^*, \dots, \phi_n^*)$ is a solution to (43) and there exists i and j such that $\phi_i^* \neq \phi_j^*$. Because π is weakly concave ([Assumption 1](#)) and c is strictly convex ([Assumption 2](#)), setting each $\phi_i = \frac{1}{n} \sum_{j=1}^n \phi_j^*$ would not change the objective function, but would ensure that (44) holds with slack (because of Jensen's Inequality). But then each constraint in (43) must also hold with slack. In that case, each ϕ_i can be increased slightly, which strictly improves the objective function and contradicts $(\phi_1^*, \dots, \phi_n^*)$ being optimal.

Because a symmetric solution is optimal, substitute $\phi_i = \phi$ for every i into [Eq. \(44\)](#). Setting it to hold with equality yields

$$\frac{c(\phi)}{\pi(\phi)} = \frac{\frac{\lambda}{r} + (n-1) \left(\frac{\lambda}{r} - X_{n+1}\right)}{1 + \frac{r}{\lambda} + n + (n-1)X_{n+1}}.$$

[Assumption 2](#) ensures that this equation has a unique solution; we call that solution ϕ_n^R . Let $w_n^R = \frac{\lambda+r}{\lambda} c(\phi_n^R)$ be the wage associated with ϕ_n^R that makes [Agent's IC](#) hold with equality.

Step 2: A BRP Equilibrium that supports ϕ_n^R on the equilibrium path. We construct a strategy profile $\tilde{\sigma}$ in

¹⁶Extending the following argument to mixed strategy equilibria is straightforward.

which each agent plays according to the following partial automaton:

1. Global Cooperation: Exert effort ϕ_n^R .
2. Global Punishment: Exert effort ϕ_n^R .
3. Bilateral Cooperation: Exert effort ϕ^B .
4. Bilateral Punishment: Exert effort ϕ^B .
5. Reward: Exert effort 0.

Transitions from the bilateral cooperation, bilateral punishment, and reward phases are the same as in $\hat{\sigma}$ in [Lemma 10](#). In the global cooperation and punishment phases, if the agent learns from another agent that the principal has deviated in another partnership, then the agent transitions immediately to the reward phase, and remains in it until the principal pays wage w^B , after which the agent transitions to bilateral cooperation. If the agent deviates while the principal pays w_n^R defined below, then the agent transitions to the global punishment phase. Otherwise transition to the global cooperation phase.

The principal plays according to a different partial automaton in his relationship with each agent i :

1. Global Cooperation: Pay wage w_n^R .
2. Global Reward: Pay wage 0.
3. Contagion: Pay wage 0.
4. Bilateral Cooperation: Pay wage w^B .
5. Bilateral Punishment: Pay wage w^B .
6. Reward: Pay wage 0.

Transitions from the bilateral cooperation, bilateral punishment, and reward phases are the same as in $\hat{\sigma}$. Transitions from the other states are a bit more complicated:

- Global Cooperation: If the agent deviates while the principal pays w_n^R , then transition to the reward phase. If the principal pays agent i any wage other than w_n^R , then transition to the bilateral punishment phase. If the principal deviates in the partnership with any agent $j \neq i$ or if any agent $j \neq i$ deviates, then transition to the contagion phase. Otherwise remain in the global cooperation phase.
- Contagion phase: If the principal pays any wage other than w_n^R , then transition to the punishment phase. If the principal deviates to pay wage w_n^R and the agent exerts effort ϕ_d^R , then remain in the contagion phase. If the principal deviates to pay wage w_n^R and the agent exerts any effort other than ϕ_d^R , then transition to the reward phase.
- Global Reward phase: Remain in the global reward phase until agent i exerts effort ϕ_d^R , then transition to the global cooperation phase.

We now verify incentives. By construction of ϕ_n^R and w_n^R , the equilibrium path incentive constraints hold for the principal and each agent. In the Global Punishment phase, an agent's expected payoff, beginning with the next interaction, is 0, and no action offers the agent a strictly positive payoff, so the agent has no incentive to deviate. Of course, in the Global Reward phase, the principal is acting in ways that maximize both her myopic and long-term payoff so there is no incentive to deviate. Furthermore, each agent i always weakly prefers to communicate truthfully. Communicating truthfully with other agents has no effect on her payoff. A agent who deviates obtains a payoff of 0. Finally, we verify that once the principal has reneged on one agent, he strictly prefers to renege on other agents. This incentive constraint is identical to the contagion phase

incentive constraint studied by [Kandori \(1992\)](#), [Ellison \(1994\)](#), and [Ali and Miller \(2013\)](#), and follows from adapting the proofs of Lemmas 4 and 5 in [Ali and Miller \(2013\)](#).

Since this equilibrium maximizes equilibrium-path effort, by [Lemma 3](#) if $\bar{\sigma}$ is not BRP then there must exist a BRP equilibrium with the claimed properties. □

B Supplementary Appendix: Definition of Plain PBE

A “plain” *perfect Bayesian equilibrium* (PBE; [Watson 2016](#)) is a strategy profile that is sequentially rational for a *conforming appraisal system* that satisfies *plain consistency*. An appraisal system represents the players’ beliefs about strategy profiles conditional on their information sets. A *conforming appraisal profile* is a conditional probability system over the space of pure strategy profiles, where conditioning with respect to information sets in the extensive form correctly represents the (behavior) strategy profile. An appraisal system satisfies *plain consistency* if each player updates as a Bayesian when he moves from a predecessor information set to a successor information set, both on and off the path of play. In particular, his updating respects independence properties that were present in his beliefs at the start of the game and have not been contradicted by evidence.

