



HAL
open science

Hierarchy, Opportunism in Teams

Eline van Der Heijden, Jan Potters, Martin Sefton

► **To cite this version:**

Eline van Der Heijden, Jan Potters, Martin Sefton. Hierarchy, Opportunism in Teams. Journal of Economic Behavior and Organization, 2008, 69 (1), pp.39. 10.1016/j.jebo.2008.09.007. hal-00655021

HAL Id: hal-00655021

<https://hal.science/hal-00655021>

Submitted on 25 Dec 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accepted Manuscript

Title: Hierarchy, Opportunism in Teams

Authors: Eline van der Heijden, Jan Potters, Martin Sefton

PII: S0167-2681(08)00184-4
DOI: doi:10.1016/j.jebo.2008.09.007
Reference: JEBO 2258



To appear in: *Journal of Economic Behavior & Organization*

Received date: 1-9-2006
Revised date: 17-9-2008
Accepted date: 18-9-2008

Please cite this article as: van der Heijden, E., Potters, J., Sefton, M., Hierarchy, Opportunism in Teams, *Journal of Economic Behavior and Organization* (2008), doi:10.1016/j.jebo.2008.09.007

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Hierarchy and Opportunism in Teams*

Eline van der Heijden^a

Jan Potters^b

Martin Sefton^c

^a Department of Economics, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands,
eline.vanderheijden@uvt.nl

^b Department of Economics, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands,
j.j.m.potters@uvt.nl (corresponding author)

^c School of Economics, Nottingham University, University Park, NG7 2RD, Nottingham, United Kingdom,
martin.sefton@nottingham.ac.uk

August, 2006

This revision December, 2007

Abstract: We use an experiment to compare two institutions for allocating the proceeds of team production. Under revenue-sharing, each team member receives an equal share of team output; under leader-determined shares, a team leader has the power to implement her own allocation. Both arrangements are vulnerable to opportunistic incentives: under revenue-sharing team members have an incentive to free ride, while under leader-determined shares leaders have an incentive to seize team output. We find that most leaders forego the temptation to appropriate team output and manage to curtail free riding. As a result, compared to revenue-sharing, the presence of a team leader results in a significant improvement in team performance.

JEL codes: C9, D2, L2

Keywords: leadership, team production, experiment

* This paper was partly written while Sefton was visiting CentER at Tilburg University. We are grateful to the Netherlands' Organization for Scientific Research (NWO) for financial support. We thank Sigrid Suetens for assistance in conducting the experiment, three referees, the editor (Claude Menard), Spiros Bougheas, Robin Cubitt, Doug DeJong, Simon Gächter, Sebastian Kube, Doug Nelson, Elke Renner, Daniel Seidmann, Joep Sonnemans, seminar participants at the Universities of Copenhagen, Exeter, Iowa, Leicester, Lausanne and Nottingham, and conference participants at the Economic Science Association International Meeting, Montreal 2005, for helpful comments.

Hierarchy and Opportunism in Teams*

August, 2006

This revision, December 2007

Abstract: We use an experiment to compare two institutions for allocating the proceeds of team production. Under revenue-sharing, each team member receives an equal share of team output; under leader-determined shares, a team leader has the power to implement her own allocation. Both arrangements are vulnerable to opportunistic incentives: under revenue-sharing team members have an incentive to free ride, while under leader-determined shares leaders have an incentive to seize team output. We find that most leaders forego the temptation to appropriate team output and manage to curtail free riding. As a result, compared to revenue-sharing, the presence of a team leader results in a significant improvement in team performance.

INTRODUCTION

Activities in group settings are prone to free rider problems when the actions of individual group members are not subject to enforceable contracts. Team production, public good provision, collective action, and common pool resource extraction are well-known guises of the same basic problem. One classic solution by which groups may try to attenuate the free-rider problem is to install some kind of hierarchy within the group, for example, by appointing a “leader”. By giving the leader a stake in the total benefits produced by the group (for example, a residual claimant status) she will have an incentive to reduce free riding, and by giving her the right to determine rewards to other group members she will also have the means to do so (Alchian and Demsetz 1972).

A hierarchy, though, can have its own disadvantages. When the leader is a residual claimant she will be tempted to reduce the rewards to the other group members as this will increase the residual that accrues to her. This has been referred to as the central dilemma in a hierarchy: “how to constrain the self-interest of those with a stake in the inevitable residual generated by an efficient incentive scheme” (Miller 1992, p. 155). As long as the actions of individual team members are non-contractable (i.e., not verifiable by a court of law), it will be impossible for a leader to commit to a scheme for rewarding team members. Therefore, the willingness of the leader to enforce a particular rule is subject to incentive problems, just as in the absence of a leader, the willingness of group members to behave in accordance with the common interest is subject to incentive problems (Bianco and Bates 1990, Binmore 1998, Calvert 1995). In addition, social psychological research on leadership has shown that installing a leader may be perceived as a threat to the decisional freedom of the group members. Depending on the extent to which individuals lose freedom of choice, which depends on the leadership style, this may be viewed more or less negatively (Van Vugt and De Cremer 1999).

In this paper we present an experiment designed to examine whether installing a leader can curtail free riding and improve welfare in groups. To do this we consider a team production setting under two institutional arrangements. In our revenue-sharing institution, there is no leader, and each team member receives an equal share of team output. In our leader-determined shares institution, a leader can monitor

the effort of team members and has discretion to allocate the proceeds from team production. Here, the leader can condition the allocation on the observed effort levels of the individual members, including her own, even though she cannot commit to a certain allocation. In other words, individual efforts are perfectly and costlessly observable, but not contractable.¹ Thus, the key feature of our leader treatment is that the leader is a residual claimant with discretionary reward power.

The classic example of the type of leader we consider is the owner of a small firm who has the power to set the wages of the employees (Alchian and Demsetz). Since the profits accrue to the owner she will try to use her reward power to induce the employees to work hard. At the same time, she has an incentive to curb wages in order to increase profits. Alchian and Demsetz discuss the implications of these incentives in detail and, like us, compare them with those in a team governed by equal profit sharing. It is not only in market-like settings though that we can find leaders with residual claimancy status and reward power. Similar incentives are sometimes imputed on leaders of divisions in larger (non-profit) organizations.² For example, the dean of a faculty is often accredited with substantial discretionary rights over the allocation of the faculty's budget. A benevolent dictator as dean would use this discretion to enhance the faculty's performance. However, this may require rewarding other faculty members in ways that do not serve the personal utility of the dean. An opportunistic dean will have an incentive to curtail these rewards in order to retain a larger budget to spend on those things she personally values (e.g., her own research group, personal salary, office space, support staff, or a sabbatical).

¹ This assumption accords well, for example, with the situation in some academic departments. In these, many heads will assert that they know which faculty members contribute more to the performance of the department while at the same time they will find it impossible to write complete contracts to reward faculty on the basis of these contributions.

