

Cooperation with Network Monitoring*

Alexander Wolitzky

Microsoft Research and Stanford University

February 2012

Abstract

This paper studies the maximum level of cooperation that can be sustained in perfect Bayesian equilibrium in repeated games with *network monitoring*, where players observe each other's actions either perfectly or not at all. The foundational result is that the maximum level of cooperation can be robustly sustained in grim trigger strategies. If players are equally well monitored, comparative statics on the maximum level of cooperation are highly tractable and depend on the monitoring technology only through a simple statistic, its *effective contagiousness*. Typically, cooperation in the provision of pure public goods is greater in larger groups, while cooperation in the provision of divisible public goods is greater in smaller groups, and making monitoring less uncertain in the second-order stochastic dominance sense increases cooperation. For fixed monitoring networks, a new notion of network *centrality* is developed, which determines which players cooperate more in a given network, as well as which networks support greater cooperation.

1 Introduction

How can groups sustain as much cooperation as possible? Should they rely exclusively on punishing individuals who are caught shirking, or should they also reward those who are “caught working?”

*This paper was previously circulated under the title, “Repeated Public Good Provision.” I thank the editor, Marco Ottaviani, and four anonymous referees for helpful comments. I thank my advisors, Daron Acemoglu, Glenn Ellison, and Muhamet Yildiz, for detailed comments and suggestions and for extensive advice and support; and thank Nageeb Ali, Abhijit Banerjee, Alessandro Bonatti, Gabriel Carroll, Matt Jackson, Anton Kolotilin, Parag Pathak, Larry Samuelson, Juuso Toikka, Iván Werning, and seminar participants at MIT and the 2010 Stony Brook International Conference on Game Theory for additional helpful comments. I thank the NSF for financial support.

Relatedly, what kinds of groups can sustain the most cooperation? Large ones or small ones? Ones where “who observes whom” in the group is known, or where it is uncertain?

These are fundamental questions in the social sciences (Olson, 1965; Ostrom, 1990; Coleman, 1990; Putnam, 1993; Greif, 2006). In economics, existing work on the theory of repeated games provides a framework for answering these questions when individuals can perfectly observe each other’s actions (e.g., Abreu, 1988), but provides much less explicit answers in the more realistic case where monitoring is imperfect. This weakness is particularly acute in settings where public signals are not very informative about each individual’s actions and high quality—but dispersed—private signals are the basis for cooperation. Consider, for example, the problem of maintaining a school in a small village in the developing world. Every year, say, different villagers must contribute different inputs to running the school: some provide money, some provide labor to maintain the building, some volunteer in other capacities, etc. These inputs are not publicly observable, and different villagers observe each other’s actions with different probabilities. The overall quality of the school is very hard to observe directly, and indeed one might not be able to infer much about it until one sees how well the students do years down the road, by which time the entire system of providing education in the village may have changed. This problem was studied theoretically and empirically (using data on schools and wells in rural Kenya) by Miguel and Gugerty (2005), under the assumption that each household’s contribution is publicly observable, but this assumption is often unrealistic; for example, Miguel and Gugerty emphasize the importance of ethnic divisions in the villages they study, so a natural assumption would be that a household is more likely to be monitored by households from the same ethnic group. A second example is the problem of cooperation in long-distance trade, argued by Greif and others to be an essential hurdle to the development of the modern economy. Here, a key issue is often how sharing information through network-like institutions like trading coalitions (Greif 1989, 1993), trade fairs (Milgrom, North, and Weingast, 1990), and merchant guilds (Greif, Milgrom, and Weingast, 1994) facilitates cooperation. Thus, it is certainly plausible that local, private monitoring plays a larger role than public monitoring in sustaining cooperation in many interesting economic examples, and very little is known about how cooperation is best sustained under this sort of monitoring.

This paper studies cooperation in repeated games with *network monitoring*, where in every period a network is independently drawn from a (possibly degenerate) known distribution, and players perfectly observe the actions of their neighbors but observe nothing about any other player’s action. The model covers both monitoring on a fixed network (as when a household’s actions are

always observed by its geographic neighbors, or by households in the same ethnic group), and random matching (as when traders randomly meet in a large market). Each player’s action is simply her level of cooperation, in that higher actions are privately costly but benefit others. The goal is to characterize the maximum level of cooperation that can be sustained *robustly* in equilibrium, in that it can be sustained for any information that players may have about who has monitored whom in the past. This robustness criterion captures the perspective of an outside observer, who knows what information players have about each other’s actions, but not what information players have about each other’s information about actions (or about their information about others’ information about actions, and so on), and who therefore must make predictions that are robust to this higher-order information.¹

A first observation is that for any given specification of players’ higher-order information, the strategies that sustain the maximum level of cooperation can depend on players’ private information in complicated ways that involve a mix of rewards and punishments, and determining the maximum level of cooperation appears intractable. In contrast, my main theoretical result is that the *robust* maximum level of cooperation is always sustained by simple grim trigger strategies, where each player cooperates at a fixed level unless she ever observes another player fail to cooperate at his prescribed level, in which case she stops cooperating forever. Thus, robust cooperation is maximized through strategies that involve punishments but not rewards. In addition, grim trigger strategies also maximize cooperation when players have perfect knowledge of who observed whom in the past (as is the case when the monitoring network is fixed over time, for example); interestingly, it is when players have *less* higher-order information that more complicated strategies can do better than grim trigger. A rough intuition for these results is that when players know who observed whom in the past there is a kind of “strategic complementarity” in which a player is willing to cooperate more at any on-path history whenever another player cooperates more at any on-path history, because—with network monitoring and grim trigger strategies—shirking makes every on-path history less likely; but this strategic complementarity breaks down when players can disagree about who has observed whom.

This result about *how* groups can best sustain cooperation has implications for *what* groups can sustain the most cooperation. For these more applied results, I focus on two important special cases of network monitoring: *equal monitoring*, where in expectation players are monitored “equally

¹There are of course other kinds of robustness one could be interested in, and strategies that are robust in one sense can be fragile in others. See the conclusion of the paper for discussion.

well”; and *fixed monitoring networks*, where the monitoring network is fixed over time.

With equal monitoring, I show that the effectiveness of a monitoring technology in supporting cooperation is completely determined by one simple statistic, its *effective contagiousness*, which is defined as

$$\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\textit{number of players who learn about a deviation within } t \textit{ periods}].$$

This result formalizes the simple idea that more cooperation can be sustained if news about a deviation spreads throughout the network more quickly. It implies that cooperation in the provision of pure public goods (where the marginal benefit of cooperation is independent of group size) is increasing in group size if the expected *number* of players who learn about a deviation is increasing in group size, while cooperation in the provision of divisible public goods (where the marginal benefit of cooperation is inversely proportional to group size) is increasing in group size if the expected *fraction* of players who learn about a deviation is increasing in group size. Hence, cooperation in the provision of pure public goods tends to be greater in larger groups, while cooperation in the provision of divisible public goods tends to be greater in smaller groups. In addition, making monitoring more “uncertain” in a certain sense reduces cooperation.

With fixed networks, I develop a new notion of network *centrality* that determines both which players cooperate more in a given network and which networks support more cooperation overall, thus linking the graph-theoretic property of centrality with the game-theoretic property of robust maximum cooperation. For example, adding links to the monitoring network necessarily increases all players’ robust maximum cooperation, which formalizes the idea that individuals in better-connected groups cooperate more.

The results of this paper may bear on questions in several fields of economics. First, a literature in public economics studies the effect of group size and structure on the maximum equilibrium level of public good provision. One strand of this literature studies repeated games, but characterizes maximum cooperation only with perfect monitoring. Papers in this strand have found few unambiguous relationships between group size and structure and maximum cooperation.² A second

²Pecorino (1999) shows that with perfect monitoring public good provision is easier in large groups, because shirking—and thus causing everyone else to start shirking—is more costly in large groups. Haag and Lagunoff (2007) show that with heterogeneous discount factors and a restriction to stationary strategies, maximum cooperation is increasing in group size. Bendor and Mookherjee (1987) consider imperfect public monitoring, and present numerical evidence suggesting that higher payoffs can be sustained in small groups when attention is restricted to trigger strategies. In a second paper, Bendor and Mookherjee (1990) allow for network structure but return to the

strand studies one-shot games of public good provision in networks (Ballester, Calvó-Armengol, and Zenou, 2006; Bramoullé and Kranton, 2007a; Bramoullé, Kranton, and D’Amours, 2011), where the network determines local payoff interactions—and, in particular, incentives for free-riding—rather than monitoring. These papers find that more central players (measured by Bonacich centrality or a modification thereof) cooperate less and receive higher payoffs, due to free-riding, and that adding links to a network decreases average maximum cooperation, by increasing free-riding. In contrast, my model, which combines elements from both strands of the literature, makes the following predictions, which are made precise later:

1. Cooperation in the provision of pure public goods is greater in larger groups, while cooperation in the provision of divisible public goods is greater in smaller groups.
2. Less uncertain monitoring increases cooperation.
3. More central players cooperate more (unlike in the public goods in networks literature) but still receive higher payoffs with local public goods (like in that literature).
4. Adding links to a monitoring network increases all players’ cooperation.

Second, several seminal papers in institutional economics study the role of different institutions in sharing information about past behavior to facilitate trade (Greif, 1989, 1993; Milgrom, North, and Weingast, 1990; Greif, Milgrom, and Weingast, 1994). Ellison (1994) notes that the models underlying these studies resemble a prisoner’s dilemma, and shows that cooperation is sustainable in the prisoner’s dilemma with random matching for sufficiently patient players, which suggests that information-sharing institutions are not always necessary. The current paper contributes to this literature by determining the maximum level of cooperation in a prisoner’s dilemma-like game at any fixed discount factor for any network monitoring technology. Thus, it allows one to determine the exactly how much more cooperation can be sustained in the presence of a given information-sharing institution.

Third, a young and very active literature in development economics studies the impact of network structure on different kinds of cooperation, such as favor exchange (Karlan et al, 2009; Jackson, Rodriguez-Barraquer, and Tan, 2011) and risk-sharing (Ambrus, Möbius, and Szeidl, 2010; Bramoullé and Kranton, 2007b; Bloch, Genicot, and Ray, 2008). The predictions of this paper enumerated above can be suggestively compared to some early empirical results in this literature,

assumption of perfect monitoring, and find an ambiguous relationship between group size and maximum cooperation.

although clearly much empirical work remains to be done. For example, Karlan et al (2009) find that indirect network connections between individuals in Peruvian shantytowns support lending and borrowing, consistent with my finding that more central players cooperate more. More subtly, Jackson, Rodriguez-Barraquer, and Tan (2011) find that favor-exchange networks in rural India exhibit high *support*, the property that linked players share at least one common neighbor. While it seems natural that support (which is the key determinant of cooperation in Jackson, Rodriguez-Barraquer, and Tan’s model) should be correlated with robust maximum cooperation in my model, I leave studying the precise empirical relationship between the two concepts for future research.

A few final comments on related literature: It should be noted that the aforementioned paper of Ellison (1994), along with much of the related literature (e.g., Kandori, 1992; Deb, 2009; Takahashi, 2010) focuses on the case of sufficiently high discount factors and does not characterize efficient equilibria at fixed discount factors, unlike my paper. In addition, a key concern in these papers is ensuring that players do *not* cooperate off the equilibrium path. The issue is that grim trigger strategies may provide such strong incentives to cooperate on-path that players prefer to cooperate even after observing a deviation. Ellison resolves this problem by introducing a “relenting” version of grim trigger strategies tailored to make players indifferent between cooperating and shirking on-path, and then noting that cooperation is more appealing on-path than off-path (since off-path at least one opponent is already shirking). This issue does not arise in my analysis because, with continuous action spaces, players must be just indifferent between cooperating and shirking on-path in the most cooperative equilibrium, as otherwise they could be asked to cooperate slightly more. By essentially the same argument as in Ellison, this implies that players weakly prefer to shirk off-path. Hence, the key contribution of this paper is showing that grim trigger strategies provide the strongest possible incentives for (robust) cooperation on-path, *not* that they provide incentives for shirking off-path.³

The most closely related paper is contemporaneous and independent work by Ali and Miller (2011). Ali and Miller study a network game in which links between players are recognized according to a Poisson process. When a link is recognized, the linked players play a prisoner’s dilemma with variable stakes, and can also make transfers to each other. Like my model, Ali and Miller’s features smooth actions and payoffs, so that, with grim trigger strategies, binding on-path incentive constraints imply slack off-path incentive constraints. The most important difference

³Another difference is that it is important for the current paper that in each period the monitoring network is observed after actions are chosen, whereas this timing does not matter in most papers on community enforcement.

between Ali and Miller’s paper and mine is that they do not show that grim trigger strategies always maximize cooperation in their model. Ali and Miller also do not emphasize strategic complementarity or robustness to higher-order information. They do however discuss network formation and comparisons among networks, developing insights that are complementary to mine.

Finally, this paper is related more broadly to the study of repeated games with private monitoring. Most papers in this literature study much more general models than mine, and either prove folk theorems or study robustness to small deviations from public monitoring (Mailath and Morris, 2002, 2006; Sugaya and Takahashi, 2011).⁴ However, to my knowledge this is the first paper that characterizes even a single point on the Pareto frontier of the set of perfect Bayesian equilibrium payoffs in a repeated game with imperfect private monitoring where first-best payoffs are not attainable. I make no attempt to characterize the entire set of perfect Bayesian equilibria, or any large subset thereof. Instead, I use the strategic complementarity discussed above to derive an upper bound on each player’s maximum cooperation, and then show that this bound can be attained with grim trigger strategies. It would be interesting to see if similar indirect approaches, perhaps also based on strategic complementarity, can be useful in other classes of repeated games with private monitoring of applied interest.

The paper proceeds as follows: Section 2 describes the model. Section 3 presents the key result that maximum cooperation is robustly sustained in grim trigger strategies. Section 4 derives comparative statics in games with equal monitoring. Section 5 studies games with fixed monitoring networks. Section 6 concludes and discusses directions for future research. Major omitted proofs and examples are in the appendix, and minor ones are in the online appendix.

2 Model

There is a set $N = \{1, \dots, n\}$ of players. In every period $t \in \mathbb{N} = \{0, 1, \dots\}$, every player i simultaneously chooses an action (“level of cooperation,” “contribution”) $x_i \in \mathbb{R}_+$. The players have common discount factor $\delta \in (0, 1)$. If the players choose actions $x = (x_1, \dots, x_n)$ in period t , player i ’s period- t payoff is

$$u_i(x) = \left(\sum_{j \neq i} f_{i,j}(x_j) \right) - x_i,$$

where the functions $f_{i,j} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfy

⁴Of the many private monitoring folk theorem papers, the most related are probably Ben-Porath and Kahneman (1996) and Renault and Tomala (1998), which assume a fixed monitoring network.

- $f_{i,j}(0) = 0$, $f_{i,j}$ is non-decreasing, and $f_{i,j}$ is either strictly concave or identically 0.
- $\lim_{x_1 \rightarrow \infty} \left(\sum_{j \neq i} f_{i,j}(x_1) \right) - x_1 = \lim_{x_1 \rightarrow \infty} \left(\sum_{j \neq i} f_{j,i}(x_1) \right) - x_1 = -\infty$.

The assumption that $f_{i,j}$ is non-decreasing for all $i \neq j$ is essential for interpreting x_j as player j 's level of cooperation. Note that the stage game is a prisoner's dilemma, in that playing $x_i = 0$ ("shirking") is a dominant strategy for player i in the stage game. The second assumption states that the cost of cooperation becomes infinitely greater than the benefit for sufficiently high levels of cooperation. Concavity and the assumption that $u_i(x)$ is separable in (x_1, \dots, x_n) play important roles in the analysis, and are discussed below.

