

Peer Discipline and Incentives Within Groups[☆]

David K. Levine¹, Salvatore Modica^{2,*}

Abstract

We investigate how a collusive group can sustain non-Nash actions by enforcing internal discipline through costly peer punishment. We give a simple and tractable characterization of schemes that minimize discipline costs while preserving incentive compatibility. We apply the model to a public goods contribution problem. We find that if the per-capita benefit from the public good is low, then regardless of whether peer discipline is feasible or not only small groups will contribute to the good. If the public good benefit is significant but peer discipline is infeasible it remains the case that only small groups contribute. On the other hand, if the public good benefit is significant but peer discipline is feasible then full contribution takes place regardless of group size. We reconcile this result with Olson's idea that small groups are more effective by considering the case where the per-capita benefit of the public good varies with group size.

JEL Classification Numbers:

C72 - Noncooperative Games

D7 - Analysis of Collective Decision-Making

D72 - Political Processes: Rent-Seeking, Lobbying, Elections, Legislatures, and Voting Behavior

Keywords: Group Incentives, Peer Discipline, Organization, Group

[☆]First Version: November 26, 2012. We are especially grateful to Michele Boldrin, Michael Chwe, Drew Fudenberg, Zacharias Maniadis, Eric Maskin, Andy Postelwaite, Debraj Ray and seminar participants at the St. Louis Federal Reserve Bank, the Venice Conference on Economic Theory, LUISS, Helsinki, and Warwick. We are grateful to the EIEF, to NSF Grant SES-08-51315 and to MIUR PRIN 20103S5RN3 for financial support.

*Corresponding author Salvatore Modica, Facoltà di Economia, Viale delle Scienze, 90128 Palermo, Italy

¹Department of Economics, WUSTL and European University Institute; *email*: david@dklevine.com

²Università di Palermo; *email*: salvatore.modica@www.unipa.it

1. Introduction

Groups do not act as individuals - as Olson (1965) and others have emphasized incentives within groups matter. Here we study self-sustaining group discipline that overcomes free-riding problems through costly peer punishment. We investigate schemes that might be adopted by a collusive group to minimize the cost of enforcing actions which are not Nash equilibria in the absence of punishment. We use the model to study how the strength of a group, measured by its members' contributions to a public good, depends on the size of the group.

We describe a group which plays an underlying base game - this might involve contribution to a public good, voting, or any other sort of interaction. Following the game, discipline occurs through a series of audit rounds during which signals of past play are observed and punishment may be issued for observed violations of social norms. The design of the audit rounds is endogenous and is treated as a mechanism design problem for the group. We investigate the type of play that can be sustained in the primitive round and the corresponding enforcement costs.

In this model punishment is intrinsically recursive, for failure to punish is itself punishable. This condition has been found crucial also in field work such as that conducted by Elinor Ostrom and reported for instance in Ostrom (1990) and Ostrom, Walker and Gardner (1992), or the empirical work of Coleman (1988). Indeed, from a conceptual point of view punishment cannot have a definite end. As Juvenal asked in the 2nd Century CE, “Quis custodiet ipsos custodes?” - who will guard the guardians? The answer is that they must guard each other, and for that to be possible the game should not have a definite ending.³ To take this into account the free-riding problem is often embedded in a repeated game model where group actions and (possibly) punishments occur each period, as for example in Pecorino (1999), Wolitzky (2013) or in the experimental work of Fehr and Gächter (2000) and monitoring and communication has often been studied in this setting, for example in Ben-Porath and Kahneman (2003). The repeated game approach has generally been taken in the network literature, for example by Vega-Redondo (2006) or Balmaceda and Escobar (2014).⁴ In practice, however, punishment is usually distinct from the actions of a group, and often it is also the case that the basic actions are not repeated, as for instance when members of a group should provide effort to induce government to implement a policy measure favorable to the group. This leads us to consider the setting of initial base game in which group members take a one-shot action that has both individual and group costs and benefits, followed by a repeated and open-ended game of auditing and punishment. The latter is effectively a “closure of social network”, as described by Coleman (1988). In the peer punishment game group members repeatedly audit each other and determine punishments for signals of bad behavior.⁵ While this does not describe

³A recent paper by Rahman (2012) also raises the issue of who will guard the guardians. However his model is a static model with a principal who has unlimited commitment power and rewards agents for monitoring each other. It is very unlike the Kandori (1992) setting we adopt which also answers the question of who will guard the guardians, but, because players are peers, in a way that addresses the issue as originally raised by Juvenal.

⁴The latter paper is oriented in a somewhat different direction than the others being aimed in part at understanding Burt (1995)'s intuition about structural holes.

⁵Our auditing procedure is simpler than that considered in the auditing literature, for example Border and Sobel

the method used to provide incentives in all groups at all times, the setting is a relevant one for many of groups of practical interest: depending on the group, these punishments for violating social norms may range from social ostracism, as in the case of unions, to physical abuse as in the case of criminal gangs. We should emphasize that while the simple version of the model is a highly structured model of auditing, neither the audits nor the way in which they are conducted needs to be highly structured. When we consider extensions we show that the basic results are robust to the way in which the audits are conducted, and in many practical circumstances audits may be of a relatively informal nature - perhaps more akin to gossip than to audits conducted by trained auditors.

The use of punishments and rewards in one game to induce desired behavior over actions in a different and earlier game is not a new idea. It is the basis of the efficiency wage model of Shapiro and Stiglitz (1984), and also of models of collusion proofness as described, for example, by Laffont (1999). Generally speaking, these models have not had costs associated with enforcement - in the efficiency wage model there is generally no punishment on the equilibrium path, while in the Laffont model punishments and rewards take the form of transfer payments so that there is no net cost. In practice however, there is punishment on the equilibrium path - as in Laffont - but practical forms of punishment, such as exclusion, generally have a net cost associated with them. Moreover it is generally costly to observe the signals that are needed to trigger punishment. In our setup we allow positive punishment costs on the equilibrium path, along which enforcement of first stage punishments is carried out through subsequent rounds of costly auditing rather than through a commitment to carry out punishments.

Our model also formalizes the type of community enforcement through ostracism and other measures as documented in many studies such as Ostrom (1990) and Ostrom, Walker and Gardner (1992) or Baron (2010). Another class of models has been studied theoretically in which cooperation in a group is induced not through observation of other members but through contagion: if anybody deviates from the social norm this can trigger a chain of punishments that results in a social collapse. The idea originates in Ellison (1994) and plays a key role in the analysis of Ali and Miller (2010), Deb and González-Díaz (2010) and in the community enforcement of Acemoglu and Wolitzky (2015).⁶ While no doubt part of social norms are enforced through the threat of social collapse, in practice it seems that targeted punishments are more common. We should also point out the implication originally noted by Ellison that with noise when the group is large a high enough probability of punishment to deter deviation leads very quickly to collapse. This we do not generally observe in large organizations or societies.

The model of repeated auditing that we analyze is a simple adaptation of the Kandori (1992) model of social norms in which the punishers are distinct from the aggrieved.⁷ We adapt the

(1987) and Di Porto, Persico and Sahuguent (2013), because no signal is observed prior to conducting the audit. While those papers focus on a one-shot audit, our focus here is not on optimal methods of conducting audits, but rather how group members monitor each in a repeated setting.

⁶They also consider enforcement by specialists.

⁷See also Takahashi (2010) for a more recent variation of the Kandori model with a more limited information

model in two dimensions: first we allow for noisy signals - this leads to essentially the strongly symmetric computation of Abreu, Pearce and Stachetti (1986).⁸ Second, we view the problem as one of mechanism design for the group: it has control over the basic design of the punishment game, which consists in determining the probability with which the game ends and no further punishments take place. In effect the discount factor becomes endogenous and subject to group choice. The basic trade-off is that if the game ends too quickly insufficient incentives to avoid free-riding are provided; on the other hand if it ends too slowly too much effort is spent in auditing and too much utility is lost from punishments. The other parameters of the peer-punishment game - the basic technology of punishment and monitoring - we view as economic fundamentals over which the group has no control. In this setting our main results characterize the optimal choice for the group. We show first that the optimal discipline scheme maximizes the probability of the game ending each period subject to incentive compatibility. We then study primitive actions, and show that if auditing costs are small the group will accept the possibility of large gains to deviations in exchange for high utility in the base game.

We use the model to examine how peer monitoring may work in the context of a group that faces a linear public good contribution problem. We show that if the social benefit of the public good is low, then - regardless of the availability of peer discipline technology - in small groups full effort is provided but in large groups no effort is provided.⁹ If, on the other hand, the benefit is large then peer discipline makes all the difference. That is, when peer discipline is not feasible it remains the case that only small groups provide effort. However when peer discipline is feasible full effort provision is induced regardless of how large the group is. Thus peer discipline is needed to make large groups effective. This is important because it explains the fact that even in very large groups such as farm lobbies and voting constituencies we observe substantial contributions to public goods. This result is more subtle than the simple Olson (1965) theory that group effectiveness diminishes with size. We discuss further the relation of our results to Olson's insights at the end of Section 4.