² Even though in practice contractual factors may limit the discretion of a leader, in many organizations she has substantial power to reward team members. This may include an ability to allocate simple pecuniary rewards, such as bonuses, but also the allocation of fringe benefits and task assignments, influence on career progression, as well as less immediate or tangible rewards. The essence of the dilemma we study is unchanged under the more moderate assumption that the leader has discretion to allocate *some* (rather than *all*) of the revenues that accrue to the group.

In our environment, standard economic theory, based on the assumption that individuals maximize own earnings, predicts that teams will perform just as inefficiently under revenue-sharing and leader-determined sharing. In the former case inefficiency is driven by the classic free-rider problem, while in the latter case concerns that the leader will appropriate all team output drive inefficiency.

Our experimental results show that free-riding incentives indeed undermine performance in revenue-sharing teams. These results are qualitatively similar to those found in previous experiments on the voluntary provision of public goods (see Ledyard 1995 for a survey). In contrast, in the teams with leaders we find that most leaders forego the temptation to appropriate all of the team output for themselves and manage to induce high team effort by giving nothing to shirkers and rewarding workers with an equal share of team output. Thus, we find that the presence of a leader results in a significant and substantial improvement in team performance.

Our experiment adds to the experimental economics literature that examines the force of institutional remedies to free riding. These include investigations of pre-play communication (e.g., Isaac and Walker 1988, Ostrom et al. 1992, Frey and Bohnet 1996, Bochet et al. 2006), group incentive contracts (e.g., Nalbantian and Schotter 1997, van Dijk et al. 2001), and mutual monitoring (e.g., Ostrom et al. 1992, Fehr and Gächter 2000, Bowles et al. 2001, Sefton et al. 2007). Previous experiments on leadership have focused on leading-by-example (e.g., Meidinger and Villeval 2002, Moxnes and van der Heijden 2003, Gächter and Renner 2004, Potters et al. 2007). Leading-by-example refers to a setting in which one team member decides about his contribution to the team output before the others do. By setting a good example, the leader can try to induce others to follow him. This is an informal type of leadership in which the leader lacks formal authority or sanctioning power. Generally speaking, leadership can attain many forms and consist of many elements. The element on which we concentrate corresponds to what is sometimes referred to as "transactional leadership" (Bass 1985). The leader engages in an instrumental exchange relationship with subordinate team members by supplying contingent rewards in return for desired behavior. This is a more formal type of leadership embedded in a hierarchical relationship.

To our knowledge, the only experimental paper examining central monitoring is Vyrastekova and van Soest (2003). They study a common pool resource extraction game in which a ‘police officer’ can try to collect fines from resource users who over-extract. The police officer can keep the fines for himself and thus faces incentives to monitor the users (just like a residual claimant). They find that the presence of the police officer curtails over-extraction substantially. An important feature of their environment is that the police officer cannot fine users who do not over-extract. Thus, in their experiment there is no possibility that police officers abuse their power. Central to the present paper is that leaders are free to implement *any* allocation of team output; thus their ability to discourage free riding relies crucially on their willingness to forego the opportunity to appropriate all the team output for themselves.

There is also a large social psychological literature that considers leadership as a possible solution to social dilemmas (e.g. Messick and Brewer 1983) and studies various types and aspects of leadership such as group identification (Van Vugt and De Cremer), preferences for leadership (van Dijk et al. 2003), and leader commitment and leader fairness (De Cremer and Van Vugt 2002). For instance, Van Vugt and De Cremer report that people prefer to adopt a leader with a legitimate power base (a democratic, elected, internal leader) and that contributions to a public good are raised when groups are managed by a leader who imposes fines on noncontributing group members (i.e. an instrumental leader) rather than by a leader who only encourages members to contribute voluntarily (i.e. a relational leader).

These papers from the experimental social psychology literature differ from the current one in a number of ways. First, incentives provided in these experiments or scenario studies (no monetary payments, or a flat fee) usually do not meet the standard requirements for salience in economic experiments, which may affect the reliability and variance of the decisions. Second, most experiments have only a very limited number of periods (often one) and stop after subjects have decided to install a leader or a sanctioning system. That is they focus on the process that may lead to a change rather than on subjects’ behavior after a change or in different treatments. Third, although many different leader styles have been considered we do not know of a study that uses the type of leader and the environment we consider. Finally, and most importantly, there are no real subjects as leaders and messages by the ‘leader’

are standardized. Not having real subjects as leaders obviously excludes the possibility of studying the endogenous behavior of leaders and how it evolves in light of the incentive problem. This is the main focus of our study.

A SIMPLE MODEL OF TEAM PRODUCTION

Our experiment is based on the following simple model of team production. Each team consists of 4 members, and each member chooses either to contribute to team production (“work”) or not (“shirk”). Let $e_i = 1$ if member i works and $e_i = 0$ if i shirks. Each team member has an endowment of 120 which he loses in case he decides to contribute. This can be interpreted as the disutility of work or as the return from some alternative activity. We suppose that team output is determined by the convex team production function $Q = 60 (\sum e_i)^2$. That is, we assume that efforts of individual team members are complementary: an individual’s marginal product increases with the efforts of others.³ Team output is then shared among team members so that i ’s payoff is

$$\pi_i = q_i + 120(1-e_i),$$

where q_i is the share of team output given to individual i , $0 \leq q_i \leq Q$, and $\sum q_i = Q$.

Consider first what efficiency entails. Total social surplus, as measured by aggregate payoffs, is

$$\sum \pi_i = 60 (\sum e_i)^2 + 120 \sum (1-e_i).$$

Since this function is convex, either all must work or all must shirk for efficiency to be attained. If all team members work aggregate payoffs are 960, while if all shirk aggregate payoffs are 480. Thus, it is efficient for all team members to work.

³ Alchian and Demsetz take this as a defining feature of team production.

Equilibrium under revenue-sharing

How will output be allocated in the absence of a leader? A natural benchmark to consider is that team output is divided equally; that is, $q_i = Q/4$. Under this revenue-sharing rule a team member's payoff is

$$\pi_i = 15 (\sum e_i)^2 + 120(1 - e_i).$$

Standard game theory supposes that each team member maximizes his payoff, taking as given the other team members' choices. Suppose that the rest of the team shirks. Then obviously the remaining team member should also shirk. Thus there is an inefficient equilibrium where all team members shirk. The only other candidate for an equilibrium is where all team members work. If the rest of the team works, the remaining team member earns 240 when working and 255 when shirking. This implies that under revenue sharing shirking is a dominant strategy and complete free riding is the unique equilibrium.

Equilibrium under leader-determined sharing

Suppose now that a leader has the authority to divide up the team output in any way she sees fit. Thus, the game consists of two stages. In the first stage team members (including the leader) simultaneously decide whether to work or shirk, and in the second stage the leader determines how team output will be divided up, where the leader can condition the division on the decisions observed in the first stage.⁴

Again, standard theory is easy to apply to this situation. A selfish leader will take any output for herself in the second stage, giving nothing to other team members. Anticipating this, the other team members should not work, as it is costly to do so and they expect to get no benefit. Since other team members do not work, the leader can either shirk and get a payoff of 120, or work and get at most a payoff of 60. Hence, the leader is better off shirking as well. In the unique subgame-perfect equilibrium each team member shirks, and team output is zero.⁵

⁴ In the terminology of Bianco and Bates, our leader has "enhanced capabilities" rather than "limited capabilities" since not just total output, but also individual efforts, are observable.

