Coordinated deterrence from a network perspective

Lawrence R. De Geest*

Abstract

Poaching is a major threat to common-pool resources (CPRs), and coordinating the deterrence of poachers, a public good, is a major challenge. Existing studies on CPR management and protection focus largely on average outcomes, but these can gloss over important details about how groups coordinate to those outcomes. Using data from a CPR experiment I study how insiders coordinate punishment to deter poaching by outsiders. Punishment in each period creates a network of insiders (nodes) sending sanctions (directed, weighted edges) to outsiders or insiders. I recover these networks in two treatments that vary the monitoring of outsiders (Partial Monitoring or Full Monitoring) and then run simulations that explore the emergence and stability of these "deterrence networks". In both treatments I show that non-random coordination emerges, but coordination is still below optimal. While average deterrence is similar across treatments, the network approach further shows that there was greater and more stable coordination in Full Monitoring. Understanding how groups managing CPRs coordinate their actions is crucial to designing effective conservation policies. My result show that coordination can vary with external conditions (e.g., the monitoring of outsiders), even if average outcomes do not.

Keywords: Common-pool resources, poaching, punishment, networks, social dilemmas
JEL Codes: C9, C63, D85, Q2

*Department of Economics, Suffolk University. Correspondence to: `ldegeest@suffolk.edu`
Thank you to Abdul Kidwai, John Miller, Scott Page, John Stranlund, Bill Gibson, John Spraggon, Javier Burroni, Charles Mason and Antonio Sirrianifor helpful comments.

1 Introduction

Poaching is a threat to many common-pool resources (CPR). In Chile for instance, where benthic resources like abalaone are managed by decentralized fishing unions known as Territorial User Rights Fisheries (TURFs), poaching has led to violent conflict, and is regarded as the biggest problem facing conservation (Chávez et al., 2018; Gelcich et al., 2017). TURFs members must coordinate their actions (e.g. monitoring and punishment) to deter poachers. But deterrence is a public good, and since there is an incentive to freeride, groups struggle to coordinate, and deterrence is under-provided (Davis et al., 2017; Stevens et al., 2015).[1]

Laboratory experiments are a good way to study coordinated behavior in CPR management and protection (Ostrom, 2006). De Geest et al. (2017) were the first to study how TURFs ("insiders") deter poachers ("outsiders") under imperfect and imperfect monitoring of poachers. Other studies explore complimentary punishment by a central authority (Chávez et al., 2018) and earned versus randomly assigned property rights (De Geest et al., 2020). In all these studies insiders coordinate the use of punishment to deter outsiders, and by and large, all these studies find that insiders do not deter enough. However, these studies mainly focus on average outcomes, like average punishment, and the average may obscure useful details. We are not only interested in the outcome, we are also interested in *how* insiders got there through coordination or miscoordination.

In this paper I explore coordinated deterrence from a network perspective. The basic idea is that punishment is a network: subjects are nodes, and punishments are directed, weighted edges. This network evolves over the course of an experiment. Since networks have structural properties like centrality and clustering, we can derive a benchmark deterrence network that captures optimal insider coordination, and then compare its properties to those in the empirical networks using novel visualization and simulation techniques.

To explore this idea, I re-analyze the data from the experiment by De Geest et al. (2017). Insiders could either imperfectly monitor (the Partial Monitoring treatment) or perfectly monitor (Perfect Monitoring) outsiders, and only outsiders who were monitored could be punished.[2] The punchline from De Geest et al. (2017) is that insiders in both monitoring treatments on average under-provide deterrence. Here I show that a network approach gives new insights into how groups in both treatments coordinated to deter outsiders, and how conditions like the monitoring of outsiders affect the emergence and stability of coordinated deterrence.

[1] Coordination failures in CPRs with poaching are also documented in (Muchapondwa et al., 2014), Afflerbach et al. (2014) and Ovando et al. (2013).

[2] There is also a zero monitoring treatment, where insiders could not monitor and therefore could not punish outsiders.

Simulations show that deterrence networks in both Partial Monitoring and Full Monitoring are different from random, showing that insiders were able to coordinate their actions to some degree, even if they were also different from the optimal benchmark deterrence network. However, visualizations show less clustering in these networks in Full Monitoring, meaning more insiders participated in deterrence when outsiders were fully monitored. Moreover, simulations of the long-run dynamics of deterrence networks suggest that coordination in Full Monitoring was more stable. In addition, deterrence networks in Full Monitoring were less likely to be in states where insiders were punishing each other, further suggesting that there was better coordination in Full Monitoring, even if that coordination did not lead to the deterrence of outsiders.

The main contribution of this paper is to the growing experimental literature on insider deterrence of outsiders in CPR management. Valuable details about coordination lost in average outcomes can be recovered using the network approach I present in this paper. Punishment is among the most studied institutions in CPR experiments (Anderies et al., 2011; Ostrom, 2006), and the methods presented here can be easily used in any experiment that involves punishment.

Moreover, the methods presented here are general enough to be applied to any kind of exchange between subjects in a CPR experiment (e.g., punishment , communication). Therefore, this paper contributes to the broader CPR and social dilemma literature. Other papers on CPRs study the effects social networks (Barnes et al., 2016; Stevens et al., 2015; Bodin and Crona, 2009), communication networks (Mantilla, 2015), and monitoring networks (Shreedhar et al., 2020), and there is a large literature that studies networks in public goods games Leibbrandt et al. (2015); Bramoullé et al. (2014); Jackson et al. (2008). However, the networks in all those papers are taken as exogenous: they are time invariant during the period of study (e.g., a subject assigned to a monitoring network cannot form a new monitoring edge during the experiment). This paper is the first of its kind to study the *use* of monetary sanctions from a network perspective. To that end, this paper is also the first to study the emergence of endogenous networks in a social dilemma experiment.

The rest of this paper proceeds as follows. Section 2 briefly reviews the experiment from De Geest et al. (2017) before discussing how coordination to deter outsiders can be analyzed as a network. Section 3 visualizes deterrence networks in both monitoring treatments, and then conducts two simulations that explore the emergence and stability of these networks. Section 4 concludes.

2 Data and methods

This paper analyzes punishment data from the CPR game with poaching introduced by De Geest et al. (2017). First I give brief overview of the experiment. Then I explain coordinated deterrence from a network perspective.

2.1 The experiment design from De Geest et al. (2017)

Subjects were randomly sorted into groups of eight and then randomly sorted into types: five subjects were assigned to be "insiders" with the remaining three "outsiders". The main difference between insiders and outsiders in our experiment is that a) insiders could communicate with other insiders while outsiders could communicate with no one and b) insiders could sanction behavior by insiders and outsiders with monetary punishment. To avoid framing effects I used neutral language to describe types ("Group 1" for insiders and "Group 2" for outsiders). Similarly, harvesting and poaching decisions were simply referred to as investments into an account called "The Account". Groups and types remained the same for the duration of the experiment.

Each subject in each period received an endowment to harvest or poach from The Account. The remainder of the endowment was invested in a private account with a fixed return. Payoffs were:

$$\pi_{hk} = T + c(e - g_{hk}) + g_{hk} \left(a - b \left(G_i + G_o\right)\right) \tag{1}$$

where h is the subject's type (insider i or outsider o) and k is her group, $e = 12$ is her endowment, g_{hk} is her appropriation (harvests for insiders, poaching for outsiders) from the CPR, $G_i = \sum_{h=1}^{n_i} g_{hi}$ are aggregate insider harvests by n_i insiders, $G_o = \sum_{h=1}^{n_o} g_{ho}$ are aggregate outsider poaching by n_o outsiders, $T > 0$ is a constant to avoid bankruptcy, c is the payoff to the private account, a influence b returns from harvesting or poaching. The parameters a, b and c satisfy a social dilemma ($a > c > b > 0$ and $0 < b < 1$).