The theory we describe is relatively easy to apply to more complicated settings involving the interaction of several groups. The general conceptual issues of equilibrium with interacting collusion groups have been examined in Dutta, Levine and Modica (2015), while a simplified version of the peer monitoring described here has been used in Levine and Mattozzi (2015) to study the classical Palfrey and Rosenthal (1985) model of voter participation.

We initially present a stripped down version of the model and derive the optimal schemes for the group. The stripped down model has a number of limitations. After studying the application we remove many of these restrictions. We show how to incorporate renegotiation, how to allow for general matching technologies, how to introduce social costs of punishment, how to endogenize the

structure.

⁸The structure of peer games is not subject to the types of complications found in more general repeated games as described in Fudenberg Levine and Maskin (1994) and Sugaya (2011).

⁹As will be clear from the model, peer discipline is feasible if adequate punishments and signals about members behavior are available and auditing costs are sufficiently low.

costs of auditing and punishment, and how to incorporate general signaling technologies.

2. The Basic Peer Discipline Model

We first introduce a streamlined model, deferring extensions to Section 5. There are $N > 2$ identical members $i = 1, \dots, N$ of a collusive group. The group plays a primitive game at $t = 0$ which members choose actions $a^i \in A$, where A is finite, individual payoffs being $u(a^i, a^{-i})$. Given a common action of the other members a^R we let $u(a^i, a^R) = u(a^i, a^R, \dots, a^R)$. To avoid triviality and since we have allowed only a finite number of (possibly mixed) actions, we assume that there is at least one symmetric Nash equilibrium of this game, that is, there exists an $a^R \in A$ such that for all $a^i \in A$ we have $u(a^i, a^R) \leq u(a^R, a^R)$.

The focus of this paper is on sustainability of common actions a^R taken at time $t = 0$ through an incentive compatible peer monitoring scheme. We assume that the group has access to a *peer discipline* technology - based on Kandori (1992)'s information systems approach - in which members may audit each others behavior, the audit possibly resulting in a punishment for the auditee. The model accounts for the self-referential nature of punishment equilibria by supposing that the group plays a potentially unlimited number of audit rounds $t = 1, 2, \dots$. Peer monitoring takes place over a simple circular network in which each group member is connected with the member to his left and member 1 is connected to member N . Each auditor chooses whether or not to conduct the audit.

The super-game starts with an initial “meeting” in which group members agree on a scheme to maximize the utility of group members - this makes sense since members are *ex ante* identical, so there is no conflict of interest. At this meeting the group agrees on a common action a^R to be played at $t = 0$, and probabilities δ_t , $t = 0, 1, \dots$ that an audit round will take place at $t + 1$ (with probability $1 - \delta_t$ the game ends after round t). We assume that auditing rounds take place sufficiently quickly so that there is no discounting beyond that induced by δ_t . Initially and for concreteness we assume a simple stylized network structure governing auditing - we show subsequently that our results are robust to the specification of the auditing network. Specifically, if an audit round takes place members are matched in pairs with *auditor* i matched according to the network structure with *auditee* $j = i - 1$. In the initial audit round an auditor who carries out an audit assesses whether the auditee played the chosen action a^R . In subsequent audit rounds an auditor who carries out an audit assesses whether the auditee conducted the audit in the previous round. These assessments are based on a noisy signal seen only by the auditor about the auditees choice in the previous round.

The signal technology is as follows. Actions in the primitive game generate a signal $z_0^j \in \{0, 1\}$ about the play of member j where we interpret 0 as “bad” and 1 as “good.” The probability of bad signal about j is $\pi_0(a^j, a^R)$. In subsequent rounds, depending on whether j in the role of auditor conducted the audit or not, another bad-good signal $z_t^j \in \{0, 1\}$ is generated. If the audit was conducted then the bad signal is generated with probability π ; otherwise with probability $\pi^p \geq \pi$.

When an audit is conducted and the signal $z_{t-1}^j \in \{0, 1\}$ of the behavior of the auditee in the previous round is 0 (interpreted as “bad” behavior) the auditee is punished. This punishment inflicts an exogenously determined cost on the auditee of $P > 0$ while the cost to the auditor of conducting the audit is $\theta_t P \geq 0$, for $t \geq 1$; we assume that $\theta_t = \theta$ for $t > 1$.¹⁰

To emphasize how this works: in the initial audit round the auditor assesses whether the auditee played the “correct” action. In subsequent audit rounds the auditor assesses whether the auditee conducted the audit in the previous round.

Since the group is bound by incentive constraints only incentive compatible plans can be chosen.

Definition 1. A plan $a^R, \delta_t|_{t=0}^\infty$ is *peer feasible* if the individual strategies of playing a^R in the primitive round and always conducting an audit in the audit rounds is a Nash equilibrium¹¹ of the super-game induced by the continuation probabilities δ_t .

At the initial meeting the group may either choose a peer feasible plan, or it may choose a static Nash equilibrium of the primitive game together with $\delta_0 = 0$. Among these plans the group chooses the plan that maximizes the *ex ante* expected utility of the members (who, recall, are *ex ante* identical).

3. Enforceability, Peer Feasibility, and Optimal Group Plans

We study the optimality of group plans in three parts. First, we analyze which actions can be enforced by some punishment scheme. Second, we analyze which actions are peer feasible and the corresponding cost minimizing continuation probabilities. Finally we use these results to characterize the optimal group plan.

3.1. Enforceability

Recall that in the initial primitive round the probability of a “bad” signal 0 is $\pi_0(a^i, a^R)$ and utility is $u(a^i, a^R)$. Following the repeated game literature such as Fudenberg Levine and Maskin (1994) we say that the initial common action a^R is *enforceable* with respect to punishment \bar{P} if punishing the bad signal makes a^R incentive compatible. This means that for all a^i we have $u(a^R, a^R) - \pi_0(a^R, a^R)\bar{P} \geq u(a^i, a^R) - \pi_0(a^i, a^R)\bar{P}$. If for all a^i we have $u(a^i, a^R) - u(a^R, a^R) \leq 0$ we reiterate that a^R is *static Nash*. This case is not terribly interesting since no peer discipline is required to implement it as an outcome.

¹⁰We allow the costs of the initial audit θ_1 to be different than that of subsequent audits θ since determining compliance in the primitive game is different than determining whether an audit was conducted. Note that there is no loss of generality in taking the audit cost proportional to the punishment cost as θ_1, θ are arbitrary.

¹¹We use Nash equilibrium because this is an infinite horizon game with private information, where refinements such as subgame perfection have no bite and the definition and analysis of more suitable refinements such as sequentiality is complicated. However for this class of games refinements do not matter in the generic case in which both signals have positive probability. In this case every information set is reached with positive probability. Hence the Nash equilibrium problem of “off the equilibrium path play” does not arise. It follows that sequential equilibrium or even stronger refinements such as extensive form trembling hand perfect equilibrium are identical to Nash equilibrium.

We now characterize enforceability. Let $\sigma_0(a^i, a^R) \equiv \pi_0(a^i, a^R) - \pi_0(a^R, a^R)$ be the *bad signal increase*. Observe that if $\sigma_0(a^i, a^R) > 0$ we can rewrite the incentive constraint as

$$\tilde{G}(a^i, a^R) \equiv \frac{u(a^i, a^R) - u(a^R, a^R)}{\sigma_0(a^i, a^R)} \leq \bar{P}$$

so that the *gain function* $\tilde{G}(a^i, a^R)$ measures the smallest punishment needed for incentive compatibility. For $\sigma_0(a^i, a^R) < 0$ the inequality is reversed and the gain function measures the largest punishment compatible with incentive compatibility. In case $\sigma_0(a^i, a^R) = 0$, so that a^i is *indistinguishable* from a^R , and $u(a^i, a^R) = u(a^R, a^R)$ we define the gain function to be $\tilde{G}(a^i, a^R) = 0$ since there can be no gain from deviating. If a^i is indistinguishable from a^R and $u(a^i, a^R) \neq u(a^R, a^R)$ the gain function is defined as $\tilde{G}(a^i, a^R) = [u(a^i, a^R) - u(a^R, a^R)] \cdot \infty$ - if the utility increase is negative incentive compatibility is not an issue; if it is positive incentive compatibility is impossible. It is then easy to check that

Lemma 1. *The group action a^R is enforceable with respect to \bar{P} if and only if*

$$\max_{a^i: \sigma_0(a^i, a^R) \geq 0} \tilde{G}(a^i, a^R) \leq \bar{P} \leq \min_{a^i: \sigma_0(a^i, a^R) < 0} \tilde{G}(a^i, a^R)$$

If $\max\{0, \max_{a^i: \sigma_0(a^i, a^R) \geq 0} \tilde{G}(a^i, a^R)\} \leq \min_{a^i: \sigma_0(a^i, a^R) < 0} \tilde{G}(a^i, a^R)$ define

$$G(a^R) \equiv \max\{0, \max_{a^i: \sigma_0(a^i, a^R) \geq 0} \tilde{G}(a^i, a^R)\}$$

otherwise set $G(a^R) = \infty$.¹²

3.2. Peer Feasibility

We first work out the optimal decision making of the group concerning implementation of a particular action a^R . In Section 3.3 below we address the question of which a^R to choose. Define the *audit signal increase* $\sigma = \pi^p - \pi$ and notice that by assumption this is non-negative. The next result says that optimal implementation of action a^R requires minimizing the probabilities δ_t while preserving incentive compatibility.

