Efficient coordination in weakest-link games

Arno Riedl, Ingrid M.T. Rohde, Martin Strobel[1]

March 14, 2011

Preliminary and incomplete.

Please do not quote and circulate without permission of the authors.

Comments highly welcome.

[1]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. We thank Jakob Goeree and Friederike Mengel for helpful comments.

Abstract

Problems of coordination are at the core of most economic and social decision situations. Theory and evidence from laboratory experiments show that when people are forced to interact with other people efficiency strongly decays over time. This leaves us with the puzzling uncertainty if all behavior we observe outside theoretical models and outside the lab also has this tendency towards inefficiency. In a series of laboratory experiments we show that in weakest-link (minimum effort games) maximal efficiency can be achieved when participants are allowed to choose their interaction neighborhood. This holds for groups of up to 24 members, which is the largest group size ever investigated in laboratory weakest-link games. The results suggest that the possibility of choosing interaction partners is a key mechanism to achieve efficient outcomes and to sustain them over time. This has important consequences for the design of efficient organizations.

1 Introduction

Societies continuously face incidents in which coordination plays a crucial role. Many of these situations have a weakest link characteristic, i.e. the weakest player determines the outcome for all involved parties.

An intriguing example is given by Hirshleifer (1983). He describes Anarchia, a fictitious island on which each citizen owns a wedge-shaped sector. In order to prevent the island from occasional floods, each citizen builds a dike along the coastline of his plot of land. The topography of Anarchia is flat. This means that if the sea overcomes one dike the island will completely be flooded. In the absence of a governing body the citizens have to coordinate on the height and strength of the dikes. The lowest and weakest dikes determine the level of flood protection which Anarchia's citizens may expect.

Another example that has recently attracted attention of the media, is the fight against infectious diseases such as SARS, avian influenza, swine influenza, or AIDS. To prevent the diseases from breaking out or spreading, involved parties have to invest into precautionary means.[1] The party which exerts the lowest effort determines to a large extent the chances of outbreak or eradication of a disease. In the optimal case all parties would choose the highest precaution level. The same reasoning transfers to computer networks. Hackers or viruses try to enter weakly protected computers and use them as stepping stones to move further into the network.

An example from business is the situation of an airport ground crew (see for example Camerer 2003, p.381). While some teams are responsible for re-fueling and maintaining the aircraft, other teams de-board and board passengers or unload and load luggage. The aircraft can only take off again after all processes are completed. The time is determined by the team which is slowest.[2]

In all examples players face the problem to coordinate on high efforts, which is maximizing both, welfare as well as individual payoffs. Hence, in contrast to typical dilemma

[1] Involved parties could be individual persons or countries. Precautionary means could be for example vaccination on the individual level or the provision of medication or health education on the country level.

[2] Various other examples can be found in Sandler (1998).

situations, everyone providing high effort is an equilibrium but it is characterized by high strategic uncertainty. One single player deviating downwards is sufficient to cause substantial for each individual as well as for the whole group.

A closer look to our examples, however, yields a crucial structural difference. The inhabitants of Anarchia and the teams of the airport ground crew are exogenously bound to their neighborhood by the geography of the island and the organizational decisions of the airline's management. In contrast, the actors in the disease example and the computer network, be it countries or individuals, have the possibility to choose their neighborhood to interact with. If for example a network is repeatedly hacked via the same weakly protected computer, system administrators of other computers might deny access from this specific computer to theirs. Similarly governments may restrict, or at least discourage, free travel from and to countries with repeated outbreaks of infectious diseases.

The structural difference between endogenously chosen and exogenously given neighborhoods opens up a new perspective on coordination problems and might turn out as a crucial determinant when it comes to the observation of behavior in such situations.

For the case of exogenously given neighborhoods an appropriate game-theoretic model is the so-called minimum effort game which was introduced by Bryant (1983).[3] Van Huyck, Battalio and Beil (1990, VHBB henceforth) were the first to examine it experimentally. Subjects in their experiments repeatedly chose an effort level between 1 and 7 and received round payoffs that decreased in their own chosen level, but more strongly increased in the minimum level of the group. Any strategy profile where players play the same effort level constitutes an equilibrium. Moreover these equilibria can be ranked according to the criteria of risk dominance and payoff dominance (see Harsanyi and Selten, 1988).[4] VHBB were interested whether subjects could in the course of several rounds tacitly coordinate on high efforts.[5] The results were rather disappointing from the efficiency perspective. Only

[3]Bryant suggested the game as a simple, parsimonious micro-foundation to describe how coordination failure in macroeconomics could lead to bad equilibria.

[4]Although the concept of risk dominance is not transitive in general, it turns out to be for this particular type of game.

[5]Strictly speaking, VHBB separated two types of coordination problems. Firstly, players may fail to

in very small groups of two participants and in the case of costless effort (i.e. high effort was riskless) subjects were able to coordinate on the payoff dominant equilibrium.

These results have led to a wave of experiments in which their robustness was tested and conditions for the emergence of Pareto-efficient behavior were explored (see Devetag and Ortmann 2007 for an excellent overview). Almost all attempts focus on the reduction of strategic uncertainty. The most obvious are the potential efficiency gains and the size of the group. Brandts and Cooper (2006) showed that the increase in the potential efficiency gains mitigates the tendency to coordinate on inefficient equilibria. VHBB already demonstrated that small groups more easily coordinate on high efforts than large groups (see also Based on this result Weber (2006) showed that starting out from small groups and adding participants one by one also enables coordination on high effort levels in larger groups although coordination eventually collapsed in many groups. Other ways to reduce strategic uncertainty and to induce higher effort levels are the admission of cheap talk before the game (Blume and Ortmann 2007), the reduction of exploration costs by an increase of the number of rounds (Berninghaus and Ehrhart 1998), or (almost) common knowledge, unanimous "positive" advice of previous subjects (Chaudhuri et al. 2009).

