



Excludability: A laboratory study on forced ranking in team production [☆]



Rachel Croson ^a, Enrique Fatas ^b, Tibor Neugebauer ^{c,d,e,*}, Antonio J. Morales ^f

^a University of Texas at Arlington, United States

^b University of East Anglia, United Kingdom

^c University of Luxembourg, Luxembourg

^d UJI, Spain

^e UCSB, United States

^f University of Malaga, Spain

ARTICLE INFO

Article history:

Received 9 October 2012

Received in revised form 10 February 2015

Accepted 7 March 2015

Available online 14 March 2015

JEL classification:

C72

C92

H41

Keywords:

Cooperation

Competition

Exclusion

Linear public goods games

Best-shot mechanism

Weakest-link mechanism

ABSTRACT

Exclusion has long been employed as a common disciplinary measure against defectors, both at work and in social life. In this paper, we study the effect of excludability – exclusion of the lowest contributor – on contributions in three different team production settings. We demonstrate theoretically and experimentally that excludability increases contributions. Excludability is particularly effective in production settings where the average or maximum effort determines team production. In these settings, we observe almost immediate convergence to full contribution. In settings where the minimum effort determines team production, excludability leads to a large increase in contributions only if the value of the excluded individual's contribution to the public good is redistributed among the included individuals.

© 2015 Published by Elsevier B.V.

1. Introduction

The social and economic success of organizations and societies depends on the cooperative interactions of motivated individuals. In organizations, teams are often employed in traditional management functions because they can execute tasks better, learn faster, and change more easily than traditional structures. In societies, cooperation in groups can yield

[☆] Enrique Fatas gratefully acknowledges the financial support of the ESRC Network for Integrated Behavioural Science (NIBS). The scientific research presented in this publication has also been given financial support by the National Research Fund of Luxembourg (F2R-368 LSF-PMA-13SYSB). We also thank participants at the Economic Science Association meetings for helpful suggestions and comments.

* Corresponding author at: University of Luxembourg – FDEF, Luxembourg School of Finance, 4, rue Albert Borschette, L-1246, Luxembourg. Tel.: +352 466644 6285; fax: +352 466644 6811.

E-mail addresses: croson@uta.edu (R. Croson), e.fatas@uea.ac.uk (E. Fatas), tibor.neugebauer@uni.lu (T. Neugebauer), a.morales@uma.es (A.J. Morales).

efficiency and flexibility. However, teams and groups face the free rider problem: individual incentives are often at odds with efficient actions. Much research has focused on how to overcome or alleviate this problem.¹

In this paper, we focus on a novel institution designed to alleviate the free rider problem: *excludability*.² Excludability combines two incentives that have been identified in the literature as being crucial for motivation on the job: *competition* and *exclusion*. Nalebuff and Stiglitz (1983) note that “one of the dominant characteristics of modern societies is the important role played by competition; competition is the force providing work incentives. Rewards within a firm (...) are at least partially based on relative performance” (p. 21). Similarly, exclusion is a common disciplinary measure against defectors both at work and in social life. For example, shirking workers are fired (Shapiro and Stiglitz, 1984); uncooperative neighbors are not invited to neighborhood parties and other social events; societal defectors are incarcerated or expelled (Hirshleifer and Rasmusen, 1989); and countries that violate international conventions are boycotted.

A combination of competition and exclusion is utilized in many organizations as an implicit or explicit incentive mechanism. Jack Welch of GE famously fired the bottom 10% of employees each year, thus implementing *competition* among employees to stay in the top 90% and *exclusion* of the bottom 10%. An estimated 20% of US firms utilize some sort of forced ranking, including Ford, Sun and Microsoft. Although common, this method has met with much controversy, and the evidence supporting its practice is somewhat mixed. The *stack ranking* mechanism employed by Microsoft is a notorious example. By utilizing this incentive system, managers are requested to rank their employees in three categories and distribute bonuses accordingly. Given that the proportion of workers in each category is fixed, the system relies entirely on the relative performance of employees rather than on absolute levels of productivity.

While the benefits of competition as suggested in Nalebuff and Stiglitz (1983) have been repeatedly documented (see Knoeber and Thurman, 1994; Nalbantian and Schotter, 1997; Blanes i Vidal and Nossol, 2011), the flip side of competition in organizations has received attention only more recently. Competition may discourage teamwork and become detrimental in very different ways. Charness et al. (2013) and Berger et al. (2013) observe the emergence of disreputable behavior in two experiments in which participants sabotage others' work to increase their chances of winning the competition. Bandiera et al. (2013) find strategic partner selection when rank incentives are introduced, as workers choose to be part of teams with other workers of similar ability to avoid competition, leading to substantial drops in performance.

We design an experiment to examine exclusion of the lowest-contributor under three production functions: the standard voluntary contribution mechanism (VCM) where individual contributions are averaged to create team production, the weakest link mechanism (WLM) where the lowest contribution determines the team production level, and the best-shot mechanism (BSM) where the highest contribution determines the team output.^{3,4} Our experiment includes a baseline treatment without and two treatments with excludability. In the treatments with excludability, if all players contribute the same amount, then no one is excluded. In the first excludability treatment, the excluded party's value from the team production is simply lost from the perspective of the team. This corresponds to a situation in which, for example, in a social setting of the neighborhood, the value of the noncontributing neighbor from attending the party is not captured. In the second excludability treatment, the value of the excluded party's consumption of the public good is redistributed among the included members. For example, when a low-contributing employee is excluded from the bonus pool, the remaining members get larger bonuses. In an organizational perspective, *exclusion* generates savings for the employer (as she keeps some team benefits or bonuses), while *redistribution* is neutral relative to the baseline condition in the sense that all of the team output remains within the team, and incentives for those not excluded increase.

Excludability is an attractive incentive institution for at least two reasons. First, it involves lower informational requirements than does exclusion without competition and with externally fixed threshold level. The mechanism designer does not need to determine in advance the threshold below which contributors will be excluded (how low is too low?). In addition, it involves lower information requirements for implementation. Participants do not need to know exactly how much each of their team members has contributed (a cardinal measure), only the ordering of contributions (an ordinal measure). These lower informational requirements are most likely the reason that excludability has been observed in the field. Second, it taps into the forces of competition and allows these competitive forces to work in favor of increasing contributions. In contrast, exclusion without competition has more of a contractual structure; everyone knows in advance how much they

¹ Laffont (1987) is the classic reference reviewing theoretical proposals. Ledyard (1995) summarizes early findings of the experimental literature in his well-known review. More recent surveys are offered in Keser (2002), Zelmer (2003), Kosfeld and Riedl (2004) and Chaudhuri (2011).

² The term excludability is a reference to the public goods literature to which our study links. Public goods are characterized by non-excludability from consumption and non-competition in consumption. Our institution excludability implies the possibility of immediate exclusion of the worst free rider from the consumption of the public good. If contributions are the same across contributors, however, no exclusion takes place. Our institution redistribution that we explain in detail below implies both a possibility of exclusion and also a competition in consumption.

³ Classic examples of the WLM involve meetings that can begin only when all participants arrive or joint production tasks in which each member's contribution is critical to producing the output. A typical example of the BSM is the volunteer's dilemma, where one individual's contribution is sufficient to create joint benefit, such as one employee stopping the assembly line to prevent the firm from producing more defective goods. An extreme example occurs when soldiers in a trench at wartime face a live grenade, and one soldier jumps on the grenade and loses their own life but saves the lives of all of their comrades.

⁴ Hirshleifer (1983) and Hirshleifer and Harrison (1989) first analyzed these production functions in a two-player setting. While many studies examined the weakest link game (Van Huyck et al., 1993; Cacho and Camerer, 1996; Bornstein et al., 2002), the best-shot mechanism has previously only been studied as a sequential game in a two-player setting (Hirshleifer and Harrison, 1989; Prasnikar and Roth, 1992; Duffy and Feltovich, 1999; Carpenter, 2002). Thus, we will over-sample this treatment in our experimental design.