Every period t , a *monitoring network* $L_t = (l_{i,j,t})_{i,j \in N \times N}$, $l_{i,j,t} \in \{0, 1\}$, is drawn independently from a fixed probability distribution μ on $\{0, 1\}^{n^2}$. In addition, *higher-order information* $y_t = (y_{i,t})_{i \in N}$, $y_{i,t} \in Y_i$ is drawn independently from a probability distribution $\pi(y_t | L_t)$, where the Y_i are arbitrary finite sets. At the end of period t , player i observes $h_{i,t} = \{z_{i,1,t}, \dots, z_{i,n,t}, y_{i,t}\}$, where $z_{i,j,t} = x_{j,t}$ if $l_{i,j,t} = 1$, and $z_{i,j,t} = \emptyset$ if $l_{i,j,t} = 0$. That is, player i observes the action of each of her out-neighbors and also observes the signal $y_{i,t}$, which may contain information about who observes whom in period t (as well as information about others' information about who observes whom, and so on).⁵ The special case of *perfect higher-order information* is when $y_{i,t} = L_t$ with probability 1 for all $i \in N$; this is the case where who observes whom is common knowledge (while monitoring of actions remains private). Assume that $\Pr(l_{i,i} = 1) = 1$ for all $i \in N$; that is, there is perfect recall. A repeated game with such a monitoring structure has *network monitoring*, the distribution μ is the *monitoring technology*, and the pair $(Y = Y_1 \times \dots \times Y_n, \pi)$ is the *higher-order information structure*. Let $h_i^t \equiv (h_{i,0}, h_{i,1}, \dots, h_{i,t-1})$ be player i 's private history at time $t \geq 1$, and denote the null history at the beginning of the game by $h^0 = h_i^0$ for all i . A (behavior) strategy of player i 's, σ_i , specifies a probability distribution over period t actions as a function of h_i^t .

Many important repeated games have network monitoring, including random matching (as in Kandori (1992) and Ellison (1994)) and monitoring on a fixed network (where L_t is deterministic and

⁵As to whether players observe their realized stage-game payoffs, note that $f_{i,j}(x_j)$ can be interpreted as player i 's *expected* benefit from player j 's action, and player i may only benefit from player j 's action when $l_{i,j,t} = 1$. However, some combinations of assumptions on $f_{i,j}$ and μ are not consistent with this interpretation, such as monitoring on a fixed network with global public goods, where $\Pr(l_{i,j,t} = 1) = 0$ but $f_{i,j} \neq 0$ for some i, j . An alternative interpretation is required in these cases: for example, the infinite time horizon could be replaced with an uncertain finite horizon without discounting, with payoffs revealed at the end of the game and δ viewed as the probability of the game's continuing. The former interpretation is appropriate for the long-distance trade example, while the alternative interpretation is appropriate for the school example.

constant, see Section 5). For random matching, by changing the higher-order information structure the model can allow for the case where players learn nothing about who matches with whom outside their own matches ($Y_i = \emptyset$ for all i), the case where who matches with whom is common knowledge ($y_{i,t} = L_t$ with probability 1 for all i), or any intermediate case. For monitoring on a fixed network, however, players already know who matches with whom, so higher-order information is irrelevant (although technically higher-order information could still act as a correlating device in this case). To fix ideas, note that a repeated game in which players observe the actions of their neighbors on a random graph that is determined in period 0 and then fixed for the duration of play does *not* have network monitoring, because the monitoring network is not drawn independently every period (e.g., player i observes player j 's action in period 1 with probability 1 if she observes it in period 0, but she does not observe player j 's action with probability 1 in period 0).

Throughout, I study weak perfect Bayesian equilibria (PBE) of this model with the property that, for every player i , time t , and monitoring network $L_{t'}$, for $t' < t$, the sum $\sum_{\tau=t'}^{\infty} \delta^{\tau-t} \mathbb{E} \left[u_i \left(\left(\sigma_j \left(h_j^\tau \right) \right)_{j=1}^n \right) | L_{t'} \right]$ is well-defined; that is, $\lim_{s \rightarrow \infty} \sum_{\tau=t}^s \delta^{\tau-t} \mathbb{E} \left[u_i \left(\left(\sigma_j \left(h_j^\tau \right) \right)_{j=1}^n \right) | L_{t'} \right]$ exists.⁶ This technical restriction ensures that players' continuation payoffs are well-defined, conditional on any past realized monitoring network. Fixing a description of the model other than the higher-order information structure—that is, a tuple $(N, (f_{i,j})_{i,j \in N \times N}, \delta, \mu)$ —let $\Sigma_{PBE}(Y, \pi)$ be the set of PBE strategy profiles when the higher-order information structure is (Y, π) . Player i 's *level of cooperation* under strategy profile σ is defined to be $(1 - \delta) \sum_{t=0}^{\infty} \delta^t \mathbb{E} [\sigma_i(h_i^t)]$. The main object of interest is the highest level of cooperation for each player that can be sustained in PBE for *any* higher-order information structure.

Definition 1 *Player i 's maximum cooperation with higher-order information structure (Y, π) is*

$$x_i^*(Y, \pi) \equiv \sup_{\sigma \in \Sigma_{PBE}(Y, \pi)} (1 - \delta) \sum_{t=0}^{\infty} \delta^t \mathbb{E} [\sigma_i(h_i^t)].$$

Player i 's robust maximum cooperation is

$$x_i^* \equiv \inf_{(Y, \pi)} x_i^*(Y, \pi).$$

Player i 's robust maximum cooperation is the highest level of cooperation that is sure to be sustainable in PBE for a given stage game, discount factor, and monitoring technology. Put

⁶Recall that a weak perfect Bayesian equilibrium is a strategy profile and belief system in which, for every player i and private history h_i^t , player i 's continuation strategy is optimal given her beliefs about the vector of private histories $(h_j^t)_{j=1}^N$, and these beliefs are updated using Bayes' rule whenever possible.

differently, it is the highest level of cooperation that an outside observer who does not know the higher-order information structure can be sure is sustainable. This seems reasonable for applications like local public good provision or long-range trade, where it seems much more palatable to make assumptions only about the probability that players observe each other's actions (the monitoring technology), rather than also making assumptions about what players observe about each other's observations, what they observe about what others observe about this, and so on.⁷

One more definition: a strategy profile σ is *higher-order information free* if $\sigma_i(h_i^t)$ does not depend on $(y_{i,\tau})_{\tau=0}^{t-1}$ for all $i \in N$. A higher-order information free strategy profile can naturally be viewed as a strategy profile in the game corresponding to any higher-order information structure (Y, π) . So the following definition makes sense.

Definition 2 *For any player $i \in N$ and level of cooperation x_i , a higher-order information free strategy profile σ robustly sustains x_i if $x_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \mathbb{E}[\sigma_i(h_i^t)]$ and $\sigma \in \Sigma_{PBE}(Y, \pi)$ for every higher-order information structure (Y, π) .*

This definition is demanding, in that a strategy profile can robustly sustain a level of cooperation only if it is a PBE for *any* higher-order information structure. However, my main theoretical result (Theorem 1) shows that there exists a grim trigger strategy profile that robustly sustains all players' robust maximum cooperation simultaneously (and the applied analysis in Sections 4 and 5 then focuses on this equilibrium). The resulting equilibrium is particularly important when it is also the PBE that maximizes social welfare. This is the primary case of interest in the literature on public good provision, where the focus is on providing incentives for sufficient cooperation, rather than on avoiding providing incentives for excessive cooperation. For example, the grim trigger strategy profile that simultaneously robustly sustains each player's maximum cooperation also maximizes utilitarian social welfare if x_i^* is below the first-best level (Lindahl-Samuelson benchmark) for every $i \in N$. Letting $f'_{j,i}$ denote the left-derivative of $f_{j,i}$ (which exists by concavity of $f_{j,i}$), this sufficient condition is

$$\sum_{j \neq i} f'_{j,i}(x_i^*) \geq 1 \text{ for all } i \in N.$$

This condition can be checked easily using the formula for $(x_i^*)_{i=1}^n$ given by Theorem 1.⁸ As a consequence, when this condition holds, all of the comparative statics on robust maximum cooper-

⁷However, I have implicitly assumed that the higher-order information structure is common knowledge among the players. But relaxing this would not affect the results.

⁸It would of course be desirable to characterize the entire set of payoffs that can be robustly sustained in PBE, or at least the entire Pareto frontier of this set, rather than only the equilibrium that robustly sustains maximum

ation derived below are also comparative statics on inefficiency relative to the Lindahl-Samuelson benchmark. I also present some quantitative examples of the relationship between inefficiency and network structure in Sections 4 and 5.

Before beginning the analysis, let me remark briefly on the motivation for studying this model. The model is intended to capture the essential features of cooperation in settings like those discussed in the introduction. Consider again the example of maintaining a school in a small village. In this setting, it is natural to think that villagers sometimes observe each other's contributions quite accurately but sometimes do not observe them at all (e.g., a villager might usually know how hard her friends have been working on the school, and might occasionally see someone else working, or learn that someone else has contributed money), and that it is very hard to observe the school's overall quality (e.g., because school quality might be best measured by students' labor market outcomes in the distant future). This suggests that repeated game models with (possibly imperfect) pure public monitoring are not well-suited for studying cooperation in this setting. My model instead makes the opposite assumption of pure network monitoring, and this leads to predictions that are very different from those that would emerge with imperfect public monitoring; for example, none of the four predictions enumerated in the introduction have been made in the literature on repeated games with imperfect public monitoring, and those predictions that relate a player's location in a monitoring network to her maximum cooperation cannot possibly be made in such models. It will become clear that my model is also very tractable: given a monitoring technology, it is easy to calculate each player's robust maximum cooperation. Of course, allowing players to access both network monitoring and noisy public monitoring—which is certainly more realistic than either pure public monitoring or pure network monitoring—remains a very interesting direction for future research. I discuss this possibility further in the conclusion.

3 Characterization of Robust Maximum Cooperation

This section presents the main theoretical result of the paper, which shows that all players' robust maximum cooperation can be robustly sustained in grim trigger strategies. To further motivate the focus on robustness, Section 3.1 presents an example showing that, with a given higher-order information structure, maximum cooperation may be sustained by complicated strategies that seem

cooperation. However, this problem appears intractable, just as it seems intractable in general repeated games with imperfect private monitoring (for fixed δ , rather than in the $\delta \rightarrow 1$ limit).




Figure 1: An Example where Complex Strategies are Optimal

“non-robust.” Section 3.2 then presents the main theoretical result.

3.1 Optimality of Complex Strategies with Imperfect Higher-Order Information

This section shows by example that for some higher-order information structures a player’s maximum level of cooperation cannot be sustained in (stationary) grim trigger strategies. I sketch the example here and defer the details to the appendix.

There are three players, arranged as in Figure 1. Player 1 is observed by player 2 with probability $1/2$ and is never observed by player 3. Players 2 and 3 always observe each other. Player 1 observes nothing. The realized monitoring network (drawn independently every period) is unobserved; in particular, player 3 does not observe when player 2 observes player 1 and when he does not (formally, $Y_i = \emptyset$ for all i). For each player i , $u_i((x_j)_{j=1}^3) = (\sum_{j \neq i} \sqrt{x_j}) - x_i$, and $\delta = 1/2$. It is straightforward to show that player 1’s maximum cooperation in grim trigger strategies equals 0.25 (see the appendix). I now sketch a strategy profile in which player 1’s maximum cooperation equals 0.2505 .

Player 1 always plays $x_1 = 0.2505$ on-path. Players 2 and 3 each have two on-path actions, denoted x_2^L , x_2^H , x_3^L , and x_3^H , with $x_2^L < x_2^H$ and $x_3^L < x_3^H$. Player 2 plays $x_2 = x_2^H$ in period 0. At subsequent odd-numbered periods t , player 2 plays x_2^H with probability 1 if he observed player 1’s period- $t - 1$ action, and otherwise plays each of x_2^H and x_2^L with probability $1/2$. At subsequent even-numbered periods t , player 2 plays x_2^H with probability 1 if he observed player 1’s period- $t - 2$ action, and otherwise plays each of x_2^H and x_2^L with probability $1/2$. Thus, if player 2 observes player 1’s action in even-numbered period t , he then plays x_2^H with probability 1 in *both* periods $t + 1$ and $t + 2$. Finally, player 3 plays $x_3 = x_3^H$ in period 0, and in every period $t \geq 1$ he plays x_3^H if player 2 played x_2^H in period $t - 1$, and plays x_3^L if player 2 played x_2^L in period $t - 1$. If any player i observes a deviation from this specification of on-path play (i.e., if any player deviates herself; if player 2 observes $x_1 \neq 0.2505$ or observes player 3 failing to take her prescribed action; or if player 3 observes $x_2 \notin \{x_2^L, x_2^H\}$), she then plays $x_i = 0$ in all subsequent periods. In the

appendix, I specify x_2^L , x_2^H , x_3^L , and x_3^H , and verify that the resulting strategy profile is a PBE.

Why can strategies of this form sustain greater maximum cooperation by player 1 than grim trigger strategies can? The key is that the difference between player 1's expectation of player 3's average future cooperation when player 1 cooperates and when player 1 shirks, conditional on the event that player 2 observes player 1 (which is the only event that matters for player 1's incentives), is larger than with grim trigger strategies. To understand this, consider what happens after period 2 sees player 1 play 0.2505 in period $t - 1$, for t odd. Conditional on this event, player 1's expectation of player 3's action in both periods $t + 1$ and $t + 2$ equals x_3^H ; but player 3's expectation of his own action in period $t + 2$ after seeing player 2 play x_2^H in period t is less than x_3^H , because he is not sure that player 2 observed player 1 in period $t - 1$. Indeed, if player 3 were sure that player 2 had observed player 1 in period $t - 1$, he would not be willing to play x_3^H (as he would have to play x_3^H in period $t + 2$ in addition to $t + 1$). Thus, the disagreement between player 1's expectation of player 3's average future cooperation (conditional on player 2 observing player 1) and player 3's (unconditional) expectation of his own average future cooperation improves player 1's incentive to cooperate without causing player 3 to shirk.

Note that all this example directly proves is that player 1's maximum cooperation is not sustainable in grim trigger strategies. However, it is not hard to show that any strategies that sustain more cooperation than is possible with grim trigger must involve "rewards," in that on-path actions must sometimes increase from one period to the next. This observation places a lower bound on how "complicated" the strategies that do sustain player 1's maximum cooperation in the example must be, even though actually computing these strategies seems intractable.⁹

3.2 Robust Optimality of Grim Trigger Strategies

This section shows that all players' robust maximum cooperation can be robustly sustained in grim trigger strategies, defined as follows.

Definition 3 *A strategy profile σ is a grim trigger strategy profile if there exist actions $(x_i)_{i=1}^n$ such that $\sigma_i(h_i^t) = x_i$ if $z_{i,j,\tau} \in \{x_j, \emptyset\}$ for all $z_{i,j,\tau} \in h_{i,\tau}$ and all $\tau < t$, and $\sigma_i(h_i^t) = 0$ otherwise.*

In a grim trigger strategy profile player i 's action at an off-path history h_i^t does not depend on the identity of the initial deviator. In particular, by perfect recall, player i plays $x_i = 0$ in every

⁹It is also trivial to modify this example to show that a player's *payoff* need not be maximized by grim trigger strategies: simply add a fourth player, observed by no one, who only values player 1's contributions.

period following a deviation by player i herself. Also, if a grim trigger strategy profile σ sustains each player's robust maximum cooperation, then under σ each player i plays x_i^* at every on-path history. Finally, grim trigger strategy profiles are clearly higher-order information free.