Theorem 1. *If the action a^R is not static Nash it is peer feasible for some $\delta_t|_{t=0}^\infty$ if and only if $P \geq G(a^R)$, $\theta_1/\sigma \leq 1$ and $\theta/\sigma < 1$, in which case the group optimally chooses the termination probabilities*

$$\delta_0 = G(a^R)/P, \delta_t = \theta_t/\sigma \text{ for } t > 0$$

The corresponding utility attained by each member is

$$U = u(a^R, a^R) - \left(\pi_0(a^R, a^R) + \theta_1 + \frac{\theta_1(\theta + \pi)}{\sigma - \theta} \right) G(a^R).$$

¹²Note that $\min_{a^i: \sigma_0(a^i, a^R) < 0} \tilde{G}(a^i, a^R) \geq 0$ if and only if $\sigma(a^i, a^R) \geq 0$ for all a^i with $u(a^i, a^R) - u(a^R, a^R) > 0$.

Proof. Let $\pi_t = \pi_0(a^R, a^R)$ if $t = 0$ and $\pi_t = \pi$ otherwise. When the representative member plays a^R initially and then always chooses to perform the audit he gets utility

$$U = u(a^R, a^R) - \sum_{t=1}^{\infty} \left(\prod_{\tau=0}^{t-1} \delta_{\tau} \right) (\theta_t + \pi_{t-1})P,$$

which is strictly decreasing in δ_t for each t .

Incentive compatibility in the primitive round is covered by Lemma 1 with $\bar{P} = \delta_0 P$; this gives $G(a^R) \leq \delta_0 P \leq \min_{\sigma_0(a^i, a^R) < 0} \tilde{G}(a^i, a^R)$. From the first inequality and $\delta_0 \leq 1$ we then get the condition $P \geq G(a^R)$; and since the group utility requires us to minimize δ_0 , we see that $\delta_0 = G(a^R)/P$. This also implies $\delta_0 P \leq \min_{\sigma_0(a^i, a^R) < 0} \tilde{G}(a^i, a^R)$.

Consider next the decision by auditor i not to audit in round t . The only consequences of this decision are the saving of the cost $\theta_t P$ and the increased probability of punishment in the subsequent round $\delta_t \sigma P$. The incentive constraint is therefore $\delta_t \sigma \geq \theta_t$. Since $\delta_t \leq 1$ this gives the condition $\theta_t / \sigma \leq 1$. If the incentive constraint holds with strict inequality, then we should lower δ_t , so the optimum requires this constraint hold with exact equality, that is, $\sigma \delta_t = \theta_t$.¹³ Hence $\delta_1 = \theta_1 / \sigma$ and for $t > 1$ we have $\delta_t = \theta / \sigma$. If the cost of auditing is to be finite, the latter must be less than one, giving the condition $\theta / \sigma < 1$. Substituting in we find

$$U = u(a^R, a^R) - \delta_0 P \left((\pi_0 + \theta_1) + (\theta_1 / \sigma) \left[\sum_{t=1}^{\infty} (\theta / \sigma)^{t-1} (\pi + \theta) \right] \right),$$

adding up the geometric series, and substituting for $\delta_0 P$ using $\delta_0 = G(a^R)/P$ from above we get the result. \square

Remark. Observe that we can always increase δ slightly during the audit stage and obtain an equilibrium that is strict in all the audit rounds - the price is a small reduction in group welfare. Such an equilibrium can be more robust as it does not require individuals to “make the right choice” when indifferent. In the first period, by Theorem 1 we can achieve strict incentive compatibility if and only if $\max\{0, \max_{\sigma_0(a^i, a^R) \geq 0} \tilde{G}(a^i, a^R)\} < \min_{\sigma_0(a^i, a^R) < 0} \tilde{G}(a^i, a^R)$.

3.3. The Optimal Group Plan and Comparative Statics

To analyze the choice of primitive group action a^R observe that for given a^R that satisfies $P \geq G(a^R)$, $\theta_1 / \sigma \leq 1$ and $\theta / \sigma < 1$, so is peer feasible, Theorem 1 tells us that the optimum peer-feasible utility U has a very simple form. Define $v(a^R) = u(a^R, a^R) - \pi_0(a^R, a^R)G(a^R)$ to be the *initial net utility* of each member from playing a^R net of the minimum cost of punishment required to make a^R incentive compatible. Define the *unit cost of auditing* as

$$C = \theta_1 + \frac{\theta_1(\theta + \pi)}{\sigma - \theta}$$

¹³As is usual in the optimal punishment literature the key step is identifying the incentive constraint that binds.

Then the optimum peer feasible utility U attainable from action a^R for each member is

$$U = v(a^R) - CG(a^R)$$

In addition to choosing for each a^R an optimal audit procedure, the group must also choose an optimal a^R . One possibility is always to choose the static Nash equilibrium of the primitive game that maximizes group utility with corresponding audit cost of zero. If C is very large this will certainly be optimal. For smaller C the group may instead wish to choose a U -maximizing peer feasible plan. The next result says that this entails larger $v(a^R)$ and $G(a^R)$ as C declines. That is, as the unit cost of auditing declines, the optimal action in the primitive game has larger gains to deviation but higher net utility.

Theorem 2. *As C increases the optimized values of $v(a^R)$ and $G(a^R)$ weakly decrease.*

Proof. While this follows from the general monotone selection theorem of Milgrom and Shannon (1994) the direct proof is simple and we include it for completeness. Suppose $C' > C$ and that v', v and G', G are the respective values of $v(a^R)$ and $G(a^R)$ corresponding to optimal a^R for C' and C respectively. Then $v - CG \geq v' - CG', v' - C'G' \geq v - C'G$. From the first inequality $v - v' \geq C(G - G')$ and from the second inequality $v - v' \leq C'(G - G')$. Since $C' > C$ these inequalities imply $G \geq G'$. Rewriting the inequalities as $(1/C)(v - v') \geq G - G'$ and from the second inequality $(1/C')(v - v') \leq G - G'$ shows that $v \geq v'$ as well. \square

Finally, we can do comparative statics by analyzing the properties of C : it is evident that C is increasing in $\theta_1, \theta, \pi, 1/\sigma$. In particular if the group could choose between different auditing technologies, it would prefer technologies that have low cost θ_1, θ , a low failure rate on the equilibrium path π , and a high ability to discriminate cheating as measured by the auditing signal increase σ . Note that even with perfect monitoring - $\pi = 0, \pi^p = 1$ - the unit cost $C = \theta_1/(1 - \theta)$ is still positive provided that there is a cost of conducting the initial audit, and this is true even if the cost of subsequent audits θ is zero. On the other hand, if the cost of conducting an initial audit θ_1 is zero then the unit cost of auditing is also zero. In the case $C = 0$ the group will simply choose the action a^R that maximizes $v(a^R)$, the initial net utility.

The message here is the same as that in much of the moral hazard literature: more precise monitoring leads to higher-powered incentives, an explicit statement of which can be found, for example, in Demougin and Fluet (2001).

4. Group Size and the Strength of Groups

A particular focus of the study of groups is their ability to overcome public goods problems, as for example in Olson (1965)'s work on individuals who voluntarily participate in special interest groups. For example the group might be attempting to corrupt a politician as in Ades and DiTella (1997) or Slinko and Yakovlev (2005), or it could be a consortium bidding on a contract. The more participation a group can get from its members the stronger it will be. Here we apply the model

of peer discipline to a prototypical model of a group facing a public goods contribution problem. In this context we look at the relation between group size and effectiveness - a classical problem studied by Olson (1965) and others - and we show how the size of a group interacts with peer discipline to determine group strength.