Table 1
Experimental design.

	Condition		
	Baseline	Excludability (EX)	Redistribution (RE)
Production function			
VCM (average)	VCM (6)	VCM-EX (6)	VCM-RE (8)
WLM (minimum)	WLM (6)	WLM-EX (6)	WLM-RE (12)
BSM (maximum)	BSM (12)	BSM-EX (6)	BSM-RE (6)

Note: The number of four-subject teams is recorded in parentheses. Because no previous researchers have investigated the BSM as a simultaneous-move game, we oversampled in this treatment, recruiting 12 independent teams of four participants, twice as many as the other baseline treatments, which had been studied extensively.

need to contribute to be included, and many contributions at that minimum level are observed. A negative side of our mechanism though is that it might exclude a relatively high contributor, yielding a sizable loss if there is no redistribution and compromising the willingness to participate.

Our paper contributes to the experimental work studying the effects of exclusion (without competition) on contributions to a public good. In some studies, exclusion is explicit in the sense that individuals who contribute below a certain known threshold are excluded (Swope, 2002; Kocher et al., 2005). In other studies, exclusion can occur because individuals are voted out of the group (Masclot, 2003; Cinyabuguma et al., 2005; Maier-Rigaud et al., 2010) or because they are not part of an endogenously formed team (Brosig et al., 2005; Ehrhart and Keser, 1999; Ahn et al., 2008; Cabrera et al., 2013) or both (Charness and Yang, 2010; Charness et al., 2011). These studies find that the possibility of exclusion increases contributions in a significant way, decreasing free riding behavior and increasing efficiency, although the efficiency gains depend on the particular institutional arrangements. Those who are voted out of the group are typically the lowest contributors. In general, the level of cooperation when the possibility of endogenous group formation exists is higher by the awareness that free riding behavior will result in exclusion.⁵

Our experiment also relates to a work studying how competition to avoid punishment increases contributions. In Falkinger et al. (2000), Orrison et al. (2004) and Harbring and Irlenbusch (2005), participants compete to avoid exogenous sanctions or to gain exogenous rewards.⁶ Recently, a number of papers have studied how targeting the lowest contributor to linear public goods increases contributions, for example, the punishment mechanisms investigated by Xiao and Houser (2011) and Andreoni and Gee (2012). While the former keep punishment small enough to maintain the incentives to free ride, the latter carefully manipulates sanctions so that the person would have rather been the second least contributor. Both papers demonstrate that targeting the lowest contributor to a linear public good game is a way of achieving larger contributions.⁷

We find that excludability helps to mitigate the free riding problem in team production. It dramatically increases contributions and obtains sustained and very high efficiency levels in the linear public good (reaching almost 100%). Contributions are also increased without full efficiency in the BSM (also reaching full contribution). However, excludability fails to promote similarly dramatic efficiency gains in the WLM unless combined with redistribution.

The intuition is that excluding low contributors is enough in the VCM and the BSM because top contributors continue to benefit from a team output partially or fully generated by their own efforts. Excludability has a smaller impact on contributions in the WLM because contributors are still paid by a team output determined by low contributors. This generates an intense and almost immediate contribution reduction by some top contributors and traps the whole team in a low-effort equilibrium.

The rest of the paper is organized as follows. In Section 2, we outline the details of the experimental design, and in Section 3 we summarize the theoretical predictions. Section 4 reports the experimental results, Section 5 contains explains and Section 6 concludes.

2. Experimental design and procedures

We employ a three (production function)-by-three (exclusion conditions) experimental design, as depicted in Table 1. Equilibrium predictions and efficiency properties for each treatment of the design are described in the respective subsection of Section 3. Our participants consisted of 272 economics undergraduate students. We utilized a between-subject design: each participant faced only one production function and one excludability condition. Participants were assigned to teams of size four ($N = 4$), and played a finitely repeated (10-round) game with fixed teams. We included a surprise restart (Andreoni,

⁵ Riedl et al. (2012) is the only paper dealing with exclusion in a non-linear environment. Subjects can choose the members with whom they want to interact in a WLM. Endogenous exclusion generates large efficiency gains because high performers exclude low performers from future interactions, who end up learning to become high performers, too.

⁶ All of these studies find a large initial boost to cooperation, which diminishes over time.

⁷ Other papers considering exclusion in different forms are Güth et al. (2007) and Levati et al. (2007). While the former finds that exclusion empowers leadership and increases contributions, the latter does not observe a significant effect. Fatas et al. (2010) find that probabilistic exclusion generates large cooperation gains, even when full efficiency is never observed.

1998; Croson, 1996; Croson et al., 2005; Andreoni and Croson, 2008) followed by a second ten-round game. The number of independent teams is noted in parentheses in each cell of Table 1.

The examined production functions involve three well-known games: the voluntary contribution mechanism (hereafter VCM) where production is a linear function of the average of the contributions, the weakest link mechanism (hereafter WLM) where production is a linear function of the minimum contribution, and the best-shot mechanism (hereafter BSM) where production is a linear function of the maximum contribution.

In the baseline condition, all participants receive the same return from team production regardless of their contribution. The individual payoff to player i takes the form

$$\pi_i^{\text{Baseline}}(x) = (e - x_i) + 2F(x) \quad (1)$$

where π denotes payoff, x is the vector of individual contributions, $(e - x_i)$ is the level of resources kept by i for her individual consumption and $F(x)$ is the team production. The three production functions corresponding to the three games are as follows: $F(x) = \frac{\sum x_i}{N}$ for the VCM, $F(x) = \min\{x_1, \dots, x_N\}$ for the WLM and $F(x) = \max\{x_1, \dots, x_N\}$ for the BSM.

In the two other conditions (excludability [EX] and redistribution [RE]), the participant who contributes the least is excluded from enjoying the benefits of the team production. In the EX treatments, this excluded contributor's share of team production is lost, while in the RE treatments, it is redistributed among the included contributors.

The payoff function to player i under the excludability condition [EX] is as follows:

$$\pi_i^{\text{EX}}(x) = \begin{cases} (e - x_i) & \text{if } \max\{x_1, \dots, x_N\} > x_i = \min\{x_1, \dots, x_N\} \\ \pi_i^{\text{Baseline}}(x) & \text{otherwise} \end{cases} \quad (2)$$

Each individual whose contribution is the lowest within the team (i.e., $\max\{x_1, \dots, x_N\} > x_i = \min\{x_1, \dots, x_N\}$) is excluded from the benefits from the team production and enjoys only the benefits of their own consumption $e - x_i$. All non-excluded individuals receive the same payoff as in the baseline treatment. In the extremes, up to $N - 1$ individuals can be excluded or, if all contribute the same amount, no one is excluded. It is important to note that in all symmetric profiles, no one is excluded and therefore, every individual receives the same payoff as in the baseline treatment.

Finally, in the redistribution condition [RE], the payoff to player i is as follows

$$\pi_i^{\text{RE}}(x) = \begin{cases} (e - x_i) & \text{if } \max\{x_1, \dots, x_N\} > x_i = \min\{x_1, \dots, x_N\} \\ \pi_i^{\text{Baseline}}(x) + \frac{n}{N - n} 2F(x) & \text{otherwise} \end{cases} \quad (3)$$

where n denotes the number of excluded contributors. Note that excluded contributors receive the same payoff in the RE condition as in the EX condition, but included contributors receive the payoffs as in the baseline condition plus an extra payment, which can be interpreted as a transfer from the excluded to the included. The extra payment is the return from the team production, which is not paid to the excluded contributors $2nF(x)$, and is divided among the $(N - n)$ included contributors. As in the EX condition, for symmetric contribution profiles, no individual is excluded and subjects receive the same payoff as in the baseline.

