Efficient coordination in weakest link games

Citation for published version (APA):

Riedl, A. M., Rohde, I. M. T., & Strobel, M. (2016). Efficient coordination in weakest link games. *Review of Economic Studies*, 83(2), 737-767. <https://doi.org/10.1093/restud/rdv040>

Document status and date:

Published: 01/04/2016

DOI:

[10.1093/restud/rdv040](https://doi.org/10.1093/restud/rdv040)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Efficient Coordination in Weakest-Link Games

ARNO RIEDL

Maastricht University

INGRID M. T. ROHDE

Maastricht University

and

MARTIN STROBEL

Maastricht University

First version received January 2013; final version accepted June 2015 (Eds.)

Coordination problems resembling weakest-link games with multiple Pareto ranked equilibria are ubiquitous in the economy and society. This makes it important to understand if and when agents are able to coordinate efficiently. Existing research on weakest-link games shows an overwhelming inability of people to coordinate on efficient equilibria, especially in larger groups. We show experimentally that freedom of neighbourhood choice overcomes the problem and leads to fully efficient coordination. This implies substantial welfare effects with achieved welfare being about 50% higher in games with neighbourhood choice than without it. We identify exclusion of low effort providers who in response start providing high effort as the simple but effective mechanism enforcing efficient coordination. A variety of other treatments show that the efficiency result as well as the identified mechanism are robust to changes in the information condition, payoff specification, and a substantial increase in group size. Moreover, we find that neighbourhood choice boosts efficiency even when exclusion does not materially affect the excluded agent. Our results are widely applicable on the societal and organizational level, *e.g.* containment of diseases, fight against terrorism, and co-authorship networks.

Key words: Efficient coordination, Weakest-link, Minimum effort, Neighbourhood choice, Experiment

JEL Codes: C72, C92, D02, D03, D85

1. INTRODUCTION

Societies continuously face incidents in which efficient coordination of economic agents is crucial for social welfare (Schelling, 1960; Cooper, 1999). Many of these situations can be described as weakest-link problems that are characterized by a multiplicity of Pareto-rankable strict Nash equilibria (Nash, 1950).¹ In such problems, each agent has an incentive to coordinate

1. The classical example of the fictitious island Anarchia can be found in Hirshleifer (1983). Sandler (1998) and Nordhaus (2006) argue that many global public goods have weakest-link characteristics. Camerer (2003, pp. 381–382) describes weakest-link situations in professional organizations.

on high efforts, implying high individual and group welfare, but also faces considerable strategic uncertainty, because one single “trembling” player suffices to cause substantial losses for all.

Traditionally, weakest-link problems are studied for groups where agents are forced to interact with each other. These studies have produced important insights on the role of, *e.g.* incentives, communication, leadership, social identity, and culture on efficiency in such problems (see, *e.g.* Brandts and Cooper, 2006a, 2007; Hamman *et al.*, 2007; Chaudhuri *et al.*, 2009; Engelmann and Normann, 2010, for some more recent studies).

We lack, however, empirical evidence about the behaviour of economic agents in weakest-link games where they are not bound to interact with everybody else, but instead can freely choose their interaction neighbourhood.² In this article, we provide such evidence and experimentally demonstrate that the freedom to choose interaction neighbourhoods boosts efficiency and welfare in weakest-link games.

Field examples resembling important features of weakest link problems with the possibility of neighbourhood choice come readily to mind. For instance, in the global public good of preventing the outbreak and spread of infectious diseases the country with the poorest preventive measures determines the likelihood of an outbreak. Governments with higher standards can respond not only by lowering their own costly preventive measures but also by restricting trade or traveling to and from countries with low standards. Similar arguments can be made for the fight against international terrorism. Further, in computer and other infrastructure networks hackers or viruses can easily enter through weakly protected network nodes. In such cases, network system administrators may not respond by lowering their own costly security standards but by restricting access to sufficiently safe network nodes. Further, in groups of co-authors the slowest member determines how fast the paper is finished and the sloppiest its overall quality. In response, co-authors may either lower their own costly efforts leading to a poorer paper or terminate cooperation even if it comes at the cost of losing complementary expertise. An important aspect common to these examples is that agents may have overlapping interaction neighbourhoods and do not have to choose between mutually exclusive groups and tasks.

For weakest-link games played in fixed groups the seminal papers by Harrison and Hirshleifer (1989) and Van Huyck *et al.* (1990, VHBB, henceforth) indicated that when played in pairs substantial coordination on the Pareto-dominant equilibrium occurs, whereas efficient coordination breaks down completely when groups grow large enough. The latter result was replicated in many experimental studies in a variety of settings building on VHBB. In particular, when groups reached sizes of 8 or larger only the least efficient equilibria were observed (see, Devetag and Ortmann, 2007, for an overview).³

Weber (2006) exogenously manipulates group size in weakest-link games. In a smart design, he exploits the fact that very small groups relatively easily coordinate on the efficient equilibrium. In his study, players start in groups of two and subsequently the size of the group is exogenously increased by slowly adding more players. With this mechanism, groups up to size 12 regularly

2. Theoretical studies in evolutionary game theory have been looking into the long-run effects of allowing interaction in neighbourhoods and/or location choice in 2×2 coordination games (see, *e.g.* Blume, 1993; Ellison, 1993; Oechssler, 1999; Ely, 2002). The latter two suggest that free neighbourhood choice can be crucial for achieving the Pareto-dominant equilibrium in 2×2 stag-hunt games.

3. Specifically, studies have shown that more—but not fully—efficient outcomes can be reached for small groups (size 5) when increasing the incentives for playing the more efficient action, introducing leaders and communication, increasing the length of the game, allowing for sorting, and inducing social identity (Cooper *et al.*, 1992; Knez and Camerer, 1994; Cachon and Camerer, 1996; Berninghaus and Ehrhart, 1998; Charness, 2000; Berninghaus and Ehrhart, 2001; Goeree and Holt, 2001; Weber *et al.*, 2001; Brandts and Cooper, 2006a; Weber, 2006; Brandts and Cooper, 2007; Brandts *et al.*, 2007; Hamman *et al.*, 2007; Chaudhuri *et al.*, 2009; Chen and Chen, 2011; Kogan *et al.*, 2011; Cooper *et al.*, 2014).

coordinate on efforts higher than the lowest one. However, larger groups only rarely achieve full efficiency and in most cases efficiency eventually declines with group size. This indicates that exogenously growing initially small high effort groups is not a remedy for decreasing effort in large groups.

In this article, we show that the possibility of freely choosing ones interaction neighbourhood has an overwhelmingly positive impact on achieved efficiency in weakest-link games, even in very large groups and under various information and payoff settings. Specifically, we experimentally test behaviour of subjects in various versions of repeated weakest-link games, without and with neighbourhood choice, in groups of size 8 and 24. In our BASELINE TREATMENT (BT) 8 players are forced to play with each other and to simultaneously choose an integer number (effort) between 1 and 7. All choosing 1 is the least efficient Nash equilibrium and all choosing 7 the most efficient Nash equilibrium.

Our NEIGHBOURHOOD TREATMENT (NT) differs from the BT only in that, interaction between any two players is endogenous and requires mutual consent. Consequently, a player's benefit is determined by the minimum effort in her neighbourhood. Further, in line with our field examples, which exhibit returns from increasing interaction neighbourhood size, we provide incentives to endogenously form larger neighbourhoods. Importantly, our design guarantees, that when each player chooses to be in the interaction neighbourhood of each other player, the incentives in NT coincide with those of BT.

In addition to these two basic treatments, we run four more treatments in order to check the robustness of results in respect to three important aspects of the interaction environment. First, we test the role of access to information about past behaviours. In NT subjects can look up all past behaviour of other group members irrespective of whether they have actually interacted with them or not. In a treatment with information loss, NT-IL, a subject knows another group member's behaviour only if she has interacted with that other subject. Second, to check the robustness of results with respect to the payoff specification we run a treatment, NT-AP, where costs are constantly high, but benefits change with neighbourhood size. Third, to test robustness with regard to group size, we run two treatments where we implement groups with twenty-four subjects, a size not studied before in weakest-link game experiments. These two treatments, BT-XL and NT-XL, are replicates of BT and NT, respectively, except that groups are three times as large in size.^{4,5}

Our findings in the treatments with fixed interaction structure, BT and BT-XL, are in accordance with the results reported in the literature. Efficient coordination is rarely observed in groups of size 8 and never in groups of size 24. In stark contrast, in all of our treatments with neighbourhood choice, subjects quickly coordinate on the fully efficient equilibrium, where virtually each subject chooses the highest effort level and interacts voluntarily with everybody

4. There are a few other experimental studies, which allow for endogenous interaction structures in coordination games. Corbae and Duffy (2008) investigate the effect of payoff shocks in a four players two-stage stag-hunt game. Corten and Buskens (2010) explore groups of size 8 playing pair-wise stag-hunt games in a dynamic setting where players can remove or add at most one interaction link in each round. Chen (2011) creates group identity by allowing for endogenous group formation via a ranking mechanism in three-person minimum effort games. Yang *et al.* (2013) use a complex voting mechanism to allow for exit and merger of small groups in a weak-link game. Salmon and Weber (2015) allow players to migrate from low performing (initially) large groups to high performing (initially) small groups.

5. Some other related papers look into the effect of group formation on contribution rates in social dilemma or public good games. An exemplary paper of this research is Charness and Yang (2014). They investigate linear public goods games where groups can endogenously form through voting on exit and entry of group members and merging of subgroups. Other important studies allowing for endogenous group formation in public goods experiments are Ehrhart and Keser (1999), Coricelli *et al.* (2004), Cinyabuguma *et al.* (2005), Page *et al.* (2005), Güth *et al.* (2007), Ahn *et al.* (2008) and Ahn *et al.* (2009a). Some studies also explore variations of prisoners' dilemma games when groups are endogenously formed (Hauk and Nagel, 2001; Riedl and Ule, 2002; Ule, 2008; Ahn *et al.*, 2009b; Brandts *et al.*, 2009).

else. Furthermore, over time there is no tendency towards less efficient equilibria. We identify a simple robust mechanism that drives this result: subjects who choose low effort levels are frequently excluded from the neighbourhood of those who choose higher effort levels and the former increase effort in response. Over time, when eventually all subjects are choosing the highest effort, exclusion becomes obsolete, which in turn boosts overall welfare.

The possibility of neighbourhood choice sorts two main effects. First, exclusion of low effort providers reduces strategic uncertainty and thus allows for self-protection of the excluding player. Second, as exclusion is, on average, costly for excluding as well as excluded subjects it resembles to some extent a costly punishment mechanism (Fehr and Gächter, 2000). In order to see if the punishment element is necessary for the identified mechanism to be effective we conducted another treatment. In this treatment, NT-SP, subjects can protect themselves from low effort providers via neighbourhood choice but those who are “excluded” are materially unaffected, effectively switching off any punishment opportunities. We find that in NT-SP behaviour is basically identical to behaviour observed in NT. This shows that free neighbourhood choice is very effective in boosting efficient coordination even in the absence of a punishment element.