Next, I introduce an important piece of notation: define $D(\tau, t, i)$ recursively by

$$\begin{aligned} D(\tau, t, i) &= \emptyset \text{ if } \tau < t \\ D(t, t, i) &= \{i\} \\ D(\tau + 1, t, i) &= \{j : z_{j,k,\tau} = x_{k,\tau} \text{ for some } k \in D(\tau, t, i)\} \text{ if } \tau \geq t. \end{aligned}$$

That is, $D(\tau, t, i)$ is the set of players in period τ who have observed a player who has observed a player who has observed... player i since time t . By perfect recall, $D(\tau + 1, t, i) \supseteq D(\tau, t, i)$ for all τ, t , and i . The set $D(\tau, t, i)$ is important because $j \in D(\tau, t, i)$ is a necessary condition for player j 's time τ history to vary with player i 's actions at times after t . In particular, if players are using grim trigger strategies and player i shirks at time t , then $D(\tau, t, i)$ is the set of players who shirk at time τ . Note that the probability distribution of $D(\tau, t, i)$ is the same as the probability distribution of $D(\tau - t, i) \equiv D(\tau - t, 0, i)$, for all i and $\tau \geq t$.

I now state the main theoretical result of the paper.

Theorem 1 *There is a grim trigger strategy profile σ^* that robustly sustains each player's robust maximum cooperation. Furthermore, the vector of players' robust maximum cooperation $(x_i^*)_{i=1}^n$ is the (component-wise) greatest vector $(x_i)_{i=1}^n$ such that*

$$x_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \Pr(j \in D(t, i)) f_{i,j}(x_j) \text{ for all } i \in N. \quad (1)$$

Given that grim trigger strategies sustain each player's robust maximum cooperation, equation (1) is almost immediate: the left-hand side of (1) is the cost to player i of conforming to σ^* ; and the right-hand side of (1) is the benefit to player i of conforming to σ^* , which is that, if player i deviated, she would lose her benefit from player j 's cooperation whenever $j \in D(t, i)$. Thus, (1) states that the vector of robust maximum cooperation is the greatest vector of actions that equalizes the cost and benefit of cooperation for each player. In addition, it is easy to compute the vector $(x_i^*)_{i=1}^n$, as discussed in footnote 21 in the appendix.

Thus, the substance of Theorem 1 is showing that grim trigger strategies sustain each player's robust maximum cooperation. As shown above, grim trigger strategies do not sustain each player's maximum cooperation with every higher-order information structure. However, if one shows that

a grim trigger strategy profile σ sustains each player i 's maximum cooperation x_i with *some* higher-order information structure, then this implies that both $x_i^* \leq x_i$ (by definition of x_i^*) and $x_i^* \geq x_i$ (because σ must robustly sustain x_i),¹⁰ so Theorem 1 follows. Hence, the following key lemma implies Theorem 1.

Lemma 1 *The grim trigger strategy profile with on-path actions given by (1) sustains each player's maximum cooperation with perfect higher-order information.*

Lemma 1 is also of interest in its own right, as it shows that grim trigger strategies maximize cooperation when higher-order information is perfect. For example, Lemma 1 implies that grim trigger strategies always maximize cooperation for fixed monitoring networks, as with fixed monitoring networks who observes whom is always common knowledge. Since grim trigger strategies are higher-order information free, Lemma 1 also implies that each player's maximum cooperation with perfect higher-order information is weakly less than her maximum cooperation with any other higher-order information structure.

The key idea behind Lemma 1 is that a player is willing to cooperate (weakly) more at any on-path history if any other player cooperates more at any on-path history, because the first player is more likely to benefit from this increased cooperation when she conforms than when she deviates.¹¹ Thus, there is a kind of strategic complementarity between the actions of any two players at any two on-path histories. This suggests the following “proof” of Lemma 1: Define a function ϕ that maps the vector of all players' on-path actions at every on-path history, \vec{x} , to the vector of the highest actions that each player is willing to take at each on-path history when actions at all other on-path histories are as in \vec{x} , and players shirk at off-path histories. Let \bar{X} be an action greater than any on-path PBE action, and let \vec{X} be the vector of on-path actions \bar{X} . By complementarity among on-path actions, iterating ϕ on \vec{X} yields a sequence of vectors of on-path actions that are all constant across periods and weakly greater than the greatest fixed point of ϕ , and this sequence converges monotonically to the greatest fixed point of ϕ . Therefore, the greatest fixed point of ϕ is constant across periods, and it provides an upper bound on each player's maximum cooperation. Finally, verify that the grim trigger strategy profile with on-path actions given by the greatest fixed

¹⁰It is not difficult to show that if a grim trigger strategy profile sustains a player's maximum cooperation x_i with some higher-order information structure then it robustly sustains x_i . See the appendix.

¹¹This observation relies on the assumption of network monitoring, since otherwise a deviation by the first player may make some on-path histories more likely.

point of ϕ is a PBE.¹²

The problem with this “proof” (and there must be a problem, because the “proof” does not mention perfect higher-order information) is that, while the highest action that a player is willing to take at any on-path history is non-decreasing in every *other* player’s on-path actions, it is decreasing in *her own* future on-path actions. That is, a player is not willing to cooperate as much today when she knows that she will be asked to cooperate more tomorrow. Hence, the function ϕ as defined in the previous paragraph is not isotone, and thus may not have a greatest fixed point. This problem may be addressed by working not with players’ stage-game actions $\sigma_i(h_i^t)$, but rather with their “continuation actions” $X_i^t \equiv (1 - \delta) \sum_{\tau \geq t} \delta^{\tau-t} \sigma_i(h_i^\tau)$. Indeed, it can be shown that

$$\mathbb{E} [X_i^t | h_i^t] \leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} [X_j^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] \right),$$

for every player i and on-path history h_i^t . The intuition for this inequality is that, if player i shirks at time t , then player j starts shirking at time τ with probability $\Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i))$, and this yields lost benefits of at least $f_{i,j} \left(\mathbb{E} [X_j^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] \right)$ to player i . This inequality yields an upper bound on player i ’s expected continuation action, $\mathbb{E} [X_i^t | h_i^t]$, in terms of her expectation of *other* players’ continuation actions only. This raises the possibility that the function ϕ could be isotone when defined in terms of continuation actions X_i^t , rather than stage-game actions. For an approach along these lines to work, however, one must be able to express $\mathbb{E} [X_j^\tau | j \in D(\tau, t, i) \setminus D(\tau-1, t, i)]$ in terms of $\mathbb{E} [X_j^\tau | h_j^\tau]$ for player j ’s private histories h_j^τ . With perfect higher-order information (but not otherwise),

$$\mathbb{E} [X_j^\tau | j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] = \mathbb{E} [\mathbb{E} [X_j^\tau | h_j^\tau] | j \in D(\tau, t, i) \setminus D(\tau-1, t, i)],$$

so such an approach is possible.¹³

4 Equal Monitoring

This section imposes the assumption that all players’ actions are equally well-monitored in a sense that leads to sharp comparative statics results. In particular, assume throughout this section:

¹²For this last step, one might be concerned that grim trigger strategies do not satisfy off-path incentive constraints, as a player might want to cooperate off-path in order to slow the “contagion” of defecting, as in Kandori (1992) and Ellison (1994). As discussed in the introduction, this problem does not arise with continuous actions and payoffs.

¹³The assumptions that payoffs are concave and separable are also necessary. Without concavity, PBE actions could be scaled up indefinitely. Without separability, higher cooperation may be sustained when players take turns cooperating (see Example A1 in the online appendix).

- *Parallel Benefit Functions:* There exists a function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and scalars $\alpha_{i,j} \in \mathbb{R}_+$ such that $f_{i,j}(x) = \alpha_{i,j}f(x)$ for all $i, j \in N$ and all $x \in \mathbb{R}_+$.
- *Equal Monitoring:* $\sum_{t=0}^{\infty} \delta^t \sum_{k \neq i} \Pr(k \in D(t, i)) \alpha_{i,k} = \sum_{t=0}^{\infty} \delta^t \sum_{k \neq j} \Pr(k \in D(t, j)) \alpha_{j,k}$ for all $i, j \in N$.

Parallel benefit functions imply that the importance of player j 's cooperation to player i may be summarized by a real number $\alpha_{i,j}$. With this assumption, equal monitoring states that the expected discounted number of players who may be influenced by player i 's action, weighted by the importance of their actions to player i , is the same for all $i \in N$. To help interpret these assumptions, note that if $\alpha_{i,j}$ is constant across players i and j then, for generic discount factors δ , equal monitoring holds if and only if $\mathbb{E}[\#D(t, i)] = \mathbb{E}[\#D(t, j)]$ for all $i, j \in N$ and $t \in \mathbb{N}$; that is, if and only if the expected number of players who find out about shirking by player i within t periods is the same for all $i \in N$.

Section 4.1 derives a simple and general formula for comparative statics on robust maximum cooperation under equal monitoring. Sections 4.2 and 4.3 apply this formula to the leading special case of *(global) public good provision*, where $\alpha_{i,j} = \alpha$ for all $i \neq j$; that is, where all players value each other's actions equally. Section 4.2 studies the effect of group size on public good provision, and Section 4.3 considers the effect of "uncertainty" in monitoring on public good provision.

Finally, the higher-order information structure plays no role in this section or the following one, because these sections study comparative statics on players' maximum robust cooperation, which is independent of the higher-order information structure by definition.

4.1 Comparative Statics Under Equal Monitoring

The section derives a formula for comparative statics on robust maximum cooperation under equal monitoring. The first step is noting that each player's robust maximum cooperation is the same under equal monitoring (proof in appendix).

Corollary 1 *With equal monitoring, $x_i^* = x_j^*$ for all $i, j \in N$.*

Thus, under equal monitoring each player has the same robust maximum cooperation x^* . I wish to characterize when x^* is higher in one game than another, when both games satisfy equal monitoring and have the same underlying benefit function f . Formally, a *game with equal monitoring* $\Gamma = (N, (\alpha_{i,j})_{i,j \in N \times N}, \delta, \mu)$ is a model satisfying the assumptions of Section 2 as well as equal

monitoring. For any game with equal monitoring Γ , let $x^*(\Gamma)$ be the robust maximum cooperation in Γ , and let

$$B(\Gamma) \equiv (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \Pr(j \in D(t, i)) \alpha_{i,j}$$

be player i 's benefit of cooperation (i.e., the right-hand side of (1)) when $f(x_j) = 1$ for all $j \in N$, which is independent of the choice of $i \in N$ by equal monitoring. The comparative statics result for games with equal monitoring is the following:

Theorem 2 *Let Γ' and Γ be two games with equal monitoring. Then $x^*(\Gamma') \geq x^*(\Gamma)$ if $B(\Gamma') \geq B(\Gamma)$, with strict inequality if $B(\Gamma') > B(\Gamma)$ and $x^*(\Gamma') > 0$.*

Proof. Since $x_i^* = x^*$ for all $i \in N$, (1) may be rewritten as

$$x^* = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \Pr(j \in D(t, i)) \alpha_{i,j} f(x^*) = B(\Gamma) f(x^*).$$

Hence, $x^*(\Gamma)$ is the greatest zero of the concave function $B(\Gamma) f(x) - x$. If $B(\Gamma') \geq B(\Gamma)$, then $B(\Gamma') f(x^*(\Gamma)) - x^*(\Gamma) \geq B(\Gamma) f(x^*(\Gamma)) - x^*(\Gamma) = 0$, which implies that $x^*(\Gamma') \geq x^*(\Gamma)$. If $B(\Gamma') > B(\Gamma)$ and $x^*(\Gamma') > 0$, then either $x^*(\Gamma) = 0$ (in which case $x^*(\Gamma') > x^*(\Gamma)$ trivially) or $x^*(\Gamma) > 0$, in which case $B(\Gamma') f(x^*(\Gamma)) - x^*(\Gamma) > B(\Gamma) f(x^*(\Gamma)) - x^*(\Gamma) = 0$, which implies that $x^*(\Gamma') > x^*(\Gamma)$. ■

Theorem 2 gives a complete characterization of when $x^*(\Gamma)$ is greater or less than $x^*(\Gamma')$, for any two games with equal monitoring Γ and Γ' . In particular, robust maximum cooperation is greater when the expected discounted number of players who may be influenced by a player's action, weighted by the importance of their actions to that player, is greater. For example, in the case of *global* public good provision (where all players value all other players' actions equally), robust maximum cooperation is greater when the sets $D(t, i)$ are likely to be larger; while if each player only values the actions of a subset of the other players (her geographic neighbors, her trading partners, etc.), then robust maximum cooperation is greater when the *intersection* of the sets $D(t, i)$ and the set of players whose actions player i values is likely to be larger. Hence, Theorem 2 characterizes how different monitoring technologies sustain different kinds of cooperative behaviors.

4.2 The Effect of Group Size on Global Public Good Provision

This section uses Theorem 2 to analyze the effect of group size on robust maximum cooperation in the leading special case of global public good provision, where $\alpha_{i,j} = \alpha$ for all $i \neq j$.

In the case of (global) public good provision,

$$B(\Gamma) = \alpha \sum_{t=0}^{\infty} \delta^t (\mathbb{E}[\#D(t, i)] - 1).$$

Thus, for public goods, all the information needed to determine whether changing the game increases or decreases the (robust) maximum per capita level of public good provision is contained in the product of two terms: the marginal benefit to each player of public good provision, α , and $(1/(1 - \delta))$ less than) the *effective contagiousness* of the monitoring technology, $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t, i)]$. Information such as group size, higher moments of the distribution of $\#D(t, i)$, and which players are more likely to observe which other players are not directly relevant. In particular, the single number $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t, i)]$ —the effective contagiousness—completely determines the effectiveness of a monitoring technology in supporting public good provision.

This finding that comparative statics on the per-capita level of public good provision are determined by the product of the marginal benefit of the public good to each player and the effective contagiousness of the monitoring technology yields useful intuitions about the effect of group size on the per capita level of public good provision. In particular, index a game Γ by its group size, n , and write $\alpha(n)$ for the corresponding marginal benefit of contributions and $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t, n)]$ for the effective contagiousness (I use this simpler notation for the remainder of this section). Normally, one would expect $\alpha(n)$ to be decreasing in n (a larger population reduces player i 's benefit from player j 's contribution to the public good) and $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t, n)]$ to be increasing in n (a larger population makes it more likely that player i 's action is observed by more individuals), yielding a tradeoff between the marginal benefit of contributions and the effective contagiousness. Consider again the example of constructing a local infrastructure project, like a well. In this case, $\alpha(n)$ is likely to be decreasing and concave: since each individual uses the well only occasionally, there are few externalities among the first few individuals, but eventually it starts to become difficult to find times when the well is available, water shortages become a problem, etc.. Similarly, $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t, n)]$ is likely to be increasing, and may be concave if there are “congestion” effects in monitoring. Thus, it seems likely that in typical applications $\alpha(n) \sum_{t=0}^{\infty} \delta^t (\mathbb{E}[\#D(t, n)] - 1)$, and therefore per capita public good provision, is maximized at an intermediate value of n .

Theorem 2 yields particularly simple comparative statics for the leading cases of pure public goods ($\alpha(n) = 1$) and divisible public goods ($\alpha(n) = 1/n$), which are useful in examples below.

Corollary 2 *With pure public goods ($\alpha(n) = 1$), if $\mathbb{E}[\#D(t, n')] \geq \mathbb{E}[\#D(t, n)]$ for all t then $x^*(n') \geq x^*(n)$, with strict inequality if $\mathbb{E}[\#D(t, n')] > \mathbb{E}[\#D(t, n)]$ for some $t \geq 1$ and $x^*(n') >$*

0.

With pure public goods, $x^*(n)$ is increasing unless monitoring degrades so quickly as n increases that the expected *number* of players who find out about a deviation within t periods is *decreasing* in n , for some t . This suggests that $x^*(n)$ is increasing in n in many applications.