We consider as a base game a simple linear¹⁴ public good contribution game: each group member chooses between two actions $a^i \in A = \{0, 1\}$ representing the utility cost of contributing to the public good. If a contribution is made, that is $a^i = 1$, this results in a benefit to the group of $s > 1$ divided equally among all N members. In this case $a^R = 0$ means that nobody except for i contributes while $a^R = 1$ means that everybody except for i contributes. So $u(a^i, a^R)$ is given by

	$a^R = 0$	$a^R = 1$
$a^i = 0$	0	$s - (s/N)$
$a^i = 1$	$(s/N) - 1$	$s - 1$

Our central interest is to establish the conditions under which the full contribution action $a^R = 1$ is sustained. Notice that the condition for full contribution in the primitive game is $N \leq s$, while for $N > s$ there is no contribution. We assume that neither $\pi_0(a^i, a^R)$, π nor π^p depend on the size of the group - in other words we assume that auditors are close to the auditees regardless of size. The general picture is described in the following

Theorem 3. Abbreviate $\sigma_0 = \sigma_0(0, 1) = \pi_0(0, 1) - \pi_0(1, 1)$. Define

$$\bar{N}(s, P) = \begin{cases} s / \left(1 - \frac{\sigma_0(s-1)}{\pi_0(1,1)+C}\right) & \text{for } s \leq 1 + [\pi_0(1, 1) + C] \cdot \min\{P, 1/\sigma_0\} \\ s / (1 - \sigma_0 P) & \text{for } s \geq 1 + [\pi_0(1, 1) + C]P, P < 1/\sigma_0 \\ \infty & \text{for } s \geq 1 + [\pi_0(1, 1) + C]/\sigma_0, P \geq 1/\sigma_0 \end{cases}.$$

For $N \leq s$ the group contributes full effort, requires no costly auditing, and achieves utility $U = u(1, 1) = s - 1$. For $s < N \leq \bar{N}(s, P)$ and $\theta_1/\sigma \leq 1, \theta/\sigma < 1$ the group employs costly auditing, contributes full effort and achieves utility

$$U = s - 1 - [\pi_0(1, 1) + C][1 - (s/N)]/\sigma_0.$$

For $N > \bar{N}(s, P)$ or $\theta_1/\sigma > 1$ or $\theta/\sigma \geq 1$ the group contributes no effort and achieves utility $U = 0$.

Proof. The gain to cheating is $u(0, 1) - u(1, 1) = 1 - (s/N)$ from which we see that if $s \geq N$ there is no gain to cheating and voluntary contributions sustain an equilibrium with utility $u(1, 1) = s - 1$. Otherwise we have $G^R \equiv G^R(a^R) = (1 - (s/N))/\sigma_0$ and require $P \geq G^R$ and $\theta_1/\sigma \leq 1, \theta/\sigma < 1$. The first of these conditions can be written as $s \geq N(1 - \sigma_0 P)$. In addition if the group is to wish

¹⁴The idea that value generated by group effort is linear in group effort is consistent with the continuous approach to special interest groups argued by Becker (1983) and Tullock (2001), but is not appropriate for a setting such as a winner take all election where there is a sharp distinction between winning and losing.

to cooperate then it must be that the optimum peer feasible utility is U as given in the statement. Since $s < N$, the condition for willingness to audit is $U \geq u(0, 0)$, which then reads

$$s - 1 - [\pi_0(1, 1) + C][1 - (s/N)]/\sigma_0 \geq 0.$$

Starting with the constraint $U \geq u(0, 0)$: if $s \geq 1 + [\pi_0(1, 1) + C]/\sigma_0$ this holds for all N . Otherwise it holds with equality when

$$N = N_U(s) \equiv s / \left(1 - \frac{\sigma_0(s-1)}{\pi_0(1, 1) + C} \right),$$

and we may define $N_U(s) = \infty$ when $s \geq 1 + [\pi_0(1, 1) + C]/\sigma_0$. For smaller N the constraint holds strictly, and for larger N it is violated. The function $N_U(s)$ is non-decreasing in s .

Turning to the constraint $P \geq G^R$: if $P \geq 1/\sigma_0$ this always holds. Otherwise it holds with equality when

$$N = N_G(s, P) \equiv s / (1 - \sigma_0 P),$$

and we may define $N_G(s, P) = \infty$ when $P \geq 1/\sigma_0$. For smaller N the constraint holds strictly, and for larger N it is violated. The function $N_G(s, P)$ is non-decreasing in s .

It follows from this that peer discipline is feasible and optimal exactly when

$$s \leq N \leq \min\{N_U(s), N_G(s, P)\} = \bar{N}(s, P).$$

Since $N_U(s), N_G(s, P)$ are weakly increasing in s so is $\bar{N}(s, P)$. It remains to explicitly compute the minimum.

Take first the case where $P \geq 1/\sigma_0$ so that $N_G(s, P) = \infty$. In this case $\bar{N}(s, P) = N_U(s)$. This covers the cases $s \geq 1 + [\pi_0(1, 1) + C]/\sigma_0, P \geq 1/\sigma_0$ and $s \leq 1 + [\pi_0(1, 1) + C] \min\{P, 1/\sigma_0\}, P \geq 1/\sigma_0$.

Finally suppose that $P < 1/\sigma_0$. Observe that for $s = \bar{s} = 1 + [\pi_0(1, 1) + C]P$ we have $N_U(s) = N_G(s, P)$, and for $s \leq \bar{s}$ it is $N_U(s) \leq N_G(s, P)$. This covers the remaining cases. \square

What does this theorem tell us? If peer discipline is not available because $\theta_1/\sigma > 1$ or $\theta/\sigma \geq 1$ there is a pure public goods problem, and the group contributes full effort as long as individuals have adequate incentive to provide effort - that is $N \leq s$. Once the group becomes too large the group ceases to provide effort. If peer discipline is available, that is $\theta_1/\sigma \leq 1, \theta/\sigma < 1$, then the group continues to provide full effort in the range $s < N \leq \bar{N}(s, P)$. In case $\bar{N}(s, P)$ is finite, qualitatively this is similar to the pure public goods case - peer discipline can merely sustain contribution with a larger group size. The comparative statics of $\bar{N}(s, P)$ have the monotonicity properties we expect: lower cost of peer discipline as measured by smaller $\pi_0(1, 1) + C$ and larger σ_0 increase the size of group that can sustain effort.

Of particular interest is the case in which $\bar{N}(s, P) = \infty$. This requires that the punishment be adequately large for the given initial signal quality - $P \geq 1/\sigma_0$ - and that s be sufficiently large:

$s \geq 1 + [\pi_0(1, 1) + C]/\sigma_0$. This case is useful to understand what happens with very large groups. For back of the envelope purposes it might be useful to think of the large groups of farmers across countries. Here we have a problem similar to the paradox of voting: it is not very plausible that the individual lobbying efforts of a single farmer increase the chances of farm subsidies enough to be individually worthwhile. Moreover, in countries of different sizes the absolute number of farmers varies considerably. If the peer discipline technology and the benefit per farmer of farm subsidies s are roughly the same in the different countries, and if $\bar{N}(s, P)$ is finite, then in countries with few farmers $N \leq \bar{N}(s, P)$ we should find lobbying effort and farm subsidies, while in countries with many farmers $N > \bar{N}(s, P)$ we should find no lobbying and no farm subsidies. However, we find that countries of wildly different sizes often have similar levels of subsidies. For example (OECD 2010 data): the EU and Canada have roughly similar per capita GDP, similar democratic political systems and both have about a 2% share of agriculture in GDP. As a percentage of output value farm subsidies in Canada are 16% and the EU 20%. Yet as measured by either GDP or population the EU is more than ten times the size of Canada. The case where peer discipline leads to $\bar{N}(s, P) = \infty$ covers this facts: full effort is provided independent of group size, so no matter the number of farmers or size of country, the amount of per capita public good achieved should be roughly similar - as it is.

This argument seems to go against the Olson (1965) idea that larger groups should be less effective. Can we reconcile our peer discipline model with the Olsonian observation that a small group is often more effective than a large group? Here we think of farmers as a small group, compared to the large group of non-farmers. So far we have been considering increasing group size while keeping the per-capita size of the benefit s fixed. This makes sense if we want to compare, say, farm lobbies between different countries of different sizes. The Olson experiment though compares different lobbies of different sizes within the same country competing for the same prize. To wit: if the entire group contributes effort, the total benefit to the group of size N is $S = Ns$. Here we are asked to keep S fixed - that is, the value to all farmers of getting subsidies is (roughly) equal to the value to all non-farmers of not paying subsidies. Hence S is fixed and the same for groups of different sizes, so the per capital benefit $s = S/N$ depends on group size. For the few farmers the s_F corresponding to receiving farm subsidies is large; for the many non-farmers the s_{NF} corresponding not paying for farm subsidies is small. Hence, even if both groups have access to exactly the same peer discipline technology, we can have $s_F \geq 1 + [\pi_0(1, 1) + C]/\sigma_0$ for farmers and $s_{NF} < 1 + [\pi_0(1, 1) + C]/\sigma_0$ for non-farmers. In this case, farmers will be effective and contribute full effort, but non-farmers will be ineffective and not contribute effort.

We should remark on another implication of the model with respect to group size: the total effort of the group increases with the size of the group unless $\bar{N}(s, P)$ is finite and the group becomes too large. This is because a larger group has more resources - that is, there are more members to provide full effort. This anti-Olsonian implication has been noted in other models such as Esteban and Ray (2001) in which there is a mix of private and public incentives.