2. THE WEAKEST-LINK COORDINATION GAME

The games we investigate are based on the minimum effort game of VHBB. The **BASELINE GAME** (BG) is the same as in VHBB and players can only choose between different effort levels. In the **NEIGHBOURHOOD GAME** (NG) we allow players, in addition to their effort choice, to choose their interaction neighbourhood by mutual consent. This extension requires an adjustment of the payoff function for the cases where not all group members interact with each other. The BG is a special case of the NG, where the neighbourhood of each player is fixed to be the entire group.

2.1. *The baseline game*

Let $N = \{1, 2, 3, \dots, n\}$ be a group of players and $E = \{1, 2, \dots, 7\}$ be a set of effort levels. In the BG, each player simultaneously chooses an effort level $e_i \in E$. Let $s = (e_i)_{i \in N}$ be the strategy profile of the players in the group, b the marginal cost of effort, and a the marginal return from the lowest effort in the group, with $a > b > 0$. The payoff of player i is given by

$$\pi_i(s) = a \min_{j \in N} \{e_j\} - be_i + c, \quad (1)$$

where $c > 0$ ensures non-negative payoffs for all strategy profiles, and $a > b > 0$ implies that every player has a monetary incentive to align her effort level with the minimum level chosen by the other players. Therefore, the strategy profiles $(\tilde{e})_{i \in N}$ where all players choose the same effort level $\tilde{e}$ are the pure strategy Nash equilibria,⁶ which can be Pareto-ranked from the highest to the lowest effort equilibrium. Further, the strategy profile where every player chooses the lowest effort is pairwise risk-dominating any other equilibrium.⁷

6. In line with the experimental literature, we focus on pure strategy equilibria. Equilibria in mixed strategies do exist, however.

7. There is no generally accepted definition of risk-dominance beyond 2×2 games and, in general, risk-dominance is also not transitive (Harsanyi and Selten, 1988). However, in our particular game it can be easily checked that any lower effort equilibrium risk dominates any higher effort equilibrium. For a recent concept of risk-dominance beyond 2×2 games, see Peski (2010).

2.2. The neighbourhood game

As in BG each player i chooses simultaneously an effort level $e_i \in E$. Additionally, each player i simultaneously chooses a set of players $I_i \subseteq N \setminus \{i\}$ with whom she would like to interact. For an interaction to actually take place mutual consent is required. That is, two players i and j interact with each other if and only if $i \in I_j$ and $j \in I_i$. Let $s = (s_1, s_2, \dots, s_n)$ with $s_i = (e_i, I_i)$ be a strategy profile in NG. The interaction neighbourhood of player i is given by $N_i(s) = \{j | j \in I_i \wedge i \in I_j\}$. The size of i 's neighbourhood, *i.e.*, the number of players i is interacting with, is denoted by $n_i = |N_i(s)|$.

In contrast to BG, in NG players do not necessarily interact with all other players in the group. In such situations, the payoff function in NT needs to have the following two properties. First, it should be comparable to the payoff function in BG and, second, it should reflect the trade-off players face when choosing between more or less interaction partners.

The first property is achieved by applying the same payoff function as in the BG for any *given* interaction neighbourhood size. The second property regards the trade-off between increasing strategic uncertainty due to increasing neighbourhood size, on the one hand, and increasing potential benefits due to returns from larger interaction neighbourhoods, on the other hand. A natural way to implement this trade-off is to let a player's benefit depend on the minimum effort in the player's interaction neighbourhood and to make the payoff proportional to the neighbourhood size. We, therefore, introduce the relative neighbourhood size $\frac{n_i}{n-1}$ into the payoff function of BG, such that a player's payoff in NG is given by

$$\pi_i(s) = \frac{n_i}{n-1} \left[a \left(\min_{j \in N_i(s) \cup \{i\}} \{e_j\} \right) - b e_i + c \right]. \quad (2)$$

This specification guarantees that isolated players always earn zero, implying that it is always beneficial to interact. Further, all players, except isolated ones, are engaged in a weakest-link game with their neighbourhood. Hence, irrespective of the neighbourhood size incentives are weakest-link game incentives. Importantly, for the case where each player's interaction neighbourhood consists of all other players (*i.e.* all interact with each other: for all i , $I_i = N \setminus \{i\}$ and $i \in I_j$ for all $j \neq i$) the payoff functions in NG and BG coincide. This guarantees equivalent incentives between games, when all players of a group interact.⁸

3. EXPERIMENT DESIGN

Our basic experiment design comprised a BASELINE TREATMENT (BT) where we implemented the BG and a NEIGHBOURHOOD TREATMENT (NT) where we implemented the NG, both with groups of size 8. In each treatment, subjects played the corresponding game repeatedly for thirty rounds in fixed matching groups. To assure anonymity, subjects did not get to know the identity of the other group members. They were referred to themselves as *me* and to the others in their group with capital letters *A*, *B*, *C*, etc. For each subject these identifiers remained fixed throughout the experiment. The parameters of the payoff function were the same as in VHBB: $a=20$, $b=10$, and $c=60$. Table 1 shows the corresponding payoff table of a player i .

In BT, in each round each subject interacted with all other group members and chose, simultaneously and independently, an effort level. In NT, subjects in each round did not only

8. An alternative payoff function in NG would have the proportionality factor applied only to the benefits and not to the costs. While in applications both scenarios are conceivable, there are good reasons to assume that costs increase with neighbourhood size (*e.g.* due to communication costs with neighbours, maintenance costs of infrastructure, costs of diplomacy with trading partners, etc.). Nevertheless, we have run an additional treatment where we implemented the alternative payoff function (see, Section 5.5.2).

TABLE 1
Payoffs of Player i

	$\min_{j \in N_i(s) \cup \{i\}} \{e_j\}$						
	7	6	5	4	3	2	1
e_i	7	130	110	90	70	50	30
	6	–	120	100	80	60	40
	5	–	–	110	90	70	50
	4	–	–	–	100	80	60
	3	–	–	–	–	90	70
	2	–	–	–	–	–	80
	1	–	–	–	–	–	–

Note: In the table neighbourhood size is ignored. In BT the neighbourhood coincided with all other group members and actual payoffs coincided with the table entries. In NT the payoffs in the table were multiplied by the fraction of players i interacted with relative to all other members in the group.

choose an effort level, but also decided simultaneously whether or not to propose an interaction to any of the other group members. For any pair of subjects an interaction took place only if both proposed to interact with each other. Figure 1 shows an example screen-shot for NT. (For an example screen-shot of BT, see Figure C.1 of the experiment instructions provided in Online Supplementary Appendix C.) Note, that two subjects could interact with each other while having been involved in only partially overlapping interaction neighbourhoods. For instance, in the situation depicted in the screen-shot, A and G interact with each other. Yet, A 's interaction neighbourhood comprises next to G only subject me , while G interacts also with subjects B and C . In all treatments, when making their decisions, subjects had access to the complete history of their own and other group members' actions.⁹

4. EXPERIMENT PROCEDURES

All sessions were conducted at the Behavioral and Experimental Economics laboratory (BEElab) at Maastricht University. A majority of subjects were students of business (53%) and economics (25%) while the rest came from other programs (22%). Participants were recruited through announcements by email and students' intranet or ORSEE (Greiner, 2015). Each subject only participated in one experiment session.

The experiment was computerized using z-Tree experiment software (Fischbacher, 2007). To ensure anonymity and avoid communication during the experiment subjects were seated in sight-shielded cubicles. Thereafter, subjects received written instructions which they could study at their own pace. For clarifications they could ask questions in private. The experiment did not start before all subjects had correctly answered a series of comprehension questions (see Online Supplementary Appendix C for the instructions).

We conducted nine sessions with two groups of size 8 in each session. In four sessions we implemented the BT and in five sessions NT, with in total 64 and 80 subjects. Subjects were

9. This design element resembles the actual information structure in the network infrastructure examples presented in the Introduction. For other examples, like co-authors, it is unlikely that past effort is known, although probably some statistics on past performance are available. Therefore, we do not consider the implemented information structure as more realistic than alternative ones, but as a good starting point. In Section 5.5.1 we report on a treatment with a radically coarser information structure.

FIGURE 1

Decision screen in the neighbourhood treatment.

Note: On the left side of the screen subjects could browse through the previous outcomes. Thick lines connecting two letters indicate that the corresponding pair of subjects interacted in that round. Thin non-connecting lines between two letters mean that only one subject proposed to interact and interaction did not take place (*e.g.* subject *C* wanted to interact with subject *A*, but *A* did not wish to interact with *C*). In order to save some time and effort costs, the interaction decisions from the previous round were used as default for the current round. The decision screens in the BT looked similar except that there were interaction lines between all pairs of letters shown on the left side and the decision was only about the effort level

paid according to their performance in the experiment. In the experiment, all earnings were calculated in points which were converted into cash (20 points = € 0.10) and confidentially paid out immediately after the experiment. A typical session in BT and NT lasted 80 and 99 minutes, respectively, and participants earned on average € 11.80 in BT and € 17.12 in NT.

5. RESULTS

We start with the presentation and comparison of effort choices in BT and NT, followed by a discussion of the endogenously formed interaction neighbourhoods and the role of exclusion in NT. Thereafter, we present results on achieved welfare in both treatments.¹⁰

10. Online Supplementary Appendix A discusses benchmark predictions for both treatments based on stochastic stability (Young, 1993, 1998) and Ely's (2002) evolutionary model of coordination with neighbourhood choice.

FIGURE 2

Cumulative distribution of efforts over rounds (BT & NT). (a) BT; (b) NT.

5.1. Effort levels without and with neighbourhood choice

In the experiment, we collected data from eight groups in BT and ten groups in NT. Each group forms an independent observation and, if not stated otherwise, our statistical tests are based on aggregated measures of these independent groups. All tests are two-sided.

Figures 2(a) and (b) show for both treatments how the cumulative distribution of effort levels develop over time. In the first round, we observe little difference between treatments. The average effort level is 5.66 in BT and 5.99 in NT. A Mann–Whitney (MW) test applied to the individual first round effort choices does not reject the hypotheses of equality ($p = 0.8919, n = 144$). Despite their similarity in the first round, the chosen effort levels show very different dynamics in the two treatments. In BT the frequency of the lowest effort (11% in round 1) is strongly increasing over time and becomes the most frequent choice as of round 19. The frequency of the highest effort level deteriorates over time from 64% to about 30% in the last few rounds. In NT, in contrast, the frequency of the lowest effort never reaches above 4%, while the frequency of the highest effort level strongly increases over time from about 60% in the first round to almost 100% in later rounds. In fact, as of round 4, its frequency is mostly above 96% and never falls below 88%.