Corollary 3 *With divisible public goods ($\alpha(n) = 1/n$), if $(\mathbb{E}[\#D(t, n')] - 1)/n' \geq (\mathbb{E}[\#D(t, n)] - 1)/n$ for all t then $x^*(n') \geq x^*(n)$, with strict inequality if $(\mathbb{E}[\#D(t, n')] - 1)/n' > (\mathbb{E}[\#D(t, n)] - 1)/n$ for some $t \geq 1$ and $x^*(n') > 0$.*

With divisible public goods, $x^*(n)$ is increasing only if the expected *fraction* of players (other than the deviator herself) who find out about a deviation within t periods is non-decreasing in n , for all t . This suggests that, with divisible public goods, $x^*(n)$ is decreasing in many applications.

The following two examples demonstrate the usefulness of Theorem 2 and Corollaries 2 and 3. An earlier version of this paper (available upon request) contains additional examples.

4.2.1 Random Matching

Monitoring is *random matching* if in each period every player is linked with one other player at random, and $l_{i,j,t} = l_{j,i,t}$ for all $i, j \in N$ and all t . This is possible only if n is even.

It can be show that, with random matching, $\mathbb{E}[\#D(t, n)]$ is non-decreasing in n and is increasing in n for $t = 2$. Therefore, Corollary 2 implies that, with pure public goods, robust maximum cooperation is increasing in group size.

Proposition 1 *With random matching and pure public goods, if $n' > n$ then $x^*(n') \geq x^*(n)$, with strict inequality if $x^*(n') > 0$.*

However, it can also be shown that $\sum_{t=0}^{\infty} \delta^t (\mathbb{E}[\#D(t, n')] - 1)/n' < \sum_{t=0}^{\infty} \delta^t (\mathbb{E}[\#D(t, n)] - 1)/n$ whenever $n' > n$, n' and n are sufficiently large, and $\delta < 1/2$. In this case, Theorem 2 implies that, with divisible public goods, robust maximum cooperation is decreasing in group size.

Proposition 2 *With random matching and divisible public goods, if $\delta < \frac{1}{2}$ then, for any $\gamma > 0$, there exists \bar{N} such that $x^*(n') \leq x^*(n)$ if $n' > (1 + \gamma)n \geq \bar{N}$, with strict inequality if $x^*(n') > 0$.*

4.2.2 Monitoring on a Circle

Monitoring is *on a circle* if the players are arranged in a fixed circle and there exists an integer $d \geq 1$ such that $l_{i,j,t} = 1$ if and only if the distance between i and j is at most d .

It is a straightforward consequence of Corollary 2 that robust maximum cooperation is increasing in group size with monitoring on a circle and pure public goods.

Proposition 3 *With monitoring on a circle and pure public goods, if $n' > n$ then $x^*(n') \geq x^*(n)$, with strict inequality if $x^*(n') > 0$.*

In contrast, Corollary 3 implies that robust maximum cooperation is decreasing in group size with monitoring on a circle and divisible public goods.

Proposition 4 *With monitoring on a circle and divisible public goods, if $n' > n$ then $x^*(n') \leq x^*(n)$, with strict inequality if $d < n'/2$ and $x^*(n') > 0$.*

Finally, monitoring on a circle is a simple test case in which to compare robust maximum cooperation with the first-best (Lindahl-Samuelson) benchmark, for various discount factors and group sizes. Assume that n is odd, and consider first the case of pure public goods, with $f(x) = \sqrt{x}$. Then first-best cooperation is given by

$$(n-1) \frac{1}{2\sqrt{x^{FB}}} = 1,$$

or

$$x^{FB} = \left(\frac{n-1}{2}\right)^2.$$

Robust maximum cooperation is given by equation (1) (or by the simpler equation (2) of Section 5.1), which yields

$$x^* = 2\delta \left(\frac{1 - \delta^{\frac{n-1}{2}}}{1 - \delta}\right) \sqrt{x^*},$$

or

$$x^* = \left(2\delta \left(\frac{1 - \delta^{\frac{n-1}{2}}}{1 - \delta}\right)\right)^2.$$

The following table displays robust maximum cooperation for various combinations of δ and n .

		n			
		11	31	101	1001
δ	.5	3.75	4.00	4.00	4.00
	.7	15.07	21.57	21.78	21.78
	.9	54.33	204.3	320.7	324.0
	.99	94.17	767.8	6,116.6	38,691

For comparison, first-best cooperation does not depend on δ , and equals 25 when $n = 11$, 225 when $n = 31$, 2500 when $n = 101$, and 250,000 when $n = 1001$. Several remarks are in order here. First, robust maximum cooperation is less than first-best cooperation for all combinations of δ and n other than $(\delta = .9, n = 11)$ and $(\delta = .99, n \in \{11, 31, 101\})$. Thus, the equilibrium that sustains robust maximum cooperation also maximizes utilitarian social welfare, unless players are very patient and the group is relatively small. Second, robust maximum cooperation falls far short of the first-best benchmark in large groups, unless players are very patient. Third, network structure matters: for example, if the monitoring network is a clique rather than a circle (i.e., if monitoring is perfect), then robust maximum cooperation is given by $x^* = (\delta(n - 1))^2$, and hence first-best cooperation can be sustained whenever $\delta \geq 1/2$. Finally, all of these conclusions also hold for the case of divisible public goods, as with divisible public goods both x^{FB} and x^* are given by dividing the corresponding quantities with pure public goods by n^2 .

4.3 The Effect of Uncertain Monitoring on Global Public Good Provision

This section provides a result comparing monitoring technologies in terms of the maximum level of (robust) global public good provision they support, for a fixed group size. As discussed in the previous subsection, a monitoring technology supports greater robust maximum cooperation in global public good provision if and only if it has greater effective contagiousness, $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t)]$, where the parameter n is omitted because it is held fixed in this subsection. I compare “less certain” monitoring, where it is likely that either a large or small fraction about the population finds out about a deviation, with “more certain” monitoring, where it is likely that an intermediate fraction of the population finds out about it, in the sense of second-order stochastic dominance. Under fairly broad conditions, more certain monitoring supports greater robust maximum cooperation.

The analysis of this subsection relies on the following assumption, which states that the distribution over $\#D(t + 1)$ depends only on $\#D(t)$.

- There exists a family of functions $\{g_k : \{0, \dots, n\} \rightarrow [0, 1]\}_{k=1}^n$ with such that, whenever $\#D(t) = k$, $\Pr(\#D(t+1) = k') = g_k(k')$, for all t, k , and k' .

This assumption is satisfied by random matching, for example, but not by monitoring on a circle, because with monitoring on a circle the distribution of $\#D(t+1)$ depends on the identities of the members of $D(t)$.

Given a probability mass function g_k , define the corresponding distribution function $G_k(k') \equiv \sum_{s=0}^{k'} g_k(s)$. Recall that a distribution \tilde{G}_k *strictly second-order stochastically dominates* G_k if $\sum_{s=0}^n \eta(s) \tilde{g}_k(s) > \sum_{s=0}^n \eta(s) g_k(s)$ for all increasing and strictly concave functions $\eta : \mathbb{R} \rightarrow \mathbb{R}$. The following result compares monitoring under $\{\tilde{g}_k\}_{k=1}^n$ and $\{g_k\}_{k=1}^n$.

Theorem 3 *Suppose that $\tilde{G}_k(k')$ and $G_k(k')$ are decreasing and strictly convex in k for $k \in \{0, \dots, k'\}$ and $k' \in \{0, \dots, n\}$, and that \tilde{G}_k strictly second-order stochastically dominates G_k for $k \in \{1, \dots, n-1\}$. Then robust maximum cooperation is strictly greater under a monitoring technology corresponding to $\{\tilde{g}_k(\cdot)\}_{k=1}^n$ than under a monitoring technology corresponding to $\{g_k(\cdot)\}_{k=1}^n$.*

The intuition for Theorem 3 is fairly simple: If \tilde{G}_k strictly second-order stochastically dominates G_k for all k , then under \tilde{G}_k it is more likely that an intermediate number of players find out about an initial deviation each period. Since $G_k(k')$ and $\tilde{G}_k(k')$ are decreasing and convex in k , the expected number of players who find out about the deviation within t periods increases in t more quickly when it is more likely that an intermediate number of players find out about the deviation each period. Hence, $\sum_{t=0}^{\infty} \delta^t \mathbb{E}[\#D(t)]$ is strictly higher under a monitoring technology corresponding to $\{\tilde{g}_k(\cdot)\}_{k=1}^n$ than under a monitoring technology corresponding to $\{g_k(\cdot)\}_{k=1}^n$, and the theorem then follows from Theorem 2.

5 Fixed Monitoring Networks

This section studies both global and local public good provision with network monitoring when the monitoring network is fixed over time. That is, throughout this section I make the following assumption on the (deterministic) monitoring technology.

- *Fixed Undirected Monitoring Network:* There exists a network $L = (l_{i,j})_{(i,j) \in N \times N}$ such that $l_{i,j,t} = l_{i,j} = l_{j,i}$ for all t .

I also assume that the stage game satisfies one of the following two properties, where $N(i)$ is the set of player i 's neighbors in L .

- *Global Public Goods:* $u_i(x) = \left(\sum_{j \neq i} f(x_j)\right) - x_i$.
- *Local Public Goods:* $u_i(x) = \left(\sum_{j \in N(i)} f(x_j)\right) - x_i$.

The extensions of all of the results in this section to directed networks is straightforward. I discuss below where the assumption of global or local public goods can be relaxed.

Section 5.1 introduces a new definition of centrality in networks, and uses Theorem 1 to show that more central players have greater robust maximum cooperation. Section 5.2 shows that centrality can also be used to determine when one network “dominates” another in terms of supporting cooperation. Finally, Section 5.3 remarks on the stability of monitoring networks, emphasizing differences between the cases of global and local public goods.

5.1 Centrality and Robust Maximum Cooperation

Theorem 1 provides a general characterization of players’ robust maximum cooperation as a function of the discount factor and benefit functions. Here, I provide a partial ordering (“centrality”) of players *in terms of their network characteristics* under which higher players have greater robust maximum cooperation *for all discount factors and benefit functions*. Intuitively, player i is “more central” than player j if i has more neighbors (within distance t , for all $t \in \mathbb{N}$) than j , i ’s neighbors have more neighbors than j ’s neighbors, i ’s neighbors’ neighbors have more neighbors than j ’s neighbors’ neighbors, and so on. Formally, let $d(i, j)$ be the distance (shortest path length) between players i and j , with $d(i, j) \equiv \infty$ if there is no path between i and j . The definition of centrality is the following.¹⁴

Definition 4 *Player i is 1-more central than player j if, for all $t \in \mathbb{N}$, $\#\{k \in N : d(i, k) \leq t\} \geq \#\{k \in N : d(j, k) \leq t\}$. Player i is strictly 1-more central than player j if in addition $\#\{k \in N : d(i, k) \leq t\} > \#\{k \in N : d(j, k) \leq t\}$ for some t .*

For all integers $s \geq 2$, player i is s -more central than player j if, for all $t \in \mathbb{N}$, there exists a surjection $\psi : \{k \in N : d(i, k) \leq t\} \rightarrow \{k \in N : d(j, k) \leq t\}$ such that, for all k with $d(j, k) \leq t$, there exists $k' \in \psi^{-1}(k)$ such that k' is $s - 1$ -more central than k . Player i is strictly s -more central than player j if in addition k' is strictly $s - 1$ -more central than k for some t , ψ , k , and k' .

Player i is more central than player j if i is s -more central than j for all $s \in \mathbb{N}$. Player i is strictly more central than player j if in addition i is strictly s -more central than j for some $s \in \mathbb{N}$.

¹⁴This seems like a natural notion of centrality, but I am not aware of any references to it in the literature.




Figure 2: A Five-Player Example

As a first example, consider five players arranged in a line (Figure 2). Player 3 is strictly more central than players 2 and 4, who are in turn strictly more central than players 1 and 5. To see this, note first that player 3 is strictly 1-more central than players 2 and 4, who are in turn each strictly 1-more central than players 1 and 5. For example, player 2 is strictly 1-more central than player 5 because player 2 has 3 neighbors within distance 1 (including player 2 herself), 4 neighbors within distance 2, and 5 neighbors within distance 3 or more; while player 5 has 2 neighbors within distance 1, 3 neighbors within distance 2, 4 neighbors within distance 3, and 5 neighbors within distance 4 or more. Next, suppose that player 3 is s -more central than players 2 and 4, and that players 2 and 4 are both s -more central than players 1 and 5. Then it is easy to check that player 3 is also $s + 1$ -more central than players 2 and 4, who in turn are both $s + 1$ -more central than players 1 and 5; for example, one surjection $\psi : \{k \in N : d(2, k) \leq 2\} \rightarrow \{k \in N : d(5, k) \leq 2\}$ that satisfies the terms of the definition is given by $\psi(1) = \psi(2) = 5$, $\psi(3) = 3$, $\psi(4) = 4$ (noting that a player is always more central than herself, because in this case ψ can be taken to be the identity mapping). Thus, by induction on s , player 3 is strictly more central than players 2 and 4, who are in turn strictly more central than players 1 and 5.

The main result of this section states that, with either global or local public goods, more central players have greater robust maximum cooperation, regardless of the discount factor δ and benefit function f . The result can easily be generalized to allow for utility functions intermediate between global and local public goods, where a player's benefit from another player's action is a decreasing function of the distance between them.¹⁵ The proof uses a monotonicity argument similar to that in the proof of Lemma 1, which shows that more central players cooperate more at every step of a sequence of vectors of actions that converges to the vector of robust maximum cooperation.

Theorem 4 *With either global or local public goods, if player i is more central than player j , then $x_i^* \geq x_j^*$. The inequality is strict if player i is strictly more central than player j and $x_k^* > 0$ for all $k \in N$.*

The proof of the strict inequality in Theorem 4 uses the following lemma.

¹⁵Formally, Theorem 4 holds whenever there exist a function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and constants $\alpha_d \in \mathbb{R}_+$ such that $\alpha_d \geq \alpha_{d+1} \geq 0$ for all $d \in \mathbb{N}$ and $u_i(x) = \left(\sum_{j \neq i} \alpha_{d(i,j)} f(x_j)\right) - x_i$ for all $i \in N$.

Lemma 2 *Player i is more central than player j if and only if for all $t \in \mathbb{N}$ there exists a surjection $\psi : \{k \in N : d(i, k) \leq t\} \rightarrow \{k \in N : d(j, k) \leq t\}$ such that, for all k with $d(j, k) \leq t$, there exists $k' \in \psi^{-1}(k)$ such that k' is more central than k .*

Proof of Theorem 4. I prove the result for global public goods. The proof for local public goods is similar.

Let $\phi : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$, be defined as in the proof of Lemma 1; with a fixed monitoring network and global public goods, this simplifies to

$$\phi_i \left((x_j)_{j=1}^n \right) = \sum_{j \neq i} \delta^{d(i,j)} f(x_j) \text{ for all } i \in N.$$

As in the proof of Lemma 1, define x_i^m recursively by $x_i^1 = \bar{X}$ (a large constant defined in Step 1a of the proof of Lemma 1) and $x_i^{m+1} = \phi_i \left((x_j^m)_{j=1}^n \right)$. The proof of Lemma 1 shows that $x_i^* = \lim_{m \rightarrow \infty} x_i^m$.