⁵ Since individual efforts are not contractable the leader cannot commit to rewarding team members conditional on their efforts. If team output were contractable, in principle the leader could write a forcing contract in which she

Repeated Game Considerations

In real-life settings teams interact repeatedly. For example, a commonplace situation is where a manager evaluates effort of a relatively fixed group annually. This motivated our use of an experimental design where subjects interact repeatedly in fixed groups.⁶ In our experiment the game will be repeated only a finite number of periods. Strictly speaking there is no room for cooperation: in the unique subgame perfect equilibrium all team members shirk in every period for both environments. However, many experimental studies have found that players often manage to cooperate even in finitely repeated games with a unique stage game equilibrium (see, e.g., Selten and Stoecker 1986 or Engle-Warnick and Slonim 2006).

In our revenue-sharing environment contributing to team production is analogous to contributing to a public good. Numerous experiments have shown that many individuals are willing to contribute to a public good, even when it is not in their narrowly defined self-interest to do so (Ledyard). The typical pattern is for contributions to start somewhere halfway between the individually rational level and the collectively efficient level. Contributions then gradually decline with repetition. Thus the evidence from public goods experiments is that incentives to free ride undermine cooperation. We see no reason to expect a different pattern for our revenue-sharing environment. Are there reasons to expect that cooperation will be easier to sustain in the presence of a leader?

promises to reward each team member with 120 (plus some small amount) if output is equal to 960 and to give them nothing if output is below 960. For such a contract only output needs to be verifiable. It is easily checked, however, that with such a contract the leader would have an incentive to free ride and make sure that the maximum output level is not reached so that she need not pay the other team members (Holmstrom 1982). Thus, the best output level a contract would be able to achieve is 540. However, we do not consider such explicit contracts in the sequel.

⁶ Since the main purpose of our experiment is to compare institutions, we use a ‘partners protocol’ (i.e. subjects interact repeatedly in fixed groups) that parallels the pertinent institutional context. If the purpose were to test whether behavior conforms to the equilibrium predictions of the static model, one might consider designs that minimize possibilities for players to use repeated game strategies, either by limiting attention to a one-shot game or

A frequently used measure to examine the scope for cooperation is the maximum discount rate, \bar{r} (or alternatively, the minimum discount factor $\underline{\delta} = 1/(1 + \bar{r})$), that allows ‘grim’ trigger strategies to support cooperation as a subgame perfect equilibrium in the infinitely repeated game (Friedman 1971). The strategies specify cooperating as long as no player has ever deviated from cooperation, but punishing any deviation by reverting to the stage game equilibrium for the remainder of the game. Letting π^N denote the payoff from the Nash equilibrium of the stage game, π^C the stage game payoff from cooperation, and π^D the stage game payoff from the best reply to cooperation (i.e., optimal defection), the maximum discount rate is

$$\bar{r} = \frac{\pi^C - \pi^N}{\pi^D - \pi^C}.$$

In both our treatments the Nash equilibrium payoff in the stage game is unique: $\pi^N = 120$. Under revenue-sharing, mutual cooperation gives each team member a stage payoff of $\pi^C = 240$, and if the other team members cooperate optimal defection yields $\pi^D = 255$. Thus, under revenue sharing $\bar{r} = 8$. Under leader-determined shares the analysis is a bit more involved as we have to determine the leader’s sharing strategy. The leader’s sharing strategy must give the other members at least 120 each. Note that as long as other members get at least 120, they have no myopic incentive to shirk, whatever their discount rates. Thus, as long as the sharing strategy meets this condition, only the team leader has a myopic incentive to deviate. If the leader decides to deviate, the best thing she can do is still to work, just like the others, and then to take all the team output for herself: $\pi^D = 960$. Finally, the maximum discount rate that will prevent the leader from deviating will be when the leader’s gains from cooperation are greatest. This implies a sharing strategy that gives each other member 120, leaving $\pi^C = 600$ for the leader (i.e. $960 - 3 \times 120$). Thus, the threshold discount rate for the leader is $\bar{r} = (600 - 120)/(960 - 600) = 1.33$. This threshold is smaller than that in case of revenue-sharing, implying a more strict condition. Thus, from this perspective we should expect cooperation to be *more difficult* to sustain with a leader than without a leader. The

by using a ‘strangers protocol’ (i.e., randomly changing group composition between repetitions). We suspect that under a strangers protocol cooperation would be more difficult in both institutions.

reason is that defecting from the cooperative outcome is a very tempting option for the leader and much more attractive than “plain” free riding under revenue-sharing.

However, there is a counter-balancing force. For cooperation to be sustainable with revenue sharing, the discount rate of *all* members has to be below the threshold discount rate. The support of all team members is essential. With leader-determined shares, only the leader’s discount rate matters. She should be (very) patient, but the discount rates of the other team members do not matter. Note also the different ways the benefits of cooperation are shared: under revenue-sharing all share equally, whereas under leader-determined shares all the benefits of cooperation accrue to the leader. If players care about the distributions of earnings as well as own earnings, the leader may not be able to induce other team members to work by giving them 120 each. Indeed, if players are sufficiently averse to disadvantageous inequality (c.f. Fehr and Schmidt 1999), the leader will have to offer each member 240, and, as a result, the incentives for the leader to share team output become even weaker.

In summary, it is not clear that the introduction of a leader will help to improve team performance. Under either revenue-sharing or leader-determined shares, standard theory predicts that all team members will shirk. Even in a repeated game setting, installing a leader is not predicted to have an unambiguously positive effect.

EXPERIMENTAL DESIGN

The experiment was conducted in Spring 2004 using subjects recruited by e-mail from a pool of undergraduate students at Tilburg University. Five sessions were run, and 80 subjects participated in total, with no subject participating in more than one session.

All sessions were computerized and used an identical protocol. Upon arrival, subjects were randomly assigned to a group of four subjects and randomly assigned an identification number: subject 1, 2, 3 or 4; these groups and identification numbers were kept fixed throughout the session. Subjects were

then given a written set of instructions which the experimenter read aloud.⁷ Subjects were allowed to ask questions by raising their hands and speaking to the experimenter in private. Subjects were not allowed to communicate with one another throughout the session, except via the decisions they entered on their terminal.

The decision-making phase of the session consisted of 15 rounds. At the beginning of each round the four subjects in each group were prompted to choose between two actions: A or B. The subjects made these choices simultaneously, and each choice of B increased group output as described in Table 1.