This impression of opposite dynamics in effort choices in the two treatments is corroborated by Jonckheere–Terpstra (JT) tests.¹¹ The tests show that in BT the frequency of effort level 1 is significantly increasing over rounds ($p = 0.0002, n = 8$), while the frequencies of effort levels equal to or larger than 4 are significantly decreasing ($p \leq 0.0018, n = 8$) over time. In NT, in contrast, only the frequency of effort level 7 is significantly increasing ($p < 0.0001, n = 10$), while the frequencies of all other effort levels decrease or do not change significantly.

Table 2 reports descriptive statistics for respectively average and minimum effort at the group level, across rounds. It shows that both measures are significantly larger in NT than in BT. The mean of group average effort is 4.00 in BT and 6.85 in NT. The difference is statistically highly significant according to an MW test ($p = 0.0010, n = 18$).¹² Similarly, the mean of group minimum

11. The JT test is a non-parametric test for ordered differences of a response variable among classes (Pirie, 2006). Here it tests the null hypothesis that the distribution of the frequency of a given effort level does not differ among rounds. The alternative hypothesis is that there is an ordered difference among rounds. That is, if $f(e)_i^t$ denotes the frequency of effort level i in round t then $f(e)_i^1 \leq f(e)_i^2 \leq \dots \leq f(e)_i^{29} \leq f(e)_i^{30}$ (or $f(e)_i^1 \geq f(e)_i^2 \geq \dots \geq f(e)_i^{29} \geq f(e)_i^{30}$) with at least one strict inequality.

12. If not otherwise indicated, we use non-parametric MW tests for comparing treatments, using group averages across all thirty rounds as units of observations. This guarantees independence but also ignores a large amount of

TABLE 2
Descriptive statistics — average and minimum effort

	Average effort		Minimum effort		No. of obs.
	Mean	St.dev.	Mean	St.dev.	
BT	4.00	2.09	2.93	2.37	8
NT	6.85	0.18	6.27	0.69	10

Note: No. of obs. refers to the average across rounds per group.

FIGURE 3

Group means of average and minimum efforts over rounds (BT & NT).

effort is with 6.27 significantly larger in NT than in BT, where it amounts only to 2.93 ($p=0.0033$, $n=18$).

Figure 3 shows that average and minimum effort also exhibit very different dynamics between treatments and that differences become more pronounced over time. JT tests confirm that average efforts are significantly decreasing in BT ($p=0.0001$, $n=8$) but significantly increasing in NT ($p<0.0001$, $n=10$). Similarly, minimum effort levels significantly increase in NT ($p<0.0001$, $n=10$) but stay low in BT ($p=0.6030$, $n=8$).¹³

At the individual group level we find that in BT three groups manage to achieve the highest effort for a number of rounds. In the remaining five groups, effort levels quickly deteriorate and stay at or close to the lowest effort of 1 for most remaining rounds. Although chosen efforts do not fully converge to the lowest possible one in all investigated groups, the observed effort levels and the clear downward dynamics of minimum efforts are within the sampling variation

information. In Online Supplementary Appendix B.1 we also report the results of random effects regressions with group averages per round as unit of observation and treatment dummies as independent variables.

13. The observed statistically significant upward dynamics in NT hold despite a so-called endgame effect frequently observed in repeated games. Note, also that the effect for minimum effort looks strong but that in fact only 1–2 subjects per group deviate in the rounds close to the end.

FIGURE 4

Frequencies of interactions and exclusions over rounds (NT).

in previous experiments with a similar information feedback structure.¹⁴ In NT, in nine of the ten investigated groups, minimum effort converges to the highest possible level of 7 within the first ten rounds. Also in the remaining group the average effort level is never below 6 from round 3 onwards.¹⁵ We report the development of average and minimum effort over time for each group in each treatment in Online Supplementary Appendix B.3.

5.2. Size and development of interaction neighbourhoods

In NT it is important to know whether interactions actually take place, because only then high effort levels also have a high economic return. For instance, if no interaction takes place all group members earn zero points, irrespective of the effort choice, while if all 8 subjects interact with each other and choose effort 7 they together earn 1040 points.

For any pair of subjects, an interaction actually takes place only if both sides propose an interaction. Thus, if a pair of subjects does not interact this can be due to either mutual or unilateral exclusion (*i.e.* either none or only one of both involved subjects proposes to interact with each other). Figure 4 depicts the frequencies of these three possible situations over time. Averaged over all rounds the frequency of interactions amounts to 93%. In the first round, on average, 99% of all possible interactions are proposed and 78% actually take place. Thereafter, there is a slight decrease in the frequency of interactions which reaches its minimum of 74% in round 3. From there onwards, this frequency is almost monotonically and significantly increasing over

14. With full information feedback, Brandts and Cooper (2006b) and Berninghaus and Ehrhart (2001) report relatively high minimum effort levels in groups of size 4 and 8, respectively, while Devetag (2005) and VHBB (1990) report convergence to low effort levels.

15. As other studies (*e.g.* Brandts and Cooper, 2006b) reporting non-minimal effort, we observe an endgame effect. It is observed in eight of the eighteen groups (two in BT, six in NT). In all groups (except one) the deviation to lower levels is caused by exactly one subject. In one NT group four players deviated simultaneously.

FIGURE 5

Frequencies of interaction neighbourhood sizes over rounds (NT).

rounds ($p < 0.0001$, JT test, $n = 10$) and as of round 9 it never drops below 94%. This increase in interactions is accompanied by a simultaneous decrease in the frequency of unilateral exclusions ($p < 0.0001$, JT test, $n = 10$), which drops to practically zero.¹⁶

Although suggestive, the above is not fully informative about the development of the size of interaction neighbourhoods. Figure 5 provides this information. It shows the evolution of the cumulative frequencies of the neighbourhood sizes in NT over time. Recall that in our definition neighbourhood size does not include the subject itself. Thus, a neighbourhood of size 7 is maximal and means that a subject interacts with all other group members, and a size of 0 indicates an isolated group member. The figure shows that in the first round almost 90% of the neighbourhoods were of at least size 4 and almost 60% of at least size 6. Consistent with the already discussed development of interaction frequencies, up to round 3 the frequencies of neighbourhoods of size smaller than 7 slightly increase at the expense of neighbourhoods of maximal size. Thereafter, we see a sharp increase in the frequency of the largest possible neighbourhood and as of round 10 about 90% of all individual neighbourhoods are maximal (just as in BT) and virtually all neighbourhoods are at least of size 5. These visual impressions are corroborated by statistical tests. Over rounds, only neighbourhoods of maximal size significantly increase in frequency ($p < 0.0001$, JT test, $n = 10$), whereas the frequency of neighbourhood sizes 6 to 1 significantly decrease ($p \leq 0.0022$, JT test, $n = 10$). The negligible frequency of isolated group members does not change over time ($p = 0.6284$, JT test, $n = 10$).

5.3. Exclusion as efficiency enforcement

Here we examine if low effort providers are excluded by higher effort providers and, if yes, whether being excluded makes the former to increase subsequently chosen effort levels. If that

16. In fact, as of round 3, in all but one group all possible interactions take place in virtually each round. The one group where sometimes not all members interact with each other is the same group that do not fully settle on everybody always choosing the highest effort of 7. Recall, however, that also in this group average effort never falls below 6.

TABLE 3
Exclusion rates and responses to exclusion

$t-1$	Effort of i relative to effort of j and efforts in j 's neighbourhood					
	$e_i^+ : e_i \geq e_j$		$e_i^- : e_i < e_j$ but $e_i > \min_{k \in N_j} \{e_k\}$		$e_i^{--} : e_i < e_j$ and $e_i = \min_{k \in N_j} \{e_k\}$	
t	Exclusion rates (in%)					
	0.6 (84/14,738)		23.6 (21/89)		38.5 (105/273)	
$t+1$	i 's response (in%)					
	$j \in I_i$ $j \notin I_i$		$j \in I_i$ $j \notin I_i$		$j \in I_i$ $j \notin I_i$	
$e_i \uparrow$	11.8 (9)	2.6 (2)	71.4 (15)	9.5 (2)	61.6 (61)	10.1 (10)
$e_i =$	68.4 (52)	14.5 (11)	4.8 (1)	4.8 (1)	18.2 (18)	0.0 (0)
$e_i \downarrow$	1.3 (1)	1.3 (1)	9.5 (2)	0.0 (0)	10.1 (10)	0.0 (0)

Note: In panel t , the number of cases where exclusion takes place and the total number of cases (interactions in $t-1$) are in parentheses. In panel $t+1$, $j \in I_i$ ($j \notin I_i$) indicates the instances where a subject i excluded in t by j proposes (does not propose) in $t+1$ an interaction link to her excluder j , with the number of cases (of exclusion) in parentheses. The sum of cases in round $t+1$ can be lower than in round t due to exclusion in $t=30$ for which no further round exists.

turns out to be the case it would explain the observed parallel dynamics of effort levels and interaction frequencies.

Specifically, we look at the changes in dyadic relationships for all pairs of group members, i and j , and their effort choices in all three-rounds intervals from $t-1$ to t and $t+1$. We do so in three steps. First, we compare the effort levels of i and j in $t-1$. Second, we analyse whether j excludes i from her interaction neighbourhood in t and how this depends on the chosen effort levels in $t-1$. Third, we check how the chosen effort level of i in $t+1$ compares with her effort level in $t-1$, and how a change in effort depends on having been a low effort provider in $t-1$ and on being excluded in t .

For the first step we categorize all dyadic relations into three distinct effort classes in dependence of the relative efforts of i and j in $t-1$. The first class includes all cases where i provided at least as high an effort as j ($e_i \geq e_j$). The second class consists of the cases where i provided a lower effort than j but a higher effort than the lowest effort in j 's neighbourhood ($e_i < e_j$ but $e_i > \min_{k \in N_j} \{e_k\}$). The third class includes the cases where i 's effort was lower than j 's and also the lowest effort in j 's neighbourhood ($e_i < e_j$ and $e_i = \min_{k \in N_j} \{e_k\}$). For brevity we will refer to the first effort class as e_i^+ , the second as e_i^- , and the third as e_i^{--} .