Our study reports the results of computerized experiments conducted at LINEEX, the experimental laboratory at the University of Valencia.⁸ By participating in the experiment, a participant earned an average of €16; experiments took less than one hour to run. At the beginning of a session, written instructions were read aloud. Thereafter, participants went through a test including four control questions.⁹ The instructions were re-read until everybody had answered all questions correctly. After the experiment, subjects were debriefed in a questionnaire, and they were asked individually whether they had understood the instructions. Given their replies and the procedure, we are confident that the tasks and the incentives were understood.

Participants were assigned to one session and played in only one of the nine treatments. The experiments entailed ten rounds of play (original game) with another ten-round surprise restart game in a partner's setting. Upon arriving at the lab, participants were randomly and anonymously arranged in teams of four in the first round and remained together throughout both the original and the restart game. Participants received a record of all individual contributions in increasing order after each repetition; individual contributions were not identified with their contributor. Additionally, participants were informed about their own earnings both in total and subdivided by the return from the endowment kept and the return from team production. Only neutral language was used; for example, participants allocated Eurocents (rather than contributing them) to a group project.

⁸ Fischbacher's (2007) zTree was used for the computer program. Experiments were run in 2002 and 2005.

⁹ The translated instructions and tests are included in Appendix D.

3. Theory

In this section, we analyze the equilibrium predictions for the different team productions and conditions. All proofs are contained in Appendix A.

3.1. Nash predictions

3.1.1. VCM: the voluntary contribution mechanism

The baseline condition corresponds to the standard VCM; one of the most frequently studied social dilemma environments in experimental economics. It is well known to have only one equilibrium contribution level that corresponds to free riding. Players have a dominant strategy to free ride because every unit contributed toward team production yields half a unit payoff to each individual in the team (because $N = 4$ and therefore, $2N^{-1} = 0.50$), while each unit contributed toward personal consumption yields one unit payoff to the contributor only. Collective free riding is similarly predicted for the finitely repeated game because it is the unique Nash equilibrium. This result is socially inefficient because any contribution to team production returns twice as much to the team as a contribution to individual consumption; in fact, if players contributed all of their resources to the team production, they would achieve the socially efficient outcome.

When excludability is considered, the dominance of the free-riding equilibrium is eliminated. In fact, the game now has a multiplicity of equilibria because every symmetric profile of contributions is a Nash equilibrium. This result hinges on the fact that no one is excluded if all contribute the same amount. In these symmetric profiles, there are incentives neither to increase one's own contribution (we are dealing with a public good game) nor to decrease it (because the player is excluded). When redistribution is included, incentives to decrease one's own contribution from any symmetric profile remain unchanged, but incentives to increase improve because of the extra bonus of getting the team production of the excluded members. As a result, the possibility of excluding all others players by unilaterally increasing one's own contribution is so profitable that players seek increasingly larger contributions until they reach the full contribution (it is a case of inverse price competition). [Proposition 1](#) formalizes these intuitions.

Proposition 1. For the VCM, excludability [EX] turns the game into a coordination game where all symmetric profiles are Nash equilibria. Under redistribution [RE], full contribution is the unique Nash equilibrium.

3.1.2. WLM: the weakest link mechanism

In the WLM, the *minimum* individual contribution determines the team production; thus $F(x) = \min\{x_1, \dots, x_N\}$. This game is a coordination game in which all symmetric profiles are Nash equilibria. To see it, note that in any symmetric profile, there are incentives neither to increase one's own contribution (because the team output would remain unchanged), nor to decrease it (because the opportunity cost is twice the return from individual consumption). These symmetric equilibria are Pareto ranked from collective zero contribution to collective full contribution of the entire endowment, where the latter represents the Pareto-efficient and payoff-dominant equilibrium. Excludability has no effect on the equilibrium structure of the game because deviating by contributing less will induce exclusion, which will reduce profits, and deviating by contributing more will reduce private consumption without increasing the return from the team production because the team output is determined by the minimum individual contribution. Under redistribution, the possibility of receiving the output share of the excluded members gives incentives to increase contributions (provided that the team output is positive), and as a result, full contribution and no contribution are the only pure NE. These intuitions are collected in [Proposition 2](#).

Proposition 2. For the WLM, all symmetric profiles are NE in the Excludability [EX] condition. Under Redistribution [RE], full contribution and zero contribution are the only pure Nash equilibria of the game.

3.1.3. BSM: the best-shot mechanism

The team production function in the BSM is the maximum contribution $F(x) = \max\{x_1, \dots, x_N\}$. The NE structure of the one-shot BSM involves four pure asymmetric Nash equilibria¹⁰ in which one subject contributes her entire endowment and the other three team-members contribute nothing. These pure-strategy equilibria are also efficient while increased contributions by the other players would result in social inefficiencies. Excludability prevents any asymmetric profile – such as the efficient NE of the baseline – from being an equilibrium. The incentives to avoid exclusion are powerful. As a consequence, the unique Nash equilibrium in the BSM-EX is full contribution. This equilibrium is inefficient, as there are multiple individuals contributing the maximum level toward team production, when only one such contribution is necessary to reach the same output level.¹¹ The addition of redistribution adds nothing, as exclusion on its own is enough to drive contributions to the largest levels. [Proposition 3](#) formalizes these intuitions.

¹⁰ a mixed strategy equilibrium exists in which each subject contributes their endowment to team production with a probability of $1 - p^{1/(N-1)} = 0.2063$, where $p = 1/2$ for our parameters (see [Croson et al., 2005](#)). Several other (less-efficient) mixed strategy equilibria exist as well.

¹¹ The social payoff in the BSM-EX would be greatest if one subject were to contribute their entire endowment (creating the maximum team payoff), one subject were to contribute nothing (and would thus be excluded) and the other two subjects were to contribute the smallest possible amount (and would thus avoid exclusion). Unfortunately, this efficient allocation is not an equilibrium; the excluded subject has an incentive to increase her contribution to enable her to be included. Thus, this game has properties of a rent-seeking game where individuals (inefficiently) contribute too much.

Table 2
Theoretical predictions for contributions in percentage of endowment.

Treatment	NE	RDE	QRE
VCM			
Baseline	0	0	0
VCM-EX	Any symmetric	100	100
VCM-RE	100	100	100
WLM			
Baseline	Any symmetric	Any symmetric	0
WLM-EX	Any symmetric	0 or 100	0
WLM-RE	0 or 100	0 or 100	0
BSM			
Baseline	Asymmetric (0, 0, 0, 100) or mixed strategy	Mixed strategy	Mixed strategy
BSM-EX	100	100	100
BSM-RE	100	100	100

Proposition 3. For the Best-Shot Mechanism, full contribution is the unique NE under both the excludability [EX] and the redistribution [RE] condition.

3.2. Refinements

Our previous analysis reveals that the Nash concept makes unique predictions in only four out of the nine treatments (VCM, VCM-RE, BSM-EX and BSM-RE). This weakness is especially serious in the WLM, where the NE offers no unique point prediction in any of the three treatments, and in two of them (WLM and WLM-EX) every contribution level is part of a NE profile.

In this subsection, we seek to refine our predictions pursuing two different avenues. The first one applies the equilibrium selection criteria proposed by [Harsanyi and Selten \(1988\)](#): payoff dominance and risk dominance. Payoff dominance selects full contribution whenever full contribution is a NE. Risk dominance allows no refinement of our prediction in the WLM (see [Fatas et al., 2006](#) for details) because no NE risk-dominates the others (hence, it can be said that all symmetric profiles are risk dominant). Risk dominance selects full contribution in the VCM-EX and the symmetric mixed strategy equilibrium in the BSM. Finally, full contribution and zero contribution are risk-dominant equilibria (RDE) of the WLM-EX and WLM-RE.