5. Extensions and Limitations

We have analyzed a very spare and stripped down model, which necessarily has many limitations. Here we discuss these limitations and the relevant extensions of the basic model. We should emphasize that we do not explore changing the basic structure of the model: *ex ante* identical members and sequential auditing. Without *ex ante* identical members we do not have a clear theory of what agreement they might reach, and with more complicated dynamic auditing we cannot apply the simple Kandori (1992) repeated game social norm computations. We also do not consider more complicated dynamic games a group might play: for example repeated primitive games and reputational effects. In addition to the extension considered in this section, the Web Appendix shows how to extend to model to multiple signals in place of 0-1 signals.

5.1. Renegotiation

If the group can collude initially, why can it not do so later? If the group can costlessly collude after the primitive game has been played it will always cancel the audits, since *ex post* these are costly and the actions in the primitive round can no longer be changed. This makes peer discipline impossible *ex ante* since everybody will know that the agreement to audit is not credible. Hence for peer discipline to take place it must be costly to collude in order that the peer discipline mechanism has commitment value.

It is sensible to assume that collusion is costly because in practice it is. In the simple model we described the initial collusion as taking place in a “meeting.” This might be a single face-to-face meeting of the entire group, or some more decentralized or less personal means of reaching an agreement. We know that explicit discussions are important in practice since they are explicitly forbidden under anti-trust law yet never-the-less people are sometimes caught engaging in these discussions. From a theoretical point of view we know that common knowledge is central to Nash equilibrium, and again from practice we know that common knowledge is reinforced through discussions and meetings - looking the other person in the eye both figuratively and actually.

This leads us to introduce a simple model of costly collusion: we assume that collusion can take place only if a “meeting” is held, and this meeting has a per participant cost - the cost of both organizing and attending the meeting. There are two types of meetings that are important: the initial meeting to decide the collusive scheme and subsequent meetings after the primitive round in which the original scheme may be renegotiated. These subsequent meetings may be held at the end of any round. We assume that each type of meeting has a cost associated with it, and that the meeting is held if it is in the best interest of the group to meet at that time.

Formally we assume that the initial meeting has a cost per member κ_0 associated with it, and that if the meeting does not take place, a default outcome a_0^R occurs. Since there is no meeting to agree on collusion, it is natural to assume that a_0^R is static Nash, and if there are several static Nash outcomes, for simplicity we assume that the group can - even without a meeting - coordinate on a most favorable static Nash equilibrium.

If a meeting is held subsequent to the initial primitive round the agenda is clear: the remaining audits and punishments should be dropped and the game ended. *Ex ante* at the initial meeting and prior to the initial round it is in the interest of the group to avoid this outcome, since it will undermine the collusive peer discipline arrangement. Hence the group may undertake activities to make subsequent meetings more costly and difficult - to take an extreme example, at the end of the initial meeting everybody may be required to smash their cell phones with a hammer as part of the agreed upon action a^R . To model this, we assume that the cost per member of subsequent meetings may depend upon the action, that is, has the form $\kappa(a^R)$. In other words, some actions ("smash the cell phones") may make subsequent meetings more costly.

We say that an action a^R is *renegotiation feasible* if it is peer feasible, if the optimal member utility from implementing a^R is at least κ_0 and if at the end of any round the expected cost of auditing and punishment is no greater than $\kappa(a^R)$. In other words, it pays to hold the initial meeting and does not pay to hold subsequent meetings.

In effect our original simple model assured renegotiation feasibility by implicitly assuming that $\kappa_0 = 0$ and that $\kappa(a^R) = \infty$. We now consider the general case. In the Appendix we prove

Theorem 4. *If the action a^R is not static Nash it is renegotiation feasible if and only if $P \geq G(a^R)$, $\theta_1/\sigma \leq 1$ and $\theta/\sigma < 1$ (as in Theorem 1), and additionally the three following inequalities are satisfied:*

$$\begin{aligned} u(a^R, a^R) - (\pi_0(a^R, a^R) + C) G(a^R) &\geq u(a_0^R, a_0^R) + \kappa_0 \\ (\pi_0(a^R, a^R) + C) G(a^R) &\leq \kappa(a^R) \end{aligned}$$

$$\frac{\theta + \pi}{\sigma - \theta} (\max\{\theta_1, \theta\} P) \leq \kappa(a^R).$$

The first condition is that the benefit of introducing a peer enforcement scheme must be sufficiently large that it is worth holding a meeting in the first place. The second condition is that it must not pay to hold a meeting after the primitive round, the third condition is that it not pay to hold a meeting after the first or subsequent audit rounds.

To understand this result, note that the third inequality is different in nature than the other two. In particular the left-hand-side $\frac{\theta + \pi}{\sigma - \theta} (\max\{\theta, \theta_1\} P)$ does not depend upon a^R so that no matter how good $u(a^R, a^R)$ and how small a positive number is $G(a^R)$ this condition may none-the-less fail. If $\kappa(a^R) = \kappa$ independent of the action taken, so that different actions do not have different commitment value, then this condition either fails for all a^R or is satisfied for all a^R and is independent as well of N . In this sense the third inequality represents an absolute requirement similar to the peer feasibility conditions $\theta_1/\sigma \leq 1$ and $\theta/\sigma < 1$. Notice also that the inequality seems also to restrict the size of P - but this is deceptive, because it really restricts $\theta P, \theta_1 P$ which is to say, not the size of the punishment, but rather and more intuitively, the cost of auditing. If this is too great relative to the cost of having a meeting then no renegotiation feasible arrangement is possible.

The first inequality embodies two conditions. The first is that $G(a^R)$ should not be too large, the second is that $u(a^R, a^R) - u(a_0^R, a_0^R)$ should be fairly large. Both of these strengthen an already

existing requirement for a peer discipline implementation to be desirable in the sense of being better than static Nash.

The fact that $\kappa(a^R)$ depends on a^R obviously introduces a bias in favor of actions for which $\kappa(a^R)$ is large: all other things equal such actions are more likely to be renegotiation feasible.

The Public Good Contribution Case

A good way to appreciate the content of the theorem is to apply it to our public good contribution example. Assume for simplicity that $P > 1/\sigma_0$ and $s \geq 1 + [\pi_0(1, 1) + C]P$ so that $\bar{N}(s, P) = \infty$. As a benchmark case suppose that all meetings have the same cost: that is, $\kappa_0 = \kappa(a^R) = \kappa$. For contribution enforced by peer discipline to be the group optimum we need as before the peer feasibility conditions $\theta_1/\sigma \leq 1$ and $\theta/\sigma < 1$. The third condition for renegotiation feasibility is

$$\frac{\theta + \pi}{\sigma - \theta}(\max\{\theta, \theta_1\}P) \leq \kappa,$$

while since $G^R(a^R) = (1 - s/N)/\sigma_0$ the second condition is

$$[\pi_0(1, 1) + C][1 - s/N]/\sigma_0 \leq \kappa$$

for which in turn a sufficient condition is $[1 + C]/\sigma_0 \leq \kappa$. Both of these conditions are absolute in the sense that they depend neither on a^R nor on N . If κ is large enough both are satisfied.

On the other hand, large κ goes against the first constraint which may be written as

$$s - 1 - [\pi_0(1, 1) + C][1 - s/N]/\sigma_0 \geq \kappa.$$

In the original model this condition with $\kappa = 0$ was exactly the one needed for willingness to audit, so that this condition simply strengthens the existing condition. It is no longer enough for the benefit of auditing to be positive: it must be larger than κ , the cost of calling the meeting. We can also give a sufficient condition for this constraint that does not depend on N : $s - 1 - (\pi_0(1, 1) + C)/\sigma_0 \geq \kappa$. If s is sufficiently large then this condition is satisfied and the group colludes on full contribution for all N . However, without the availability of peer punishment they make the full contribution only if $N \leq s$. The key qualitative feature of the peer discipline model of public goods contribution is preserved in the face of renegotiation even when the cost of holding all meetings is the same.

5.2. Generalized Matching Procedures

The basic model supposes a very simple procedure in which each member j is audited by the member $i = j + 1$ to his right and in which all the audits end at the same time as determined by $1 - \delta_t$. However, more elaborate procedures are possible. We now describe a broader class of procedures and show that the results based on simple matching are robust.

Let $I = \{1, 2, \dots, N\}$ be the set of members. In general we can consider a matching of auditees to auditors described by maps $m_t : I \rightarrow I$ with the convention that if a member is assigned to audit himself $m_t(j) = j$ then no audit takes place. Moreover, we require that if $i \notin m_t(I)$ then

$m_{t+1}(i) = i$, that is, once a player does not conduct an audit, he cannot be audited next period as there is nothing to audit. A generalized matching process is then a random choice of m_t where the probability distribution over matchings depends on the history of previous matchings. A generalized matching process thus incorporates the random rule for ending the game. The simple base case is that if the game has not ended $m_t(j) = j + 1$ while if the game has ended $m_t(j) = j$. With generalized matching the group in its initial meeting does not choose ending probabilities but rather a generalized matching process - which necessarily also incorporates the ending procedure.

Notice that we allow the possibility that an auditor is assigned to audit more than one auditee. In this case we assume that each audit has a separate cost, that the auditor makes a separate decision about whether to conduct each audit, and that in the next period if he is audited his auditor must separately audit each of his audits. This means that from the perspective of incentive compatibility each audit is a separate unit and it makes no difference if two audits are conducted by the same or by different members.