In the second step, we examine for each of these classes the frequency of exclusion of i by j in round t . The upper panels, $t-1$ and t , of Table 3 report the results. Panel $t-1$ restates the introduced classes of relative efforts between i and j . Panel t reports for each class the frequency (in %) of severed links, *i.e.* the exclusion rate, together with the number of cases in parentheses. From the leftmost column it can be seen that, when i chooses an effort level that is not lower than the effort level of j , the exclusion rate is negligible (0.6%). In stark contrast, when i chooses a lower effort level than j , but is not the lowest effort provider in j 's interaction neighbourhood, the risk of being excluded from j 's neighbourhood is quite high (23.6%) and further increases to 38.5% when

i is the lowest effort provider in j 's neighbourhood (middle and rightmost columns, respectively). To test whether these differences in exclusion rates across effort classes are statistically significant, we calculate the exclusion rates for each class and each independent matching group separately and apply a Wilcoxon signed-rank test. All three pairwise comparisons are significant ($0.6 < 23.6$, $p = 0.039$, $n = 8$; $23.6 < 38.5$, $p = 0.016$, $n = 8$; $0.6 < 38.8$, $p = 0.002$, $n = 10$; exact tests). Hence, higher effort providers indeed frequently exclude lower effort providers from their interaction neighbourhoods.¹⁷

In the third step, we examine whether the observed exclusion actually affects subsequently chosen effort levels of excluded subjects. To this end, we investigate the change in chosen effort levels from period $t-1$ to $t+1$ for those cases where i was excluded in round t . An excluded subject can react in two dimensions. First, she may still propose an interaction to j ($j \in I_i$) or avoid interaction with j ($j \notin I_i$) in $t+1$. Second, she may not change the effort level ($e_i =$), increase it ($e_i \uparrow$) or decrease it ($e_i \downarrow$).

Panel $t+1$ of Table 3 reports the results. In 81.5 to 89.9% of the cases, an excluded subject proposes an interaction to the excluding subject (see panel $t+1$ in Table 3, columns $j \in I_i$). Hence, excluded subjects overwhelmingly do not break interactions with those who excluded them. Further, low effort providers (classes e_i^- and e_i^{--}) strongly respond to exclusion with an increase in effort levels. Specifically, in 80.9 and 71.7% of the cases an excluded low effort provider increased her effort in response to being excluded (see Table 3, panel $t+1$, row $e_i \uparrow$: $80.9 = 71.4 + 9.5$ and $71.7 = 61.6 + 10.1$). If a subject was excluded although she did not provide less effort than the excluder (class e_i^+) then effort did not change in 82.8% of the cases (see row $e_i =$: $82.8 = 68.4 + 14.4$).¹⁸ Thus, exclusion indeed makes low effort providers subsequently increase their effort levels and does not discourage high effort providers.

An important question regards the incentives a subject considering to exclude a low effort provider may have. Exclusion of low effort providers, while keeping own effort level high, can have different short-term effects on the excluding subject's earnings. On the one hand, it may increase earnings because exclusion can increase the minimum effort in the remaining neighbourhood, which may overcompensate the loss from fewer interactions. On the other hand, an excluding subject may incur short-run losses, due to the reduced neighbourhood size. In the first case exclusion is consistent with myopic earnings maximization, whereas the second case is reminiscent of costly punishment in public goods games (see, e.g. Fehr and Gächter, 2000).

As effort and neighbourhood choices are made simultaneously, costs from exclusion can be calculated either in expected terms or in experienced terms. We opt for the latter. That is, for a subject j who excludes a subject i in round t , we calculate (given the efforts actually chosen in t) the counter-factual payoff in t if j would not have excluded this subject and compare it to the actual payoff of j in t . This returns the cost (or benefit when negative) for an excluding subject j in round t per excluded subject i .

Overall, in 37.6% of all cases where a subject drops an interaction link to any of her neighbors it leads to an increase (or no decrease) in the excluding subject's short-term earnings. In the remaining 62.4% of the cases exclusion is costly for the excluding subjects. On average, benefits

17. We assume that a subject j is mainly concerned with effort provisions in her own neighbourhood and thus analyse how i 's effort ranking in j 's neighbourhood affects the likelihood that j excludes i . One may also ask whether j also is concerned with how i behaves relative to i 's own neighbourhood. We looked into this and find that the effort rankings of i in her own neighbourhood and in j 's neighbourhood are highly correlated (Spearman's $\rho = 0.87$) and, hence, the results do not substantially differ.

18. Ideally, we could provide a statistical test showing that in comparison to subjects not excluded, those who face exclusion increase effort levels significantly more. However, there are too few cases where low effort providers are not excluded by any neighbour. (Note, that the 23.6 and 38.5% in Table 3 refer to the excluding subject j and not to the excluded subject i .)

of exclusion amount to 15.2 points and costs amount to 13.6 points, per exclusion. Taken together, for the excluding subject exclusion is costly (2.8 points) because of the relatively larger number of costly exclusions.

The experienced costs for excluded subjects are calculated in a similar way as for excluding subjects. For each round t , given the actually chosen efforts, we calculate the counter-factual payoff of a subject i excluded by a subject j if the latter did not exclude the former and compare it to the actual payoff of i in round t . This returns the cost (or benefit when negative) for an excluded subject i in round t per excluding subject j . Overall, in 27.1% of all cases excluded subjects experienced no costs or even a benefit but in 72.9% of all cases they experienced a loss.¹⁹ On average, per exclusion, benefits amount to 6.3 points whereas costs amount to 13.0 points. Overall, being excluded is clearly costly for excluded subjects as they lose on average 7.8 points per exclusion.²⁰

In summary, high effort providers often exclude lower effort providers from their interaction neighbourhood, even though in half of the cases this implies short-term costs. In response, lower effort providers increase their effort levels and are eventually included again. This explains the strong dynamics towards the highest possible effort level of 7 where everybody interacts with everybody else (cf. Sections 5.1 and 5.2).

5.4. *Welfare*

Failure of (efficient) coordination can be the source of large welfare losses. VHBB distinguish two types of coordination failure. First, players may fail to predict the effort levels of other players and, therefore, fail to coordinate on any equilibrium (individual coordination problem). Second, players may coordinate, but do so on inefficient equilibria (collective coordination problem). In the previous sections, we have seen that in both treatments the individual coordination problem is solved over time. However, the collective coordination problem is only solved in NT and, as shown above, the mechanism behind this is exclusion. Yet, we have also seen that exclusion is on average costly for both, the excluding and the excluded subject. Therefore, at the outset it is not clear whether the overall earnings in NT are higher than in BT.

To examine that, we calculate a group's welfare as the sum of earnings of all group members. In addition, we calculate the maximally possible welfare where every group member chooses the highest effort level of 7 and interacts with all other group members (optimal benchmark). Figure 6 shows the average welfare levels over time for both treatments as well as the optimal benchmark.

Despite an endgame effect triggered by a few subjects, in both treatments average welfare levels are significantly increasing over rounds ($p < 0.0001$, JT tests, $n = 8$ and $n = 10$, respectively). In BT, this is mainly due to the fact that subjects learn to overcome the individual coordination problem and coordinate on the same (low) effort level (cf. Section 5.1 and Figure 3). In NT, in contrast, the stark increase in welfare is induced by overcoming both the individual and the collective coordination problem. As of round 10 up to almost the last round welfare is basically identical to optimal welfare. Taken over all rounds, average welfare in NT is economically and statistically significantly larger than in BT (BT: 629.5, NT: 913.1, $p = 0.0129$, $n = 18$). Hence, endogenous neighbourhood choice not only increases effort levels but also welfare.

19. Being excluded can lead to a benefit if the excluding subject provides a lower effort than the excluded subject. Note, that the reported cases include those where both sides dropped a link simultaneously which leads to zero benefits (and costs).

20. For the whole distribution of costs and benefits of excluding as well as excluded subjects, see the histograms reported in Online Supplementary Appendix B.5.

FIGURE 6
Welfare levels over rounds (BT & NT).

5.5. Robustness of results: information, payoff function, group size

The results presented in the previous section clearly show that neighbourhood choice can boost coordination in groups of intermediate size to full efficiency levels. In this section we check the robustness of this result with three additional treatments exploring first, the effect of information loss in case of exclusion, second, an alternative payoff function, and third, a large group size.

First, in the setting investigated until now, subjects had access to the full history of past links and effort choices of all other subjects. This could be a crucial assumption because subjects may be less willing to exclude low effort providers when they know that this will have the consequence of losing information about the behaviour (effort choices and linking decisions) of the excluded. As a result, exclusion of low effort providers may not take place and effort may decrease to low levels similar to the BT. Moreover, even if exclusion takes place, the fact that no information about behaviour of the excluded is available may make high effort providers hesitant to propose to interact with them again. This could lead to neighbourhood structures where subjects only interact in small isolated cliques. In both scenarios, the possibility of neighbourhood choice thus might be ineffective in achieving efficiency.

Second, we have assumed a payoff function where both costs and benefits, change proportionally with neighbourhood size (see equation (2)). While we believe that this is a natural choice, interesting alternative specifications are conceivable. Specifically, it seems interesting to study a case where benefits are generated in proportion to neighbourhood size while costs are independent of it (cf. Footnote 8). This implies that in small neighbourhoods costs have a relatively high weight, which may make it more difficult for small neighbourhoods to sustain high effort levels. If subjects anticipate this, high effort providers could be reluctant in excluding low effort providers. In that case, the possibility of neighbourhood choice may not be effective in facilitating high efficiency, because non-exclusion of low effort providers disrupts the mechanism leading to high effort in NT.

Third, the experimental literature discussed in the Introduction unambiguously shows that for fixed groups of size larger than 12, eventual convergence of behaviour towards inefficient

FIGURE 7

Decision screen in the NT-IL treatment.

Note: As in NT thick lines connecting two letters indicate that the corresponding pair of subjects interacted in that round and thin non-connecting lines between two letters mean that only one subject proposed to interact and interaction did not take place. In contrast to NT, in NT-IL, the subject did not receive information—about chosen effort levels and interactions (proposals) with other group members—of group members and the subject did not interact within a given round (e.g. subject *me* neither saw *C*'s effort level nor any interaction (proposal) involving *C* and another group member). Moreover, for interacting subjects no information was provided about the interaction with *other* subjects (e.g. subject *me* saw *G*'s effort level but nothing about *G*'s interactions (proposals) with other group members).

equilibria is unavoidable. In the previous section, we have seen that allowing for neighbourhood choice provides an effective countervailing force in groups of intermediate size. It is a natural question to ask if neighbourhood choice is equally effective when groups are considerably larger.

5.5.1. Neighbourhood choice with information loss. To test whether neighbourhood choice remains an effective mechanism for inducing efficient effort provision, even if exclusion of neighbors entails a significant information loss, we ran treatment NT-IL. In this treatment, information on past behaviour depends on neighbourhood choices in the following way. In any round t , a subject i received information on the chosen effort level of another group member j in any previous round t' ($t' < t$) only if j was in the neighbourhood of i in round t' . In other words, a subject could access information on the effort levels of other group members only for those rounds in which these subjects had interacted with each other. Moreover, even when a group member j was in the neighbourhood of subject i in a given round, the latter did not receive any information about the other neighbors of j . Naturally, each subject knew to whom she proposed to interact with and what her actual neighbourhood was. Each subject also got to know when other subjects proposed to interact with her. All other aspects were identical to treatment NT.

Figure 7 shows an example screen shot of treatment NT-IL. Experiment instructions for this treatment can be found in Online Supplementary Appendix C.3. In three sessions, we collected

FIGURE 8

Cumulative distribution of efforts over rounds (NT-IL and NT-AP). (a) NT-IL; (b) NT-AP.