Table 1. Team Production Function

Number of group members choosing option B	0	1	2	3	4
Group Output	0	60	240	540	960

This table appeared in the instructions given to subjects, and its simplicity stems from our design choices of four-person groups and binary actions. Obviously, explaining the non-linear team production function to subjects becomes more difficult as the set of feasible actions expands. We opted for a design that made explaining the incentive structure as simple as possible. One implication of this is that a subject's action has a very clear interpretation: the subject is either contributing to group output or she is not. This element of the design may make it easier for subjects to interpret free riding and act on any concerns they may have about fairness.

Subjects' earnings were determined as their share of group output plus an additional 120 points if they chose option A. At the end of each round, subjects were informed of all group members' choices and earnings. Thus, in terms of the model presented above, choosing A corresponds to shirking ($e_i = 0$), choosing B corresponds to working ($e_i = 1$).

⁷ Reading the instructions aloud made the information and move structure public knowledge. A copy of the instructions for the experiment can be found in Appendix A.

We employed two experimental treatments. In our revenue-sharing treatment, all subjects received an equal share of their group's output irrespective of their own decision. In our leader-determined shares treatment subject 1 was informed of the decision of each group member and then decided how to allocate the group output among the group members. Subjects were informed that subject 1 could "choose any division he or she wishes as long as no subject receives a negative share and the shares of the four subjects add up to group output." Both treatments were run with ten different groups of four subjects each.

At the end of the experiment subjects were privately paid €0.004 per point, based on their accumulated point earnings from all 15 rounds. Sessions lasted between 40 and 70 minutes, and subjects earned on average €10.75 (with a minimum of €4.75 and a maximum of €18).⁸

RESULTS

Treatment Effects

We first present the results from treatment comparisons. Table 2 presents averages and standard deviations of team effort, team production, and team earnings for each treatment. For each group we averaged over all fifteen rounds and then report the average and standard deviations of the resulting two sets of ten observations. We also report p-values for tests of treatment effects in the last row.⁹ Relative to the revenue-sharing treatment, introducing a leader to allocate team output results in significant increases in average team effort ($p = 0.038$), output ($p = 0.034$), and earnings ($p = 0.031$). Under revenue-sharing, on average 1.75 team members work each round producing 314 units of output, while under leader-determined shares 3.03 team members work producing 682 units of output. Thus, these results clearly show that teams in which the leaders determine the shares work harder and produce more output than teams in which the shares are equally distributed. Because team earnings are not monotonically increasing in the amount of effort, it is not immediately evident that higher effort supply translates into changes in

⁸ At the time of the experiment the exchange rate was approximately €1 = \$1.21.

⁹ Unless otherwise noted, tests are based on (strictly independent) group level observations.

earnings.¹⁰ However, as the last column shows, the presence of a leader also enhances team performance from a welfare perspective: introducing a leader results in a significant increase in team earnings from 583.6 to 798.0 ($p = 0.031$).

Table 2. Treatment Comparisons

	Average Team Effort (Standard Deviation)	Average Team Output (Standard Deviation)	Average Team Earnings (Standard Deviation)
Revenue-Sharing	1.75 (1.08)	314.0 (276.9)	583.6 (151.2)
Leader-determined Shares	3.03 (1.10)	682.0 (301.3)	798.0 (169.9)
p-value	0.038	0.034	0.031

Note: p-values are based on Wilcoxon ranksum tests using the group-level data reported in Appendix B.

Standard theory predicts that teams will attain 50% of maximum attainable earnings in both treatments. In fact we find that on average teams attain 61% of maximum possible earnings in our revenue-sharing treatment. Thus, on average, team performance is a little better than predicted by standard theory, but still falls far short of efficient levels. In fact, three of ten teams earned less on average than the equilibrium level of 480 points per round, and the hypothesis that a team is equally likely to earn more or less than equilibrium predicts cannot be rejected ($p = 0.344$, using a two-sided Binomial test). In contrast, in teams where the leader determines how team output is allocated, all ten teams earned more than predicted by equilibrium, and average earnings are 83% of the maximum possible.

In each treatment 150 stage games were played, each resulting in team earnings of either 420 (if one team member worked), 480 (if none or two members worked), 660 (if three members worked), or 960

¹⁰ With complete free-riding team earnings are 480, whereas a single effort decision lowers team earnings to 420. For the team to earn more than in equilibrium, either 3 or 4 team members must supply effort.

(if all worked) points. Table 3 gives the percentages of games resulting in each outcome. While in the revenue-sharing treatment most games resulted in payoffs at or below the equilibrium level, in the leader-determined shares treatment most games resulted in the efficient outcome. We next look for the sources of these differences by examining choices in each treatment separately.

Table 3. Stage Game Outcomes.

	Percentage of rounds in which number of workers equals				
	0	1	2	3	4
Revenue-Sharing	27	24	16	14	19
Leader-determined Shares	14	5	7	11	63

Effort and Earnings under Revenue-Sharing

Figure 1 shows how average team effort levels, aggregated over all ten groups of our revenue-sharing treatment and ten groups of the leader-determined shares treatment, developed across rounds. In the first round the proportions of subjects choosing to work do not differ significantly across treatments ($p=0.252$, based on Fisher's exact test applied to individual level data). Afterwards, differences between treatments emerge as subjects defect more rapidly from cooperation in revenue-sharing than in the leader-determined shares treatment.

--- Figure 1 here ---

We will focus first on the revenue-sharing treatment. In the first round 30 of 40 subjects chose to work. Four groups attained the efficient outcome, so that in these groups each team member earned 240 points. The other six groups included shirkers, and these shirkers earned 120 points more than fellow team members who chose to work.

Effort declined over subsequent rounds: after round seven, most subjects were shirking, and in the last round of the session all but four subjects shirked. This pattern is similar to the pattern of declining voluntary contributions observed in public goods experiments. As shown in Figure 2, this pattern of declining effort induces a similarly declining pattern in team earnings.

--- Figure 2 here ---

The averages shown in Figures 1 and 2 may mask substantial variation across groups. To examine this we consider Figure 3, which shows how effort supply develops in each individual team. The four teams that attained efficient levels of production in round one also attained efficient production in round two (and two of these teams, labeled Equal_3 and Equal_4, managed to sustain efficient production levels for most of the experiment). In contrast, in five of the six teams where a team member shirked in round one, production fell in round two. Although there appear to be several attempts to renew cooperation in these teams, the general trend of declining production and earnings is evident. Recall that in order for a team to earn more than equilibrium earnings, at least three team members must work. Across all ten teams, this occurred in 56% of cases in the first five rounds, 34% in the second five, and only 10% in the last five. Moreover, a similar pattern is observed in every team: the number of times the team attained greater than equilibrium earnings was at least as high in the first five rounds as the second five rounds, which in turn was at least as high as in the last five rounds. Thus, although there is some heterogeneity across teams and some teams are more successful than others, the general picture from our revenue-sharing treatment is that teams are unable to overcome the underlying free-riding incentives.

--- Figure 3 here ---

Effort and Earnings under Leader-Determined Shares

Turning now to our leader-determined shares treatment, Figure 1 shows that initial effort supply is high, remains high for most of the experiment, and then declines abruptly in the later rounds. Figure 2 shows that the same holds for average team earnings, which stay close to 840 until round 12.