The second approach we consider involves the quantal response equilibrium (QRE) ([McKelvey and Palfrey, 1995](#)), which always converges to an NE when the noise parameter goes to 0, and therefore, it will offer unique predictions for all treatments.¹² It selects full contribution in VCM-EX but the inefficient NE for any treatment in the WLM, where the NE multiplicity problem is more severe.¹³ Why these contrasting predictions? QRE assumes that the choice probabilities are proportional to expected payoffs thus incorporating the intuitive notion that it is more likely that a player will make “mistakes” in the direction of higher expected payoffs. Excludability makes more costly those mistakes associated with small contributions in the VCM-EX (because they increase exclusion probability), favoring large contributions and resulting in QRE selecting full contribution.

This is not the case for the WLM because when compared to the VCM, a large contribution submitted to avoid exclusion is not as profitable (because the team output is not affected), and low contributions that avoid exclusion are less costly because by definition, they do not affect the team output (the output is determined by the minimum contribution that has been excluded). As a consequence, QRE picks zero contribution. The bonus for large contributions in the redistribution case is not large enough (typically only one member is excluded and her team production share is to be divided among three members) to overcome the attraction of low contributions, and QRE also picks zero contribution in WLM-RE.¹⁴ [Table 2](#) collects all of the theoretical predictions for contributions.

We are primarily interested in the effectiveness of exclusion on contribution behavior, equilibrium behavior and efficiency levels, rather than testing these theoretical predictions per se. Thus, we will compare within a production function across excludability conditions. However, for the theoretically minded reader, we will also compare experimental outcomes with these equilibrium predictions. Note that the RDE and QRE models predict that excludability will imply full contribution in VCM and BSM. In WLM, QRE predicts zero contribution whereas RDE predicts zero or full contribution.

¹² Obviously, it will coincide with RDE and NE whenever there is a unique NE.

¹³ It picks the mixed strategy NE in the BSM baseline, as RDE does.

¹⁴ QRE has been successful in organizing experimental data when intermediate values of the noise parameter are considered. For these cases, the intuition is correct that exclusion and redistribution increase contributions. For example, for a noise parameter $\mu = 10$, the mean value of the QRE distribution, as a fraction of the total endowment, is 22%, 24% and 30% for WLM, WLM-EX and WLM-RE, respectively. However, the QRE prediction will never be above 50% of the total endowment for any intermediate value of the noise parameter (in the limit as noise goes to infinity, QRE converges to the uniform distribution with an average contribution equal to the midpoint of the interval.).

Table 3
Average contributions in percentage of endowment.

Treatment	Round				Game		
	1	10	11	20	Original	Restart	Overall
VCM							
Baseline	40	18	40	10	32	29	31
VCM-EX	70	93	92	97	82	93	88
VCM-RE	84	99	96	100	94	98	97
WLM							
Baseline	42	21	34	17	24	21	22
WLM-EX	62	30	51	34	42	45	44
WLM-RE	54	61	79	81	68	86	77
BSM							
Baseline	62	29	42	25	40	30	35
BSM-EX	94	100	99	100	98	100	99
BSM-RE	81	100	93	100	93	95	94

Table 4
Average efficiency and team production.

Treatment	Efficiency (payoff/max. payoff) ^a	Team production ^b
VCM		
Baseline	65%	31%
VCM-EX	83%	88%
VCM-RE	99%	97%
WLM		
Baseline	52%	13%
WLM-EX	52%	26%
WLM-RE	72%	61%
BSM		
Baseline	81%	79%
BSM-EX	86%	100%
BSM-RE	75%	100%

^a Note that this efficiency measure is linearly related to contribution levels only for VCM.

^b Recall that team production coincides with the average contribution level only for VCM.

4. Experimental results

Table 3 provides descriptive statistics of contributions in the nine treatments over the different time horizons of the experiment. We focus on the level of contributions and their comparison in the baseline, excludability and redistribution treatments within production functions and also the efficiency and team production levels that result from these contribution decisions. Parallel tables describing outcomes for those variables can be found in the online supplement (Tables B3.1 and B3.2).

We begin by analyzing the impact of excludability on contributions under the three production functions. Our statistical analyses utilize two-tailed Mann–Whitney nonparametric tests unless it is stated otherwise. In comparing first round contributions, we conduct our statistical tests on the individual contributions; it is reasonable to think these are independent, as no participant has seen the previous contributions of any other participant. All other tests are conducted on team contributions, as individual contributions no longer represent independent observations.

In addition to the contribution as a measure of team performance, we are also interested in two alternative (derived) dependent measures, the efficiency of the outcome and the level of team production. Economists are generally concerned with the efficiency of institutions. We will measure efficiency in our context as the team's payoff relative to the maximum possible team payoff.¹⁵ This analysis will enable us to answer the question of how efficient the various institutions are in practice and can aid in institutional design questions of interest to the social planner. The efficiency levels for each treatment averaged over the entire game are recorded in Table 4 (the interested reader should see Table B3.1 in the online supplement for more disaggregated descriptions and Table B1.2 for a time trend analysis of efficiency across the three mechanisms). We statistically compare these levels utilizing the groups as independent observations.

Finally, while a social planner or institution-designer may be interested in efficiency, from the perspective of a principal or an organization, performance may be measured by the team production generated rather than by efficiency. The second

¹⁵ In the VCM, efficiency is positively (and linearly) correlated with the contribution amounts. But in the other production functions, there is no linear mapping from contributions to efficiency. For example, unanimous full contribution in the BSM is inefficient. Note, however, that full contribution by all team members still achieves a rather high efficiency level of 84.21%.




Fig. 1. Average contributions in VCM in percentage of endowment across treatments.

column of Table 4 records for the nine treatments the average team production levels as a proportion of the maximum possible production level. Further details can be found in Tables B1.3 and B3.2 and Fig. B3.3 in the online supplement.

4.1. VCM

Fig. 1 plots contributions in the three treatments: baseline, excludability and redistribution (graphs of each team's contribution are provided in the online supplement). Consistent with previous experimental results (see e.g. Croson, 2007; Neugebauer et al., 2009), we find contributions in the VCM starting at just below one half the endowment and decreasing over the course of the game to approximately 10–20% by the last round. The decline is significant in both the original and the restart game (see Table B1.1 in the online supplement). Again replicating previous work, we observe a significant restart effect in the VCM ($p=0.0277$, $N=6$, Wilcoxon signed rank test). Contributions double between these rounds from 18% in round 10–40% in round 11.

Excludability significantly increases contribution levels for the VCM production function in every round, consistent with the theoretical predictions, and generally increases the efficiency in terms of subjects' payoff ($p=0.002$ for the entire game).¹⁶ Overall, the proportion of the endowment contributed averages 31% in the VCM and 88% in the VCM-EX, implying that the average contribution increases by 184% if one adds excludability. Excludability also eliminates the pattern of declining contributions; instead, in the VCM-EX treatment, contributions increase from 70% in the first round to 97% in round 20 (Table B.1.1. in the online supplement statistically examines these time-trends). Excludability also eliminates the restart effect ($p=0.5688$, $N=6$, Wilcoxon signed rank test).¹⁷

The addition of redistribution has an additional effect. Contributions in the VCM-RE are significantly larger than in the baseline and marginally higher than the VCM-EX treatment ($p=0.059$ for the entire game).¹⁸ In the VCM-RE treatment, we observe a marginal restart effect ($p=0.0854$; $N=8$) and almost full contribution (overall 97%).

Result 1a. Excludability is an effective mechanism for increasing and sustaining contributions in the VCM (approximately 88% of the endowment overall). The addition of redistribution has an additional positive effect, increasing contributions to almost full levels.