TABLE 4
Descriptive statistics — average and minimum effort

	Average effort		Minimum effort		No. of obs.
	Mean	St.dev.	Mean	St.dev.	
BT	4.00	2.09	2.93	2.37	8
NT	6.85	0.18	6.27	0.69	10
NT-IL	6.87	0.09	6.32	0.46	9
NT-AP	5.59	1.63	4.08	2.19	9

Note: No. of obs. refers to the average across rounds per group.

data from nine groups of size 8 with in total seventy-two subjects. A typical session lasted 79 minutes and average earnings of a participant were €16.78.

Our main interest is, first, to explore if we see a similar effort choice and neighbourhood pattern and, second, if we see the same interaction dynamics as in NT. Figure 8(a) shows the development of the cumulative distribution of effort over time for the new treatment NT-IL. It is clear from the figure that the dynamics of the different effort levels in NT-IL and NT are almost identical (cf. Figure 2(b)). The frequencies of all effort levels smaller than 7 are significantly decreasing ($p < 0.0075$, JT tests, $n = 9$) whereas the frequency of effort level 7 is significantly increasing over rounds ($p < 0.0001$, JT tests, $n = 9$).

Consequently, average and minimum effort levels are close to the maximum effort level of 7 (cf. row NT-IL Table 4) and do not differ significantly between NT-IL and NT (average effort: $p = 0.6242$; minimum effort: $p = 0.8381$; $n = 19$). In contrast, compared to BT, both measures are significantly larger in NT-IL (average effort: $p = 0.0029$; minimum effort: $p = 0.0122$; $n = 17$).

Despite the high effort levels in NT-IL, there is still the possibility of inefficiency because of small neighbourhoods, as the loss of information about other group members behaviour may induce fragmentation. Figure 9(a) shows the development of frequencies of interactions and mutual and unilateral exclusions over rounds for NT-IL. The figure shows that the dynamics are very similar to those in NT (cf. Figure 4). The frequency of mutual exclusion is generally negligible and the frequency of unilateral exclusion is significantly decreasing ($p < 0.0001$, JT test, $n = 9$), while the frequency of interactions is significantly increasing over rounds ($p < 0.0001$, JT test, $n = 9$). Specifically, as in NT, also in NT-IL there is a convergence to full interaction. Overall, between treatments there is no difference in interaction frequencies ($p = 0.2883$, $n = 19$). Hence, in NT-IL there is no evidence for inefficiency due to small neighbourhoods.

FIGURE 9

Frequency of interactions and exclusions over rounds (NT-IL and NT-AP). (a) NT-IL; (b) NT-AP.

The high effort levels and high interaction frequencies already suggest that welfare in NT-IL will be higher than in BT and similar to NT. This is indeed the case. Taken over all rounds, in NT-IL average welfare is 894.9 and similar to NT (913.1) ($p=0.5676$, $n=19$) but significantly larger than in BT (629.5) ($p=0.0209$, $n=17$).

The presented results clearly show that the high effort and welfare levels observed in NT are robust to severe reduction in information about other group members' behaviour, in case of exclusion.

5.5.2. Alternative payoff function. In this section, we explore if the results will be affected by a different specification of the payoff function. In NT, both costs and benefits from providing effort were proportionally affected by the size of the neighbourhood of a player (cf. equation (2)). Here, in our alternative payoff treatment, NT-AP, we explore the specification

$$\pi_i(s) = \frac{n_i}{n-1} \left[a \left(\min_{j \in N_i(s) \cup \{i\}} \{e_j\} \right) \right] - be_i + c, \quad (3)$$

where benefits from effort provision are still proportional to neighbourhood size but costs of providing effort are independent of it. All other elements were exactly the same as in the original NT.²¹ A consequence of making costs independent of neighbourhood size is a higher costs-to-benefits ratio in smaller neighbourhoods in comparison to larger neighbourhoods. We know from other studies (Goeree and Holt, 2001; Brandts and Cooper, 2006a) that higher cost-to-benefit ratios make it considerably more likely that effort converges to low levels even when groups are small. Hence, in this treatment the reduction in neighbourhood size through exclusion of low effort providers has two effects. On the one hand, it decreases strategic uncertainty that can induce the choice of higher effort levels but, on the other hand, it may also induce the choice of lower efforts through the increase in the cost-to-benefit ratio of effort provision. In fact, with our parameter specifications it can be shown that for neighbourhoods of size smaller than 4, the weakest-link game turns into a social dilemma game.²² That is, in such small neighbourhoods it

21. In order to avoid negative payoffs, we increased the value of parameter c from 60 to 70. The strategic structure of the game remains the same.

22. This can be easily seen by checking the marginal benefits and marginal costs of effort provision, $mb(e) = \frac{n_i}{n-1}a$ and $mc(e) = b$, respectively. It is easily seen that $mb(e) < mc(e)$ iff $n_i < \frac{7}{2}$ for our parameters, $a=20$, $b=10$ and $n=8$.

becomes a dominant strategy to choose the lowest effort level 1. In that sense this payoff function provides a rather harsh test of the robustness of the results observed in NT.

In treatment NT-AP, we collected data from nine groups of size 8 in three sessions comprising in total seventy-two subjects. A typical session lasted 105 minutes and average earnings of a participant were €15.20. Experiment instructions for this treatment are provided in Online Supplementary Appendix C.4.

Figure 8(b) shows the development of the cumulative distribution of effort over time for NT-AP. It shows that, similar to NT (cf. Figure 2(b)), over time there is a clear development towards a high frequency of the maximum effort level of 7. It increases from around 35% in round 1 to around 60% by the end of the game. This increase in frequency is highly significant ($p=0.0165$, JT test, $n=9$). Also similar to NT, the effort levels 6 to 2 significantly decrease over time ($p<0.0001$, JT test, $n=9$). However, in contrast to NT, the dynamics are not strong enough to wipe out these lower effort levels. Until the end of the game there is a non-negligible fraction (around 40%) of effort levels 5 and lower. Interestingly, the frequency of the lowest effort level of 1 is slightly and marginally significantly increasing over time ($p=0.0652$, JT test, $n=9$). This is consistent with results from fixed neighbourhood weakest-link and stag-hunt games, showing that higher costs (and lower benefits) of effort provision make high effort levels more difficult to achieve (Goeree and Holt, 2001; Brandts and Cooper, 2006a). The incentive effect appears to carry over to weakest-link games with neighbourhood choice. Importantly, however, although the dynamics towards full efficiency are weaker than in NT they are still in clear contrast to those observed in BT (cf. Figure 2(a)).

Row NT-AP of Table 4 reports descriptive statistics for group average and minimum effort levels. They reflect the observed differences in frequencies of effort levels between NT and NT-AP. The mean of average effort as well as minimum effort is somewhat lower in NT-AP than in NT (average: $p=0.0037$; minimum: $p=0.0113$, $n=19$). Compared to BT, these measures are economically considerably larger in NT-AP, despite the unfavorable specification of the payoff function (average effort: $p=0.1489$; minimum effort: $p=0.2881$, $n=17$).²³

Figure 9(b) shows the development of exclusion and interactions over time in NT-AP. Overall the dynamics are very similar to NT (cf. Figure 4) with initially a decreasing interaction frequency, fueled by unilateral exclusion, which recovers soon and approaches full interaction. Statistical tests corroborate this visual impression. Over time the frequency of unilateral exclusion is significantly decreasing ($p<0.0001$, JT test, $n=9$) and the frequency of interactions is significantly increasing ($p<0.0001$, JT test, $n=9$). Over all rounds the interaction frequency in NT-AP is with 84% somewhat lower than in NT where it is 93% ($p=0.0090$, $n=19$). However, as of round 12 it never falls below 88.9% and the difference with NT vanishes in the last 10 rounds ($p=0.4627$, $n=19$).²⁴

Taken over all rounds, in NT-AP average welfare amounts to 810.6 which is a bit lower than in NT (913.1) ($p=0.2207$, $n=19$). Compared to BT (629.5), welfare is about 30% higher and the difference is statistically significant ($p=0.0543$, $n=17$).²⁵

In sum, the higher costs-to-benefits ratio of effort provision in NT-AP makes it somewhat more difficult to achieve the highest effort level of 7. However, in comparison to NT the differences are small, whereas in comparison to BT they are considerable. Moreover, the welfare effects of

23. When using random effects regression analysis, effort levels are significantly different at $p=0.027$ and $p=0.012$, respectively.

24. That there is no perfect convergence to full interaction is due to only two groups which reach stable interaction frequencies of 75%, while the other seven groups stabilize at (almost) full interaction from round 21 onwards.

25. A random effects regression returns $p=0.026$.

FIGURE 10

Cumulative distribution of efforts over rounds (BT-XL and NT-XL). (a) BT-XL; (b) NT-XL.

neighbourhood choice in NT-AP are indistinguishable from NT but statistically and economically significant in comparison to BT without neighbourhood choice.

5.5.3. Large groups. For fixed groups the literature on weakest-link games undoubtedly shows that there is a strong negative relation between group size and the likelihood to observe high effort levels. Specifically, in groups of size 12 or larger almost solely the least efficient effort level of 1 is chosen after a few repetitions (cf. Van Huyck *et al.*, 1990; Weber, 2006). Therefore, increasing the group size is a crucial test for the robustness of the effectiveness of neighbourhood choice in guaranteeing (close to) maximal efficiency.

To test this we conduct treatments with groups of size 24, which is, to the best of our knowledge, the largest group size ever investigated for the weakest-link game in the laboratory. In comparison to our BT, a NT with group size 24 differs in two dimensions: the possibility of neighbourhood choice and the size of the group. Therefore, in order to provide a clean comparison, we conducted two treatments with large groups. In one treatment, BT-XL, we replicate with large groups the basic treatment without neighbourhood choice. In the other treatment, NT-XL, we implement the possibility of neighbourhood choice in the same way as in NT. Hence, BT-XL and NT-XL are identical to BT and NT, respectively, except for group size and, consequently, a scaling in the payoff function. Experiment instructions for these treatments can be found in Online Supplementary Appendix C.2.

For the large groups with 24 members we gathered data from six groups, three BT-XL and three NT-XL groups, run in six sessions. Hence, both treatments comprised seventy-two subjects. For matters of comparison, we kept the marginal incentives the same as in BT and NT, but paid an extra lump-sum of €5.— to compensate for the longer duration of about 30 minutes. A typical session in BT-XL and NT-XL lasted 103 and 128 minutes, respectively. Average earnings (including the lump-sum) were €14.33 in BT-XL and €20.27 in NT-XL.

Figure 10 shows the cumulative distributions of effort levels over rounds for BT-XL (Figure 10(a)) and NT-XL (Figure 10(b)). Similar as observed in BT and NT, in the first round, the lowest effort level of 1 is chosen equally infrequent in BT-XL and NT-XL (6.9%). The highest effort level of 7 is more often chosen in the NT-XL (65.3%) than in the BT-XL (44.4%). This translates into a significantly higher average effort in NT-XL than in BT-XL, already in round 1 ($p = 0.0238$, MW test, $n = 144$).