Interestingly, it is not just the leaders, but also the other team members who benefit from the presence of a leader. Leaders earn 238.8 points on average, which is significantly more than the average earnings of 186.4 points of the other team members (two-sided Wilcoxon matched-pairs signed-ranks test $p = 0.016$). The latter earnings, however, are significantly higher than the average earnings (145.9) in the revenue-sharing treatments (two-sided Wilcoxon ranksum test $p = 0.089$). The reason for this is that leaders allocate a larger share of team output to the other working team members than what would be required (120) to compensate them for the cost of effort. On average, the leader allocates 187.3 points to another team member who chose to work (and only 3.2 points to another team member who chose to shirk). Notably, when all team members work, the average share that leaders allocate to another team member is 224.6 (i.e. 23.4% of team output). To see that the leader's behavior gives an incentive to other team members to work, consider what happened when three team members (including the leader) worked, and one team member shirked. In this situation, which occurred at least once in seven groups, the leader allocated herself 170 points, the other two workers 165 points each, and the shirker 26 points on average.

To appreciate fully the different strategies employed, it is useful to discuss the behavior of some teams in detail. In the first round 35 of 40 subjects chose to work (including nine of ten leaders), and all ten teams earned more than the equilibrium level of earnings. How did leaders distribute team output? In seven of the ten teams, output was evenly divided among those team members who chose to work. That is, leaders allocated 180 points to a worker if output was 540 and allocated 240 points to a worker if output was 960. Shirkers were allocated zero, so they earned 120 points for the round. In all these teams output either remained the same or increased in round two. In two of the other three teams the leader kept a disproportionately large share of output, and in the one remaining team the leader allocated 120 to one worker, 360 to another worker and 240 to himself and the last worker. In each of these three teams output fell in round two.

Figure 4 shows how effort levels developed in each team over the course of the experiment in the leader-determined shares treatment. The figure shows that even though teams with a leader performed reasonably well on average, installing a leader is not a sufficient condition for success. Some teams performed very well, while others did very poorly. What is important is how the leader used her reward power. The three teams where leaders did not divide first-round output equally among workers are shown in the first column; in the other columns are the seven teams whose leaders divided output equally among workers. To understand the further development of effort choices it is useful to consider these teams separately.

--- Figure 4 here ---

Consider the seven teams whose leaders divided output equally among workers in the first round and allocated nothing to shirkers. Five of these teams managed to keep earnings at the maximum (i.e. at 960) from round 2 until at least round 13. Another team attained 91% of maximum earnings until round 11, and the remaining team attained 76% of maximum earnings until round 14. Now consider the three teams in which output was not evenly divided among the workers. For two of these teams earnings were at or below equilibrium levels from round 3 onwards. Finally, in one of the three teams (labeled Leader_1 in Figure 4), the leader distributed output equally in round 2 (giving 60 points to all team members, including shirkers), and then in round 3 began using the strategy of dividing output equally among workers and allocating nothing to shirkers. This leader seems to have eventually learned a successful strategy since her team output stayed at efficient levels for most of the experiment.

Table 4 summarizes the decisions taken by leaders. In 21% (32/150) of the stage games the leader shirked, and if there was any output to divide, the leader almost always divided it in a way that favored herself. In 79% (118/150) of the stage games the leader worked, and in 82% (97/118) of these the leader divided output equally among workers.

Table 4. Leader Decisions (number of choices)

	Division of Output					Total
	No output to divide	Equal share among workers	Equal share among group	Favor self	Other	
Leader Shirks	21	0	1	10	0	32
Leader Works	0	97	1	11	9	118

The simple strategy whereby leaders divide output equally among workers was employed by seven of ten leaders for most of the experiment.¹¹ Given this leader behavior, other team members had no incentive to free ride. As already seen in Figure 1, team effort levels remained high until near the end of the experiment, indicating that the leader was successful in encouraging team members to work. Figure 2 shows the earning implications: in contrast to the revenue-sharing treatment, earnings in the leader-determined shares treatment were substantially above the equilibrium level, though they declined somewhat toward the end of the session.

DISCUSSION

What is the explanation for this effect of leadership? In a previous section we discussed how repeated game considerations might enable a team to perform efficiently. However, in that model a selfish leader allocates just enough to the other team members to induce them to work while our leaders are typically more generous. On average our leaders allocate more than enough output to workers to compensate them for their effort, and indeed successful leaders typically give workers as much as they give themselves. This might suggest that leaders have an intrinsic preference for “fair” distributions. A widely used model

¹¹ An eighth leader used this strategy five times, but followed a somewhat more resentful and selfish strategy in other rounds: when a subject chose to start working again after having shirked in the previous round, this leader ‘punished’ the subject by assigning less than an equal share (even in rounds thereafter). This happened to two subjects. The other team member did cooperate in all rounds, but the leader assigned less to this subject than to himself in rounds 7 and 8. These decisions, and other decisions that are difficult to summarize, fall in the “Other” category of Table 4.

where payoffs refer to distributional considerations as well as to own-earnings is Fehr and Schmidt's model of inequality aversion. However, according to their model a sufficiently inequality averse leader is predicted to use her reward power to equalize earnings; instead, our leaders typically give very little to shirkers and equalize the earnings of the working team members. Thus, to the extent that the allocations reflect fairness considerations, our results suggest that leaders reward team members in accordance with their inputs, which is more in line with equity theory (Homans 1974).

Even so, closer scrutiny of endgame effects suggests that most of the apparently "fair" leaders may be only "acting fairly" in earlier rounds. Of the eight leaders who encouraged team members to work by rewarding them, only two of these pursued this policy through to the end of the session. The other six decided to take all team output for themselves in some round before the end of the session. Thus, only two out of ten leaders seem intrinsically motivated by fairness considerations.

Nevertheless, the mere possibility that a leader is truly motivated by fairness may also induce an opportunistic leader to act fairly in the early rounds in order to increase the belief held by the other team members that she is a truly fair leader and thereby induce them to work. To examine this we analyze the equilibrium of an incomplete information game along the lines of the reputation model by Kreps and Wilson (1982).¹² Here we outline only the main features of the game. Further details are given in Appendix C.

The central assumption is that there are two possible types of leaders: an opportunistic type who tries to maximize her expected payoffs, and a fair type who is intrinsically motivated to share team output equally among workers (but not shirkers). Let q_1 denote the prior probability that a leader is of the fair type. This belief is updated after each round in accordance with Bayes' rule, and members maximize their

¹² We note that Hollander's (1958) theory of idiosyncratic credits could be an alternative explanation for the fact that some leaders take more than an equal share of team output after having successfully managed the group for a substantial period of time. According to this (social psychological) theory credits are earned through the demonstration of competency in helping the group achieve goals and conforming to the group norms.

expected payoffs given their current beliefs. An opportunistic type of leader maximizes her expected payoffs taking into account the impact of her strategy on the beliefs of the other team members.