Figure 1 illustrates the stages and treatments. Subjects played the game for fifteen periods. In odd-numbered periods, insiders could communicate with each other through a chat box. Insiders and outsiders then simultaneously chose levels of harvest and poaching. Next, each insider could punish any other insider, and punish any outsider they monitored. I used a 1:3 punishment technology: an insider could pay one experimental dollar (ED) to reduce the payoffs of another individual, insider or outsider, by 3 EDs. Spending on punishment was constrained by initial payoffs, and a target's payoffs could not be reduced below zero.

The two treatments varied the monitoring of outsiders. In Partial Monitoring, the experiment randomly chose a single outsider each period to monitored by (and possibly sanctioned) by insiders. Insiders also observed the *total* level of poaching. In Full Monitoring, insiders could monitor all three outsiders. Insiders always had full monitoring of each other.

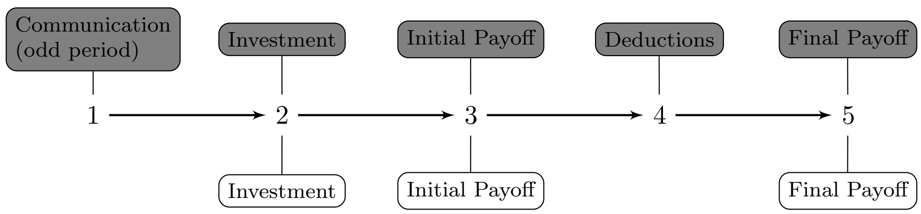

(a) Experiment stages

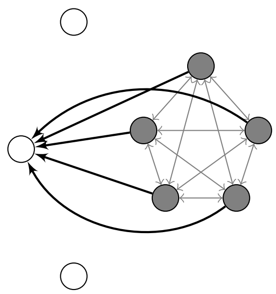

(b) Treatment: Partial Monitoring

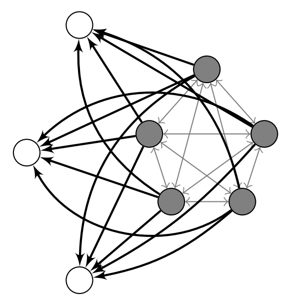

(c) Treatment: Full Monitoring

Figure 1: Experiment stages (a) and treatments (b)-(c). Insiders are in gray, outsiders white. In (b) and (c) paths between nodes (subjects) describe information (monitoring) on appropriations by insiders and poaching by outsiders.

2.2 Coordinated deterrence as a network

De Geest et al. (2017) show that average deterrence of outsiders, across all levels of poaching, was similar in Partial Monitoring and Full Monitoring. Figure 2 breaks down average deterrence by poaching and underscores the similarity across treatments. In both treatments, insiders deter higher levels of poaching, but deterrence is still below one hundred percent, and there are scant differences between the two monitoring treatments. De Geest et al. (2017) conclude from their experiment that insiders do not sufficiently deter outsiders, just like in real-world case studies (Davis et al., 2017).

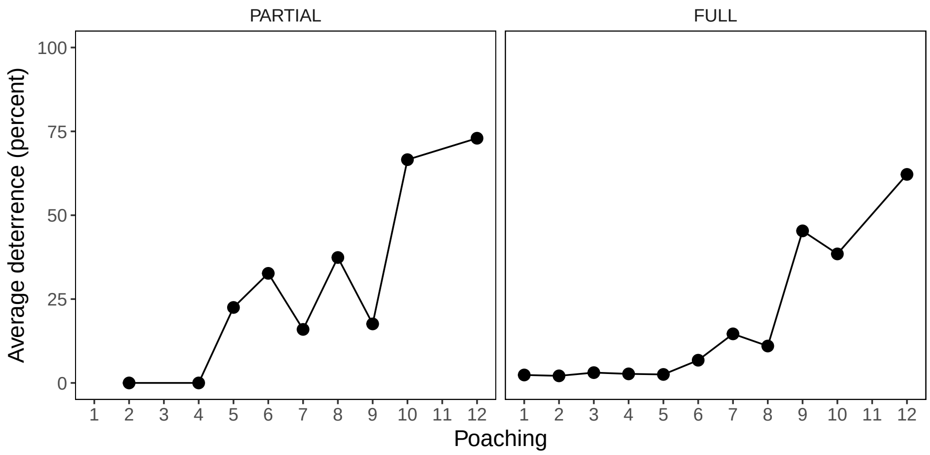

Figure 2: Average deterrence by levels of poaching in De Geest et al. (2017).

However, average deterrence may hide useful details. We are not only interested in the outcome, we are also interested in how insiders coordinated to that outcome. This paper shows how to study coordinated deterrence from a network perspective.

Networks are nodes and edges. Each node in Figure 1 is a subject, and each edge shows would could monitor who, and by definition, who could punish who. In each period, an insider chose whether to punish a target (i.e. form an edge) and how much (i.e. the weight of that edge). So, we can think of deterrence a problem of network formation: insiders coordinate to find the minimum configuration of edges (punishment instances) and weights (punishment values) to deter outsiders.[3]

The benefit of studying coordination as a network is that we can derive a simple benchmark that describes optimal coordination and compare it to what insiders actually did.

From Equation 1 we can write an outsider's payoffs from poaching as $\pi_o - (T + ce)$. The experiment treatments vary how many outsiders can be monitored. When $n_m \leq n_o$ outsiders are monitored, deterrence means choosing punishment $\overline{p}$ that eliminates all gains from poaching, $\overline{p} = \frac{n_{ho}(\pi_{ho} - (T+ce))}{n_m}$. Coordinating deterrence therefore means insiders coordinating to sanction outsiders such that $\overline{p} = \sum_{i=1}^{n_1} p_i$.

Notice that deterrence is a linear public good. This means that the distribution of deterrence contributions that maximizes joint insider payoffs is $\frac{\overline{p}}{n_i}$. We can represent this as a network: a directed star with n_i nodes and $n_i \times n_o$ edges, shown in Figure 3. This is our benchmark.

[3] Here is a more formal description. In each period, a punishment network is realized from a hypergraph (Berge and Minieka, 1973) – the set of feasible edges (sanctions) between a fixed set of nodes (subjects). For a given treatment let $G_k = (V, E, w)$ be the hypergraph of group k, where V is the set of nodes (subjects) indexed by insiders $i \in \{1, \ldots, 5\}$ and outsiders $o \in \{1, \ldots, 3\}$, E is the set of edges (sanctioning events) indexed by sanctions on insiders e_i and outsiders e_o, and w are the corresponding edge weights (sanction sizes). A punishment network in period t is simply a realization of $G_{kt} \in T$.

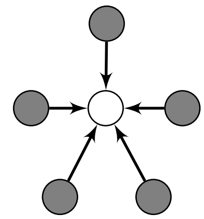

Figure 3: An optimal deterrence network with one outsider. If π_j is the outsider's payoff from poaching, then the value of each edge e_{ij} is $\frac{\pi_j}{n_i}$. This network has a reciprocity and transitivity of zero, and $n_i = 5$ distinct communities. Note that these properties also hold for Full Monitoring, since each targeted outsider will essentially be the core of a distinct star network.

All networks have structural properties (Jackson et al., 2008). I use properties from the benchmark network to study coordination by insiders to deter outsiders.

For starters, some nodes may have more edges and thus more influence; this is network centrality. The benchmark deterrence network in Figure 3 shows that under optimal deterrence, insiders all contribute the same to deterrence, meaning no insider has a high centrality score. Outsiders can also be central in the deterrence network if they are on the receiving end of edges (sanctions). These outsiders should be the largest poachers. In the analysis I plot the aggregate deterrence networks (summed up across periods) to examine centrality.