Suppose that player i is more central than player j . I claim that $x_i^m \geq x_j^m$ for all $m \in \mathbb{N}$, which proves the weak inequality. Trivially, $x_i^1 = \bar{X} \geq \bar{X} = x_j^1$. Now suppose that $x_{k'}^m \geq x_k^m$ whenever player k' is more central than player k , for some $m \in \mathbb{N}$. Since player i is $m+1$ -more central than player j , for any $t \in \mathbb{N}$ there exists a surjection $\psi : \{k \in N : d(i, k) \leq t\} \rightarrow \{k \in N : d(j, k) \leq t\}$ such that, for all k with $d(j, k) \leq t$, there exists $k' \in \psi^{-1}(k)$ such that k' is m -more central than k . Since $x_{k'}^m \geq x_k^m$, this implies that $\sum_{k: d(i,k) \leq t} f(x_k^m) \geq \sum_{k: d(j,k) \leq t} f(x_k^m)$. This holds for all t , which implies that $x_i^{m+1} = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k: d(i,k) \leq t} f(x_k^m) \geq (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k: d(j,k) \leq t} f(x_k^m) = x_j^{m+1}$. It follows by induction that $x_i^m \geq x_j^m$ for all $m \in \mathbb{N}$.

To prove the strict inequality, suppose that player i is strictly more central than player j and that $x_k^* > 0$ for all $k \in N$. Rewrite (1) as

$$x_i^* = \sum_{j \neq i} \delta^{d(i,j)} f(x_j^*). \quad (2)$$

Suppose that i is more central than j and strictly 1-more central than j , let $\underline{x}^* \equiv \min_k x_k^*$ (which is positive by assumption), let $\bar{x}^* \equiv \max_k x_k^*$, and let \bar{d} be the diameter of L (i.e., the maximum distance between any two path-connected nodes in L). Then, by Lemma 2 and (2), $x_i^* \geq x_j^* + \delta^{\bar{d}-1} \min\{\delta, 1 - \delta\} f(\underline{x}^*)$, as player i has at least one more distance- t neighbor than player j for some $t \in \mathbb{N}$.¹⁶ Therefore, there exists $\varepsilon_1 > 0$ such that $x_i^* - x_j^* \geq \varepsilon_1 > 0$ whenever i is more

¹⁶The $\min\{\delta, 1 - \delta\}$ term corresponds to the possibility that player i may have one more distance- \bar{d} neighbor than player j , or may have one more distance- $\bar{d} - 1$ neighbor and the same number of distance- \bar{d} neighbors.

central than j and strictly 1-more central than j . Now suppose that there exists $\varepsilon_s > 0$ such that $x_i^* - x_j^* \geq \varepsilon_s > 0$ whenever i is more central than j and strictly s -more central than j . Suppose that i is more central than j and strictly $s + 1$ -more central than j . Then $x_i^* \geq x_j^* + \delta^{\bar{d}-1} \max\{\delta, 1 - \delta\} (f(\bar{x}^* + \varepsilon_s) - f(\bar{x}^*))$, by Lemma 2 and (2), which implies that there exists $\varepsilon_{s+1} > 0$ such that $x_i^* - x_j^* \geq \varepsilon_{s+1} > 0$. By induction on s , it follows that $x_i^* > x_j^*$ whenever i is strictly more central than j . ■

Four remarks on Theorem 4 are in order. First, the conclusion of Theorem 4 would still hold for local public goods (but not global public goods) if the definition of centrality was weakened by specifying that player i is 1-more central than player j whenever $\#N(i) \geq \#N(j)$ (by essentially the same proof). Thus, players' robust maximum cooperation can be ordered for more networks with local public goods than with global public goods. Second, the fixed point equation (2)—which is substantially simpler than the general fixed point equation (1)—orders players' robust maximum cooperation for *any* fixed monitoring network. So comparing players' robust maximum cooperation for any fixed monitoring network is not difficult, even if the centrality ordering is very incomplete.

Third, Theorem 4 provides a new perspective on the Olsonian idea of the “exploitation of the great by the small.” Olson (1965) notes that small players may free ride on large players if larger players have greater private incentives to contribute to public goods. Theorem 4 illustrates an additional reason why larger players might contribute disproportionately to public goods: larger players may be more central, in which case they may be punished more effectively for shirking. While this “exploitation” implies that more central players receive lower payoffs than less central players with global public goods, Corollary 7 below implies that more central players receive *higher* payoffs than less central players with local public goods, which shows that with local public goods the benefit of having more neighbors more than offsets the cost of contributing more.

Fourth, my definition of centrality is related to Bonacich centrality (Bonacich, 1987). My definition of centrality is a partial order, as it ranks players in a way that is invariant to the benefit function and discount factor, so a more direct comparison with Bonacich centrality results from comparing players' robust maximum cooperation for a fixed benefit function, f , and discount factor, δ ; in this case, δ is analogous to the decay factor in the definition of Bonacich centrality, β . Indeed, the formula for a player's robust maximum cooperation, (2), is very similar to the formula for her Bonacich centrality, with the important difference that a player's robust maximum cooperation depends on other players' robust maximum cooperation through the concave function f , while this dependence is linear for Bonacich centrality. As a consequence, the vector of players' robust




Figure 3: A Seven-Player Example

maximum cooperation is unique, while the vector of players' Bonacich centrality is determined only up to multiplication by a constant.

For general monitoring networks, it may be difficult to verify that one player is more central than another, making it hard to apply Theorem 4. Sometimes, however, symmetries in the network can be exploited to determine which players are more central than others more easily. The remainder of this section shows how this can be done. Corollary 4 states that, if player i is closer to *all* players $k \neq i, j$ than is player j , then player i is more central than player j . Corollary 5 shows that if players i and k are in “symmetric” positions in the monitoring network (in that there exists a graph automorphism ρ on L such that $k = \rho(i)$) and player k is more central than player j , then player i is more central than player j as well.¹⁷

Corollary 4 *If $d(i, k) \leq d(j, k)$ for all $k \neq i, j$, then player i is more central than player j . Player i is strictly more central than player j if in addition the inequality is strict for some $k \neq i, j$.*

Corollary 5 *If there exists a graph automorphism $\rho : N \rightarrow N$ such that $\rho(i)$ is more central (resp., strictly more central) than j , then i is more central (resp., strictly more central) than j .*

The “bow tie” network in Figure 3 illustrates the usefulness of Corollaries 4 and 5.¹⁸ First, Corollary 4 immediately implies that player 3 is more central than players 1 and 2, and that player 5 is more central than players 6 and 7. Second, observe that the following map ρ is an automorphism

¹⁷A *graph automorphism* on L is a permutation ρ on N such that $l_{i,j} = l_{\rho(i),\rho(j)}$ for all $i, j \in N$. That is, a graph automorphism is a permutation of vertices that preserves links.

¹⁸This example is the same as that in Figure 2.13 of Jackson (2008), which Jackson uses to illustrate various graph-theoretic concepts of centrality. An important difference between my definition and those discussed by Jackson is that my definition gives a partial order on nodes, while all the definitions discussed by Jackson give total orders.

of L : $\rho(1) = 7$, $\rho(2) = 6$, $\rho(3) = 5$, $\rho(4) = 4$, $\rho(5) = 3$, $\rho(6) = 2$, and $\rho(7) = 1$. Thus, Corollary 5 implies that each player in $\{3, 5\}$ is more central than each player in $\{1, 2, 6, 7\}$. Given this observation, it is not hard to show that player 4 is more central than each player in $\{1, 2, 6, 7\}$. Finally, neither of players 3 and 4 are more central than the other, as player 3 has more immediate neighbors while player 4 has more neighbors within distance 2. Therefore, Theorem 5 does not say whether player 3 or player 4 has greater robust maximum cooperation. This is reassuring, because one can easily construct examples in which $x_3^* > x_4^*$ and others in which the reverse inequality holds: for example, if $f(x) = \sqrt{x}$ (with global public goods), then $x_1^* \approx 2.638$, $x_3^* \approx 3.425$, and $x_4^* \approx 3.475$ if $\delta = 0.5$, whereas if $\delta = 0.4$ then $x_1^* \approx 1.378$, $x_3^* \approx 1.849$, and $x_4^* \approx 1.839$. Indeed, it is not surprising that player 3 contributes more relative to player 4 when δ is lower, as in this case the fact that player 3 has more immediate neighbors is more important, while player 4's greater number of distance-2 neighbors matters more when δ is higher (since δ^2 is low relative to δ when δ is low). However, there are networks in which a player i is not more central than player j but nonetheless $x_i^* \geq x_j^*$ for every concave benefit function f and discount factor δ , as shown by Example A2 in the online appendix. This implies that centrality is not necessary to order players' maximum equilibrium cooperations for every benefit function and discount factor.

As an aside, the bow tie network also provides an interesting example of how social welfare depends on network structure, and how this varies with the discount factor. When $f(x) = \sqrt{x}$, first-best cooperation is $x^{FB} = ((7 - 1)/2)^2 = 9$, so robust maximum cooperation is below first-best for $\delta = 0.5$ or 0.4 .¹⁹ Were the seven players arranged in a circle rather than a bow tie, then robust maximum cooperation would be 3.984 if $\delta = 0.5$, and would be 1.777 if $\delta = 0.4$. In particular, compared to the bow tie, utilitarian social welfare is much higher with monitoring on a circle when $\delta = 0.5$, but is similar if $\delta = 0.4$ (and is lower if $\delta = 0.3$). The intuition is that the circle supports more cooperation relative to the bow tie when δ is higher, as when δ is higher players' lower average degree in the circle is more than offset by the circle's smaller diameter. More generally, networks that have lower average degree but also have shorter average distance between nodes that are not neighbors tend to support relatively more cooperation—and thus relatively greater efficiency—when δ is higher.

¹⁹All calculations in this paragraph follow from the relevant formulas in Section 4.2.2.

5.2 Comparing Networks

This section shows that centrality is a key tool for comparing different networks in terms of their capacity to support cooperation, not just for comparing individuals within a fixed network. To see this, note that the “more central” relation can be immediately extended to pairs of players in different networks L' and L by specifying that player $i' \in L'$ is more central than player $i \in L$ if player i' is more central than player i in the network consisting of disjoint components L' and L . With this definition, the result is the following.

Theorem 5 *For any network L' and connected network L , if there exists players $i' \in L'$ and $i \in L$ such that player i' is more central than player i , then there exists a surjection $\psi : L' \rightarrow L$ such that, for all $j \in L$, there exists $j' \in \psi^{-1}(j)$ such that $x_{j'}^* \geq x_j^*$.*

Proof. Let \bar{d} be the diameter of L (which is finite because L is connected). Since player i' is more central than player i , Lemma 2 implies that there exists a surjection $\psi : \{j \in L' : d(i', j) \leq \bar{d}\} \rightarrow \{j \in L : d(i, j) \leq \bar{d}\}$ such that, for all j with $d(i, j) \leq \bar{d}$, there exists $j' \in \psi^{-1}(j)$ such that j' is more central than j . By Theorem 4, $x_{j'}^* \geq x_j^*$ for any such j' and j . Finally, $\{j \in L : d(i, j) \leq \bar{d}\} = L$, by definition of \bar{d} . ■

It is easy to see that Theorem 5 applies if $L' \supseteq L$, in which case any surjection $\psi : L' \rightarrow L$ such that $\psi(i) = i$ for all $i \in L$ satisfies the condition of the theorem. This implies the following corollary, which formalizes in a natural way the widespread idea that better-connected societies can provide more public goods.²⁰

Corollary 6 *Adding links to a network weakly increases each player’s robust maximum cooperation.*

However, Theorem 5 is much more general than this. For example, if L' is a circle with n' nodes and L is a circle with n nodes, then Theorem 5 applies whenever $n' \geq n$. Similarly, if L' is a symmetric graph of degree k' on n nodes and L is a symmetric graph of degree k on n nodes, then Theorem 5 applies whenever $k' \geq k$. Finally, the example in Figure 4 shows that Theorem 5 can even apply if L' and L have the same number of nodes *and* the same number of links (here, six and seven, respectively), because a simple application of Lemma 2 and Corollary 4 shows that players 1, 2, 5, and 6 are more central than players 7, 8, 11, and 12, and that players 3 and 4 are more central than players 9 and 10.

²⁰An earlier version of this paper proves that, in addition, adding a link to a network *strictly* increases the robust maximum cooperation of every player in the same component as the added link, if the robust maximum cooperation of every such player is strictly positive.



Figure 4: Comparing Networks with Theorem 5

5.3 Network Stability

This section briefly considers the implications of allowing players to sever links in the monitoring network before the beginning of play. I assume that the resulting equilibrium involves each player making her robust maximum contribution with respect to the remaining monitoring network. I show that, with local public good provision, no player ever has an incentive to sever a link, but that this is not true with global public good provision. Given that adding any link to a monitoring network increases all players' robust maximum cooperation (by Corollary 6), these results suggest that it may be easier to sustain monitoring networks that support high robust maximum cooperation with local public goods than with global public goods.

With local public goods, every player is made worse-off when any link in the monitoring network is severed. This implies that any monitoring network is stable, in that no individual can benefit from severing a link; if players can also add links, then only the complete network is stable. Note that a less restrictive definition of local public goods is needed for this result.

Corollary 7 *Suppose that L is a subnetwork of L' . If $u_i \left((x_j)_{j=1}^n \right) = \left(\sum_{j \in N(i)} f_{i,j}(x_j) \right) - x_i$ for all $i \in N$, then every player i 's payoff when all players make their robust maximum contributions is weakly greater with monitoring network L' than with monitoring network L .*

Proof. Note that (1) simplifies to $x_i^* = \delta \sum_{j \in N(i)} f_{i,j}(x_j^*)$. Therefore,

$$u_i \left((x_j^*)_{j=1}^n \right) = \left(\sum_{j \in N(i)} f_{i,j}(x_j^*) \right) - x_i^* = (1 - \delta) \sum_{j \in N(i)} f_{i,j}(x_j^*). \quad (3)$$

The set $N(i)$ is weakly larger in L' than in L (in the set-inclusion sense), and by Corollary 6 every player's robust maximum cooperation is weakly greater with monitoring network L' than with monitoring network L . Hence, the result follows from (3). ■

Corollary 7 does not hold with global public goods. The key difference between global and local public goods is that with global public goods a player can benefit from another player's cooperation even if she is not observed by the other player, and in this case her own robust maximum cooperation is lower. Formally, with global public goods (3) becomes

$$u_i \left(\left(x_j^* \right)_{j=1}^n \right) = \sum_{j \neq i} \left(1 - \delta^{d(i,j)} \right) f_{i,j} \left(x_j^* \right).$$

This equation clarifies the tradeoff player i faces when deciding whether to sever a link with player j : severing the link increases $d(i, k)$ for some players $k \in N$, which increases $u_i \left(\left(x_j^* \right)_{j=1}^n \right)$ (by reducing player i 's robust maximum cooperation x_i^*), but also decreases x_k^* for some players $k \in N$, which decreases $u_i \left(\left(x_j^* \right)_{j=1}^n \right)$. It is easy to construct examples where the first effect dominates.

6 Conclusion

This paper studies repeated cooperation games with network monitoring and characterizes the robust maximum equilibrium level of cooperation and its dependence on group size and structure, where the notion of robustness is independence from players' information about others' information (holding fixed players' information about others' actions). The key theorem, which underlies all the results in the paper, is that robust maximum cooperation can be sustained by grim trigger strategies. This theorem is driven by an intuitive—but subtle—strategic complementarity between any two players' actions at any two on-path histories. With equal monitoring, robust maximum cooperation is typically increasing in group size with pure public goods and decreasing in group size with divisible public goods. In general, comparative statics on robust maximum cooperation depend on the product of the marginal benefit of cooperation and the effective contagiousness of the monitoring technology, which is thus identified as the property of a monitoring technology that determines how much cooperation it can support. Less uncertain monitoring, which in some cases may be interpreted as reliable local monitoring rather than unreliable public monitoring, supports greater robust maximum cooperation. With a fixed monitoring network, a new notion of network centrality is developed, under which more central players have greater robust maximum cooperation. In addition, all players have greater robust maximum cooperation when the network is better connected, although better connected networks are more likely to be stable with local public goods than with global public goods.