The intuition for the equilibrium relies on the opportunistic type of leader behaving as if she is a fair type. The opportunistic leader knows that she can "take" only once because the members will then know that she is not of the fair type and they never work again, so the best thing for her is to take as late as possible. This induces the opportunistic leader to behave as a fair type (that is, sharing output equally) in the earlier rounds of the game. Since the other team members anticipate that both a fair and an opportunistic type will share output in the earlier rounds of the game, it is optimal for them to work in these rounds. It is only in the later rounds that they need to worry that an opportunistic type will take.

The main effect of combining repetition with incomplete information is that repetition allows cooperation to be sustained more easily (i.e. with a lower q_1) because opportunistic leaders have an incentive to mimic fair leaders. The intuition is seen most easily by comparing a one-shot game with a two-fold repetition. Note that in a one-shot game the non-leaders can get 120 by shirking and get 240 with probability q_1 by working. Thus, if $q_1 < \frac{1}{2}$, non-leaders (and opportunistic leaders) shirk in the equilibrium of the one-shot game. Now suppose there are two rounds. We show that if $q_1 > \frac{1}{4}$ there is an equilibrium in which team members work in the first round, and with positive probability team members work in the second round as well. The central features of the equilibrium are that (i) opportunistic leaders play a mixed strategy in round one, sharing with probability s , and (ii) in the event that a leader acts fairly in round one, non-leaders play a mixed strategy in round two and work with probability m .

Note that in the second round an opportunistic leader will keep any output for herself, so non-leaders could get 120 by shirking or 240 with probability q_2 by working (where q_2 is the probability that the leader is of the fair type, given the observed leader decision from round one). The non-leaders will be indifferent between working and shirking when $q_2 = \frac{1}{2}$. Of course, if the leader takes in round one then she reveals herself to be opportunistic and $q_2 = 0$. If, however, the leader acts fairly the non-leaders must

use Bayes Rule to update the prior belief that the leader is of the fair type: $q_2 = \frac{q_1}{q_1 + (1 - q_1)s}$.

Rearranging this expression gives $s = \frac{q_1}{1-q_1} \frac{1-q_2}{q_2}$. Thus if $s = \frac{q_1}{1-q_1} \frac{1-\frac{1}{2}}{\frac{1}{2}} = \frac{q_1}{1-q_1}$ this leaves non-leaders

indifferent between working and shirking in round two. In order for the opportunistic leader to share with some probability $0 < s < 1$ in round one, her expected payoff from sharing must equal her expected payoff from taking. If team output is 960, this requires that $240 + 960m + 60(1-m) = 960 + 120$, or $m = 13/15$.

Finally, for all team members to work in round one we require that the expected payoff from working is as high as that from shirking. By shirking a non-leader expects to get $120 + 120$. By working she expects to get $240(q_1 + (1-q_1)s) + 120$. With $s = \frac{q_1}{1-q_1}$ this implies that we require $q_1 \geq \frac{1}{4}$. A similar logic can be

applied to the 15 round game. If $q_1 \geq \frac{1}{2}^{15} \approx 0.00003$ there is an equilibrium in which all team members work in the very first round.

The Perfect Bayesian Equilibrium of the game has the following main properties. First, all team members work in the early rounds of the game, even if the prior probability that a leader is of the fair type is quite small. For example, with a prior probability of $q_1 = 0.2$ (i.e., 2 out of 10 leaders are of the truly fair, as seems to be the case in our experiment) all members work until round 14, when they start shirking with positive probability. Second, an opportunistic leader shares output equally among workers in the earlier rounds of the game. With $q_1 = 0.2$ for example, she does so until round 13, when she starts playing a mixed strategy. The probability that an opportunistic leader shares decreases further in round 14 and drops to 0 in round 15. Third, if in any round the leader does not share, then no member (including the leader) will ever work in a later round. Finally, the probability that output collapses increases towards the end of the game. Moreover, it is more likely that this happens in *response* to an opportunistic leader taking than it is in *anticipation* of an opportunistic leader taking. In other words, breakdown is more likely to be caused by the leaders than by the other team members.

These qualitative features of the equilibrium accord remarkably well with the data of our leader-determined shares treatment. First, the proportion of team members working is 82% in rounds 1 to 13 and is fairly constant over these rounds (see Figure 1). The proportion drops to 60% in round 14 and to 13% in

round 15. Second, there are six teams in which in some round the leader takes all output for herself. This occurs only once in rounds 1 to 12 (in round 11 to be precise), once in round 13, and four times in round 14, so typically this happens towards the end of the game. Third, in the six teams in which the leader takes all output for herself in some round, only three of the forty remaining decisions (i.e., 8%) involve working. Finally, there are nine groups in which output is at the maximum level in some round and then falls in the next round. In seven of these groups output falls in response to the leader not sharing equally in the previous round, while there are only two groups in which output falls in spite of the fact that the leader shared equally in the previous round.

Summing up, although the fit is surely not perfect, and there may be alternative ways of interpreting our data¹³, it is fair to say that several of the qualitative features of our data are nicely captured by a reputation model in the spirit of Kreps and Wilson.

CONCLUSION

Under revenue-sharing free-riding incentives are difficult to overcome in our experiment and the supply of effort decreases steadily with repetition. In contrast, in the leader-determined shares treatment free-riding incentives are effectively mitigated. Most team leaders use a simple strategy of dividing team output equally among workers and allocating nothing to shirkers. This strategy turns out to be very successful, and several teams attain a level of efficiency of 100% for a substantial period of time. The implication of this behavior is that team performance is significantly better with the introduction of a leader.

However, installing a leader is no guarantee for success. In our experiment, two out of ten teams with a leader perform poorly. Thus a critical factor is not just the availability of reward power, but how it is used. What characterizes these poorly-performing teams is that their leaders fail to send unambiguous signals of fairness. As a consequence, cooperation by the other team members breaks down.

¹³ For example, as a referee pointed out, since smart egoistic and fair leaders will both share in the early periods, the team member is possibly learning more about the short-sightedness of the leader than about his or her fairness.

While standard theory suggests that installing a hierarchy in order to overcome free riding will merely shift opportunistic incentives, our results suggest that this prediction may be too pessimistic. We find that the presence of a team leader can be very effective. In particular, the combination of repeated game considerations *and* the possibility that a leader is truly motivated by a preference for fairness can account for the success of having one team member with reward power and residual claimancy.

Obviously, the results in this paper are obtained using a very stylized set-up. Future research could study extensions and examine factors that influence the effectiveness of a hierarchy. First of all, the team leader in our experiment is well equipped. Individual actions are perfectly and costlessly observable. An obvious extension would be to examine the effectiveness of leadership in an environment with imperfect monitoring, for example, one in which individual effort levels are unobservable or only observable with noise. Second, due to the binary effort decision and the symmetric set-up of our experiment, it is quite obvious what a fair allocation of team output is. When the team members differ in their abilities or endowments, there may be competing fairness standards (e.g., equity versus equality). It will be interesting to examine whether this erodes the constraining force of morality (e.g., due to self-serving biases) and how subjects react to that. A third extension would be to examine endogenous rather than exogenous leadership. If a team has the option to install a leader, will they do so, and if so how? Furthermore, will an endogenously selected leader be more effective, or will she perhaps feel entitled to a larger share of the pie? It is likely that several of these alternative settings will make the effectiveness of a hierarchy more precarious. How it will compare to alternative arrangements in these settings is less clear *ex ante*.