In addition, nodes in a network can cluster together. Two simple ways to study clustering are reciprocity and transitivity.[4]

The basic idea behind reciprocity and transitivity is that they capture insiders punishing each other and thus devoting resources away from deterring outsiders. Reciprocity $\in [0, 1]$ captures clustering at the dyadic level by calculating the total number of reciprocal dyads (e.g. A punishes B and B punishes A). Since insiders were the only ones who could assign sanctions, reciprocity by definition only occurs between insiders, so in the benchmark network this value will be zero. Transitivity $\in [0, 1]$ captures clustering by calculating the total number of closed triangles divided by the total number of triples. Since a closed triangle must involve at least one insider punishing another, this again will equal zero in the benchmark network. To study reciprocity and transitivity for each group in each treatment I use simulations to see where insider coordination falls between random and optimal.

Finally, studying deterrence as networks allows us to simulate long-run dynamics. These simulations give us a sense of how stable coordination was across treatments. I conclude the results section with a simulation that explores the stability of deterrence networks.[5]

[4]In the appendix (Section B) I also look at community detection (clustering greater than three subjects). Communities refers to partitions of the network into unique subsets of nodes (subjects) and directed edges (sanctions). The number of distinct communities in the benchmark deterrence network is equal to the number of insiders. The results confirm what we see in terms of reciprocity and transitively.

[5]Data and replication files can be found at https://github.com/lrdegeest/CoordinatedDeterrence.

3 Results

3.1 Emergent coordination

I begin by looking at where groups in each treatment lie between total coordination and total miscoordination (i.e., random deterrence).

The basic idea is to simulate deterrence networks and compare their structure to the observed networks. Reciprocity and transitivity are simple and convenient measures of network structure. Coordinating deterrence can fail when insiders devote too much punishment at each other. Reciprocity and transitivity capture this. In the benchmark network insiders focus all punishment on outsiders, so reciprocity (two insiders sanctioning each other) and transitivity (three insiders sanctioning each other) are zero.

To start, I calculate the reciprocity and transitivity score for each group and each treatment. Then to check their statistical significance I make two simulations.

In the first simulation, I simulate 1000 random deterrence networks for Partial Monitoring (one outsider randomly monitored each sweep of the simulation) and Full Monitoring (all three outsiders monitored). Deterrence is entirely random: each insider flips a coin to decide whether to punish an outsider. This is a straightforward application of the Erdős and Rényi (1960) random graph simulation, modified only so that a) outsiders cannot form edges (i.e., punish) and b) edge formation of outsiders varies with Partial and Full Monitoring.[6] The simulation returns a Beta distribution of reciprocity and transitivity scores over the range $[0, 1]$ with an expected value of 0.5.[7] If an observed deterrence network has a transitivity or reciprocity score in the tails of this distribution then we can say it is significantly different from random.

The second simulation includes more information from observed deterrence networks. Specifically, I simulate random networks that preserve the degree distributions across groups. This accounts for the fact that some subjects are more likely to punish (in-degree) and others more likely to get punished (out-degree).

I present the results visually. In each plot there are two panels for each group, one for reciprocity and the other for transitivity. Each group has a plot for reciprocity and transitivity. The vertical black line indicates the observed reciprocity or transitivity for a group's aggregate network (Figures 6 and 7). The simulated networks that *do not* preserve degree distribution are plotted in green. The simulated networks that *do* preserve degree distribution are plotted in blue. Expected values are shown with vertical lines. These visuals makes the analysis straightforward: a significant effect will be illustrated by a black line (the

[6] My algorithm to simulate directed random graphs with constraints is in the appendix.

[7] For $Beta(\alpha, \beta), \alpha = \beta$ the expected value is $\mu = \frac{1}{1+\frac{\alpha}{\beta}}$.

observed value) in the tails of the simulated, random distribution (the shaded green area).

Figure 4 shows the results for Partial Monitoring. Nearly all groups show significant transitivity and reciprocity scores. In other words, the structure of their deterrence networks are non-random. This is evidence that groups were able to coordinate deterrence. It is also important to note the variation between groups: some exhibit zero reciprocity and transitivity (like the benchmark network), others do not.

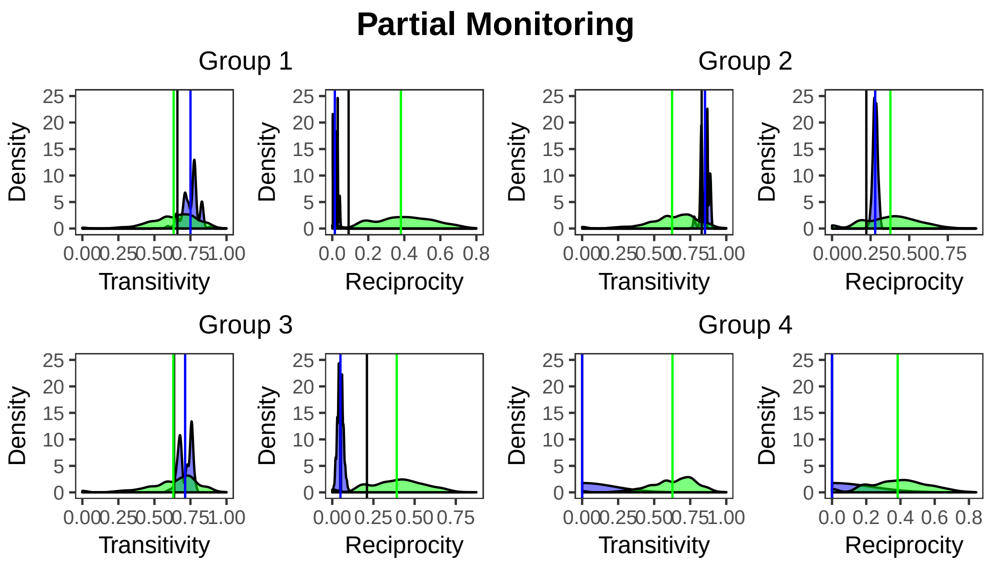

Figure 4: Simulated networks: Partial Monitoring. The black line is the observed value in the data. The green densities come from the randomly simulated networks; the green line is the expected value. The blue densities come from the simulated networks for which degree distribution is preserved; the blue line is the expected value.

Figure 4 shows the results for Full Monitoring. Once again, we see largely significant results that suggest that coordination to deter outsiders was non-random.

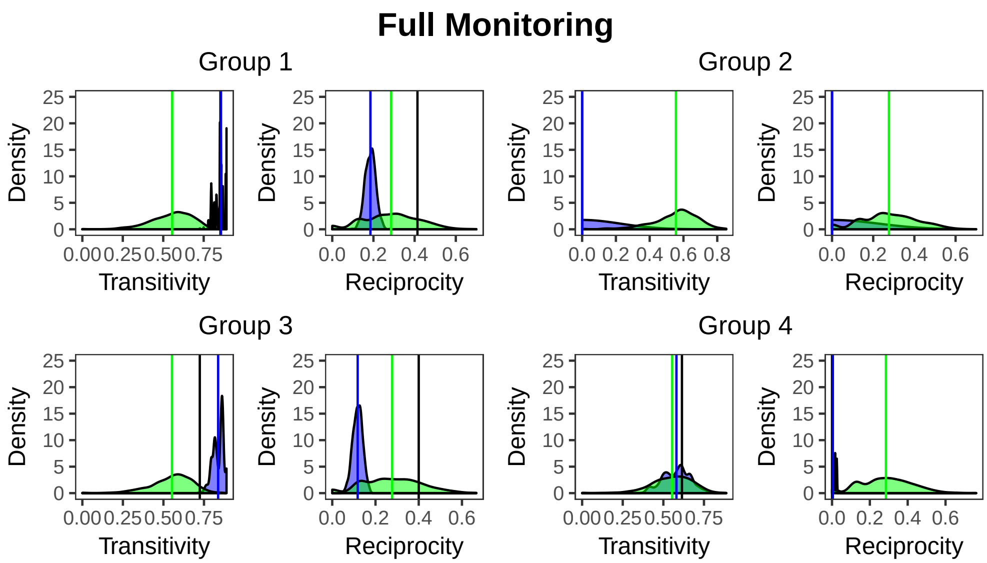

Figure 5: Simulated networks: Full Monitoring. The black line is the observed value in the data. The green densities come from the randomly simulated networks; the green line is the expected value. The blue densities come from the simulated networks for which degree distribution is preserved; the blue line is the expected value.