I conclude by discussing some limitations of the current paper and the prospects for addressing them in future research. First, it would be interesting to compare the equilibrium that sustains

robust maximum cooperation with equilibria that satisfy other desiderata, such as fairness (as payoffs in the equilibrium that sustains maximum robust cooperation can be highly asymmetric). It would also be interesting—though challenging—to try to say something about the set of payoffs that can be robustly sustained, or the set of payoffs that can be sustained with particular higher-order information structures.

Second, the assumption that players learn about actions only through network monitoring, rather than also observing public signals about aggregate outcomes, is very strong. For example, villagers certainly observe something about the quality of their schools and wells directly, in addition to observing others' contributions to their maintenance. An immediate observation here is that grim trigger strategy equilibria in games with only network monitoring continue to be equilibria when noisy public monitoring of aggregate outcomes (e.g., the sum of players' actions plus noise) is added to the game, as it is optimal for each player to ignore the public signal if everyone else does. In general, more cooperation could be sustained by using the public signal as well as information from the network, so grim trigger would no longer be optimal; however, a natural conjecture is that grim trigger is approximately optimal when the public signal is very noisy. Studying such a model in detail could lead to insights about the interaction of public monitoring and network monitoring.

Third, while grim trigger strategies are robust in terms of the higher-order information structure, they are fragile in that one instance of shirking eventually leads to the complete breakdown of cooperation. This is especially problematic in (realistic) cases where the cost of cooperation is stochastic and is sometimes prohibitively high. Hence, extending the model to allow for stochastic costs of cooperation is important for deriving yet more robust predictions about which strategies best sustain cooperation, and also seems to be an interesting and challenging problem from a theoretical perspective.

Finally, my analysis makes strong predictions about the effects of group size and structure on the level of public good provision, and on how these differ depending on whether the public good is pure or divisible and whether it is global or local. A natural next step would be to study these predictions empirically, either experimentally (as in the literature surveyed by Ledyard (1997)) or with detailed field data of the kind that is increasingly being collected by development economists (e.g., Karlan et al, 2009; Banerjee et al, 2011).

Appendix

Details of Section 3.1 Example. First, consider grim trigger strategies. As when higher-order information is perfect, player i 's maximum cooperation in grim trigger strategies equals x_i , where $(x_i)_{i=1}^3$ is the greatest vector satisfying

$$x_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \Pr(j \in D(t, i)) \sqrt{x_j} \text{ for all } i.$$

This may be rewritten as

$$\begin{aligned} x_1 &= \frac{\delta/2}{1 - \delta/2} \sqrt{x_2} + \delta \left(\frac{\delta/2}{1 - \delta/2} \right) \sqrt{x_3} \\ x_2 &= \delta \sqrt{x_3} \\ x_3 &= \delta \sqrt{x_2}. \end{aligned}$$

Solving this system of equations with $\delta = 1/2$ yields $x_1 = x_2 = x_3 = 0.25$.

Next, consider the strategy profile from Section 3.1. To define x_2^L , x_2^H , x_3^L , and x_3^H , I first define the former three numbers in terms of x_3^H . Let $\sqrt{x_3^L} \equiv \sqrt{x_3^H} - \frac{1}{10}$. This gap between x_3^L and x_3^H differentiates the resulting strategy profile from a grim trigger strategy profile. Next, I want player 2 to be indifferent among actions 0, x_2^L , and x_2^H at every on-path history, which is the case if

$$\delta \sqrt{x_3^H} - x_2^H = \delta \sqrt{x_3^L} - x_2^L = 0.$$

In order to satisfy this condition, let $x_2^H \equiv \frac{1}{2} \sqrt{x_3^H}$ and $x_2^L \equiv \frac{1}{2} \left(\sqrt{x_3^H} - \frac{1}{10} \right)$.

Given these definitions of x_2^L , x_2^H , and x_3^L in terms of x_3^H , I define x_3^H to be the number that makes player 3 indifferent between actions x_3^H and 0 after he sees player 2 play x_2^H in an odd-numbered period $t - 1$; intuitively, this is the binding incentive constraint for player 3 because the fact that player 2 plays x_2^H in period $t - 1$ is a signal that he observed player 1 in period $t - 2$, in which case he plays x_2^H with probability 1 in period t and thus requires player 3 to play x_3^H in period $t + 1$ in addition to t . To compute this number, note that player 1's future play does not depend on player 3's action, so player 3 is indifferent between playing x_3^H and 0 if and only if $(1 - \delta) x_3^H$ equals the difference in player 3's continuation value following actions x_3^H and 0, excluding player 1's actions. Clearly, this continuation value equals 0 after action 0, as players 2 and 3 play 0 in every period after player 3 plays 0. To compute this continuation value after action x_3^H , note that the probability that player 2 observed player 1's action in period $t - 2$ conditional on his playing

x_2^H in period $t - 1$ equals $\frac{1/2}{1/2+(1/2)(1/2)} = 2/3$. Therefore, player 3's assessment of the probability that player 2 plays x_2^H in period t equals

$$\frac{2}{3}(1) + \frac{1}{3}\left(\frac{1}{2}\right) = 5/6.$$

In contrast, player 3's assessment of the probability that player 2 plays x_2^H in every period $\tau \geq t + 1$ equals $\frac{1}{2}(1) + \frac{1}{2}\left(\frac{1}{2}\right) = 3/4$. Hence, since player 3's assessment of the probability that he himself plays x_3^H in period $\tau + 1$ equals his assessment of the probability that player 2 plays x_2^H in period τ , for all τ , his continuation value after playing x_3^H in t equals

$$\begin{aligned} & \delta \left((1 - \delta) \left(\frac{5}{6}(-x_3^H) + \frac{1}{6}(-x_3^L) \right) + \frac{3}{4} \left(\sqrt{x_2^H} - \delta x_3^H \right) + \frac{1}{4} \left(\sqrt{x_2^L} - \delta x_3^L \right) \right) \\ = & \frac{1}{2} \left(\begin{aligned} & \frac{1}{2} \left(-\frac{5}{6}x_3^H - \frac{1}{6} \left(\sqrt{x_3^H} - \frac{1}{10} \right)^2 \right) \\ & + \frac{3}{4} \left(\sqrt{\sqrt{x_3^H}/2} - \frac{1}{2}x_3^H \right) + \frac{1}{4} \left(\sqrt{\left(\sqrt{x_3^H} - \frac{1}{10} \right)/2} - \frac{1}{2} \left(\sqrt{x_3^H} - \frac{1}{10} \right)^2 \right) \end{aligned} \right). \quad (4) \end{aligned}$$

Define x_3^H to be the number such that $(1 - \frac{1}{2})x_3^H$ equals (4). Computing this number yields $x_3^H \approx 0.25384$, and thus $x_3^L \approx 0.16307$, $x_2^H \approx 0.25191$, and $x_2^L \approx 0.20191$.

It remains to show that this strategy profile is a PBE. The one-shot deviation principle applies, by standard arguments. Player 2 is indifferent among actions 0, x_2^L , and x_2^H at every on-path history, and clearly weakly prefers to play 0 at every off-path history, so he has no profitable one-shot deviation (as any other deviation yields a lower stage-game payoff and a weakly lower continuation payoff than does $x_2 = 0$). It is also straightforward to verify that the fact that player 3 has no profitable deviation after seeing player 2 play x_2^H in an odd-numbered period implies that he has no profitable deviation at any history; in particular, all other one-shot incentive constraints of player 3's are slack. Finally, player 1's most profitable deviation at any on-path history is playing $x_1 = 0$. If player 1 conforms in period t , for any $t \geq 1$, her expected payoff equals

$$\frac{3}{4} \left(\sqrt{x_2^H} + \sqrt{x_3^H} \right) + \frac{1}{4} \left(\sqrt{x_2^L} + \sqrt{x_3^L} \right) - 0.2505 \approx 0.71709.$$

If player 1 deviates to $x_1 = 0$ in an odd-numbered period, her expected payoff may be shown to equal

$$\begin{aligned} & (1 - \delta) \left(\frac{1}{4} \right) \left(\left(1 + \frac{\delta}{2} \right) \left(\sqrt{x_2^H} - \sqrt{x_2^L} \right) + \left(1 + \left(\frac{\delta}{2} \right)^2 \right) \left(\sqrt{x_3^H} - \sqrt{x_3^L} \right) \right) \\ & + \frac{1 - \delta}{1 - \delta/2} \left(\frac{1}{2} \sqrt{x_2^H} + \frac{1}{2} \sqrt{x_2^L} + \left(1 + \frac{\delta}{2} \right) \left(\frac{1}{2} \sqrt{x_3^H} + \frac{1}{2} \sqrt{x_3^L} \right) \right) \approx 0.71676. \end{aligned}$$

If player 1 deviates to $x_1 = 0$ in an even-numbered period $t \geq 2$, her expected payoff is strictly less than this; intuitively, this is because if player 1's period- t deviation is unobserved, player 2 plays x_2^H in period $t + 1$ with probability $3/4$ if t is odd but plays x_2^H with probability only $1/2$ if t is even. In addition, it is clear that the difference between player 1's expected payoff from conforming and from deviating to $x_1 = 0$ in period $t = 0$ is the same as the difference between her expected payoff from conforming and from deviating to $x_1 = 0$ in any other even-numbered period. Therefore, player 1 does not have a profitable deviation at any on-path history. Finally, it can be verified that deviating to $x_1 = 0.2505$ is not profitable for player 1 at any off-path history, and it is clear that no other off-path deviation is profitable. ■

Proof of Lemma 1. Let (Y^P, π^P) denote a perfect higher-order information structure. There are three steps. Step 1 shows that there exists a (component-wise) greatest vector $(\hat{x}_i)_{i=1}^n$ satisfying (1), and also makes the technical point (used in Step 2d) that there exists an upper bound $\bar{X} \in \mathbb{R}_+$ on any player's expected action, conditional on any set of monitoring realizations, at any time in any PBE. Step 2 shows that \hat{x}_i is an upper bound on player i 's maximum cooperation, $x_i^*(Y^P, \pi^P)$. Step 3 exhibits a PBE in grim trigger strategies, σ^* , such that $(1 - \delta) \sum_{t=0}^{\infty} \delta^t \mathbb{E}[\sigma_i^*(h_i^t)] = \hat{x}_i$ for all i , which proves that $x_i^*(Y^P, \pi^P) = \hat{x}_i$ for all i .

Step 1a: There exists a number $\bar{X} \in \mathbb{R}_+$ such that for every $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$, player i , time t , and set of monitoring realizations up to time t , F , $\mathbb{E}[\sigma_i(h_i^t) | F] \leq \bar{X}$ and $\sum_{j \neq i} f_{i,j}(\bar{X}) - \bar{X} < 0$.

Proof: Recall that $(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{i=1}^n \mathbb{E}\left[u_i\left(\left(\sigma_j(h_j^\tau)\right)_{j=1}^n\right) | F\right]$ is well-defined for all sets of monitoring realizations, F , by assumption. The assumptions that $f_{j,i}$ is concave for all j, i and $\lim_{x_i \rightarrow \infty} \left(\sum_{j \neq i} f_{j,i}(x_i)\right) - x_i = -\infty$ for all i imply that there exists a number $x_i^{FB} \in \mathbb{R}_+$ that maximizes $\left(\sum_{j \neq i} f_{j,i}(x_i)\right) - x_i$. For every player i , let $\bar{X}'_i \in \mathbb{R}_+$ be the number such that the sum of the players' continuation payoffs from period t onward equals 0 when player i plays \bar{X}'_i in period t , every player $j \neq i$ plays x_j^{FB} in period t , and every player j (including player i) plays x_j^{FB} in every subsequent period; that is, \bar{X}'_i is defined by

$$(1 - \delta) \left(\left(\sum_{j \neq i} f_{i,j}(x_j^{FB}) \right) - \bar{X}'_i + \sum_{j \neq i} \left(\left(\sum_{k \notin \{i,j\}} f_{j,k}(x_k^*) \right) + f_{j,i}(\bar{X}'_i) - x_j^{FB} \right) \right) + \delta \left(\sum_{j \in N} \left(\sum_{k \neq j} f_{j,k}(x_k^{FB}) \right) - x_j^{FB} \right) = 0.$$

Let $\bar{X}' \equiv \max_{i \in N} \bar{X}'_i$. Note that $(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j=1}^n \mathbb{E}\left[u_j\left(\left(\sigma_k(h_k^\tau)\right)_{k=1}^n\right) | F\right] < 0$ whenever $\mathbb{E}[\sigma_i(h_i^t) | F] > \bar{X}'$ for some player i . Now if h_i^t is a history with perfect higher-order information,

then h_j^t determines the monitoring realization up to time t , and thus whether F has occurred, so $(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [u_j ((\sigma_k (h_k^\tau))_{k=1}^n) | F] = \mathbb{E} [(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [u_j ((\sigma_k (h_k^\tau))_{k=1}^n) | h_j^t] | F]$. Therefore, if $\mathbb{E} [\sigma_i (h_i^t) | F] > \bar{X}'$ then there exists a player j such that $(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [u_j ((\sigma_k (h_k^\tau))_{k=1}^n) | h_j^t] < 0$ with positive probability, conditional on F . But $(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [u_j ((\sigma_k (h_k^\tau))_{k=1}^n) | h_j^t] \geq 0$ at any on-path history h_j^t , because $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$ and 0 is each player's minmax value. Hence, $\mathbb{E} [\sigma_i (h_i^t) | F] \leq \bar{X}'$ for every player i .

Finally, for every player i , the assumption that $\lim_{x_1 \rightarrow \infty} (\sum_{j \neq i} f_{i,j}(x_1)) - x_1 < 0$ (combined with concavity of $f_{i,j}$) implies there exists a number $\bar{X}_i \in \mathbb{R}_+$ such that $(\sum_{j \neq i} f_{i,j}(x_1)) - x_1 < 0$ for all $x_1 \geq \bar{X}_i$. Taking $\bar{X} \equiv \max \{\bar{X}', \max_{i \in N} \bar{X}_i\}$ completes the proof.

Step 1b: There exists a greatest vector $(\hat{x}_i)_{i=1}^n$ satisfying (1), and $\hat{x}_i \leq \bar{X}$ for all i .

Proof: Define the function $\phi : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ by

$$\phi_i \left((x_j)_{j=1}^n \right) \equiv (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \Pr(j \in D(t, i)) f_{i,j}(x_j) \text{ for all } i.$$

The fixed points of ϕ are precisely those vectors satisfying (1). Observe that ϕ is isotone. In addition, $\left((1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \Pr(j \in D(t, i)) f_{i,j}(\bar{X}) \right) - \bar{X} \leq 0$ for every player i , which implies that $\phi \left((\bar{X})_{j=1}^n \right) \leq (\bar{X})_{j=1}^n$. Hence, the image of the set $[0, \bar{X}]^n$ under ϕ is contained in $[0, \bar{X}]^n$. Therefore, Tarski's fixed point theorem implies that ϕ has a greatest fixed point $(\hat{x}_i)_{i=1}^n$ in the set $[0, \bar{X}]^n$. Finally, if $x_i > \bar{X}$ for some player i , then there exists a player j (possibly equal to i) such that $\left((1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k \neq j} \Pr(k \in D(t, j)) f_{j,k}(x_k) \right) - x_j < 0$. This implies that every fixed point of ϕ must lie in the set $[0, \bar{X}]^n$, and it follows that $(\hat{x}_i)_{i=1}^n$ is the greatest vector satisfying (1).²¹

Step 2a: If $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$, then for every player i and every on-path history h_i^t ,

$$(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [\sigma_i (h_i^\tau) | h_i^t] \leq (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i)) \mathbb{E} [f_{i,j}(\sigma_j (h_j^\tau)) | h_i^t, j \in D(\tau, t, i)]. \quad (5)$$

Proof: Fix strategy profile σ , player i , and on-path history h_i^t . For any player j and history h_j^τ , let $\mathbb{E} [f_{i,j}(\sigma_j (h_j^\tau)) | h_i^t, 0]$ be the expectation of $f_{i,j}(\sigma_j (h_j^\tau))$ conditional on each player $k \neq i$ following σ_k , player i following σ_i at every time $\tau < t$, history h_i^t being reached, and player i playing

²¹As an aside, note that the vector $(\hat{x}_i)_{i=1}^n$ (which by Theorem 1 equals $(x_i^*)_{i=1}^n$) may be easily computed by iterating ϕ on $(\bar{X})_{j=1}^n$. Thus, computing the vector of robust maximum cooperation is like computing the greatest equilibrium in a supermodular game (cf Milgrom and Roberts (1990)).