The key idea is the standard mechanism design procedure of standing the problem on its head. Any generalized matching procedure induces initial probabilities p_0^i that member i will be audited and depending on the history for each assigned match ij at time t probabilities p_t^{ij} that the audit decision will in turn be audited the next period.

Our basic conclusion is that allowing generalized matching procedures makes essentially no difference:

Theorem 5. *If the action a^R is not static Nash and it is peer feasible then p_0^i, p_t^{ij} is optimal if and only $P \geq G(a^R)$, $\theta_1/\sigma \leq 1$ and $\theta/\sigma < 1$ and*

$$p_0^i = G(a^R)/P, p_t^{ij} = \theta_t/\sigma.$$

The corresponding optimal utility for a member is

$$U = u(a^R, a^R) - \left(\theta_1 + \pi_0(a^R, a^R) + \frac{\theta_1(\theta + \pi)}{\sigma - \theta} \right) G(a^R).$$

Proof. This is essentially an observation about the proof of Theorem 1. Obviously the incentive constraints must be satisfied with inequality $p_0^i \geq G(a^R)/P, p_t^{ij} \geq \theta_t/\sigma$. Reducing the probability of audits reduces the expected number of audits in all future periods and so reduces the cost of audits, hence optimality requires that the incentive constraints hold with exact equality. This of course gives exactly the same expected cost of audits as in Theorem 1. \square

There is one proviso: we may wish to consider only matching procedures such as the simple one with the property that no pair of members is assigned to repeatedly audit each other. The reason is that when two members repeatedly audit each other they more or less costlessly “meet” and so will collude to avoid the audits. This is why we have assumed from the beginning of the paper that $N \geq 3$.

Notice that there are many generalized matching procedures that are not optimal, some obviously so, such as ones that result in the game ending in finite time. But besides the simple procedure, many other procedures are possible. For example we might in the initial period only audit only a fraction ϕ of the population and in subsequent periods audit all the auditors. Here the probability of being audited after the primitive round is $\delta_0\phi$. Hence the incentive compatibility condition is $\delta_0 = G^R/(\phi P)$. This means that we must require $\phi \geq G^R/P$ - that is, depending on the gain to deviating and the available punishment, we cannot choose ϕ too small.

5.3. Punishment Cost Spillover and Choice of Punishment

We have assumed that the cost of punishment is borne only by the “guilty” party. In practice, however, the cost of punishment may spill over to other group members. The most common forms of punishment - some sort of exclusion, ranging from being denied the opportunity to participate in group events to imprisonment - will generally harm group members as well as the designated target of the punishment. For example, if Tim is punished by being excluded from joining the group at the bar after work then David suffers the loss of Tim’s companionship. Or it may be that David feels sorry for Tim. We refer to this a punishment cost spillover.

We have also assumed that there is only one possible kind of punishment - generally there are degrees of punishment, for example a member may be excluded for a week or for a month, or may be partially excluded. The choice of type of punishment is closely linked to the spillover cost because different types of punishments will generally have different spillover costs - and in particular we might expect that very severe punishments may have disproportionately high spillover costs.

We first formally introduce the notion of spillover cost. In doing so we need to specify which group members suffer from the spillover. The consequence of spillover to the auditor is more costly to the group than the cost of spillover to other members: it increases the incentive of the auditor to not conduct the audit so as to avoid the spillover cost. Consequently if different group members suffer different levels of spillover costs, the group will want to appoint as auditor the member who least suffers these costs. For simplicity we assume that it is possible to appoint an auditor who suffers no spillover cost and that the spillover costs are equally divided among the remaining group members. Hence when punishment is imposed on i there is a spillover cost of ψP divided equally among group members other than i and his auditor $i - 1$. Since the punishment occurs only with probability π_t the expected cost is $\psi P\pi_t$. Since each member pays a share $1/(N - 2)$ of the cost of the $N - 2$ matches in which he is neither auditee nor auditor this is also the per capita expected spillover cost, leading to a simple change in the computation of the optimal utility:¹⁵

$$U = u(a^R, a^R) - \left(\theta_1 + \pi_0(a^R, a^R)(1 + \psi) + \frac{\theta_1(\theta + \pi(1 + \psi))}{\sigma - \theta} \right) G(a^R).$$

¹⁵It is easy to treat also the case in which the spillover costs are shared equally among all members including the auditor. In this case the cost to the auditor of conducting the audit is $(\theta_t + \psi\pi_t/(N - 1))P$, and the equation below should be adjusted accordingly. Note however that if the group is relatively large, $\psi\pi_t/(N - 1)$ will be very small compared to θ_t , so that the adjustment will not make much difference.

Note that the spillover costs do not necessarily have to be positive, although we would generally want to assume that $\psi > -1$ so that the punishment does not bring a net benefit to the group. Negative spillover cost corresponds to punishments that involve a transfer payment. The most obvious example is the payment of a fine, but there are other possibilities. For example, punishment might involve a demotion, in which case another member of the group might benefit from being promoted to fill the vacant spot. Since lower spillover costs are better, punishments involving transfer payments are highly desirable if they are feasible - in the repeated game setting with imperfect private information as in Fudenberg Levine and Maskin (1994) it is the use of transfer payment punishments that gives rise to near efficiency as the discount factor approaches one.

Having formally introduced spillover costs to the model we now examine the implications of having available many types of punishments of varying severity. If the audit costs $\theta_1 P, \theta P$ and spillover costs ψP are fixed then we see as we increase P that θ_1, θ, ψ decline and this lowers the cost of discipline and raises group utility. If, as we expect, increasing the level of punishment raises these costs then there is a trade-off that we now spell out. For simplicity we examine the case in which we may choose a punishment technology P_1 in the initial period and a technology $P_t = P$ in the subsequent periods $t > 1$, while we hold fixed the audit cost $A_t = \theta_t P_t$ and allow the spillover costs ψP_t to vary. This also reflects our intuition that the cost of auditing (for a fixed signal technology) should not much vary with the cost of punishment and that as a practical matter the reason for using more moderate punishments - community service rather than incarceration rather than the death penalty in the case of penal systems - is that the spillover cost of more severe punishments is disproportionately large.

Formally, we assume that the trade-off between punishment size and spillover cost is captured by an increasing continuous function $\psi P = f(P)$ defined on $0 \leq P \leq \bar{P}$ where \bar{P} is the worst possible punishment. In general since randomization is possible we expect f to be convex, but this is not essential. For any choice of P the audit costs A_1 and $A_t = A$, $t > 1$ are held fixed implicitly allowing θ_t to vary so that $\theta_t = A_t/P_t$. In the Appendix we prove

Theorem 6. *If there is some technology pair $\psi_1 P_1 = f(P_1)$, $\psi P = f(P)$ with corresponding $\theta_1 = A_1/P_1, \theta = A/P$ that satisfies the constraints $0 \leq P_1, P \leq \bar{P}$, $P_1 \geq G(a^R)$, $P \geq A_1/\sigma$ and $\theta/\sigma < 1$ then the problem of maximizing*

$$U = u(a^R, a^R) - \left(\theta_1 + \pi_0(a^R, a^R)(1 + \psi_1) + \frac{\theta_1(\theta + \pi(1 + \psi))}{\sigma - \theta} \frac{P}{P_1} \right) G(a^R)$$

subject to those constraints has a solution and it is an optimal peer discipline scheme.

6. Conclusion

We have developed a model of group behavior with explicit account of individual incentives inside the group. In a public goods contribution example we argue that the model gives more appropriate predictions with respect to large groups than a voluntary contribution approach.

Not all enforcement mechanisms are peer discipline mechanisms. Hierarchical schemes - such as those in the principal-agent model - are commonplace. Yet peer discipline mechanisms play a role even within hierarchical organizations. Indeed, it is not always the case that social norms within a hierarchical organization are enforced because the executive has the right to punish his subordinates, who have the right to punish their subordinates and so forth. The key point is that - regardless of contracts and laws - the “right” to punish exists because other people recognize it as a right and it is part of a social norm that itself is enforced through peer punishment. On the other hand social norms and peer enforcement mechanisms in hierarchical organizations sometimes serve to subvert rather than enhance the goals of the hierarchy. Several examples help clarify the point.

Consider first the very hierarchical military. Here is a quotation from Senator James Inhofe:¹⁶ “Army and Marines always feel that when we’re out there, we’re not doing it for the flag or the country; we’re doing it for the guy in the next foxhole.” That is - in combat - despite the hierarchical structure and severe punishments for cowardice - there are many opportunities to shirk. The enforcement mechanism - the reason soldiers risk their lives when there is little chance that their superiors will find out - is the enforcement by “the guy in the next foxhole.” In the military peer enforcement reinforces the hierarchical organization.