In addition, and even topping the results for groups of size 8, the dynamics of effort choices are dramatically different in the two treatments. In BT-XL, the average frequency of the lowest effort

TABLE 5
Descriptive statistics — average and minimum effort

	Average effort		Minimum effort		No. obs.
	Mean	St.dev.	Mean	St.dev.	
BT-XL	1.78	0.15	1.00	0.00	3
NT-XL	6.83	0.04	5.41	0.49	3

Note: No. of obs. refers to the average across rounds per group.

level of 1 quickly increases from 6.9% in the first round to at least 86.1% as of round 20, whereas the average frequency of the highest effort level 7 quickly decreases from 44.4% in round 1 to at most 4.2% as of round 16. This impression of opposite dynamics is corroborated by statistical tests. The frequency of effort level 1 is significantly increasing over rounds ($p < 0.0001$, JT test, $n = 3$), whereas the frequencies of all effort levels higher than 1 are decreasing ($p \leq 0.0010$, JT tests, $n = 3$). This is in stark contrast to what is observed in NT-XL. There the cumulative average frequency of effort levels smaller or equal 6 decreases from 34.7% in round 1 to negligible 1.4% as of round 9, while the average frequency of the highest effort level 7 increases from 65.3% in round 1 to at least 98.6% as of round 9.²⁶ These opposite trends of the highest effort level and the other lower effort levels are statistically significant (effort levels ≤ 6 decreasing: $p \leq 0.0225$; effort level 7 increasing: $p < 0.0001$, JT tests, $n = 3$).

Table 5 shows clear differences between both XL treatments in the group means of average and minimum effort. The average effort in BT-XL is only 1.78 but 6.83 in NT-XL and the minimum effort in BT-XL is only 1.00 but 5.41 in NT-XL. Inspection of the individual group data reveals that in each NT-XL group average and minimum effort levels are at or close to the highest effort level most of the time, whereas these statistics quickly converge to the lowest effort level in each BT-XL group. These differences are statistically significant ($p = 0.10$, $n = 6$; exact test).²⁷ When comparing effort levels between NT-XL and NT we see neither in average effort nor in minimum effort a significant difference (cf. Tables 5 and 2; average effort: $p = 0.161$, minimum effort: $p = 0.112$, $n = 13$; exact tests). This shows that neighbourhood choice is undoubtedly at least as effective in inducing (close to) maximum effort in large groups as it is in groups of intermediate size.

The development of actual interactions and unilateral and mutual exclusions in NT-XL is depicted in Figure 11. The picture is similar to the one in NT (cf. Figure 4), with somewhat more pronounced dynamics in the first couple of rounds. In the first round, on average 87.8% of all possible interactions take place. This frequency sharply decreases to a minimum of 58.7% in round 3, after which it almost monotonically increases. As of round 11 it never falls below 97%. The overall increasing frequencies of interactions and decreasing frequencies of unilateral exclusions are statistically significant ($p < 0.0001$, JT tests, $n = 3$).

Similar to BT and NT, in both BT-XL and NT-XL, welfare increases over rounds ($p < 0.0001$, JT tests, $n = 3$). However, the opposing dynamics in effort choices in BT-XL and NT-XL, together with the convergence towards full interaction in NT-XL, implies welfare effects that are strongly in favour of the neighbourhood treatment. As of round 6, total welfare achieved in NT-XL is above the one achieved in BT-XL. Furthermore, in BT-XL actual welfare never exceeds 52.5% (round

26. The only exception is the last round where effort level 7 is chosen with an average frequency of 90.3% and effort levels 1 and 2 are chosen with an average frequency of 8.3 and 1.4%, respectively.

27. $p = 0.10$ is the lowest achievable p -value for the two-sided non-parametric exact MW test, given that the number of our strictly independent observations is 2×3 . Random effects regression analysis returns $p < 0.0001$.

FIGURE 11

Frequency of interactions and exclusions over rounds (NT-XL).

17) of the optimally achievable welfare level, while in NT-XL it almost never falls below 94.7% of the optimal level as of round 11. Across all rounds, the achieved average welfare in BT-XL (1492.6) is only 61% of the welfare achieved in NT-XL (2443.6). The difference is statistically significant ($p=0.10$, $n=6$; exact test).²⁸

5.5.4. Exclusion in the robustness treatments. We have seen that effort choices and interactions are largely robust to changes in the information environment, payoff specification, and group size. A final question is whether exclusion behaviour and responses to it, as observed in NT, are also robust to these changes in the environment. In NT we have observed that higher effort providers in a round $t-1$ often exclude lower effort providers from their neighbourhood in round t and that the latter increase effort in response in round $t+1$ (cf. Table 3). We see qualitatively very similar behaviour in the other treatments.²⁹

Table 6 shows a summary of the exclusion rates in NT and the three robustness treatments. In all treatments, when in $t-1$ a subject i provided an effort equal to or higher than a subject j , the likelihood that j excludes i in round t is basically zero (exclusion rates: 0.2–0.6%). In contrast, when subject i provides strictly less effort than j in $t-1$, exclusion rates are substantial in all treatments. As in NT, exclusion was most pronounced when i was the subject with lowest effort in j 's neighbourhood (exclusion rates: 21.2–73.2%). Interestingly, there is also some variation, with exclusion rates being highest in NT-XL and lowest in NT-IL and NT-AP.

To test for statistical significance of differences in exclusion rates across effort classes within each treatment, we proceed in the same way as in Section 5.3. It turns out that in both robustness

28. As for effort levels this is the lowest p -value that can be achieved for the two-sided non-parametric exact MW test, given that the number of our independent matching groups is 2×3 . A random effects regression returns $p < 0.0001$.

29. We have produced tables equivalent to Table 3. For brevity and to save on space we relegate them to Online Supplementary Appendix B.2.

TABLE 6
Summary of exclusion rates (all NT's)

$t-1$	Effort of i relative to effort of j and efforts in j 's neighbourhood		
	e_i^+	e_i^-	e_i^{--}
NT	0.6	23.6	38.5
NT-IL	0.2	3.5	24.5
NT-AP	0.6	5.8	21.2
NT-XL	0.4	45.8	73.2

Note: e_i^+ denotes cases in interacting dyads in a round $t-1$ where a subject i provides at least as high an effort as subject j ; e_i^- denotes cases where a subject i provides strictly less effort than a subject j but not the lowest effort in j 's neighbourhood; e_i^{--} denotes cases where a subject i provides strictly less effort than a subject j which is also the lowest effort in j 's neighbourhood.

treatments with group size 8, all pair-wise comparisons except one are statistical significant at $p \leq 0.0357$ ($n=9$).³⁰ Similarly, in NT-XL all pair-wise comparisons are significant at the lowest possible significance level for a Wilcoxon test, given the number of independent observations ($p=0.25$, $n=3$). That is, in *each* of the groups of size 24 exclusion rates are highest for subjects i who provide the lowest effort in j 's neighbourhood, second highest for those who provide less than j but not the lowest effort in j 's neighbourhood, and basically zero for those i 's who provide at least as much as j .

Hence, although there are some quantitative differences, the pattern of exclusion in the robustness treatments resembles the one observed in NT. This holds especially regarding the exclusion of the worst performing subject in the neighbourhood of a higher performing subject.

Table 7 reports for all robustness treatments a summary of how subjects responded with their effort choices after having been excluded.³¹ As can be seen, effort responses to exclusion are very similar across treatments. In the robustness treatments, when excluded after having provided at least the same effort as the excluding subject, in an overwhelming majority of cases (63.9–81.3%) subjects keep the same effort, which is in keeping to the 82.9% observed in NT. After having been excluded, when having provided less effort than the excluding subject, in a large majority of cases (64.8–86.9%) subjects increase their effort levels. Decreasing effort is infrequent and observed only in 3.6–10.5% of cases. These numbers of the robustness treatments are similar to those in NT, where it is observed that effort increases in 73.3% and decreases in 10.0% of cases.

In sum, the responses to exclusion in treatments NT-AP, NT-IL, and NT-XL are very similar to those in NT. We conclude that the mechanism of exclusion that leads to and maintains high effort choices is robust across the investigated treatment variations.

5.6. Punishment or self-protection?

We have seen that exclusion is effective in boosting efficient coordination under various circumstances. There are two different aspects inherent in the act of exclusion that may be important in achieving that.

30. The only exception is in NT-IL when comparing exclusion rates for e_i^- with those for e_i^+ , where $p=0.9324$. This is due to the relatively low exclusion rate (3.5%) for e_i^- .

31. Tables detailing the results for e_i^- and e_i^{--} separately can be found in Online Supplementary Appendix B.2. For NT, see Table 3.

TABLE 7
Summary of responses to exclusion (all robustness NT's)

$t-1$	Effort classes	
	e_i^+	e_i^- and e_i^{--}
$t+1$	i 's response to exclusion (in %)	
	NT-IL	
$e_i \uparrow$	15.6	72.7
$e_i =$	81.3	18.2
$e_i \downarrow$	3.1	9.1
	NT-AP	
$e_i \uparrow$	21.3	64.8
$e_i =$	63.9	24.8
$e_i \downarrow$	14.8	10.5
	NT-XL	
$e_i \uparrow$	20.9	86.9
$e_i =$	75.0	9.5
$e_i \downarrow$	4.1	3.6

Note: e_i^+ denotes cases in interacting dyads in a round $t-1$ where a subject i provides at least as high an effort as subject j ; e_i^- and e_i^{--} denotes cases where a subject i provides strictly less effort than a subject j ; $e_i \uparrow$ ($e_i =$) [$e_i \downarrow$] denotes an effort increase (no change) [decrease] by i in $t+1$ after having been excluded by j in t .

First, regarding the excluded low effort provider, exclusion entails a punishment effect comparable to the one observed in public goods problems with punishment. This literature has also shown that punishment is effective when the ratio of the costs for the punishing subject to the costs for the punished subject is sufficiently small (Egas and Riedl, 2008; Nikiforakis and Normann, 2008). Our discussion at the end of Section 5.3 has shown that in NT the costs of exclusion for the excluding subject are small (sometimes exclusion is even beneficial) but relatively large for the excluded subject. Hence, punishment by exclusion may indeed be an important force behind the effectiveness of neighbourhood choice.

Second, for high effort providers, exclusion entails a self-protection element as it allows them to “remove” (potential) low effort providers from their neighbourhoods, thereby reducing strategic uncertainty. Hence, the question arises whether neighbourhood choice will still be effective in achieving efficient coordination when exclusion is ripped off its punishment aspect.