To summarize, these simulations return a straightforward and useful result. Nearly all deterrence networks in both treatments are different from the benchmark network, indicating that optimal coordination was not acheived. However, deterrence networks are also different from random, indicating that insiders were able to coordinate. In the next section I look more closely at these non-random deterrence networks to study treatment differences.

3.2 Coordination structure

Figure 6 shows the networks for Partial Monitoring and Figure 7 shows the networks for Full Monitoring. For each group in each treatment I sum E_k and w_k across all periods. Insiders are denoted by circles, outsiders by squares, and node size is determined by a subject's total appropriation from the CPR. Nodes are decorated with three standard network centrality measures to highlight the structure of the networks. Green nodes represent the insider with the highest out-degree score: the insider who punished the most. Blue nodes represent the subject with the highest in-degree score: the subject (insider or outsider) who received the most punishment. Finally, subjects with the highest eigenvector centrality are trimmed in red. This simply means subjects who receive more sanctions will receive them from subjects who send more sanctions.[8]

[8]Eigenvector centrality refers to a class of centrality measures that calculate the principle eigenvector for a specific eigenvalue problem of a network's adjacency matrix. The most common is Bonacich centrality, $\mathbf{Ax} = \lambda_1 \mathbf{x}$, where $\mathbf{Ax}$ is the adjacency matrix (the nxn matrix of weighted edges between subjects), x

Coordinating is not easy, so we should not expect the observed networks to perfectly mirror the benchmark deterrence network. Nevertheless, the first result that stands out from Figures 6 and 7 is how clustered the deterrence networks are. In the benchmark network no insider is more central than another. Yet in both treatments deterrence is largely provided by a few insiders, and these insiders tend to harvest less from the CPR: in nearly each group in both treatments the insider with the highest out-degree score (the green node) is the smallest.

Who received the most punishment? In Full Monitoring it was the outsider who poached the most. That is encouraging: insiders were able to direct punishment where it was most useful. In Partial Monitoring – where one outsider each period was randomly chosen to be monitored – it is messier. In two groups we see the same result as in Full Monitoring: the outsider who poached the most is colored blue (highest in-degree centrality) and trimmed red (highest eigenvector centrality). But that is not the case in Group 1. And in Group 2, it was an insider who bore the brunt of punishment. (Group 2 was particularly dysfunctional: the same insider also had the highest out-degree score.)

is the vector of eigenvector centralities and λ_1 is the eigenvector corresponding to the largest eigenvalue (i.e. the principal eigenvector). We use a modification of this measure known as PageRank that performs better for sparse, directed graphs by scaling the centrality measure by the subject's out-degree or provided sanctions, k_i^{out}. This avoids exaggerating the centrality of a subject simply because she is receiving sanctions by a single high-centrality subject. In the aggregate enforcement networks there is one only case where the highest out-degree subject is not the highest eigenvector central subject, Group 1 in Partial Monitoring. The leading out-degree subject (47 total sanctions, PageRank score of 0.18) has a more skewed distribution of received sanctions from each insider compared to the leading eigenvector centrality subject (43 total sanctions, PageRank score of 0.20).

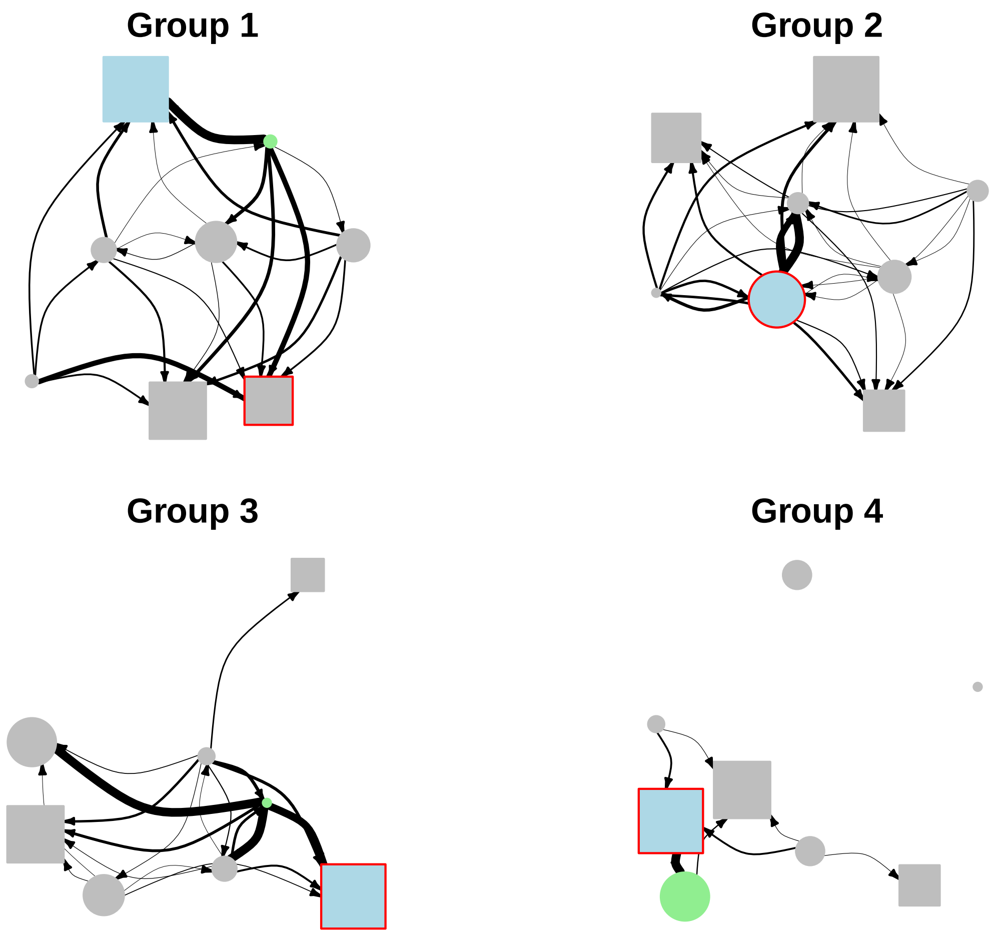

Figure 6: Aggregate group enforcement networks: Partial Monitoring. Insiders are round nodes, outsiders are square nodes. Node size represents a subject's aggregate level of harvest or poaching. Node color represents highest in-degree (blue) and out-degree (green) centrality. Red trim indicates the node with the highest eigenvector centrality. Directed edges indicate sanctions; edge size represents aggregate sanctions.