In the other direction, the “code of blue silence” is an unwritten rule that police officers do not report the misconduct of other officers. The following quotation from Frank Serpico testifying before the Knapp Commission on police corruption in New York City in 1971 reveals the enforcement scheme: “an honest police officer can[not] act...[against corruption] without fear of ridicule or reprisal from fellow officers.” Notice that peer monitoring must play an essential role in this - if an officer rats on another officer, other officers (the auditors) must report this to the police force more generally. We think we can be reasonably confident - based on our avid watching of movies about police corruption if nothing else - in asserting that failure to report a “rat” is itself subject to reprisal. In this case the peer discipline mechanism serves to subvert the hierarchical organization.

We do not believe that the simple model here sheds light on all issues involving all kinds of organizations - rather we think it provides a useful and simple tool for studying peer enforcement and one that is compatible with more elaborate models that - for example - combine hierarchical and other kinds of enforcement schemes with peer punishment schemes.

References

Acemoglu, D. and A. Wolitzky (2015): “Sustaining Cooperation: Community Enforcement vs. Specialized Enforcement D Acemoglu,” MIT.

Ali, Nageeb and David Miller (2010): “Cooperation and Collective Enforcement in Networked Societies,” UC San Diego.

Abreu, D., D. Pearce, E. Stacchetti (1986): “Optimal cartel equilibria with imperfect monitoring,” *Journal of Economic Theory* 39, pp. 251–269.

¹⁶In an interview with AFA radio on May 12, 2010 arguing against allowing gays to serve openly in the military.

Ades, A. and R. DiTella (1997): "The new economics of corruption: a survey and some new results," *Political Studies* 45, pp. 496–515.

Balmaceda, Felipe and Juan F. Escobar (2014): "Trust in Cohesive Communities," University of Chile.

Baron, David (2010): *Business and Its Environment*, 6th Edition.

Becker, G. S. (1983): "A theory of competition among pressure groups for political influence," *Quarterly Journal of Economics* 98, pp. 371–400.

Ben-Porath, Elchanan and Michael Kahneman (2003): "Communication in repeated games with costly monitoring," *Games and Economic Behavior* 44, pp. 227–250.

Border, Kim C. and Joel Sobel (1987): "Samurai Accountant: A Theory of Auditing and Plunder", *The Review of Economic Studies* 54, pp. 525–540.

Burt, Ronald S. (1995): *Structural Holes The Social Structure of Competition*, Harvard University Press.

Coleman, J.S. (1998): "Social capital in the creation of human capital", *American journal of sociology* 94, pp. S95–S120.

Deb, Joyee and Julio González-Díaz (2010): "Community enforcement beyond the prisoner's dilemma," NYU.

Demougin, Dominique and Claude Fluet (2001): "Monitoring versus incentives," *European Economic Review* 45, pp. 1741–1764.

Di Porto, Edoardo, Nicola Persico and Nicolas Sahuguet (2013): "Decentralized Deterrence, with an Application to Labor Tax Auditing", *American Economic Journal: Microeconomics* 5, pp. 35–62.

Dutta, Rohan, David K. Levine and Salvatore Modica (2015): "Collusion, Randomization, and Leadership in Groups," <http://www.dklevine.com/papers/collusion-leaders.pdf>.

Ellison, Glenn (1994): "Cooperation in the prisoner's dilemma with anonymous random matching," *The Review of Economic Studies* 61, pp. 567–588.

Esteban, Joan and Debraj Ray (2001): "Collective Action and the Group Size Paradox," *American Political Science Review* 95, pp. 663–672.

Fehr, Ernst and Simon Gächter (2000): "Cooperation and Punishment in Public Goods Experiments", *The American Economic Review* 90(4), pp. 980–994.

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

Fudenberg, Drew and David Levine (2009): "Repeated Games with Frequent Signals," *Quarterly Journal of Economics*, 124, pp. 149–195.

Kandori, Michihiro (1992): "Social norms and community enforcement," *The Review of Economic Studies* 59, pp. 63–80.

Laffont, Jean-Jacques (1999): "Political economy, information and incentives," *European Economic Review* 43, pp. 649–669.

Levine, David K. and Andrea Mattozzi (2015): "Voter Participation with Collusive Parties," <http://www.dklevine.com/archive/refs4786969000000001234.pdf>.

Levine, David K., and Salvatore Modica (2013): "Peer Discipline and the Strength of Organizations," <http://www.dklevine.com/papers/peer-discipline.pdf>.

Milgrom, Paul and Chris Shannon (1994): “Monotone Comparative Statics,” *Econometrica* 62, pp. 157-180.

Olson, Mancur (1965): *The Logic of Collective Action: Public Goods and the Theory of Groups*, Harvard University Press.

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

Ostrom, Elinor, James Walker and Roy Gardner (1992): “Covenants with and Without a Sword: Self-Governance Is Possible”, *The American Political Science Review* (86), pp. 404–417.

Palfrey, T. R. and H. Rosenthal (1985): Voter Participation and Strategic Uncertainty, *American Political Science Review* 79: 62-78.

Pecorino, Paul (1999): “The effect of group size on public good provision in a repeated game setting”, *Journal of Public Economics* 72, pp. 121–134.

Rahman, David (2012): “But Who Will Monitor the Monitor,” *American Economic Review* 102, pp. 2767–2797

Shapiro, Carl and Stiglitz, Joseph E. (1984): “Equilibrium Unemployment as a Worker Discipline Device”, *American Economic Review* 74, pp. 433–444.

Slinko, Irina, Evgeny Yakovlev, and Ekaterina Zhuravskaya (2005), “Laws for sale: evidence from Russia,” *American Law and Economics Review* 7.1, pp. 284–318.

Sugaya T. (2011): “Folk Theorem in Repeated Games with Private Monitoring,” Princeton University.

Takahashi, Sugaya (2010): “Community enforcement when players observe partners’ past play,” *Journal of Economic Theory* 145, pp. 42–62.

Tullock, G. (2001): “Efficient Rent Seeking” in *Efficient Rent Seeking*, Springer.

Vega-Redondo, Fernando (2006): “Building up social capital in a changing world,” *Journal of Economic Dynamics and Control* 30, pp. 2305–2338.

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

Appendix

Theorem. [4 in text] *If the action a^R is not static Nash it is renegotiation feasible if and only if $P \geq G(a^R)$, $\theta_1/\sigma \leq 1$ and $\theta/\sigma < 1$ and*

$$\begin{aligned} u(a^R, a^R) - (\pi_0(a^R, a^R) + C) G(a^R) &\geq u(a_0^R, a_0^R) + \kappa_0 \\ (\pi_0(a^R, a^R) + C) G(a^R) &\leq \kappa(a^R) \end{aligned}$$

$$\frac{\theta + \pi}{\sigma - \theta} (\max\{\theta_1, \theta\} P) \leq \kappa(a^R).$$

Proof. The first part simply reiterates the conditions for peer feasibility. The three displayed inequalities are the additional constraints that arise from costly meetings. There are two conditions:

it must be worth calling the initial meeting, and it should not be worthwhile to hold a subsequent meeting.

It is worth calling an initial meeting to implement a^R if and only if the utility U from the implementation exceeds the default utility $u(a_0^R, a_0^R)$ by at least κ_0 , that is,

$$u(a^R, a^R) - (\pi_0(a^R, a^R) + C) G^R - u(a_0^R, a_0^R) \geq \kappa_0.$$

Rearranging this gives the first displayed inequality of the theorem.

Second and finally it must not be strictly worth calling a meeting at the end of any round. There are three possibilities: to call a meeting at the end of the primitive round 0, at the end of the first round 1, and all subsequent rounds are identical, so if it is worth calling a meeting we can assume that it is done right away - that is, at the end of round 2. The benefit of any of these meetings is saving the expected cost of auditing and punishment conditional on the current round. In the primitive round we already know this to be

$$K_0(a^R) = (\pi_0(a^R, a^R) + C) G^R.$$

At the end of the first round we may compute

$$K_1 = \sum_{t=2}^{\infty} \left(\prod_{\tau=1}^{t-1} \delta_{\tau} \right) (\theta_t + \pi_t) P = \delta_1 \left(1 + \frac{\delta}{1 - \delta} \right) (\theta + \pi) P = \frac{\theta + \pi}{\sigma - \theta} (\theta_1 P)$$

and at the end of the second round

$$K = \sum_{t=3}^{\infty} \left(\prod_{\tau=2}^{t-1} \delta_{\tau} \right) (\theta_t + \pi_t) P = \frac{\delta}{1 - \delta} (\theta + \pi) P = \frac{\theta + \pi}{\sigma - \theta} (\theta P).$$

Hence the additional necessary and sufficient conditions are $K_0(a^R) \leq \kappa(a^R)$ which is the second displayed inequality of the theorem and $K_1 \leq \kappa(a^R), K \leq \kappa(a^R)$ which combine to form the third displayed inequality of the theorem. \square

Theorem. [6 in text] *If there is some technology pair $\psi_1 P_1 = f(P_1)$, $\psi P = f(P)$ with corresponding $\theta_1 = A_1/P_1, \theta = A/P$ that satisfies the constraints $0 \leq P_1, P \leq \bar{P}$, $P_1 \geq G(a^R)$, $P \geq A_1/\sigma$ and $\theta/\sigma < 1$ then the problem of maximizing*

$$U = u(a^R, a^R) - \left(\theta_1 + \pi_0(a^R, a^R)(1 + \psi_1) + \frac{\theta_1(\theta + \pi(1 + \psi))}{\sigma - \theta} \frac{P}{P_1} \right) G(a^R)$$

subject to those constraints has a solution and it is an optimal peer discipline scheme.