To investigate this we have run another neighbourhood choice treatment, called NT-SP, where we allow for self-protection but not for punishment by exclusion. This treatment is exactly as the original NT, except that being in a player's interaction neighbourhood does not require mutual consent. Specifically, a player j is in the neighbourhood of a player i if i proposes to interact with j , irrespective of j 's interaction decision. Importantly, this allows for asymmetric interaction as j can be in i 's neighbourhood without i being in j 's neighbourhood. If a player i proposes to interact with j but the latter does not propose to interact with the former, then i 's payoff is affected by j 's choices but not vice versa. Hence, a subject's interaction decision only affects this subject's own neighbourhood (and payoffs) but not the neighbourhood (and payoffs) of other subjects. Consequently, in NT-SP exclusion of a low effort provider by a high effort provider entails self-protection of the latter by eliminating the former from her neighbourhood, but does not lead to

FIGURE 12

Cumulative distribution of efforts over rounds (NT-SP).

punishment of the low effort provider. This allows us to disentangle the effect of self-protection from the effect of monetary punishment.

For this treatment, we ran three sessions and collected data from nine groups of size 8 with in total seventy-two subjects. On average a session lasted 91 minutes and subjects earned on average €16.83. Except for the neighbourhood choice, instructions were the same as in NT and can be found in Online Supplementary Appendix C.5.

In presenting the results, we report first on effort levels followed by the dynamics of interactions mainly in comparison to our basic neighbourhood treatment NT. Thereafter, we explore if “exclusion” of low effort providers has similar effects as in the other treatments even when there is no punishment involved.

Figure 12 reports the development over time of the cumulative distribution of effort in NT-SP. The dynamics are very similar to NT (cf. Figure 2(b)), except that in NT-SP there is a slightly higher frequency of the second highest effort level 6. Similarly to NT, statistical tests show that over time only the frequency of effort level 7 is significantly increasing ($p < 0.0001$, JT test, $n = 9$), whereas all other effort levels are either decreasing or do not change (effort 6: $p = 0.1026$, efforts 2-5: $p \leq 0.0055$, JT tests; effort 1: too few observations to conduct a JT test). Consistent with this result also the overall average and minimum efforts are similar in NT-SP and NT (cf. Table 8). Statistical tests reject the hypothesis of no differences between treatments at the 5% level, for average effort ($p = 0.0940$, $n = 19$) and minimum effort ($p = 0.4875$, $n = 19$). Hence, effort levels in NT-SP are only marginally lower, which is mainly due to the higher frequency of effort level 6. The differences are economically negligible, however. Moreover, compared to BT, average and minimum effort are significantly higher in NT-SP (average effort: $p = 0.0053$; minimum effort: $p = 0.0093$, $n = 19$).

Figure 13 shows the interaction dynamics for NT-SP, which differ slightly from NT in the beginning but show a very similar trend (cf. Figure 4). Specifically, also in NT-SP there is a clear convergence to full interaction.

In summary, neighbourhood choice is surprisingly effective in facilitating and sustaining efficient coordination even when exclusion does not entail a punishment component. This raises

TABLE 8
Descriptive statistics — average and minimum effort

	Average effort		Minimum effort		No. of obs.
	Mean	St.dev.	Mean	St.dev.	
NT-SP	6.76	0.25	6.13	0.69	9
NT	6.85	0.18	6.27	0.69	10

Note: No. of obs. refers to the average across rounds per group.

FIGURE 13

Frequency of interactions and exclusions over rounds (NT-SP).

the question if the same mechanism is at work in NT-SP as in NT. That is, does exclusion, even when it does not hurt the excluded, nevertheless make excluded subjects increasing their effort levels? To answer this question, we conduct a similar analysis as for the other treatments but need to take into account that in NT-SP mutual interaction proposals are not required for (one-sided) interaction to take place. Specifically, we say that a subject j “excludes” a subject i in round t when j , first, proposed to interact with i in $t-1$ and, second, does not propose to interact with i in t . For brevity, when looking at the effort response of an “excluded” subject i in $t+1$ we ignore whether or not that subject actually had proposed to interact with j . Again we distinguish between effort classes e_i^+ and e_i^- & e_i^{--} .³²

Table 9 shows the results and leaves little doubt that “excluded” low effort providers overwhelmingly respond to “exclusion” by increasing their effort levels. Hence, this treatment shows that neighbourhood choice boosts efficient coordination and that punishment by exclusion is not necessary to achieve this result. This, of course, raises the question why neighbourhood choice works so well even when the punishment component is absent. This is an intriguing

32. A table reporting more detailed information can be found in Online Supplementary Appendix B.2.

TABLE 9
Exclusion rates and responses to exclusion (NT-SP)

$t-1$	Effort classes	
	e_i^+	e_i^- and e_i^{--}
t	'exclusion' rates (in %)	
	1.0 (124/12,776)	17.8 (171/960)
$t+1$	i 's response to "exclusion" (in %)	
$e_i \uparrow$	10.4	67.5
$e_i =$	82.6	24.7
$e_i \downarrow$	7.0	7.8

Note: e_i^+ denotes cases in interacting dyads in a round $t-1$ where a subject i provides at least as high an effort as subject j ; e_i^- and e_i^{--} denotes cases where a subject i provides strictly less effort than a subject j ; in panel t , number of cases where 'exclusion' takes places and total number of cases (proposed interactions of j in $t-1$) in parentheses; $e_i \uparrow$ ($e_i =$) [$e_i \downarrow$] denotes an effort increase (no change) [decrease] by i in $t+1$ after having been excluded by j in t .

question which is, however, beyond the scope of this article and left for future research. We discuss some possible explanations and avenues of future research in the concluding section.

6. CONCLUSION

In this article, we study the effect of neighbourhood choice on the efficiency of coordination in repeated weakest-link (aka minimum effort) games. Theoretically, the introduction of neighbourhood choice considerably worsens the coordination problem, as it expands the strategy space and hugely increases the number of pure strategy Nash equilibria.

We test the effect of neighbourhood choice experimentally and find that it boosts efficiency in groups with eight as well as twenty-four subjects. The mechanism behind this result is that in early rounds high effort providers exclude low effort providers, who in response increase effort and are included again. The efficiency result as well as the mechanism are robust to severe changes in the information condition and the payoff specification. We also test whether the punishment aspect inherent in exclusion is necessary for the mechanism to be effective. We find that it is not and that "self-protection" from low effort providers, which reduces strategic uncertainty, is sufficient.

The effectiveness of neighbourhood choice, even when punishment is absent, is intriguing, but also leaves open a number of important questions that are beyond the scope of this article but could be tackled in future research. First, in our self-protection treatment exclusion does not hurt the excluded subject in a material sense but it may still be perceived as symbolic punishment. For public goods games it has been argued that symbolic punishment can increase cooperation (Masclet *et al.*, 2003) and it could be interesting to investigate this also for (weakest-link) coordination games. Second, exclusion in the self-protection treatment may be viewed as a means of communication as it may signal low effort providers both, that there are agents who are willing to provide high effort and that they are wrong-doing in choosing low effort. It would be interesting to investigate the relative importance of these communication effects. Third, self-protection (and exclusion) allows high effort providers to keep their effort high while being unaffected by low efforts of others. Compared to weakest-link games in fixed groups this

provides high-effort providers with more room to “lead-by-example”, which has been shown to have some positive effect in public goods games (*e.g.* Güth *et al.*, 2007; Potters *et al.*, 2007) and may also play a role in our neighbourhood treatments.

On a more general level, the behaviour observed in our experiment indicates a high potential for self-organization on efficient outcomes, provided people are given sufficient freedom in choosing their interaction partners. Our results demonstrate that once agents have coordinated on the efficient equilibrium actual exclusion will be infrequent. Hence, in the field, efficiently coordinated groups where no exclusion is observed may still rest on the threat of exclusion and/or the possibility of self-protection. This may make it difficult to directly observe the effect of neighbourhood choice on behaviour in the field. Yet, indirectly inefficient outcomes may be traced back to the lack of neighbourhood choice.

Our study can be considered as a first stepping stone towards a promising research agenda investigating the potential and limits of neighbourhood choice in coordination problems. There are many possible directions for future research.

First, in all of our treatments, a player could individually exclude other players from her interaction neighbourhood, which—except for in the self-protection treatment—caused costs for the excluded player. To investigate how crucial the individual part of such punishment by exclusion is, one could explore other mechanisms for exclusion. For instance, after each round agents could vote on whom to exclude for the next round(s). Such a mechanism contains an extra coordination problem, which may or may not make it less effective in sustaining highly efficient outcomes over time.

Second, in the field, changing the interaction neighbourhood often comes with transactions costs. In our experiment this is reflected by the opportunity costs of not interacting. The introduction of nominal costs when abandoning an interaction may cause other behavioural effects than mere opportunity costs do.

Third, we have investigated two different information structures, which can be considered as extremes in the space of all possible information structures, and find little difference in behaviour. Still it is conceivable that intermediate information loss due to exclusion affects behaviour differently than the almost complete information loss we explored as a robustness check.

Fourth, we have investigated two different payoff specifications. Naturally, many others could be explored. For instance, it could be interesting to evaluate the effect of a non-proportional increase of earnings with neighbourhood size or the effect of limits on the number of neighbors one can interact with.

Fifth, a crucial feature of our investigated environments is that those who have learned their lesson after exclusion could be integrated again. The question arises if exclusion would be similarly effective when integration after exclusion would be more difficult or even impossible.

For fixed groups, Brandts and Cooper (2006a) have shown that financial incentives can be pretty effective in achieving efficient coordination in small groups. Similarly, Weber (2006) and Salmon and Weber (2015) point out the important role of slow exogenous group growth for facilitating efficient coordination. In our experiment, slowly growing groups cannot be the main driving force behind the observed high effort choices for at least two reasons: first, we see that in all treatments with neighbourhood choice interaction frequencies are already relatively high in the beginning, and second, after a short decrease in the first few rounds interaction frequencies converge quickly towards the maximum. Hence, our experiment shows that when players have the freedom to choose their interaction neighbourhood there seems to be no need for extra incentives or slow growth, even if groups grow very large. Our study, therefore, suggests an important complementary mechanism for achieving efficiency in organizations and institutions that is probably less expensive than increasing the payoff for coordination and more effective

than exogenously growing groups, especially when there are returns to group size and groups grow large.