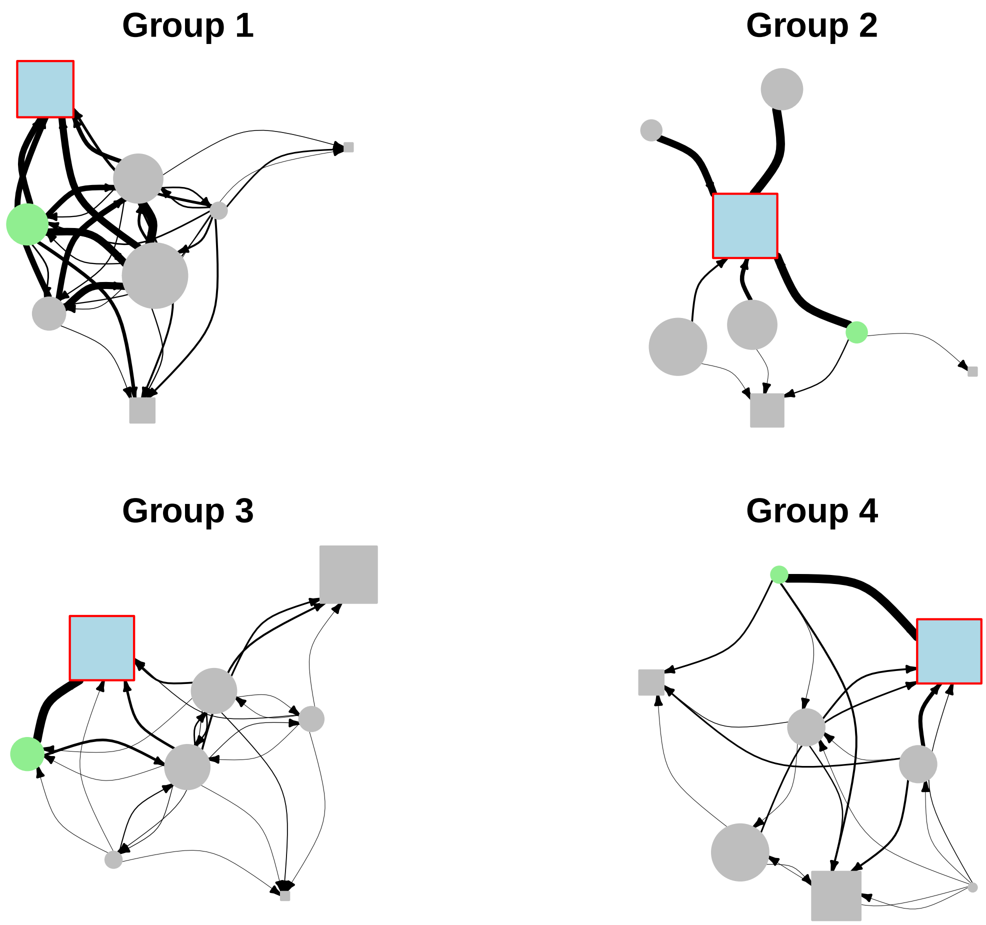

Figure 7: Aggregate group enforcement networks: Full Monitoring. Insiders are round nodes, outsiders are square nodes. Node size represents a subject's aggregate level of harvest or poaching. Node color represents highest in-degree (blue) and out-degree (green) centrality. Red trim indicates the node with the highest eigenvector centrality. Directed edges indicate sanctions; edge size represents aggregate sanctions.

The observed networks show that not all insiders contributed deterrence equally. But visual inspection does suggest that more insiders contributed in Full Monitoring. To take a closer look, Figure 8 shows how the percent each insider in each group contributed towards total deterrence. The top panels show contributions to deterrence, the bottom panels show enforcement on fellow insiders, and each bar is a subject, ordered from highest to lowest contributor. Now we see a clear pattern. Contributions to deterrence were top-heavy in Partial Monitoring: one or two insiders contributed most of it. By contrast, the distribution seems to flatten in Full Monitoring. More insiders join the deterrence effort when there is more information about outsiders. With a few exceptions, this also seems to be true in the bottom panels (sanctions on fellow insiders).

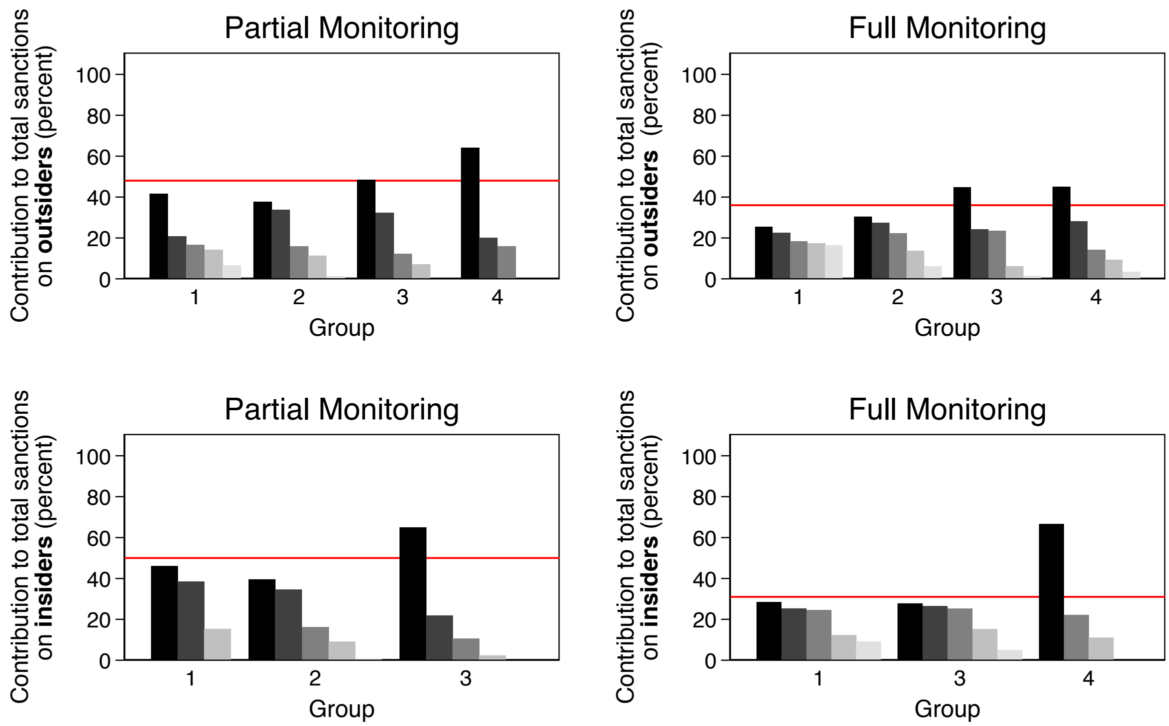

There were no sanctions on insiders in Partial Monitoring Group 4 and Full Monitoring Group 2.

Figure 8: Individual contributions to punishment. Insiders are organized by groups (1-4) across treatments. Each bar represents an insider in that group. The horizontal line in each graph indicates the average contribution by the leader insider – the insider who contributed the most.

So far we have seen clustering at the individual level: deterrence is mostly provided by a few insiders, but increased of monitoring outsiders see more insiders contribute deterrence. This suggest that more information on outsiders leads to improved coordination by insiders to deter them. In the next and final section I look at stable was coordination to deter outsiders across treatments.

3.3 Coordination stability

By studying coordinated deterrence as a problem of network formation, I can explore how stable were the networks that emerged in each group and treatment. The upshot of this simulation is that there is a clear affect of information on coordination. When insiders can perfectly monitor outsiders, deterrence networks are more stable, compared to when insiders imperfectly monitor outsiders.

To simulate stability I consider how deterrence networks transition between states. A state of the network in some period could be "punish only outsiders", "punish only insiders" or some combination, or neither. This gets to the heart of the coordination problem: insiders must coordinate to turn punishment on outsiders, but each insider can also punish group members, or simply defect from enforcement altogether. Formally, we can think of this as an ergodic Markov process, which when simulated has a unique stationary distribution describing the fraction of time spent in each state.[9] So for each group, I can derive an estimate of the probability that the punishment network is in a particular state, and thus get a sense of the stability of each group's coordination.[10]

To set-up the simulation I define the following four states that capture how insiders coordinate to allocate sanctions between in-group enforcement and out-group deterrence. The states are shown in the table below. This creates a four-state Markov chain X_k where X_{kt} denotes the position of the chain in time t.