Efficiency increases from the baseline to the treatment with excludability ($p=0.026$) and to the treatment with redistribution ($p=0.001$ versus excludability only, $p=0.001$ versus baseline). Regarding team production, it significantly increases from the baseline to excludability ($p=0.002$) and redistribution ($p=0.059$ versus excludability and $p=0.001$ versus baseline).

Result 1b. Excludability increases efficiency and team production, and adding redistribution enhances the result.

¹⁶ p -Values for the seven columns in Table 3 are 0.000, 0.002, 0.002, 0.002, 0.004, 0.002 and 0.002, respectively.

¹⁷ Increased average contributions with excludability are caused by significantly fewer free riders and more full contributors; see the last column of Tables B2.1 and B2.2 in the Appendix. While 65% of all actions in VCM-EX involve full contributions (and 2% involve zero contributions), only 3% full contributions (and 14% free riding) are observed in the VCM.

¹⁸ p -Values for two-tailed Mann-Whitney tests comparing VCM and VCM-RE are 0.000, 0.001, 0.001, 0.001, 0.001, 0.001, and 0.001 for the seven columns of Table 3, respectively. Similar comparisons between the VCM-EX and the VCM-RE treatments are 0.043, 0.852, 0.662, 0.345, 0.043, 0.491, and 0.059, respectively.

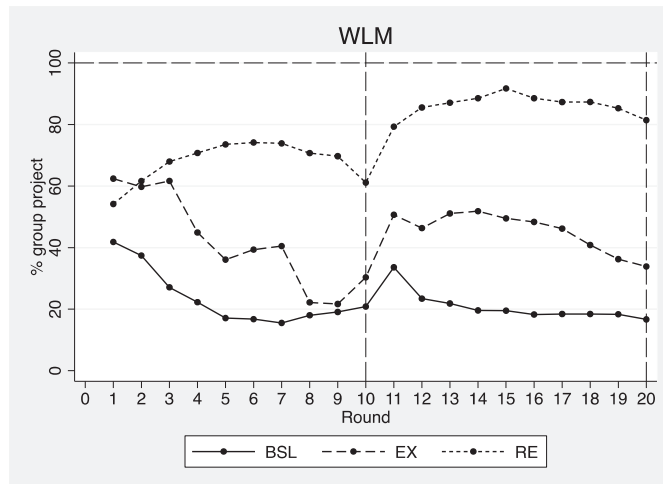


Fig. 2. Average contributions in WLM in percentage of endowment across treatments.

4.2. WLM

Fig. 2 plots contributions in the different WLM treatments (the online supplement provides graphs of each team's contribution). In the WLM baseline, contributions start at approximately 40% of the endowment and decline until ending at approximately 20%. This is consistent with previous findings in minimum effort games (Goeree and Holt, 2005), where for the treatments with high costs of effort, averaged contributions ended at approximately 20–30% of the endowment.¹⁹ Excludability has an impact in the first round, although the difference between WLM and WLM-EX treatments quickly disappears ($p = 0.240$ for the entire game).²⁰ Both treatments exhibit significantly decreasing contributions over time (see Table B.1.1) and apparent but not significant restart effects ($p = 0.1159$ and $p = 0.1400$, $N = 6$, Wilcoxon signed rank test).

Redistribution has a significant impact in the WLM. Average contributions in the WLM-RE treatment do not exhibit any decreasing trend (see Table B.1.1) and are significantly larger than in the baseline WLM ($p = 0.000$ for the entire game) and larger (marginally significantly though) than in the WLM-EX ($p = 0.067$ for the entire game) with levels at approximately 75% of the endowment. We also find a marginally significant restart effect ($p = 0.0844$, $N = 12$, Wilcoxon signed rank test).²¹

Result 2a. While excludability has only small effects on behavior in the WLM, redistribution raises the contribution levels to 75% of the endowment. This latter result cannot be rationalized by any of the models considered.

A similar picture is obtained when we analyze the impact of exclusion on the two alternate measures of team production. Whereas there is no significant difference between the baseline and the treatment with exclusion ($p = 0.394$ for efficiency and $p = 1.00$ for team production), the addition of redistribution has a positive impact on efficiency ($p = 0.032$ versus excludability, $p = 0.001$ versus baseline) and team production ($p = 0.053$ versus excludability, $p = 0.001$ versus baseline).

Result 2b. Excludability has no effect on alternative measures of team production, but the addition of redistribution improves matters significantly.

4.3. BSM

Fig. 3 displays contributions in the different BSM treatments (graphs of each team's contribution are provided in the online supplement). In the BSM, excludability increases contributions significantly for every round and over the entire game ($p = 0.000$ for the entire game).²² In the BSM-EX, overall contributions average 99% of the endowment, involving 96% full contributions. Compared to the overall average contribution in BSM of 35%, excludability increases contributions by 182%.

¹⁹ These contribution levels can be fitted by the QRE when assuming an intermediate value of the noise parameter. For $\mu = 10$, the QRE prediction in the baseline WLM is 22%. Note, however, that we rarely observe symmetric profiles at intermediate contribution levels. Thus, no equilibrium prediction is actually corroborated by the data.

²⁰ p -Values for the seven columns of Table 3 are 0.008, 0.818, 0.485, 0.937, 0.240, 0.589 and 0.240, respectively. Nevertheless, one should note that compared to the WCM, average contributions double in the WCM-EX.

²¹ p -Values for two-tailed Mann-Whitney tests comparing WLM and WLM-RE are 0.009, 0.032, 0.000, 0.000, 0.000, 0.000 and 0.000 for the seven columns of Table 3, respectively. Similar comparisons between the WLM-EX and the WLM-RE treatments are significant in some cases, being 0.123, 0.250, 0.067, 0.032, 0.083, 0.083 and 0.067, respectively.

²² p -Values for the seven columns of Table 3 are 0.008, 0.000, 0.000, 0.000, 0.000, 0.000 and 0.000, respectively.

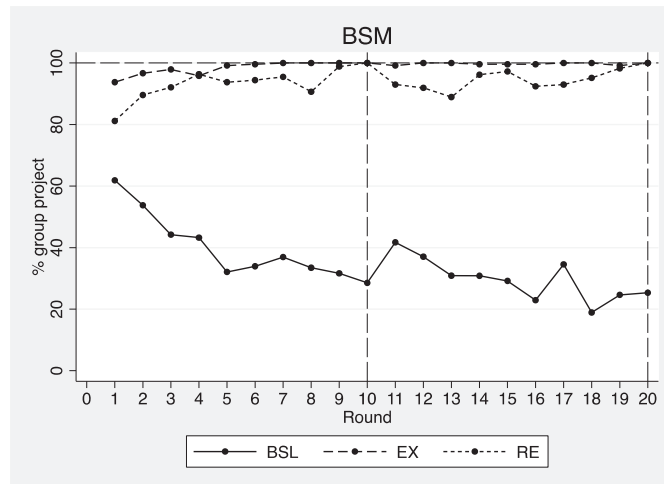


Fig. 3. Average contributions in BSM in percentage of endowment across treatments.

We see a significant declining pattern of contributions (Table B.1.1) and a significant restart effect ($p=0.0637$, $N=12$, Wilcoxon signed rank test) in the baseline BSM. However, we see no decline and no restart effect in the BSM-EX treatment ($p=0.3173$, $N=6$, Wilcoxon signed rank test). The average contribution in the baseline BSM decreases in all teams to approximately 25% of the endowment. However, efficiency does not increase over rounds in the BSM game (see Table B1.2 in the online supplement). In the BSM-EX, full contribution is reached in 98% of all outcomes, consistent with the theoretical predictions. The addition of redistribution has only a small (but negative) quantitative and non-robust additional impact on contributions, although the difference is significant over the entire game ($p=0.002$).²³ Contributions in the BSM-RE average 94% of endowment, and the negative restart effect is marginally significant ($p=0.0874$, $N=6$, Wilcoxon signed rank test).