Acknowledgments. we are grateful to Imran Rasul and four anonymous reviewers for their valuable comments. We also thank Abigail Barr, Jakob Goeree, Glenn Harrison, Friederike Mengel, Roberto Weber and participants of seminars and conferences in Amsterdam, Ann Arbor, Atlanta, Brussels, Cologne, Groningen, Innsbruck, Jena, Lyon, Mannheim, New York, Santa Fe, and Toulouse for their helpful comments. Martin Strobel thanks the Santa Fe Institute for the hospitality while working on a revision of the article. Financial support of the Oesterreichische Nationalbank (project no. 11780) is gratefully acknowledged.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- AHN, T. K., ISAAC, R. M. and SALMON, T. C. (2008), "Endogenous Group Formation", *Journal of Public Economic Theory*, **10**, 171–194.
- (2009a), "Coming and Going: Experiments on Endogenous Group Sizes for Excludable Public Goods", *Journal of Public Economics*, **93**, 336–351.
- AHN, T. K., ESAREY, J. and SCHOLZ, J. T. (2009b), "Reputation and Cooperation in Voluntary Exchanges: Comparing Local and Central Institutions", *The Journal of Politics*, **71**, 398–413.
- BERNINGHAUS, S. K. and EHRHART, K.-M. (1998), "Time Horizon and Equilibrium Selection in Tacit Coordination Games: Experimental Results", *Journal of Economic Behavior and Organization*, **37**, 231–248.
- (2001), "Coordination and Information: Recent Experimental Evidence", *Economics Letters*, **73**, 345–351.
- BLUME, L. E. (1993), "The Statistical Mechanics of Strategic Interaction", *Games and Economic Behavior*, **4**, 387–424.
- BRANDTS, J. and COOPER, D. J. (2006a), "A Change Would Do You Good ... An Experimental Study on How to Overcome Coordination Failure in Organizations", *The American Economic Review*, **96**, 669–693.
- (2006b), "Observability and Overcoming Coordination Failure in Organizations: An Experimental Study", *Experimental Economics*, **9**, 407–423.
- (2007), "Its What You Say, Not What You Pay: An Experimental Study of Manager Employee Relationships in Overcoming Coordination Failure", *Journal of the European Economic Association*, **5**, 1223–1268.
- BRANDTS, J., COOPER, D. J. and FATAS, E. (2007), "Leadership and Overcoming Coordination Failure with Asymmetric Costs", *Experimental Economics*, **10**, 269–284.
- BRANDTS, J., RIEDL, A. and WINDEN, F. V. (2009), "Competitive Rivalry, Social Disposition, and Subjective Well-being: An Experiment", *Journal of Public Economics*, **93**, 1158–1167.
- CACHON, G. P. and CAMERER, C. F. (1996), "Loss-Avoidance and Forward Induction in Experimental Coordination Games", *The Quarterly Journal of Economics*, **111**, 165–194.
- CAMERER, C. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*. (New York and Princeton: Princeton University Press).
- CHARNESS, G. (2000), "Self-serving Cheap Talk: A Test of Aumann's Conjecture", *Games and Economic Behavior*, **33**, 177–194.
- CHARNESS, G. and YANG, C.-L. (2014), "Starting Small Toward Voluntary Formation of Efficient Large Groups in Public Goods Provision", *Journal of Economic Behavior and Organization*, **102**, 119–132.
- CHAUDHURI, A., SCHOTTER, A. and SOPHER, B. (2009), "Talking Ourselves to Efficiency: coordination in Inter-generational Minimum Effort Games with Private, Almost Common and Common Knowledge of Advice", *The Economic Journal*, **119**, 91–122.
- CHEN, R. (2011), "Endogenous Groups and Cooperation", (Mimeo, University of Michigan).
- CHEN, R. and CHEN, Y. (2011), "The Potential of Social Identity for Equilibrium Selection", *American Economic Review*, **101**, 2562–2589.
- CINYABUGUMA, M., PAGE, T. and PUTTERMAN, L. (2005), "Cooperation Under the Threat of Expulsion in a Public Goods Experiment", *Journal of Public Economics*, **89**, 1421–1435.
- COOPER, R. (1999), *Coordination Games: Complementarities and Macroeconomics*. (Cambridge, UK: Cambridge University Press).
- COOPER, D. J., IOANNOU, C. A. and QI, S. (2014), "Coordination with Endogenous Contracts Incentives, Selection, and Strategic Anticipation", (mimeo).
- COOPER, R., DEJONG, D. V., FORSYTHE, R. et al. (1992), "Communication in Coordination Games", *The Quarterly Journal of Economics*, **107**, 739–771.
- CORBAE, D. and DUFFY, J. (2008), "Experiments with Network Formation", *Games and Economic Behavior*, **64**, 81–120.
- CORICELLI, G., FEHR, D. and FELLNER, G. (2004), "Partner Selection in Public Goods Experiments", *Journal of Conflict Resolution*, **48**, 356–378.

- CORTEN, R. and BUSKENS, V. (2010), "Co-evolution of Conventions and Networks: An Experimental Study", *Social Networks*, **32**, 4–15.
- DEVETAG, G. (2005), "Precedent Transfer in Coordination Games: An Experiment", *Economics Letters*, **89**, 227–232.
- DEVETAG, G. and ORTMANN, A. (2007), "When and Why? A Critical Survey on Coordination Failure in the Laboratory", *Experimental Economics*, **10**, 171–178.
- EGAS, M. and RIEDL, A. (2008), "The Economics of Altruistic Punishment and the Maintenance of Cooperation", *Proceedings of the Royal Society - B*, **275**, 871–878.
- EHRHART, K.-M. and KESER, C. (1999), "Mobility and Cooperation: On the Run", (Mimeo).
- ELLISON, G. (1993), "Learning, Local Interaction, and Coordination", *Econometrica*, **61**, 1047–1071.
- ELY, J. C. (2002), "Local Conventions", *Advances in Theoretical Economics*, **2**, Article 1.
- ENGELMANN, D. and NORMANN, H.-T. (2010), "Maximum Effort in the Minimum-Effort Game", *Experimental Economics*, **13**, 249–259.
- FEHR, E. and GÄCHTER, S. (2000), "Cooperation and Punishment in Public Goods Experiments", *American Economic Review*, **90**, 980–994.
- FISCHBACHER, U. (2007), "z-tree: Zurich Toolbox for Ready-made Economic Experiments", *Experimental Economics*, **10**, 171–178.
- GOEREE, J. K. and HOLT, C. A. (2001), "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions", *The American Economic Review*, **91**, 1402–1422.
- GREINER, B. (2015), "Subject Pool Recruitment Procedures: Organizing Experiments with Orsee", *Journal of the Economic Science Association*, 1, Forthcoming.
- GÜTH, W., LEVATI, M. V., SUTTER, M. and *et al.* (2007), "Leading by Example With and Without Exclusion Power in Voluntary Contribution Experiments", *Journal of Public Economics*, **91**, 1023–1042.
- HAMMAN, J., RICK, S. and WEBER, R. A. (2007), "Solving Coordination Failure with "All-or-none" Group-level Incentives", *Experimental Economics*, **10**, 285–303.
- HARRISON, G. W. and HIRSHLEIFER, J. (1989), "An Experimental Evaluation of Weakest Link/Best Shot Models of Public Goods", *Journal of Political Economy*, **97**, 201–225.
- HARSANYI, J. C. and SELTEN, R. (1988), *A General Theory of Equilibrium Selection in Games*. (Cambridge, MA: MIT Press).
- HAUK, E. and NAGEL, R. (2001), "Choice of Partners in Multiple Two-person Prisoner's Dilemma Games: An Experimental Study", *Journal of Conflict Resolution*, **45**, 770–793.
- HIRSHLEIFER, J. (1983), "From Weakest-link to Best-shot: The Voluntary Provision of Public Goods", *Public Choice*, **41**, 371–386.
- KNEZ, M. and CAMERER, C. (1994), "Creating Expectational Assets in the Laboratory: Coordination in 'Weakest-link' Games", *Strategic Management Journal*, **15**, 101–119.
- KOGAN, S., KWASNICA, A. M. and WEBER, R. A. (2011), "Coordination in the Presence of asset Markets", *The American Economic Review*, **101**, 927–947.
- MASCHLET, D., NOUSSAIR, C., TUCKER, S. *et al.* (2003), "Monetary and Non Monetary Punishment in the Voluntary Contributions Mechanism", *American Economic Review*, **93**, 366–380.
- NASH, Jr., J. F. (1950), "Equilibrium Points in n-person Games", *Proceedings of the National Academy of Sciences*, **36**, 48–49.
- NIKIFORAKIS, N. and NORMANN, H.-T. (2008), "A Comparative Statics Analysis of Punishment in Public Goods Experiments", *Experimental Economics*, **11**, 358–369.
- NORDHAUS, W. D. (2006), "Paul Samuelson and Global Public Goods", in Szenberg, M., Ramrattan, L. and Gottesman, A., (eds), *Samuelsonian Economics*, (Oxford, UK: Oxford University Press), 88–98.
- OECHSSLER, J. (1999), "Competition Among Conventions", *Computational & Mathematical Organization Theory*, **5**, 31–44.
- PAGE, T., PUTTERMAN, L. and UNEL, B. (2005), "Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency", *The Economic Journal*, **115**, 1032–1053.
- PESKI, M. (2010), "Generalized Risk Dominance and Asymmetric Dynamics", *Journal of Economic Theory*, **145**, 216–248.
- PIRIE, W. (2006), "Jonckheere Tests for Ordered Alternatives", *Encyclopedia of Statistical Sciences*, **9**. (New York: John Wiley).
- POTTERS, J., SEFTON, M. and VESTERLUND, L. (2007), "Leading-By-Example and Signaling in Voluntary Contribution Games: An Experimental Study", *Economic Theory*, **33**, 169–182.
- RIEDL, A. and ULE, A. (2002), "Exclusion and Cooperation in Social Network Experiments. (Mimeo, University of Amsterdam).
- SALMON, T. C. and WEBER, R. A. (2015), "Maintaining Efficiency While Integrating Entrants from Lower-Performing Environments: An Experimental Study", *Economic Journal*, forthcoming.
- SANDLER, T. (1998), "Global and Regional Public Goods: A Prognosis for Collective Action", *Fiscal Studies*, **19**, 221–247.
- SCHELLING, T. (1960), *Strategy of Conflict*. (Cambridge, MA: Harvard University Press).
- ULE, A. (2008), *Partner Choice and Cooperation in Networks: Theory and Experimental Evidence*. Number 598 in Lecture Notes in Economics and Mathematical Systems (Springer, Berlin Heidelberg).
- VAN HUYCK, J. B., BATTALIO, R. C. and BEIL, R. O. (1990), "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure", *The American Economic Review*, **80**, 234–248.

- WEBER, R., CAMERER, C., ROTTENSTREICH, Y. *et al.* (2001), "The Illusion of Leadership: Misattribution of Cause in Coordination Games", *Organizational Science*, **12**, 582–598.
- WEBER, R. A. (2006), "Managing Growth to Achieve Efficient Coordination in Large Groups", *The American Economic Review*, **96**, 114–126.
- YANG, C.-L., XU, M.-L. and TANG, F.-F. (2013), "Efficient Large-Size Coordination Via Voluntary Group Formation: An Experiment (Working paper, Research Center for Humanities and Social Sciences, Academia Sinica, Taipei).
- YOUNG, H. P. (1993), "The Evolution of Conventions", *Econometrica*, **61**, 57–84.
- (1998), *Individual Strategy and Social Structure*. (Princeton: Princeton University Press).