Proof. First with respect to the constraints we have defined $\theta_1 = A_1/P_1$. This means that $\theta_1/\sigma \leq 1$, the old form of the constraint of adequate punishment for the first period audit, would be $P_1 \geq A_1/\sigma$ which is not correct since the punishment for failure in the first period audit is now P not P_1 . Hence we explicitly write out the constraint $P \geq A_1/\sigma$. Second, the objective function is the correct one:

the group is not constrained to choose the same technology in the first period as in subsequent periods, nor since the incentive constraint is different in the initial period than in subsequent periods, will it generally wish to do so. However, we have imposed the constraint that after the initial audit of play in the primitive game a single technology be chosen for use in all audits of audits. Notice that $G(a^R) = \delta_0 P_1$ while $P_t = P$ may not equal P_1 for $t > 1$, which explains the factor of P/P_1 in the cost of auditing for $t > 1$.

The only issue in the maximization is the constraint $\theta/\sigma < 1$. Notice that the solution for $P, \psi = f(P)/P, \theta = A/P$ is independent of the rest of the problem and simply minimizes

$$J = \frac{\theta + \pi(1 + \psi)}{\sigma - \theta} P$$

subject to $0 \leq P \leq \bar{P}$, $P \geq A_1/\sigma$ and $\theta/\sigma < 1$, so the issue is the existence of a solution of this problem. Since $A_1 > 0$ we have P bounded away from zero, so J approaches $+\infty$ at a rate bounded below independent of the other parameters as $\theta/\sigma \rightarrow 1$. It follows that if there exists a feasible solution to the constraints we can find an $\epsilon > 0$ so that the optimum subject to the constraint $\theta/\sigma \leq 1 - \epsilon$ is also the optimum subject to constraint $\theta/\sigma < 1$. Basically it is not a very good idea to choose a P so low that θ ends up very close to σ .

If the probabilities of punishment on the equilibrium path $\pi_0(a^R, a^R) = \pi = 0$ then the spillover cost does not matter, and again it is best to choose the largest possible punishment to minimize θ_1, θ . Otherwise the solution may be a more moderate punishment. As we noted in the proof the solution for P, ψ is independent of the rest of the problem. This has two consequences. First, it is not a good idea to choose P very large when $\pi > 0$ since then the function J to be minimized approaches $+\infty$. Indeed, if $A = 0$ so that audits other than the initial one are costless, we should choose P as small as possible, that is, equal to A_1/σ to minimize the cost of punishing the initial auditor. Second, the choice of P, ψ, θ does not depend on a^R . By contrast the objective function determining P_1, ψ_1, θ_1 depends on $\pi_0(a^R, a^R)$ and the constraint on $G(a^R)$. That is, the solution to this problem differs from solutions we have considered previously in that the first period solution depends in general on a^R . Consequently we should write $P_1(a^R), \psi_1(a^R), \theta_1(a^R)$. This minor generalization of the theory makes perfectly good sense in any case: although we have heretofore assumed that the first period audit procedure is independent of the particular target a^R it also makes sense that different audit procedures with different costs would be used to monitor different target actions. Note, however, that if A_1 depends on a^R then in general so will P, θ, ψ because the constraint $P \geq A_1/\sigma$ now depends on a^R . \square

Web Appendix: Multiple Signals

The basic model assumes that there is a simple signal whether or not a punishment is merited. While this simplifies notation and exposition, it turns out to be without loss of generality. We focus on the case of the primitive game, since in practice the specification of the game may naturally lead to many signals. Similar considerations apply to the auditing games.

Suppose the signals $Z \in \mathcal{Z}$ can take on many values with probabilities $\pi_0(Z|a^i, a^R)$. There is a common probability $1 - \delta_0$ that no audit takes place and nobody is punished. When an audit does take place and the signal Z is observed let β_Z denote the probability that the auditee is punished. Since the probability of the punishment is arbitrary, we may continue to take the size of the punishment P to be fixed. Hence the incentive constraint in the primitive round is $u(a^i, a^R) - u(a^R, a^R) - \delta_0 \sum_{Z \in \mathcal{Z}} [\pi(Z|a^i, a^R) - \pi(Z|a^R, a^R)] \beta_Z P \leq 0$. Since we may assume the group minimizes costs of auditing in subsequent rounds, the objective function is

$$V^\beta = u(a^R, a^R) - \delta_0 P \left(\sum_{Z \in \mathcal{Z}} \pi_0(Z|a^R, a^R) \beta_Z + \theta_1 + \frac{\theta_1(\theta + \pi)}{\sigma - \theta} \right).$$

For any non-negative vector of individual punishment probabilities $\beta = (\beta_Z)_{Z \in \mathcal{Z}}$ let $|\beta|$ be the sup norm. For any such β we define a *canonical binary signal process* $z \in \{0, 1\}$ as a random function of Z by $\Pr(z = 0|Z) \equiv \beta_Z/|\beta|$, and let $\pi_0^\beta(a^i, a^R) \equiv \Pr(z = 0|a^i, a^R) = \sum_{Z \in \mathcal{Z}} \Pr(z = 0|Z) \pi_0(Z|a^i, a^R)$. For given a^R and β let $G^\beta = G(a^R)$ and $\pi_0^\beta = \pi_0(a^R, a^R)$ and let \mathcal{B} be the set of all pairs (G^β, π_0^β) generated by probability vectors β . Define

$$U^\beta \equiv \max_{(G^\beta, \pi_0^\beta) \in \mathcal{B}} u(a^R, a^R) - \left(\pi_0^\beta + \theta_1 + \frac{\theta_1(\theta + \pi)}{\sigma - \theta} \right) G^\beta.$$

The basic result is this:

Theorem 7. *For given a^R there is an incentive compatible δ_0, β if and only if a^R is enforceable with respect to the corresponding canonical binary process for some punishment $0 \leq P_1 \leq P$. In case a^R is enforceable with respect to some canonical binary process then $\max_{(G^\beta, \pi_0^\beta) \in \mathcal{B}} U^\beta$ exists and is the greatest utility achievable by any peer discipline scheme that uses the signals $Z \in \mathcal{Z}$.*

Proof. First, if there is an incentive compatible δ_0, β there is an optimal one since the incentive constraints are defined by weak inequalities. Second, observe that δ_0, β satisfy the incentive constraints if and only if $\hat{\delta}_0 = |\beta| \delta_0, \hat{\beta} = \beta/|\beta|$ does so. Moreover the cost of initial punishment

$$\delta_0 P \left(\sum_{Z \in \mathcal{Z}} \pi_0(Z|a^R, a^R) \beta_Z \right)$$

is the same for both. However, if $|\beta| < 1$ we have $\hat{\delta}_0 < \delta_0$ so that the scheme $\hat{\delta}_0, \hat{\beta}$ results in no smaller a value of the objective function, and if $\theta_1 > 0$ a strictly larger value of the objective function. Hence there is an optimal scheme in which $|\beta| = 1$.

Next fix δ_0, β . By definition of the canonical binary process we have

$$\pi_0^\beta(a^i, a^R) = \sum_{Z \in \mathcal{Z}} \pi_0(Z|a^i, a^R) \beta_Z / |\beta|$$

so that $u(a^i, a^R) - u(a^R, a^R) - \delta_0 P \sum_{Z \in \mathcal{Z}} [\pi_0(Z|a^i, a^R) - \pi_0(Z|a^R, a^R)] \beta_Z \leq 0$ if and only if $u(a^i, a^R) -$

$u(a^R, a^R) - \delta_0 |\beta| P [\pi_0^\beta(z|a^i, a^R) - \pi_0^\beta(Z|a^R, a^R)] \leq 0$. Hence if π_0^β is enforceable with respect to $\delta_0 |\beta| P$ then δ_0, β is incentive compatible. Conversely if δ_0, β is incentive compatible then π_0^β is enforceable with respect to $\delta_0 |\beta| P$.

Finally, observe that

$$U^\beta - V^\beta = \delta_0 P (1 - |\beta|) \left(\theta_1 + \frac{\theta_1(\theta + \pi)}{\sigma - \theta} \right).$$

Since there is an optimal multi-signal scheme with $|\beta| = 1$ the corresponding canonical scheme yields exactly the same value of the objective function. On the other hand if there is a feasible canonical scheme π_0^β that yields a higher utility than this then the multi-signal scheme $\delta_0 = G^\beta/P, \beta/|\beta|$ is incentive compatible and yields exactly the same utility as U^β , a contradiction. \square