Result 3a. Excludability boosts contributions to full levels, consistent with predictions. The addition of redistribution, however, has a small negative quantitative effect.

Finally, excludability significantly increases efficiency ($p=0.018$). However, including redistribution leads to a significantly lower efficiency level, worse even than the level of efficiency in the baseline ($p=0.012$ versus excludability only, $p=0.000$ versus baseline). Team production is significantly increased with excludability ($p=0.001$), but then remains high as redistribution is added ($p=0.699$ versus excludability and $p=0.001$ versus baseline).

Result 3b. Excludability increases efficiency and team production in the BSM. When redistribution is added, team production remains high while efficiency levels decrease.

In the BSM-RE, the efficient allocation requires coordination on an asymmetric contribution profile with voluntary self-exclusion. Such a profile can lead to a fair sharing of efficiency gains only if subjects rotate their positions. Players have to efficiently alternate between full contribution and zero contribution. Tacit coordination on an efficient rotation of positions between rounds is extremely difficult. The significantly lower contributions (see Figs. 3 and C3.3) in the BSM-RE treatment relative to BSM-EX indicate that subjects attempt to increase group efficiency, but the reported efficiency levels also indicate that coordination on more efficient outcomes than equilibrium outcomes largely fails.

4.4. Summary

We have compared contribution levels to team production in three team production functions with excludability and with redistribution. We find that excludability has a massive impact on contributions in the VCM and the BSM settings – with an overall contribution level of at least 90% – and that this effect is compatible with our theoretical considerations. We did not, however, anticipate the strong effect of redistribution observed in the WLM production function.

In the VCM, excludability increases efficiency, and adding redistribution enhances the result. In the WLM, excludability has no effect on efficiency, but the addition of redistribution improves matters significantly. In the BSM, excludability increases efficiency, but adding redistribution has a negative effect rather than a positive one.

Thus, our data indicate that efficiency hinges on the production function in question. In the VCM setting, excludability with redistribution achieves the highest efficiency level (99% efficiency). In the weakest link production function, excludability

²³ p -Values for two-tailed Mann–Whitney tests comparing BSM and BSM-EX-R are 0.084, 0.000, 0.000, 0.000, 0.000, 0.000, and 0.000 for the seven columns of Table 3, respectively. Similar comparisons between the BSM-EX and the BSM-RE treatments are significant in some cases, being 0.020, 1.00, 0.310, 1.00, 0.026, 0.026, and 0.002, respectively.

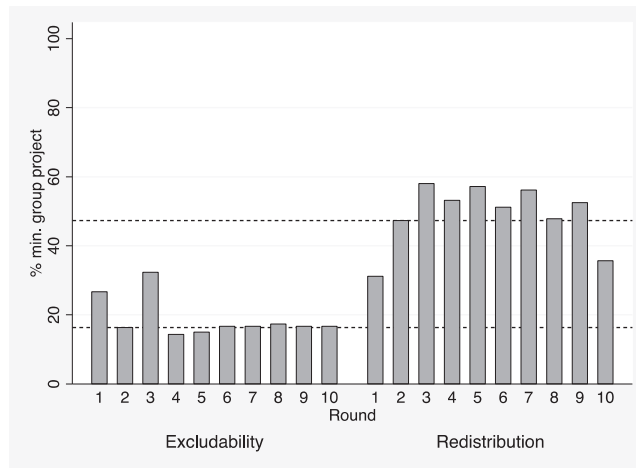


Fig. 4. Evolution of team output in WLM across conditions.

with redistribution also performs best (72%). In the best-shot setting, however, excludability leads to the highest efficiency levels (86%) with no added redistribution.

Finally, for the principal, excludability can increase team production in VCM and BSM situations, while redistribution is necessary for increased team production in WLM situations.

5. Behavioral explanation

This paper analyzes the effect of excludability and redistribution on team production for a number of team functions. The theoretical predictions made by our models (RDE and QRE) are that we should expect a massive impact of excludability in the VCM and BSM but no large impact in the WLM. Additionally, redistribution is not predicted to have any impact at all.²⁴ Our experimental results demonstrate that these predictions are basically right for the VCM and BSM,²⁵ but they are, in particular, in view of the QRE, not so successful for the WLM.²⁶ In this section, we want to understand which key feature in the production functions causes excludability and redistribution to have a differential impact.

Fig. 4 displays the evolution of the provision of the team project in the WLM under excludability and under redistribution in the first block of ten rounds. In round 1, average contributions in the WLM-EX (26.67) and WLM-RE (31.17) are not significantly different²⁷ (p -value: 0.2224, Mann–Whitney test at the group level). However, in round 2, the trajectories diverge (16.33 and 47.33, respectively, p -value: 0.0025), and after that round, the difference remains constant primarily because no trend in either treatment is observed (we have included two dashed lines at the provision levels in round 2 for the two treatments as a visual reference). Without redistribution, groups are trapped in inefficient levels of team output from round 2 on.

Fig. 4 suggests that there is a major change from the first to the second round that is strongly mediated by redistribution. To gain an understanding of why this happens, we analyze the impact of excludability in round 1. How do excluded and non-excluded subjects react to exclusion? Our answer is that they do so very differently in different conditions. Excluded participants significantly increase their contributions in round 2 with or without redistribution; on average, they go from 13.33 to 38.17 (p -value = 0.0277) in the excludability condition and from 15.58 to 29.62 (p -value = 0.0166) with redistribution. Non-excluded participants continue with similar average contributions in the redistribution condition (31.19 in round 1 and 31.21 in round 2, p -value = 0.7238), whereas average contribution levels significantly decline from 37.17 to 27.11 (p -value = 0.0464) in the excludability treatment in absence of redistribution.

In the WLM, the team output is determined by the minimum contribution in the group, not by the average contribution (Croson et al., 2005). Hence, we must analyze the individual adjustments of contributions to exclusion. Fig. 5 analyzes decisions made by those participants not excluded in the first round, across both the VCM and the WLM, and with and

²⁴ Redistribution is not predicted to have an effect on contributions because in the VCM and BSM, exclusion is already predicted to boost full contribution (and therefore, there is no room for further improvement once redistribution is allowed), while in the WLM, zero contribution is predicted in excludability and redistribution treatments.

²⁵ Point predictions are slightly disappointing, though, when the comparison is made for the overall behavior – especially for the baseline treatments. However, they capture quite nicely long-term behavior (round 20).

²⁶ The RDE predicts zero and full contribution for the WLM-EX. The contributions in the WLM-EX basically move to the extremes in the restart game. The trajectories of the four groups approach zero and of two groups approach full contribution (see Appendix C).

²⁷ All non-parametric tests employed in this section are performed at the group level.

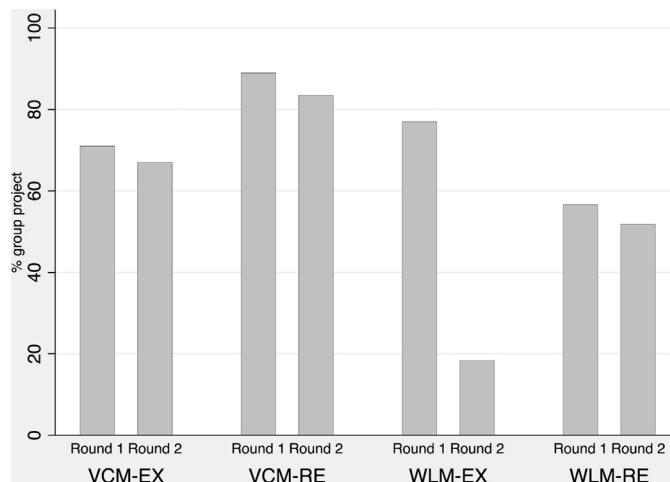


Fig. 5. Non-excluded participants' reaction to first round decisions.