Table 1: State space of the simulated enforcement networks.

State	Description
$x_1 : s_{in} = 0, s_{out} = 0$	No enforcement
$x_2 : s_{in} > 0, s_{out} = 0$	Enforcement on insiders
$x_3 : s_{in} = 0, s_{out} > 0$	Enforcement on outsiders
$x_4 : s_{in} > 0, s_{out} > 0$	Enforcement on insiders and outsiders

Transition probabilities between states come from the experiment data. I created indicator variable for each state, and in each period I mark which state the network was in. Then I counted the transitions from each state i to each state j and created a square matrix ($n \times n$, where n is the number of states), where element $\{i, j\}$ describes the number of instances the network moved from i to j across all fifteen periods of the experiment. Next I created the transition matrix M_k for group k by normalizing the rows so they summed to one. Each element of M is then defined by $P(i, j)$, the one-period probability of moving from i to j. This led to a list of four ergodic transition matrices $\mathbb{M} = \{M_1, M_2, M_3, M_4\}$ that each have a unique stationary state distribution $\Psi^* = \Psi^* M^t \ \forall \ t$ describing the long-run probability the chain is in a given state. This distribution can be estimated by simulating the marginal probabilities of X_k^t for sufficiently large t.[11]

[9] A Markov chain that is aperiodic and irreducible is ergodic. Aperiodic means the chain does not predictably cycle. Irreducibility means each state can eventually be reached every other state. See Häggström (2002) for a proof of the existence of a unique stationary state distribution in ergodic chains.

[10] Choi and Bowles (2007) use a similar approach. They statistically recover the Markov process governing an agent-based simulation to estimate the steady-state fraction of agent types in the population.

[11] The simulation begins by assigning the chain to a randomly-chosen state at $t = 0$ from distribution Ψ. For each subsequent t, the next state X_{t+1} is drawn from the distribution over the previous state, $X(S_t)$. If the simulation converges, then the average value of the time series for each state will approximate the stationary distribution. Another way to think about the steady-state is as an eigenvector of M. A steady

Table 2 shows the results for 1000 simulations for each group in each treatment.[12] Each entry represents the long-run probability a group was in a given state. Missing values indicate a group was never in a particular state.

There is a clear pattern across treatments: each group spent at least fifty-percent of the time in one state (with the exception of Group 1 in Partial Monitoring). There are also interesting differences between treatments. Two of the most visited states are (enforcement only on outsides) and $x_4 : s_{in} > 0, s_{out} > 0$ (enforcement both on insiders and outsiders). In Full Monitoring groups one a consistently stable state is $x_3 : s_{in} = 0, s_{out} > 0$. This means that networks were more focused on deterrence. By contrast, the state of zero enforcement (x_1) was fairly stable in Partial Monitoring.

Table 2: Estimated stationary states. Each entry represents long-run probability a group was in a given state.

	Partial Monitoring				**Full Monitoring**			
	Group 1	*Group 2*	*Group 3*	*Group 4*	*Group 1*	*Group 2*	*Group 3*	*Group 4*
x_1	0.18	NA	0.57	0.27	0.06	NA	NA	0.08
x_2	0.22	0.07	0.07	NA	0.20	0.21	0.07	0.08
x_3	0.23	0.59	0.07	0.58	0.52	0.71	0.36	0.64
x_4	0.36	0.34	0.29	0.15	0.22	0.07	0.57	0.20

Next I check the expected change in network structure from one period to the next. I calculated the net transition matrix $M^N = M - M^T$, where element $p^N_{ij} = p_{ij} - p_{ji}$ represents the net transition from state i to j. Where $p^N_{ij} > 0$ the transition from i to j is more likely than the other way round. Plotting a contour map of M^N provides a description of how each enforcement network will evolve. The map is like a landscape: positive values indicate troughs and negative values indicate peaks. In troughs, the system will tend to remain put; on peaks, the system will tend to move.

In Figure 9 lighter colors indicate where $p^N > 0$, increasing in brightness with larger positive values. These represent relatively stable spaces in the system, or troughs, while darker colors represent unstable spaces, or peaks. The punchline of the figure is immediately clear: across all deterrence networks there are more stable points in Full Monitoring than in Partial Monitoring. So, when we simulate long run dynamics, we see that full monitoring of poachers not only leads to better coordination among insiders, it also leads to more stable

state occurs when a probability vector multiplied by M returns the same probability vector.

[12] The time series of each simulation are shown in Figure A2 and Figure A3. Each figure plots the first 500 steps. Results show similar convergence times across groups and treatments.

coordination. By contrast, coordination is more volatile in Partial Monitoring.

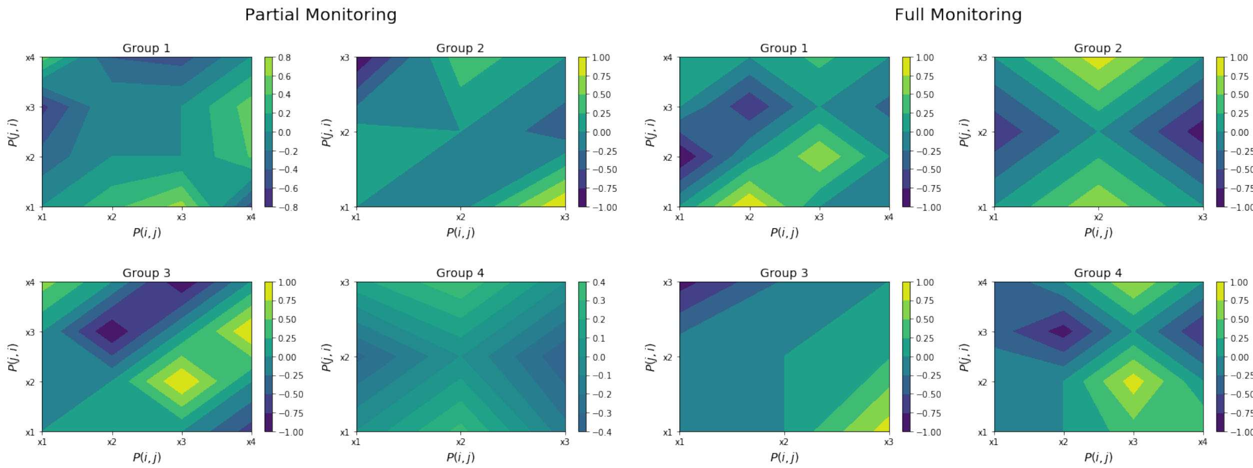

(a) Partial Monitoring

(b) Full Monitoring

Figure 9: Contour plots of net Markov transition matrices.

4 Conclusions

I study the coordinated deterrence of poachers using networks. Data come from the laboratory experiment by De Geest et al. (2017), in which groups of insiders coordinate the use of punishment to deter poaching by outsiders who are imperfectly monitored (Partial Monitoring) or perfectly monitored (Full Monitoring). I show that you deterrence can be modeled as a directed, weighted network that evolves: nodes are insiders and outsiders, edges are instances of punishment, weighted by the punishment magnitude. These networks are compared to a theoretical deterrence network that describes optimal coordination to deter outsiders. De Geest et al. (2017), like many studies, focuses on average outcomes, but we are also interested in coordination – how groups arrive to one outcome or another. Studying coordinated deterrence as a network gives us a better grasp on this question by opening doors to novel visualization and simulation techniques methods.