$x_i = 0$ at every time $\tau \geq t$. If $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$, then player i 's expected payoff from conforming to σ from h_i^t onward is weakly greater than her expected payoff from playing $x_i = 0$ at every time $\tau \geq t$. That is,

$$(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} \left[\left(\sum_{j \neq i} f_{i,j}(\sigma_j(h_j^\tau)) \right) - \sigma_i(h_i^\tau) \mid h_i^t \right] \geq (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} \left[\left(\sum_{j \neq i} f_{i,j}(\sigma_j(h_j^\tau)) \right) \mid h_i^t, 0 \right],$$

or, equivalently,

$$(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [\sigma_i(h_i^\tau) \mid h_i^t] \leq (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} (\mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t] - \mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t, 0]). \quad (6)$$

Observe that, conditional on the event $j \notin D(\tau, t, i)$, the probability distribution over histories h_j^τ does not depend on player i 's actions following history h_i^t . Therefore,

$$\mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t, j \notin D(\tau, t, i)] = \mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t, 0, j \notin D(\tau, t, i)].$$

Hence, the right-hand side of (6) equals

$$(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i)) (\mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t, j \in D(\tau, t, i)] - \mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t, 0, j \in D(\tau, t, i)]),$$

which is not more than the right-hand side of (5). Therefore, the fact that (6) holds for all players i and on-path histories h_i^t implies that (5) holds for all players i and on-path histories h_i^t .

Step 2b: For every player i , define the random variable X_i^t by

$$X_i^t \equiv (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sigma_i(h_i^\tau).$$

The right-hand side of (5) is not more than

$$\sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau - 1, t, i)) f_{i,j}(\mathbb{E} [X_j^\tau \mid h_i^t, j \in D(\tau, t, i) \setminus D(\tau - 1, t, i)]).^{22} \quad (7)$$

Proof: Fix a player j . To simplify notation, define the random variable $X_{i,j}^t$ by $X_{i,j}^t \equiv (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} f_{i,j}(\sigma_j(h_j^\tau))$; this notation is used only in this step of the proof. Note that

$$(1 - \delta) \mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) \mid h_i^t, j \in D(\tau, t, i)] = \mathbb{E} [X_{i,j}^\tau \mid h_i^t, j \in D(\tau, t, i)] - \delta \mathbb{E} [X_{i,j}^{\tau+1} \mid h_i^t, j \in D(\tau, t, i)].$$

Therefore,

$$\begin{aligned}
& (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \Pr(j \in D(\tau, t, i)) \mathbb{E} [f_{i,j}(\sigma_j(h_j^\tau)) | h_i^t, j \in D(\tau, t, i)] \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \Pr(j \in D(\tau, t, i)) \left(\mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau, t, i)] - \delta \mathbb{E} [X_{i,j}^{\tau+1} | h_i^t, j \in D(\tau, t, i)] \right) \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left(\begin{aligned} & \Pr(j \in D(\tau, t, i)) \mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau, t, i)] \\ & - \Pr(j \in D(\tau-1, t, i)) \mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau-1, t, i)] \end{aligned} \right) \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left(\begin{aligned} & \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) \mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] \\ & + \Pr(j \in D(\tau-1, t, i)) \mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau-1, t, i)] \\ & - \Pr(j \in D(\tau-1, t, i)) \mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau-1, t, i)] \end{aligned} \right) \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) \mathbb{E} [X_{i,j}^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) \mathbb{E} \left[(1 - \delta) \sum_{s=\tau}^{\infty} \delta^{s-\tau} f_{i,j}(\sigma_j(h_j^s)) | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i) \right] \\
&\leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} \left[(1 - \delta) \sum_{s=\tau}^{\infty} \delta^{s-\tau} \sigma_j(h_j^s) | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i) \right] \right) \\
&= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} [X_j^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] \right).
\end{aligned}$$

where the second equality uses the fact that $\Pr(j \in D(t-1, t, i)) = 0$, the third equality uses the fact that $D(\tau-1, t, i) \subseteq D(\tau, t, i)$, and the inequality uses concavity of $f_{i,j}$ and Jensen's inequality. Summing over $j \neq i$ completes the proof.

Step 2c: If $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$, then for every player i , time t , and subset of monitoring realizations up to time t , F ,

$$\begin{aligned}
& \mathbb{E} [X_i^t | F] \\
&\leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} [X_j^\tau | j \in D(\tau, t, i) \setminus D(\tau-1, t, i), F] \right)
\end{aligned}$$

Proof: If $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$, then (5) and Step 2b imply that, for every player i and every on-path history h_i^t ,

$$\mathbb{E} [X_i^t | h_i^t] \leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} [X_j^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] \right).$$

Thus, by concavity of $f_{i,j}$, and Jensen's inequality,

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E} [X_i^t | h_i^t] | F \right] \\ & \leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} [X_j^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] | F \right). \end{aligned}$$

Finally, perfect higher-order information implies that h_i^t (or h_j^t) determines the monitoring realization up to time t (and thus whether the events $j \in D(\tau, t, i) \setminus D(\tau-1, t, i)$ and F have occurred), so

$$\mathbb{E} \left[\mathbb{E} [X_i^t | h_i^t] | F \right] = \mathbb{E} [X_i^t | F],$$

and

$$\mathbb{E} \left[\mathbb{E} [X_j^\tau | h_i^t, j \in D(\tau, t, i) \setminus D(\tau-1, t, i)] | F \right] = \mathbb{E} [X_j^\tau | j \in D(\tau, t, i) \setminus D(\tau-1, t, i), F].$$

Step 2d: If $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$, then $\mathbb{E} [X_i^0] \leq \hat{x}_i$. In addition, if $\sigma \in \Sigma_{PBE}(Y^P, \pi^P)$ and $\mathbb{E} [X_i^0] = \hat{x}_i$ for all i , then $\sigma_i(h_i^t) = \hat{x}_i$ for every player i and on-path history h_i^t .

Proof: Define x_i^m recursively, for all $m \in \mathbb{N}$, by letting $x_i^1 \equiv \bar{X}$ and letting $x_i^{m+1} \equiv \phi_i \left(\left(x_j^m \right)_{j=1}^n \right)$ for all i . I first claim that $\mathbb{E} [X_i^t | F] \leq x_i^m$ for every player i , time t , subset of monitoring realizations up to time t , F , and number $m \in \mathbb{N}$. The proof is by induction on m . For $m = 1$, the result follows because $\mathbb{E} [\sigma_i(h_i^\tau) | F] \leq \bar{X}$ for all $\tau \geq t$, by Step 1a, and therefore $\mathbb{E} [X_i^t | F] = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathbb{E} [\sigma_i(h_i^\tau) | F] \leq \bar{X}$. Suppose the result is proved for some $m \in \mathbb{N}$. Then

$$\begin{aligned} & \mathbb{E} [X_i^t | F] \\ & \leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} \left(\mathbb{E} [X_j^\tau | j \in D(\tau, t, i) \setminus D(\tau-1, t, i), F] \right) \\ & \leq \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau-1, t, i)) f_{i,j} (x_j^m) \\ & = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i)) f_{i,j} (x_j^m) \\ & = x_i^{m+1}, \end{aligned}$$

where the first inequality follows by Step 2c and the second inequality follows by the inductive hypothesis.

Since ϕ is isotone, $\hat{x}_i \leq x_i^m$ for all $m \in \mathbb{N}$, and in addition $\hat{x}_i \leq \phi_i(\lim_{m \rightarrow \infty} x_i^m)$. Also, ϕ is continuous, which implies that $\phi_i(\lim_{m \rightarrow \infty} x_i^m) = \lim_{m \rightarrow \infty} x_i^m$. The fact that \hat{x} is the greatest fixed point of ϕ thus implies that $\hat{x} = \lim_{m \rightarrow \infty} x^m$. Therefore, the fact that $\mathbb{E}[X_i^t|F] \leq x_i^m$ for all $m \in \mathbb{N}$ implies that $\mathbb{E}[X_i^t|F] \leq \hat{x}_i$. Taking $t = 0$ and $F = \emptyset$ yields $\mathbb{E}[X_i^0] \leq \hat{x}_i$.

Step 3: Let σ^* be the strategy profile given by $\sigma_i^*(h_i^t) = \hat{x}_i$ if $z_{i,j,\tau} \in \{\hat{x}_j, \emptyset\}$ for all $z_{i,j,\tau} \in h_i^t$, and $\sigma_i^*(h_i^t) = 0$ otherwise, for all i . Then $\sigma^* \in \Sigma_{PBE}(Y^P, \pi^P)$, and $\mathbb{E}[X_i^0] = \hat{x}_i$ for all i .

Proof: It is immediate that $\mathbb{E}[X_i^0] = \hat{x}_i$ for all i . To see that $\sigma^* \in \Sigma_{PBE}(Y^P, \pi^P)$, note that the one-shot deviation principle applies, by standard arguments. I first show that no player has a profitable one-shot deviation at any on-path history, and then show that no player has a profitable one-shot deviation at any off-path history.

Fix a player i and an on-path history h_i^t . If $\hat{x}_i = 0$, then it is clear that player i does not have a profitable deviation at h_i^t . So suppose that $\hat{x}_i > 0$. Player i 's continuation payoff if she conforms to σ^* equals $\sum_{j \neq i} f_{i,j}(\hat{x}_j) - \hat{x}_i$. The most profitable deviation from σ^* is playing $x_i = 0$, as every other deviation yields the same continuation payoff and a lower stage-game payoff. I claim that player i 's continuation payoff (including period t) after such a deviation equals

$$(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \notin D(\tau, t, i)) f_{i,j}(\hat{x}_j). \quad (8)$$

Given this claim, the difference between player i 's payoff from conforming to σ^* and from playing her most profitable deviation equals

$$(1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i)) f_{i,j}(\hat{x}_j) - \hat{x}_i,$$

which equals 0 because the vector $(\hat{x}_i)_{i=1}^n$ satisfies (1). Therefore, to show that player i has no profitable deviation, it suffices to prove that player i 's continuation payoff after playing $x_i = 0$ at on-path history h_i^t equals (8).

If player i deviates from σ^* at on-path history h_i^t and $j \notin D(\tau, t, i)$ for some player j and time τ , then $\sigma_j^*(h_j^\tau) = \hat{x}_j$. Hence, the claim that player i 's continuation payoff equals (8) is equivalent to the claim that $\sigma_j^*(h_j^\tau) = 0$ whenever $j \in D(\tau, t, i)$ and $\Pr(j \in D(\tau, t, i)) > 0$. Thus, suppose that player i plays $x_i = 0$ at on-path history h_i^t , that $\Pr(j \in D(\tau, t, i)) > 0$, and that the monitoring realization up to time τ , L^τ , is such that $j \in D(\tau, t, i)$ given L^τ and $\Pr((L_s)_{s=0}^\tau = L^\tau) > 0$. I claim that $\sigma_j^*(h_j^\tau) = 0$ given L^τ . This claim is trivial if $\hat{x}_j = 0$, so assume that $\hat{x}_j > 0$. Proceed by induction on τ : If $\tau = t + 1$, then $z_{j,i,t} = 0$ given L^τ , so the fact that $0 \notin \{\hat{x}_i, \emptyset\}$ implies that

$\sigma_j^*(h_j^\tau) = 0$. Suppose that the claim holds for all $\tau \leq \tau_0$, and consider the case where $\tau = \tau_0 + 1$. Since $j \in D(\tau_0 + 1, t, i)$, player j observes the action of some player $k \in D(\tau_0, t, i)$ at time τ_0 given L^τ , and the fact that $\Pr((L_s)_{s=0}^\tau = L^\tau) > 0$ implies that $\Pr(j \in D(\tau_0 + 1, \tau_0, k)) > 0$. Since $\hat{x}_j > 0$, the fact that $\Pr(j \in D(\tau_0 + 1, \tau_0, k)) > 0$ implies that $\hat{x}_k > 0$, by the definition of $(\hat{x}_i)_{i=1}^n$. Therefore, by the inductive hypothesis, $\sigma_k^*(h_j^{\tau_0}) = 0$ given L^τ , and $0 \notin \{\hat{x}_k, \emptyset\}$. Hence, $\sigma_j^*(h_j^\tau) = 0$, completing the proof of the claim.

It remains only to show that no player has a profitable deviation at any off-path history. Intuitively, given that each player i is indifferent between playing $x_i = \hat{x}_i$ and $x_i = 0$ at every on-path history h_i^t , this follows from Ellison's (1994) observation that a player's incentive to cooperate in a grim trigger strategy profile is reduced after a shirking by another player. Formally, for any subset of players $S \subseteq N$, define $D(\tau, t, S)$ by

$$\begin{aligned} D(\tau, t, S) &= \emptyset \text{ if } \tau < t \\ D(t, t, S) &= S \\ D(\tau + 1, t, S) &= \{j : z_{j,k,\tau} = x_{k,\tau} \text{ for some } k \in D(\tau, t, S)\} \text{ if } \tau \geq t; \end{aligned}$$

note that this generalizes the definition of $D(\tau, t, i)$. Fix a player i and an off-path history h_i^t . If player i has a profitable deviation from σ^* at h_i^t , it must be playing $x_i = \hat{x}_i$, as all other actions yield the same continuation payoff as $x_i = 0$ and a strictly lower stage game payoff. By a similar argument to that in the previous two paragraphs, if $\tilde{D}(t)$ is set of players such that $z_{i,j,\tau} \notin \{\hat{x}_j, \emptyset\}$ for some $z_{i,j,\tau} \in h_i^t$, then the difference between player i 's payoff from conforming to σ^* and her payoff from deviating to $x_i = \hat{x}_i$ (and subsequently following σ^*) equals

$$\begin{aligned} & \hat{x}_i - \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr \left(j \in D(\tau, t, \tilde{D}(t)) \setminus \left(D(\tau, t, \tilde{D}(t) \setminus \{i\}) \cup D(\tau, t+1, i) \right) \right) f_{i,j}(\hat{x}_j) \\ &= \hat{x}_i - \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr \left(j \in D(\tau, t, i) \setminus \left(D(\tau, t+1, i) \cup \tilde{D}(t) \right) \right) f_{i,j}(\hat{x}_j) \\ &\geq \hat{x}_i - \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i) \setminus D(\tau, t+1, i)) f_{i,j}(\hat{x}_j) \\ &= \hat{x}_i - \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} (\Pr(j \in D(\tau, t, i)) - \Pr(j \in D(\tau, t+1, i))) f_{i,j}(\hat{x}_j) \\ &= \hat{x}_i - \frac{1}{1-\delta} \hat{x}_i + \frac{\delta}{1-\delta} \hat{x}_i = 0, \end{aligned}$$

where the last equality follows because $\sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t, i)) f_{i,j}(\hat{x}_j) = \hat{x}_i / (1 - \delta)$ and $\sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j \neq i} \Pr(j \in D(\tau, t+1, i)) f_{i,j}(\hat{x}_j) = \delta \hat{x}_i / (1 - \delta)$. Hence, player i does not have a

profitable deviation at history h_i^t for any set $\tilde{D}(t)$, and therefore player i does not have a profitable deviation at history h_i^t for any belief about the vector of private histories $\left(h_j^t\right)_{j=1}^n$. ■