APPENDIX A

The instructions given to participants are reproduced below. Text indicated by [] was given only to subjects in the leader-determined shares treatment, and text indicated by {} was given only to subjects in the revenue-sharing treatment.

Instructions

This session is part of an experiment in the economics of decision making. If you follow the instructions carefully and make good decisions, you can earn a considerable amount of money. At the end of the session you will be paid your earnings, in private and in cash.

It is important that you do not talk to any of the other participants until the session is over. If you have a question at any time raise your hand and a monitor will come to where you are sitting to answer it.

The session will consist of 15 rounds. During the session your earnings will be denoted in points. At the end of the experiment you will be paid an amount based on your total point earnings from all 15 rounds. Points will be converted to cash using an exchange rate of 1 point = 0.4 eurocents, which means 250 points = 1 euro.

Throughout the session you will be in a group with the same three other participants. The assignment of participants to groups is randomly determined by the computer at the beginning of the session. You will not be told who is in your group.

The four members of a group are labelled subject 1, subject 2, subject 3, and subject 4. Your subject number is randomly determined at the beginning of the session and will remain the same for all rounds.

Choices and earnings

In each round, each group member decides between option A and option B. If you choose option B this increases group output. Group output is determined by the following table.

Number of group members choosing option B	0	1	2	3	4
Group Output	0	60	240	540	960

As will be explained below, group output is divided among the group members. Your earnings for the round are: the share of group output you receive, plus an additional 120 points if you chose option A.

Stages and information

Every round has the same structure, and consists of a number of stages.

1. All group members decide between option A and option B.

[2. Subject 1 is informed of other group-members' decisions and group output.]

[3. Subject 1 decides how to divide group output among the four group members. Subject 1 can choose any division he or she wishes as long as no subject receives a negative share and the shares of the four subjects add up to group output.]

[4. All members of the group are informed of group output, and each group-member's decision, share of group output and earnings. If you want to retain this information for later rounds you can record it on the attached Record Sheet.]

{2. The group output is equally divided among the group members: each group member receives one-quarter (25%) of the group output.}

{3. All members of the group are informed of group output, and each group-member's decision, share of group output and earnings. If you want to retain this information for later rounds you can record it on the attached Record Sheet.}

Ending the Experiment

At the end of round 15 the experiment ends and each participant is paid his or her accumulated earnings, in private and in cash.

APPENDIX B

Table B-1 presents group-level data on average effort, output, and earnings under the two treatments.

These underlie the statistical tests reported in the text.

Table B-1. Team Effort, Output, and Earnings

Group	team effort per round	team output per round	team earnings per round	Group	team effort per round	team output per round	team earnings per round
Equal_3	3.87	912	928	Leader_3	3.93	932	940
Equal_4	3.00	652	772	Leader_10	3.87	912	928
Equal_2	1.60	296	584	Leader_4	3.67	868	908
Equal_9	2.07	332	564	Leader_2	3.67	868	908
Equal_7	2.00	320	560	Leader_8	3.47	832	896
Equal_10	1.73	248	520	Leader_1	3.53	820	876
Equal_1	1.53	212	508	Leader_7	3.20	672	768
Equal_8	0.87	92	468	Leader_6	2.93	632	760
Equal_5	0.60	60	468	Leader_9	1.33	192	512
Equal_6	0.27	16	464	Leader_5	0.73	92	484
Average	1.75	314	584	Average	3.03	682	798

APPENDIX C

The reputation model is based on the following assumptions:

1. There are two types of leader: a fair type and an opportunistic type. A fair type always shares team output equally among workers. An opportunistic type maximizes expected payoffs. The prior probability that a leader is of the fair type is q_1 .
2. The three non-leaders maximize expected payoffs. They act in concert; that is, they play correlated (mixed) strategies.¹⁴
3. Beliefs are updated in accordance with Bayes' rule.

The analysis closely resembles that of Kreps and Wilson. The reader can consult that reference or a standard textbook for further detail. We use the following notation, where $t \in \{1, \dots, 15\}$ refers to the round of the game.

- q_t : members' belief that the leader is of the fair type in round t .
- $m_t(q_t)$: probability that the three non-leaders work in round t given their belief q_t .
- $s_t(q_t)$: probability that an opportunistic leader shares output equally in round t .

The following strategies and beliefs constitute a Perfect Bayesian Equilibrium:

$$m_t(q_t) = \begin{cases} 0 & \text{if } q_t < \frac{1}{2}^{16-t} \\ \frac{13}{15} & \text{if } q_t = \frac{1}{2}^{16-t} \\ 1 & \text{if } q_t > \frac{1}{2}^{16-t} \end{cases}$$

$$s_t(q_t) = \begin{cases} \frac{q_t \cdot 1 - (\frac{1}{2})^{15-t}}{1 - q_t \cdot (\frac{1}{2})^{15-t}} & \text{if } q_t \leq (\frac{1}{2})^{15-t} \\ 1 & \text{if } q_t > (\frac{1}{2})^{15-t} \end{cases}$$

$$q_{t+1} = \begin{cases} 0 & \text{if members worked and leader did not share in } t \\ q_t & \text{if members did not work in } t \\ \max\{q_t, (\frac{1}{2})^{15-t}\} & \text{if members worked and leader shared in } t. \end{cases}$$

¹⁴ The assumption of correlated strategies is made for convenience. It is possible to allow for independent mixed strategies, but it complicates the equilibrium considerably without affecting its qualitative properties.

The general structure of the equilibrium is that for an initial phase, all team members work and opportunistic leaders act fairly. In some round t' , the opportunistic leader begins to take with positive probability, $0 < s_{t'} < 1$. If she decides to share, then in the next round, she works and non-leaders work with some probability $0 < m_{t'+1} < 1$. Thus in a second phase of the game continued cooperation rests on the outcomes of the players' mixed strategies. Finally, there can be a third phase of non-cooperation due to an endgame effect: it is possible that in some round before the last the leader takes or non-leaders shirk, after which cooperation ceases.

If $q_1 = 1/5$ (which happens to coincide with the proportion of intrinsically fair leaders in our data), the following properties, mentioned in the main text, hold along the equilibrium path:

1. $q_1 > (1/2)^{16-t}$ for $t \leq 13$. Thus all team members (including the leader) work up to and including round 13.