The network approach yields some new insights into insider coordination deter outsiders. Deterrence networks in both treatments are different from the theoretical benchmark: they are highly clustered, with most deterrence provided by a few insiders. However, the distribution of who deters outsiders flattens out in Full Monitoring, suggesting that more information about poachers increases participation in deterrence. Simulations show that the deterrence networks emerging in each treatment are non-random, a sign that insiders were able to coordinate. But coordination was more volatile in Partial Monitoring, more stable in Full Monitoring.

There are many ways to improve this work. For instance, the simulations used in this paper were fairly simple, though they reveal interesting patterns of emergence across treatments. Agent-based models (Janssen and Ostrom, 2006) could build on this work and explore richer questions about interactions between insiders and outsiders competing over a CPR. Such models could also readily explore a wider variety of monitoring networks.

It would also be interesting to apply the methods presented in this paper to other studies on insider deterrence of poachers. Because monitoring networks in De Geest et al. (2017) were fixed, they imposed tight constraints on how deterrence networks could form and grow. Different configurations of fixed networks that govern how subjects interact (e.g. monitoring networks) will influence the emergence of endogenous networks (e.g., deterrence networks). Outcomes in social dilemmas – like the provision of deterrence, a public good – are sensitive to a variety of factors, all of which could influence how insiders coordinate and arrive to those outcomes.

The main takeaway from this paper is that average outcomes can hide valuable details about coordination. To that end, the methods in this paper can be applied not just to studies about deterring poachers, but to any CPR or social dilemma experiment that involves some type of "exchange" between subjects (e.g., punishment, communication).

References

Afflerbach, J. C., S. E. Lester, D. T. Dougherty, and S. E. Poon (2014). A global survey of "TURF-reserves", Territorial Use Rights for Fisheries coupled with marine reserves. *Global Ecology and Conservation 2*, 97–106.

Anderies, J. M., M. A. Janssen, F. Bousquet, J.-C. Cardenas, D. Castillo, M.-C. Lopez, R. Tobias, B. Vollan, and A. Wutich (2011). The challenge of understanding decisions in experimental studies of common pool resource governance. *Ecological Economics 70*(9), 1571–1579.

Barnes, M. L., J. Lynham, K. Kalberg, and P. Leung (2016). Social networks and environmental outcomes. *Proceedings of the National Academy of Sciences 113*(23), 6466–6471.

Berge, C. and E. Minieka (1973). *Graphs and hypergraphs*, Volume 7. North-Holland publishing company Amsterdam.

Bodin, Ö. and B. I. Crona (2009). The role of social networks in natural resource governance: What relational patterns make a difference? *Global environmental change 19*(3), 366–374.

Bramoullé, Y., R. Kranton, and M. D'amours (2014). Strategic interaction and networks. *The American Economic Review 104*(3), 898–930.

Chávez, C. A., J. J. Murphy, and J. K. Stranlund (2018). Managing and defending the commons: Experimental evidence from TURFs in Chile. *Journal of Environmental Economics and Management 91*, 229–246.

Choi, J.-K. and S. Bowles (2007). The coevolution of parochial altruism and war. *science 318*(5850), 636–640.

Davis, K. J., M. E. Kragt, S. Gelcich, M. Burton, S. Schilizzi, and D. J. Pannell (2017). Why are fishers not enforcing their marine user rights? *Environmental and Resource Economics 67*(4), 661–681.

De Geest, L. R., A. H. Kidwai, and J. Portillo (2020). Ours, not yours: Property-rights, poaching and deterrence in common-pool resources. Working paper.

De Geest, L. R., J. K. Stranlund, and J. M. Spraggon (2017). Deterring poaching of a common pool resource. *Journal of Economic Behavior & Organization 141*(C), 254–276.

Erdős, P. and A. Rényi (1960). On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci 5*(1), 17–60.

Gelcich, S., J. Cinner, C. J. Donlan, S. Tapia-Lewin, N. Godoy, and J. Castilla (2017). Fishers' perceptions on the Chilean coastal turf system after two decades: problems, benefits, and emerging needs. *Bulletin of Marine Science 93*(1), 53–67.

Häggström, O. (2002). *Finite Markov chains and algorithmic applications*, Volume 52. Cambridge University Press.

Jackson, M. O. et al. (2008). *Social and economic networks*, Volume 3. Princeton university press Princeton.

Janssen, M. A. and E. Ostrom (2006). Empirically based, agent-based models. *Ecology and Society 11*(2), 37.

Leibbrandt, A., A. Ramalingam, L. Sääksvuori, and J. M. Walker (2015). Incomplete punishment networks in public goods games: experimental evidence. *Experimental Economics 18*(1), 15–37.

Mantilla, C. (2015). Communication networks in common-pool resource games: Field experimental evidence. *Journal of Economic Behavior & Organization 118*, 215–226.

Muchapondwa, E., K. Brick, and M. Visser (2014). Abalone conservation in the presence of drug use and corruption: implications for its management in south africa. *International Journal of Sustainable Economy 18 6*(2), 201–216.

Ostrom, E. (2006). The value-added of laboratory experiments for the study of institutions and common-pool resources. *Journal of Economic Behavior & Organization 61*(2), 149–163.

Ovando, D. A., R. T. Deacon, S. E. Lester, C. Costello, T. Van Leuvan, K. McIlwain, C. K. Strauss, M. Arbuckle, R. Fujita, S. Gelcich, et al. (2013). Conservation incentives and collective choices in cooperative fisheries. *Marine Policy 37*, 132–140.

Shreedhar, G., A. Tavoni, and C. Marchiori (2020). Monitoring and punishment networks in an experimental common pool resource dilemma. *Environment and Development Economics 25*(1), 66–94.

Stevens, K., K. A. Frank, and D. B. Kramer (2015). Do social networks influence small-scale fishermen's enforcement of sea tenure? *PloS one 10*(3), e0121431.

Zhang, X.-D. (2011). The laplacian eigenvalues of graphs: a survey. *arXiv preprint arXiv:1111.2897*.

A Algorithm to simulate directed random graphs with constraints

---

Algorithm 1: Simulate directed random graphs with constraints.

Input : Nodes (V)
Constraints ($k \mid k < V$)
Probability of forming an edge (p)

Output: Adjacency matrix $\mathbf{A}(i, j)$ forming the directed graph $G(V, E)$

begin
    $\mathbf{A}(i, j) \leftarrow \varnothing$ initialize empty $V \times V$ adjacency matrix
    $\vec{c} \leftarrow \binom{V}{k}$ set constraints vector by sampling without replacement from $\{1, 2, \ldots, V\}$
    for i **to** V **do**
        for j **to** V **do**
            if $i \neq j$ **and** $i \notin \vec{c}$ **then**
                draw $d \sim \mathcal{U}(0, 1)$
                if $d < (1 - p)$ **then**
                    $\mathbf{A}(i, j) \leftarrow 1$
                else
                    $\mathbf{A}(i, j) \leftarrow 0$
                end
            end
        end
    end
    return $\mathbf{A}(i, j)$
end

B Community structure

As a robustness check on the structure of deterrence networks I look at community detection. Communities are defined as connected components in a graph (Jackson et al., 2008). Components can be identified through spectral decomposition of the graph Laplacian (Zhang, 2011). The graph Laplacian is given by $\mathbf{L} = \mathbf{A} - \mathbf{D}$ where $\mathbf{A}$ is the adjacency matrix and $\mathbf{D}$

is a diagonal matrix of the degree sequence of the network, the list of connected, directional pairs (i.e. the degree matrix). The Laplacian matrix is singular and positive semidefinite, ensuring that its eigenvalues satisfy $\lambda_1 \geq \lambda_2 \geq \ldots \lambda_{v-1} \geq \lambda_v = 0$, where v is the number of nodes. The number of positive eigenvalues corresponds to the number of connected components, or partitions, in the graph. So if the deterrence network is optimal (Figure 3), then the number of partitions will equal the number of insiders n_i and the non-zero eigenvalues will be equal.