Fig. 6. Distribution of earnings of non-excluded team members in rounds 1 and 2.

without redistribution. More precisely, it plots the intensity of their reaction measured by the lowest contributor in the second round, in other words, decisions by non-excluded participants (in the first round) who contributed less in the second round. Fig. 5 strongly suggests that the drop in contributions is large and significant only in the WLM without redistribution. While the most extreme reaction is moderate in the VCM-EX, VCM-RE and WLM-RE (a mere 5.6%, 6.2% and 2.2%, respectively), contribution decreases by 74% in the WLM-EX and team output collapses (p -value: 0.0273; Wilcoxon signed-rank test, $N=6$).

The extreme reaction by the non-excluded participants in the second round of the original game of the WLM-EX is at odds with the fact that the theoretical prediction is identical to the VCM-EX, where excludability is very successful. As in the original WLM, all symmetric contribution profiles are NE of the game. Relative to the VCM-EX, where all contributions contribute to the team output, in the WLM-EX, low contributions drain all team benefits. Relative to the WLM-RE, where top contributors obtain an additional share of the team output (coming from excluded participants, if non-negative), there is no bonus waiting for top contributors in the WLM-EX. Fig. 6 presents the effect of this difference, presenting the distribution of earnings of non-excluded subjects in rounds 1 and 2 (across the VCM and the WLM, with or without redistribution).²⁸

Top performers in round 1 are safely above the secure payoff associated with keeping the endowment for private consumption in all games, with the exception of the WLM-EX where more than three-fourths of non-excluded subjects make less than 50. Losses associated with large contributions prevent coordination on highly ranked equilibria, and contributions collapse in the WLM-EX.

²⁸ Fig. 6 is a standard box and whisker graph. The box contains the 25–75% quartiles, the bar corresponds to the median, and whiskers include the adjacent values in each condition. Hollow diamonds are outliers.

6. Conclusions

We examine the effects of excludability in three different production functions (the voluntary contribution mechanism [VCM] where the average contribution determines production, the weakest link mechanism [WLM] where the minimum contribution determines production and the best-shot mechanism [BSM] where the maximum contribution determines production). Our results indicate that excludability with and without redistribution has an important effect on contributions. Excludability without redistribution is sufficient to raise contributions in the VCM and BSM production functions to near-total contribution levels; adding redistribution in these settings has little impact (and in the case of the BSM, a negative impact). In contrast, in the WLM, excludability by itself has only a small impact on contributions, but redistribution is also necessary to increase contributions. Increasing contributions does not necessarily lead to increased efficiency in these mechanisms, however. While adding excludability increases efficiency in the VCM and the BSM, redistribution is necessary for increased efficiency in the WLM and redistribution hurts efficiency in the BSM.

The comparative static results for the VCM and BSM production functions are mostly consistent with the Nash, risk-dominant and quantal-response equilibrium predictions in our settings, although they are unable to explain the role played by redistribution in the WLM. The general picture is that excludability successfully boosts contribution levels because the reaction to exclusion is an increase in the contribution levels, to avoid future exclusion, and the reaction of non-excluded members is to continue contributing at the same levels. This general picture has one exception: the WLM mechanism. In this mechanism, low contributions determine the team output and therefore hurt high contributors more than in the VCM and BSM. In fact, we observe that high contributors dramatically decrease contributions in round 2, trapping teams in a low performance equilibrium from which they never recover. Redistribution makes a big difference in the WLM because it gives top contributors an additional benefit: the group shares of those excluded.²⁹ This relatively small incentive is necessary for excludability to impact contributions in the weakest link mechanism, where strong complementarities between workers exist. In that sense, this paper documents how differences in production technologies, even when they do not generate changes in the theoretical analysis, may generate very different behavioral reactions.

Our policy message has to do with the positive effects of this mechanism and its limits. The possibility of exclusion from a team's benefits appears to us as a natural incentive device to sustain contribution. Many norms in firms, organizations and societies involve excludability, although its concrete realization may differ between settings. Our results predict that the incentive systems based on excludability will elicit high effort levels from the best-performing members, as well as from those individuals who want to be included in production functions in which there are no complementarities in the members' contributions. In case of complementary contributions, excludability needs to be coupled with redistribution to achieve large efficiency gains.

Like all research, this study has important limitations. For example, our results are found in symmetric settings; we induce the same opportunities and payoff functions for each participant. If contributors instead face different endowments or have different payoffs, excludability may be less effective. As with any incentive scheme, there can also be disadvantages to excludability. In general, incentive schemes are sensitive to collusion and to sabotage (see Harbring and Irlenbusch, 2011; Carpenter et al., 2010, for example); one might imagine similar problems with excludability (as has been suggested in the case of the implementation of the stack rank mechanism in Microsoft). In addition, excludability requires at least some monitoring (in our design, only ordinal information is required, but other exclusion schemes may require cardinal measurements of contribution) and may be difficult to implement. However, we believe that excludability can (and does) provide an effective, low-cost way to induce cooperation, increase efficiency and support team production in social dilemmas and related situations.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jebo.2015.03.005>.

References

- Ahn, T.K., Isaac, M., Salmon, T.C., 2008. Endogenous group formation. *J. Public Econ. Theory* 10, 171–194.
- Andreoni, J., 1998. Why free ride? Strategies and learning in public goods experiments. *J. Public Econ.* 37, 291–304.
- Andreoni, J., Croson, R., 2008. Partners versus strangers: the effect of random rematching in public goods experiments. In: Plott, C., Smith, V. (Eds.), *Handbook of Results in Experimental Economics*. North-Holland, Amsterdam.
- Andreoni, J., Gee, L.K., 2012. Gun for hire: delegated enforcement and peer punishment in public goods provision. *J. Public Econ.* 96, 1036–1046.

²⁹ Note that our results could also be consistent with social preference models à la Charness and Rabin (2002) that include considerations of subjects' efficiency concerns (see also the discussion in Engelmann and Strobel, 2004). For every positive minimum contribution, the group suffers a smaller loss in the WLM-RE than in the WLM-EX, and efficiency gains increase with an increased minimum contribution. Thus, the fact that we observe a significantly larger minimum contribution in the WLM-RE treatment could be the result of efficiency concerns among the subjects. The observed differences between the BSM-EX and BSM-RE could also be linked to efficiency concerns models. The group payoff of the BSM-RE is increased by a decreased minimum contribution for a given maximum contribution, and that is what we observe. By giving up some payoff, individuals increase the payoff to the group. Therefore, again, our data are in line with the efficiency concerns argument.