2. An opportunistic leader shares output equally as long as $q_1 > (1/2)^{15-t}$, that is, up to and including round 12. As a result $q_{13} = q_{12} = \dots = q_1$. Round 13 is the first in which an opportunistic leader plays a mixed strategy with a probability of sharing equal to $s_{13}(q_{13}) = s_{13}(q_1) = \frac{q_1 \frac{1 - 1/2^2}{1 - q_1}}{1/2^2} = 3/4$. By Bayes'

rule, we have $q_{14} = 0$ if the leader did not share in round 13 and $q_{14} = (1/2)^2$ if the leader shared in round 13. In the latter case, an opportunistic leader shares with probability

$$s_{14}(q_{14}) = \frac{q_{14} \frac{1 - 1/2}{1 - q_{14}}}{1/2} = 1/3.$$

3. Irrespective of the outcome in round 14, in the final round we have $s_{15} = 0$. If in any round t the leader does not share, we have $q_{t+1} = \dots = q_{15} = 0$. This implies that no team member will ever work after round t .

4. Breakdown of output is more likely to be caused by an opportunistic leader not sharing than by the other team members shirking. In round 13, the probability that the leader does not share is $1 - s_{13}$. This implies that the probability that output breaks down in round 14 due to an opportunistic leader is $(1 - q_1)(1 - s_{13}) = (4/5)(1/4) = 0.2$. If the leader shares in round 13, the members shirk in round 14 with

probability $1 - m_{14}(q_{14}) = 1 - 13/15 = 0.13$. Thus, the probability that output breaks down due to shirking by non-leaders is $(q_1 + (1 - q_1)s_{13})(1 - m_{14}) = ((1/5) + (4/5)(3/4))(2/15) = 0.11$. Similarly, in round 15, the probability that output breaks down due an opportunistic leader not sharing in round 14 is $(1 - q_1)(1 - s_{14}) = (4/5)(2/3) = 0.53$. The probability that output breaks down due to shirking by non-leaders is $(q_1 + (1 - q_1)s_{14})(1 - m_{15}) = ((1/5) + (4/5)(1/3))(2/15) = 0.06$.

Accepted Manuscript

REFERENCES

- Alchian, A. A., Demsetz, H., 1972. Production, information costs, and economic organization. *American Economic Review* 62, 777-795.
- Bass, B.M., 1985. *Leadership and Performance Beyond Expectations*. New York: Free Press.
- Bianco, W. T., Bates, R. H., 1990. Cooperation by design: Leadership, structure and collective dilemmas. *American Political Science Review* 84, 133-147.
- Binmore, K., 1998. *Game Theory and the Social Contract. Volume 2: Just Playing*, Cambridge MA: MIT Press.
- Bochet, O., Page, T., Putterman, L., 2006. Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior and Organization* 60, 11-26.
- Bowles, S., Carpenter, J., Gintis, H., 2001. Mutual monitoring in teams: Theory and evidence on the importance of residual claimancy and reciprocity. Working Paper, University of Massachusetts.
- Calvert, R., 1995. The rational choice theory of institutions: Implications for design. In: Weimer, D. (Ed.). *Institutional Design*, 63-94. Boston: Kluwer.
- De Cremer, D., Van Vugt, M., 2002. Intergroup and intragroup aspects of leadership in social dilemmas: A relational model of cooperation. *Journal of Experimental Social Psychology* 38, 126-136.
- Engle-Warnick, J., Slonim, R., 2006. Inferring repeated game strategies from actions: Evidence from trust games. *Economic Theory* 28, 603-632.
- Fehr, E., Gächter, S., 2000. Cooperation and punishment in public goods experiments. *American Economic Review* 90, 980-994.
- Fehr, E., Schmidt, K., 1999. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114, 817-868.
- Frey, B., Bohnet, I., 1996. Cooperation, communication and communitarianism: An experimental approach. *Journal of Political Philosophy* 4, 322-336.
- Friedman, J., 1971. Noncooperative equilibrium for supergames. *Review of Economic Studies* 28, 1-12.
- Gächter, S., Renner, E., 2004. Leading by example in the presence of free riding incentives. Working Paper, University of Nottingham.
- Hollander, E. P., 1958. Conformity, status and idiosyncrasy credit. *Psychological Review* 65, 117-127.
- Holmstrom, B., 1982. Moral hazard in teams. *Bell Journal of Economics* 13, 324-340.
- Homans, G., 1974. *Social Behavior: Its Elementary Forms (2nd Ed.)*. New York: Harcourt Brace Jovanovich.
- Isaac, R. M., Walker, J., 1988. Communication and free riding behavior: The voluntary contribution mechanism. *Economic Inquiry* 26, 585-608.

Kreps, D., Wilson, R., 1982. Reputation and incomplete information. *Journal of Economic Theory* 27, 253-279.

Ledyard, J., 1995. Public goods: A survey of experimental research. In: Kagel, J., Roth, A. (Eds.). *The Handbook of Experimental Economics*, 111-194. Princeton: Princeton University Press.

Meidinger, C., Villeval, M-C., 2002. Leadership in teams: Signaling or reciprocating? GATE Working Paper 2002-16.

Messick, D. M., Brewer, M. B., 1983. Solving social dilemmas: A review. In: Wheeler, L., Shaver, P. (Eds.). *Review of Personality and Social Psychology: Vol 4*, 11-44. Beverly Hills, CA: Sage.

Miller, Gary J. 1992. *The Political Economy of Hierarchy*. Cambridge: Cambridge University Press.

Moxnes, E., van der Heijden, E., 2003. The effect of leadership in a public bad experiment. *The Journal of Conflict Resolution* 47, 773-795.

Nalbantian, H. R., Schotter, A., 1997. Productivity under group incentives: An experimental study. *American Economic Review* 87, 314-341.

Ostrom, E., Walker, J., Gardner, R., 1992. Covenants with and without a sword: Self-governance is possible. *American Political Science Review* 86, 404-417.

Potters, J., Sefton, M., Vesterlund, L., 2007. Leading-by-example and signaling in voluntary contribution games: An experimental study. *Economic Theory* 33, 169-182.

Sefton, M., Shupp, R., Walker, J., 2007. The effect of rewards and sanctions in provision of public goods. *Economic Inquiry* 45, 671-690.

Selten, R., Stoecker, R., 1986. End behaviour in sequences of finite prisoner's dilemma supergames. *Journal of Economic Behavior and Organization* 7, 47-70.

Van Dijk, F., Sonnemans, J., van Winden, F., 2001. Incentive systems in a real effort experiment. *European Economic Review* 45, 187-214.

Van Dijk, E., Wilke, H., Wit, A., 2003. Preferences for leadership in social dilemmas: Public good dilemmas versus common resource dilemmas. *Journal of Experimental Social Psychology* 39, 170-176.

Van Vugt, M., De Cremer, D., 1999. Leadership in social dilemmas: The effects of group identification on collective actions to provide public goods. *Journal of Personality and Social Psychology* 76, 587-99.

Vyrastekova, J., van Soest, D., 2003. Centralized common pool management and local community participation. *Land Economics* 79, 500-514.

FIGURES




Figure 1. Average team effort levels




Figure 2. Average team earnings




Figure 3. Team effort levels: revenue-sharing treatment




Figure 4. Team effort levels: leader-determined shares treatment