We define PBE for the community enforcement game, following [Watson \(2016\)](#) but with some inconsequential modifications for clarity. Extension to the agents-and-principal game is straightforward. In what follows, we identify any regular private history h_i with its associated information set. In addition let each player have an additional information set h_i^0 that precedes all others, representing the beliefs that player has about the strategy profile at the start of the game; let $\bar{H}_i = H_i \cup \{h_i^0\}$. Let $\bar{H} = \bar{H}_1 \cup \dots \cup \bar{H}_n$ be the space of information sets in the game tree that are associated with regular private histories, let S be the space of pure strategy profiles, and for an information set $h \in \bar{H}$ let $S(h) \subseteq S$ be the set of pure strategy profiles for which the support of the path of play has non-empty intersection with h . Define $S(h_i^0) = S$ for all i .

Definition 5 ([Watson 2016](#), Definition 3). An *appraisal system* on S is a family of conditional probability measures $\{\theta(\cdot|E)\}_{E \subseteq S}$ on S such that

- For each $E \subseteq S$, $\theta(S \setminus E|E) = 0$;
- For any nonempty $A \subseteq B \subseteq C \subseteq S$, $\theta(A|C) = \theta(A|B)\theta(B|C)$.

An appraisal system generalizes probability distributions to allow players to update their beliefs when reaching information sets that were not in the support of the original distribution $\theta(\cdot|S)$. We are interested in an appraisal system that properly represents the strategy profile.

Definition 6 ([Watson 2016](#), Definition 8). An appraisal system θ on S *conforms* to a strategy profile σ if for each player i , for each information set $h_i \in H_i$, and each measurable set of actions $\Phi_i \subset \mathbb{R}_+$, $\theta(S(h_i, \Phi_i) | S(h_i)) = \sigma(\Phi_i|h_i)$, where $S(h_i, \Phi_i)$ is the set of pure strategy profiles that reach h_i and choose effort in Φ_i .

This definition implies that players have correct beliefs because, for each player, by definition $\theta(\cdot|S(h_i^0)) = \theta(\cdot|S)$. That is, all players start off with common, correct belief in equilibrium strategies.

Definition 7 ([Watson 2016](#), Definition 7). An appraisal system θ on S satisfies **plain consistency** if for each player i and for each measurable and non-trivial two-element partition $\{L, -L\}$ of the set of information sets, and for each pair of information sets $h_i \in \bar{H}_i$ and $h'_i \in H_i$ for player i such that h'_i is a successor of h_i , if there

exists a product set of strategy profiles $Z = Z_L \times Z_{-L} \subset S(h_i)$ such that $Z_L \subset S(h'_i)_L$, $Y \equiv S(h'_i) \cap Z = Z_L \times Y_{-L}$, and $\theta(\cdot | S(h_i))$ exhibits independence between Z_L and Z_{-L} with $\theta(Y | S(h'_i)) > 0$, then $\theta(\cdot | S(h'_i))$ exhibits independence between Z_L and Y_{-L} , and $\theta(X_L \times Y_{-L} | Y) = \theta(X_L \times Z_{-L} | Z)$ for every $X_L \subset Z_L$.

Note that the conditions on Z mean that player i , conditioning on Z at both h_i and h'_i , does not learn anything about behavior within Z_L when moving from h_i to h'_i . Then player i 's belief about any $X_L \subset Z_L$ when conditioning on Y at h'_i must be the same as her belief about X_L when conditioning on Z at h_i .

Definition 8 (Watson 2016, Definition 9). A **plain perfect Bayesian equilibrium** is a behavior strategy profile σ such that there exists a conforming appraisal θ that satisfies plain consistency, and for each player i and each information set h_i , σ satisfies **sequential rationality** with respect to θ :

$$\sigma_i(h_i) \in \Delta \arg \max_{\phi_i \in \mathbb{R}_+} \int_{s \in S(h_i, \phi_i)} \mathcal{U}_i(s | h_i) d\theta(s | S(h_i, \phi_i)). \quad (45)$$

The sequential rationality condition, incorporating the one-shot deviation principle,¹⁷ simply states that player i should choose an action at h_i that maximizes his expected payoff conditional on h_i and that action, taking the strategy profile and appraisal system as given. Together, sequential rationality and plain consistency guarantee that a PBE satisfies subgame perfection (Watson 2016, Theorem 2), because for any subgame $G \subset H$ and any $h \in H$, $\theta(\cdot | h)$ exhibits independence between $S(h)_G$ and $S(h)_{-G}$. Therefore a player updates correctly about the strategy profile within G upon entering G . Sequential rationality then implies he must be playing a best response in G . Since this argument relies on independence of strategies between $S(h)_G$ and $S(h)_{-G}$ rather than on the extensive form, it applies without modification to any portion of the game tree that is treated by the strategy profile as if it were a subgame—such as the information sets representing meetings within partnership ij that succeed a partnership history $h|_{ij}$ for which $\sigma(h|_{ij})$ is bilateral on link ij . This property guarantees that when two partners renegotiate to bilateral play along their link, they must renegotiate to a continuation strategy profile that would be a subgame perfect equilibrium if their link were isolated.

¹⁷The one-shot deviation principle holds in this environment because players are exponential discounters and stage game payoffs are uniformly bounded in any equilibrium.