Proof of Theorem 1. Let σ^* be as in Lemma 1. The proof of Step 3 of the proof of Lemma 1 applies as written to any higher-order information structure. Therefore, $\sigma^* \in \Sigma_{PBE}(Y, \pi)$ for any (Y, π) . Since $\mathbb{E}[X_i^0] = \hat{x}_i$ for all i , it follows by the definition of x_i^* that $x_i^* \geq \hat{x}_i$ for all i . On the other hand, Step 2d of the proof of Lemma 1 implies that $x_i^*(Y^P, \pi^P) \leq \hat{x}_i$ for all i , and therefore $x_i^* \leq \hat{x}_i$ for all i . Hence, $x_i^* = \hat{x}_i$ for all i , and it follows that σ^* robustly sustains each player's maximum robust cooperation. ■

Proof of Corollary 1. By Theorem 1, it suffices to show that $\hat{x}_i = \hat{x}_j$ for all $i, j \in N$. Let x_i^m be defined as in Step 2d of the proof of Lemma 1. I claim that $x_i^m = x_j^m$ for all $i, j \in N$ and $m \in \mathbb{N}$. The proof is by induction. For $m = 1$, $x_i^1 = x_j^1 = \bar{X}$. Suppose the result is proved for some $m \in \mathbb{N}$; that is, that there exists $x^m \in \mathbb{N}$ such that $x_k^m = x^m$ for all $k \in N$. Then

$$\begin{aligned}
x_i^{m+1} &= (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k \neq i} \Pr(k \in D(t, i)) \alpha_{i,k} f(x_k^m) \quad (\text{by parallel benefit functions}) \\
&= (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k \neq i} \Pr(k \in D(t, i)) \alpha_{i,k} f(x^m) \\
&= (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k \neq j} \Pr(k \in D(t, j)) \alpha_{j,k} f(x^m) \quad (\text{by equal monitoring}) \\
&= (1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{k \neq j} \Pr(k \in D(t, j)) \alpha_{j,k} f(x_k^m) \\
&= x_j^{m+1},
\end{aligned}$$

proving the claim. Finally, $\hat{x}_i = \lim_{m \rightarrow \infty} x_i^m = \lim_{m \rightarrow \infty} x_j^m = \hat{x}_j$. ■

Proof of Theorem 3. Let $g_k^t(k') \equiv \Pr(\#D(t) = k' \mid \#D(0) = k)$, let G_k^t be the corresponding distribution function, and let $\mathbb{E}_{g_k^t}[k'] \equiv \sum_{k'=0}^n k' g_k^t(k')$. By Theorem 2, it suffices to show that $\sum_{t=0}^{\infty} \delta^t \mathbb{E}_{\tilde{g}_1^t}[k'] > \sum_{t=0}^{\infty} \delta^t \mathbb{E}_{g_1^t}[k']$.

I claim that \tilde{G}_k^t strictly second-order stochastically dominates G_k^t for all $t \geq 1$ and $k \in \{1, \dots, n-1\}$, which is equivalent to $\sum_{s=0}^{k'} \tilde{G}_k^t(s) < \sum_{s=0}^{k'} G_k^t(s)$ for all $t \geq 1$ and $k' \in \{k, \dots, n-1\}$.²³ The proof is by induction on t . The $t = 1$ case is the assumption that \tilde{G}_k strictly second-order

²³Note that this claim implies that $\sum_{t=0}^{\infty} \delta^t \mathbb{E}_{\tilde{g}_1^t}[k'] \geq \sum_{t=0}^{\infty} \delta^t \mathbb{E}_{g_1^t}[k']$, but not necessarily that $\sum_{t=0}^{\infty} \delta^t \mathbb{E}_{\tilde{g}_1^t}[k'] > \sum_{t=0}^{\infty} \delta^t \mathbb{E}_{g_1^t}[k']$, which is what needs to be shown.

stochastically dominates G_k . Assume the result is proved for $t - 1$. Then

$$\begin{aligned}
\sum_{s=0}^{k'} \tilde{G}_k^t(s) &= \sum_{s=0}^{k'} \sum_{r=0}^s \tilde{g}_k^{t-1}(r) \tilde{G}_r(s) \\
&= \sum_{r=0}^{k'} \tilde{g}_k^{t-1}(r) \sum_{s=r}^{k'} \tilde{G}_r(s) \\
&= \sum_{r=0}^{k'} \tilde{g}_k^{t-1}(r) \sum_{s=0}^{k'} \tilde{G}_r(s) \\
&< \sum_{r=0}^{k'} \tilde{g}_k^{t-1}(r) \sum_{s=0}^{k'} G_r(s) \\
&< \sum_{r=0}^{k'} g_k^{t-1}(r) \sum_{s=0}^{k'} G_r(s) \\
&= \sum_{s=0}^{k'} G_k^t(s),
\end{aligned}$$

where the first line follows because $\tilde{G}_k^t(s) = \sum_{r=0}^s \tilde{g}_k^{t-1}(r) \tilde{G}_r(s)$, the second line reverses the order of sums, the third line follows because $\tilde{G}_r(s) = 0$ if $s < r$, the fourth line follows because $\tilde{G}_r(s)$ strictly second-order stochastically dominates $G_r(s)$ for all $r \in \{1, \dots, k'\}$, the fifth line follows because $\tilde{G}_k^{t-1}(r)$ strictly second-order stochastically dominates $G_k^{t-1}(r)$ (by the inductive hypothesis) and $\sum_{s=0}^{k'} G_r(s)$ is decreasing and strictly convex in r for $r \in \{0, \dots, k'\}$ (because $G_r(s)$ is decreasing and strictly convex in r for $r = \{0, \dots, s\}$ and $s \in \{0, \dots, n\}$, so the sum of such functions is decreasing and strictly convex in r for $r = \{0, \dots, k'\}$), and the sixth line follows from undoing the rearrangement of the first two lines for $G_k^t(s)$ rather than $\tilde{G}_k^t(s)$. This proves the claim.

Trivially, $\mathbb{E}_{\tilde{g}_1^0}[k'] = \mathbb{E}_{g_1^0}[k'] = 1$, and $\mathbb{E}_{\tilde{g}_1^1}[k'] \geq \mathbb{E}_{g_1^1}[k']$ because \tilde{G}_k second-order stochastically dominates G_k . I now show that that $\mathbb{E}_{\tilde{g}_1^t}[k'] > \mathbb{E}_{g_1^t}[k']$ for all $t \geq 2$. This follows because

$$\begin{aligned}
\mathbb{E}_{\tilde{g}_k^t}[k'] &= \sum_{s=0}^{k'} \tilde{g}_1^{t-1}(s) \mathbb{E}_{\tilde{g}_s^1}[k'] \\
&\geq \sum_{s=0}^{k'} \tilde{g}_1^{t-1}(s) \mathbb{E}_{g_s^1}[k'] \\
&> \sum_{s=0}^{k'} g_1^{t-1}(s) \mathbb{E}_{g_s^1}[k'] \\
&= \mathbb{E}_{g_k^t}[k'],
\end{aligned}$$

where the first line follows by the law of iterated expectation, the second line follows because

$\tilde{G}_s^{t-1}(k')$ second-order stochastically dominates $G_s^{t-1}(k')$ if $t \geq 2$ (by the claim), the third line follows because $\mathbb{E}_{g_s^\perp}[k']$ is increasing and strictly concave in s for $s \in \{0, \dots, n\}$ (since $G_k(k')$ is decreasing and strictly convex in k for $k \in \{0, \dots, k'\}$ and $k' \in \{0, \dots, n\}$) and $\tilde{G}_1^{t-1}(s)$ strictly second-order stochastically dominates $G_1^{t-1}(s)$ (by the claim), and the fourth line follows from undoing the rearrangement of the first line. Summing over t completes the proof. ■

Proof of Lemma 2. If such a surjection exists for all $t \in \mathbb{N}$, taking $t = 0$ implies that player i is more central than player j .

For the converse, I first claim that if player i is s -more central than player j , then player i is $s - 1$ -more central than player j . The proof is by induction on s . If i is 2-more central than j , then for all $t \in \mathbb{N}$ there exists a surjection $\psi : \{k \in N : d(i, k) \leq t\} \rightarrow \{k \in N : d(j, k) \leq t\}$, and therefore $\#\{k \in N : d(i, k) \leq t\} \geq \#\{k \in N : d(j, k) \leq t\}$, so i is 1-more central than j . Suppose that if i' is $s - 1$ -more central than j' then i' is $s - 2$ -more central than j' , for all $i', j' \in N$, and suppose that i is s -more central than j . Then for all $t \in \mathbb{N}$ there exists a surjection $\psi : \{k \in N : d(i, k) \leq t\} \rightarrow \{k \in N : d(j, k) \leq t\}$ such that, for all k with $d(j, k) \leq t$, there exists a $k' \in \psi^{-1}(k)$ such that k' is $s - 1$ -more central than k . By hypothesis, this implies that k' is $s - 2$ -more central than k , which, by the definition of $s - 1$ -more central, implies that i is $s - 1$ -more central than j . This establishes the claim.

The claim shows that, for any players i and k , the set of players k' such that $d(i, k') \leq t$ and k' is s -more central than k is weakly decreasing in s (in the set-inclusion sense). Since the sets $\{k \in N : d(i, k) \leq t\}$ and $\{k \in N : d(j, k) \leq t\}$ are finite, this implies that there exists $\bar{s} \in \mathbb{N}$ such that, for all k with $d(j, k) \leq t$, the set of players k' such that $d(i, k') \leq t$ and k' is s -more central than k is the same for all $s \geq \bar{s}$. Hence, if player i is more central than player j , there exists a surjection $\psi : \{k \in N : d(i, k) \leq t\} \rightarrow \{k \in N : d(j, k) \leq t\}$ such that, for all k with $d(j, k) \leq t$, there exists a player $k' \in \psi^{-1}(k)$ such that k' is s -more central than k for all $s \geq \bar{s}$. By the claim, k' is also $\bar{s} - m$ -more central than k for all m , so k' is s -more central than k for all $s \in \mathbb{N}$ and is therefore more central than k . ■

References

- [1] Abreu, D. (1988), “On the Theory of Infinitely Repeated Games with Discounting,” *Econometrica*, 56, 383-396.
- [2] Ali, S.N. and Miller, D. (2011), “Enforcing Cooperation in Networked Societies,” mimeo.

- [3] Ambrus, A., Möbius, M. and Szeidl, A. (2010), "Consumption Risk-Sharing in Social Networks," mimeo.
- [4] Ballester, C., Calvó-Armengol, A., and Zenou, Y. (2006), "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74, 1403-1417.
- [5] Banerjee, A., Chandrasekhar, A.G., Duflo, E., and Jackson, M.O. (2011), "The Diffusion of Microfinance," mimeo.
- [6] Bendor, J. and Mookherjee, D. (1987), "Institutional Structure and the Logic of Ongoing Collective Action," *American Political Science Review*, 81, 129-154.
- [7] Bendor, J. and Mookherjee, D. (1990), "Norms, Third-Party Sanctions, and Cooperation," *Journal of Law, Economics, and Organization*, 6, 33-63.
- [8] Ben-Porath, E. and Kahneman, M. (1996), "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 281-298.
- [9] Bloch, F., Genicot, G. and Ray, D. (2008), "Informal Insurance in Social Networks," *Journal of Economic Theory*, 143, 36-58.
- [10] Bonacich, P. (1987), "Power and Centrality: A Family of Measures," *American Journal of Sociology*, 92, 1170-1182.
- [11] Bramoullé, Y. and Kranton, R. (2007a), "Public Goods in Networks," *Journal of Economic Theory*, 135, 478-494.
- [12] Bramoullé, Y. and Kranton, R. (2007b), "Risk-Sharing Networks," *Journal of Economic Behavior and Organization*, 64, 275-294.
- [13] Bramoullé, Y., Kranton, R., and D'Amours, M. (2011), "Strategic Interaction and Networks," mimeo.
- [14] Coleman, J. (1990), *Foundations of Social Theory*, Cambridge: Harvard University Press.
- [15] Deb, J. (2009), "Cooperation and Community Responsibility: A Folk Theorem for Random Matching Games with Names," mimeo.
- [16] Ellison, G. (1994), "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61, 567-588.

- [17] Greif, A. (1989), “Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders,” *Journal of Economic History*, 49, 857-882.
- [18] Greif, A. (1993), “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition,” *American Economic Review*, 83, 525-548.
- [19] Greif, A., Milgrom, P., and Weingast, B. (1994), “Coordination, Commitment, and Enforcement: The Case of the Merchant Guild,” *Journal of Political Economy*, 102, 745-776.
- [20] Greif, A. (2006), *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*, Cambridge: Cambridge University Press.
- [21] Haag, M. and Lagunoff, R. (2007), “On the Size and Structure of Group Cooperation,” *Journal of Economic Theory*, 135, 68-89.
- [22] Jackson, M.O. (2008), *Social and Economic Networks*, Princeton: Princeton University Press.
- [23] Jackson, M.O., Rodriguez-Barraquer, T., and Tan, X. (2011), “Social Capital and Social Quilts: Network Patterns of Favor Exchange,” *American Economic Review*, forthcoming.
- [24] Kandori, M. and Matsushima, H. (1998), “Private Observation, Communication, and Collusion,” *Econometrica*, 66, 627-652.
- [25] Karlan, D., Möbius, M., Rosenblat, T. and Szeidl, A. (2009), “Trust and Social Collateral,” *Quarterly Journal of Economics*, 124, 1307-1361.
- [26] Ledyard, J.O. (1997), “Public Goods: A Survey of Experimental Research,” in *The Handbook of Experimental Economics*, Roth, A.E., and Kagel, J.H. eds., Princeton: Princeton University Press.
- [27] Mailath, G.J. and Morris, S. (2002), “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory*, 102, 189-228.
- [28] Mailath, G.J. and Morris, S. (2006), “Coordination Failure in Repeated Games with Almost-Public Monitoring,” *Theoretical Economics*, 1, 311-340.
- [29] Miguel, E., and Gugerty, M.K. (2005), “Ethnic Diversity, Social Sanctions, and Public Goods in Kenya,” *Journal of Public Economics*, 89, 2325-2368.

- [30] Milgrom, P. and Roberts, J. (1990), “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities,” *Econometrica*, 58, 1255-1277.
- [31] Milgrom, P.R., North, D.C., and Weingast, B.R. (1990), “The Role of Institutions in the Revival of Trade: the Law Merchant, Private Judges, and the Champagne Fairs,” *Economics and Politics*, 2, 1-23.
- [32] Olson, M. (1965), *The Logic of Collective Action*, Cambridge: Harvard University Press.
- [33] Ostrom, E. (1990), *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge: Cambridge University Press.
- [34] Pecorino, P. (1999), “The Effect of Group Size on Public Good Provision in a Repeated Game Setting,” *Journal of Public Economics*, 72, 121-134.
- [35] Putnam, R. (1993), *Making Democracy Work: Civic Traditions in Modern Italy*, Princeton: Princeton University Press.
- [36] Renault, J., and Tomala, T. (1998), “Repeated Proximity Games,” *International Journal of Game Theory*, 27, 539-559.
- [37] Sugaya, T. and Takahashi, S. (2011), “Coordination Failure in Repeated Games with Private Monitoring,” mimeo.
- [38] Takahashi, S. (2010), “Community Enforcement when Players Observe Partners’ Past Play,” *Journal of Economic Theory*, 145, 42-62.