Results are shown in Figure A1. The axes display the number of possible partitions, simply the number of nodes or subjects, and the eigenvalues of the graph Laplacian. Each line in each treatment corresponds to a group.[13] Results suggest there are significantly more partitions than would be if the networks were randomly generated. At the same time, the networks are clearly different from the benchmark. Moreover, the difference in partitioning between groups is most stark in Partial Monitoring. Consistent with Figure **??**, a few individuals contributed most enforcement, while the targets of enforcement were a subset of individuals.

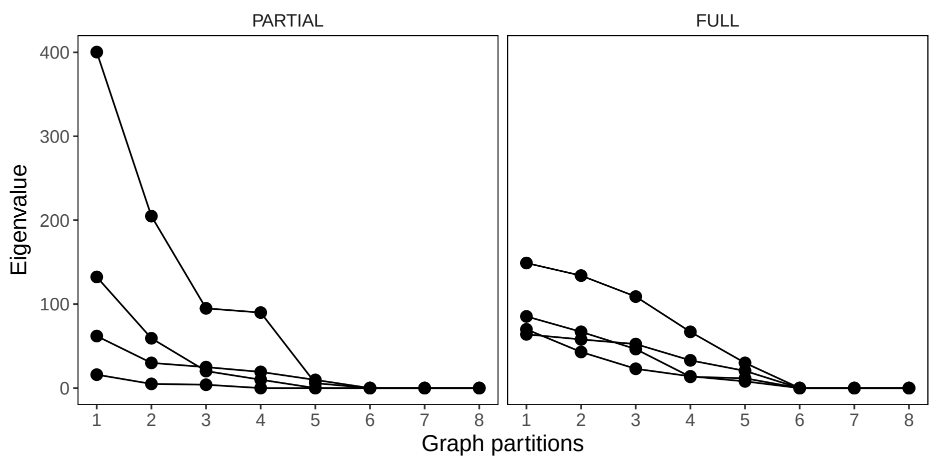

Figure A1: Community detection in aggregate enforcement networks using the Graph Laplacian.

[13]Note that the eigenvalues for indices six through eight are zero in each group. This is because the graphs are directed and constrained by the fact that outsiders could not assign sanctions.

C Markov chain simulation convergence

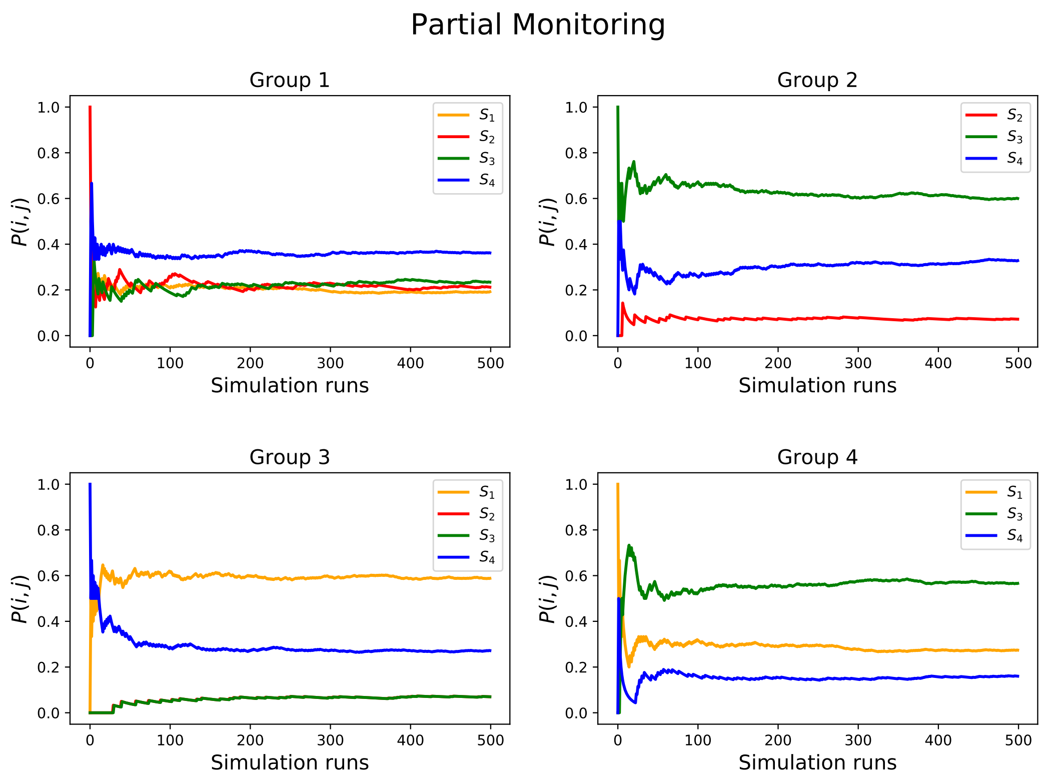

Figure A2: Simulated Markov chains for enforcement networks in Partial Monitoring.

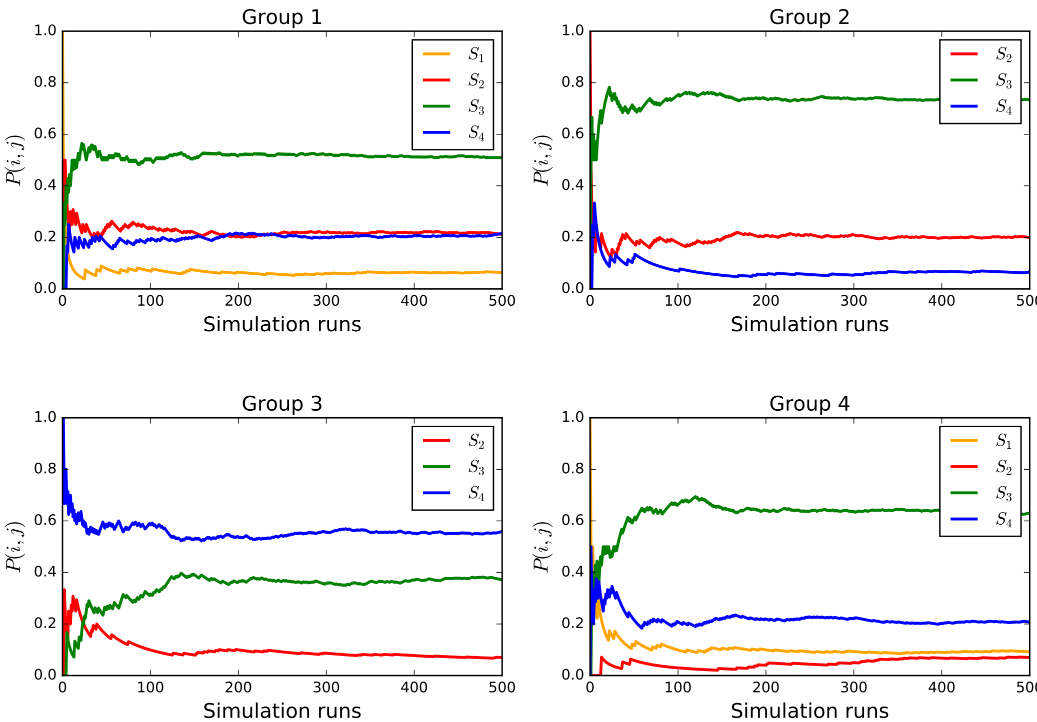

Figure A3: Simulated Markov chains for enforcement networks in Full Monitoring.