- Bandiera, O., Barankay, I., Rasul, I., 2013. Team incentives: evidence from a firm level experiment. *J. Eur. Econ. Assoc.* 11 (5), 1079–1114.
- Berger, J., Harbring, C., Sliwka, D., 2013. Performance appraisals and the impact of forced distribution. An experimental investigation. *Manag. Sci.* 59 (1), 54–68.
- Blanes i Vidal, J., Nossol, M., 2011. Tournaments without prizes: evidence from personnel records. *Manag. Sci.* 57 (10), 1721–1736.
- Bornstein, G., Gneezy, U., Nagel, R., 2002. The effect of intergroup competition on intragroup coordination: an experimental study. *Games Econ. Behav.* 41, 1–25.
- Brosig, J., Margreiter, M., Weimann, J., 2005. Endogenous Group Formation and the Provision of Public Goods: The Roles of Promises and Lies, Working paper University, Magdeburg.
- Cabrera, S., Fatas, E., Lacomba, J.A., Neugebauer, T., 2013. Splitting leagues: promotion and demotion in contribution-based regrouping experiments. *Exp. Econ.* 16, 426–441.
- Cachon, G.P., Camerer, C.F., 1996. Loss-avoidance and forward induction in experimental coordination games. *Q. J. Econ.* 111, 165–194.
- Carpenter, J., 2002. Information, fairness, and reciprocity in the best shot game. *Econ. Lett.* 75, 243–248.
- Carpenter, J., Matthews, P., Schirm, J., 2010. Tournaments and office politics: evidence from a real effort experiment. *Am. Econ. Rev.* 100, 504–517.
- Charness, G., Yang, C.L., 2010. Public Goods Provision with Voting for Exclusion, Exit, and Mergers: An Experiment, University of California at Santa Barbara Working paper.
- Charness, G., Rabin, M., 2002. Understanding social preferences with simple tests. *Q. J. Econ.* 117 (3), 817–869.
- Charness, G., Masclet, D., Villeval, M.C., 2013. The Dark Side of Competition for Status, University of California at Santa Barbara Working Paper.
- Charness, G., Cobo-Reyes, R., Jiménez, N., 2011. Efficiency, Team Building, and Identity in a Public-goods Game, University of California at Santa Barbara Working Paper.
- Chaudhuri, A., 2011. Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Exp. Econ.* 14, 47–83.
- Cinyabuguma, M., Page, T., Putterman, L., 2005. Cooperation under the threat of expulsion in a public goods experiment. *J. Public Econ.* 89, 1421–1435.
- Croson, R., 1996. Partners and strangers revisited. *Econ. Lett.* 537, 25–32.
- Croson, R., 2007. Theories of commitment, altruism and reciprocity: evidence from linear public goods games. *Econ. Inq.* 45, 199–216.
- Croson, R., Fatas, E., Neugebauer, T., 2005. Reciprocity, matching and conditional cooperation in two public goods games. *Econ. Lett.* 87, 95–102.
- Duffy, J., Feltovich, N., 1999. Does observation of others affect learning in strategics? *Int. J. Game Theory* 28, 131–152.
- Ehrhart, K., Keser, C., 1999. Mobility and Cooperation: On the Run. CIRANO Working Paper., pp. 24s–99s.
- Engelmann, D., Strobel, M., 2004. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *Am. Econ. Rev.* 94 (4), 857–869.
- Falkinger, J., Fehr, E., Gächter, S., Winter-Ebner, R., 2000. A simple mechanism for the efficient provision of public goods – experimental evidence. *Am. Econ. Rev.* 90, 247–264.
- Fatas, E., Neugebauer, T., Perote, J., 2006. Within-team competition in the minimum effort coordination game. *Pacific Econ. Rev.* 11, 247–266.
- Fatas, E., Morales, A.J., Ubeda, P., 2010. Blind justice. An experimental analysis of random punishment in team production. *J. Econ. Psychol.* 31 (3), 358–373.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10, 171–178.
- Goeree, J., Holt, C., 2005. An experimental study of costly coordination. *Games Econ. Behav.* 51, 349–364.
- Güth, W., Levat, M.V.I., Sutter, M., Van der Heijden, E., 2007. Leading by example with and without exclusion power in voluntary contribution experiments. *J. Public Econ.* 91, 1023–1042.
- Harbring, C., Irlenbusch, B., 2005. Incentives in tournaments with endogenous prize selection. *J. Inst. Theor. Econ.* 161, 636–663.
- Harbring, C., Irlenbusch, B., 2011. Sabotage in tournaments: evidence from a laboratory experiment. *Manag. Sci.* 57, 611–627.
- Harsanyi, J., Selten, R., 1988. A General Theory on Equilibrium Selection in Games. MIT Press, Cambridge.
- Hirschleifer, J., 1983. From weakest-link to best-shot: the voluntary provision of public goods. *Public Choice* 41, 371–386.
- Hirschleifer, J., Harrison, G.W., 1989. An experimental evaluation of weakest link/best shot models of public goods. *J. Polit. Econ.* 97, 201–225.
- Hirschleifer, D., Rasmusen, E., 1989. Cooperation in a repeated prisoners dilemma with ostracism. *J. Econ. Behav. Organ.* 12, 87–106.
- Keser, C., 2002. Cooperation in public goods experiments. In: Bolle, F., Lehmann-Waffenschmidt, M. (Eds.), *Surveys in Experimental Economics*. Physica-Verlag.
- Knoeber, C., Thurman, W., 1994. Testing the theory of tournaments: an empirical analysis of broiler production. *J. Labor Econ.* 12 (2), 155–179.
- Kocher, M., Matthias, S., Waldner, V., 2005. Exclusion from Public Goods as an Incentive System – An Experimental Examination of Different Institutions, Working paper. University Innsbruck.
- Kosfeld, M., Riedl, A., 2004. The design of (De) centralized punishment institutions for sustaining cooperation. In: Raith, M. (Ed.), *Procedural Approaches to Conflict Resolution*. Springer, Berlin/New York.
- Laffont, J.J., 1987. Incentives and the allocation of public goods. In: Auerbach, A., Feldstein, M. (Eds.), *Handbook of Public Economics*, vol. II. North-Holland, Amsterdam, pp. 537–569.
- Ledyard, J., 1995. Public goods: a survey of experimental research. In: Kagel, J.H., Roth, A.E. (Eds.), *Handbook of Experimental Economics*. Princeton University Press, Princeton.
- Levati, M.V., Sutter, M., van der Heijden, E., 2007. Leading by example in a public goods experiment with heterogeneity and incomplete information. *J. Conflict Resolut.* 51, 793–818.
- Maier-Rigaud, F., Martinsson, P., Staffiero, G., 2010. Ostracism and the provision of a public good: experimental evidence. *J. Econ. Behav. Organ.* 73, 387–395.
- Masclet, D., 2003. Ostracism in work teams: a public good experiment. *Int. J. Manpow.* 24, 867–887.
- McKelvey, R., Palfrey, T., 1995. Quantal response equilibrium for normal form games. *Games Econ. Behav.* 7, 6–39.
- Nalbantian, H., Schotter, A., 1997. Productivity under group incentives: an experimental study. *Am. Econ. Rev.* 87, 314–341.
- Nalebuff, B.J., Stiglitz, J.E., 1983. Prizes and incentives: towards a general theory of compensation and competition. *Bell J. Econ.* 14, 21–43.
- Neugebauer, T., Perote, J., Schmidt, U., Loos, M., 2009. Selfish-biased conditional cooperation: on the decline of contributions in repeated public goods experiments. *J. Econ. Psychol.* 30, 52–60.
- Orrison, A., Schotter, A., Weigelt, K., 2004. Multiperson tournaments: an experimental examination. *Manag. Sci.* 50, 268–279.
- Prasnikar, V., Roth, A.E., 1992. Considerations of fairness and strategy: experimental data from sequential games. *Q. J. Econ.* 107, 865–888.
- Riedl, A., Rohde, I., Strobel, M., 2012. Efficient Coordination in Weakest-Link Games, METEOR Research Memoranda Working Paper 057. Maastricht University.
- Shapiro, C., Stiglitz, J.E., 1984. Equilibrium unemployment as a worker discipline device. *Am. Econ. Rev.* 74, 433–444.
- Swope, K.J., 2002. An experimental investigation of excludable public goods. *Exp. Econ.* 5, 209–222.
- Van Huyck, J.B., Battalio, R.C., Beil, R.O., 1993. Asset markets as an equilibrium selection mechanism: coordination failure, game form auctions, and tacit communication. *Games Econ. Behav.* 5 (3), 485–504.
- Xiao, E., Houser, D., 2011. Punish in public. *J. Public Econ.* 95, 1006–1017.
- Zelmer, J., 2003. Linear public goods experiments: a meta-analysis. *Exp. Econ.* 6, 299–310.