Mixed evidence was found with respect to feedback information. While Berninghaus and Ehrhart (2001) found that average effort increases if subjects do not only get feedback about the minimum effort but also about its distribution, other studies like the original VHBB or Engelmann and Normann (2009) do not confirm this.

Not related to the reduction of uncertainty is the result of Feri et al. (forthcoming) that groups elicit more effort than individuals. Last but not least also cultural differences may play a role as being argued by Engelmann and Normann (2009).

By and large, although some studies provide gleams of hope, the picture for efficiency for the minimum effort game with exogenously given neighborhood looks quite discouraging.

predict the effort levels of other players and therefore fail to coordinate on an equilibrium (individual coordination problem). Secondly, as the equilibria can be Pareto ranked, players might coordinate but coordinate on a low payoff equilibrium (collective coordination problem). The first coordination problem is normally solvable via communication or repeated actions. Hence we concentrate on the more interesting second coordination problem, the efficiency problem.

Although coordination on the payoff-dominant equilibrium happens sometimes under some specific conditions it is not guaranteed. Moreover if happening it turns out to be quite fragile over time.

Endogenous neighborhood choice might change the picture because it introduces a kind of (sometimes) costly punishment. Instead of adjusting the effort level to the low efforts provided by some neighbors, players face the alternative to exclude them from their neighborhood. Excluding players forgo the benefits from interaction with the excluded players but they also reduces their strategic uncertainty. Hence, endogenous neighborhood choice may serve as a powerful device to achieve coordination on the payoff dominant equilibrium.

In our experiment we set out to test this hypothesis. Therefore we extend the minimum effort game and allow players to simultaneously choose their neighborhood together with their effort level. We run two treatments, a baseline treatment with a exogenously given full network (comparable to the standard minimum effort game) and a neighborhood treatment with endogenous neighborhood choice, each treatment with groups of 8 players.

We find indeed that in the long run outcomes of the neighborhood treatment converge to the maximum effort and complete neighborhoods (i.e. everybody is linked to everybody else). This is in sharp contrast to the result of the baseline treatment where most of the groups converge to the lowest effort equilibrium. The effect turned out to be remarkably stable with respect to group size. Even groups of 24 players yielded the same results.

In the remainder of the paper we will formally introduce the minimum effort game and its neighborhood extension in Section 2. Then we describe our experimental design and procedure in Sections 3 and 4. In Sections 5 and 6 we report on the results of the experiments and in Section 7 we conclude. The appendix provides the formal derivation of the theoretical benchmarks as well as experiment instructions.

2 The model

For our baseline treatment we use the same minimum effort game as in VHBB. We call it the baseline game. For the treatment with endogenous neighborhood choice we extend the players' strategy set to include a set of desired neighbors. Furthermore we adjust the payoff function (only the neighborhood is taken into account). We call this the neighborhood game.

2.1 The baseline game

Let $N = \{1, 2, 3, ..., n\}$ be a group of players and $E = \{1, 2, .., 7\}$ be a set of effort levels (1 being the lowest and 7 being the highest). Each player simultaneously chooses a strategy, i.e. an effort level $s_i = e_i$. Let $s = (s_i)_{i \in N}$ be a strategy profile of all players in the group. Further, let b denote the marginal cost of effort, let a be the marginal return on the lowest group effort, and let $a > b > 0$. The payoff of player i is determined 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.[6]

Since $a > b > 0$, every player has the incentive to choose his effort level equal to the minimum level of the other players. Therefore all profiles $s^{\tilde{e}} = (\tilde{e})_{i \in N}$ where all players play the same effort $\tilde{e}$ are Nash equilibria.[7] Furthermore $s^7 = (7)_{i \in N}$ is Pareto-dominant and $s^1 = (1)_{i \in N}$ is the risk-dominant, i.e. s^1 pairwisely risk-dominates any other pure strategy equilibrium.

2.2 The neighborhood game

As in the baseline game let $N = \{1, 2, 3, ..., n\}$ be a finite set of players and $E = \{1, 2, \ldots, 7\}$ a set of effort levels. Each player i chooses a strategy which consists of two parts: $s_i = (e_i, I_i)$. Additionally to the effort level e_i player i chooses a set of players with whom he would like to interact: $I_i \subseteq N \setminus \{i\}$.[8] An interaction requires mutual consent, i.e. i and j

[6]Negative payoffs are problematic from a practical point of view because it is not possible to legally enforce payments from subjects.

[7]We concentrate on pure strategy equilibria. Equilibria in mixed strategies exist, however.

[8]The words interaction and play can be used interchangeably. In the context of endogenous neighborhood choice we use the term interaction since whether or not to interact is a choice rather than an

interact if and only if $i \in I_j$ and $j \in I_i$. Let $s = (s_1, s_2, \ldots, s_n)$ be a strategy profile. The neighborhood of player i can now be defined as $N_i(s) = \{j | j \in I_i \wedge i \in I_j\}$ with $|N_i(s)|$ being the number of i's neighbors.

In reality, having more interactions also implies higher potential payoffs (recall the computer network and the disease example example from the introduction). An increase in the number of neighbors, however, also increases the chance of facing a low effort level in the own neighborhood. In our model players should face the same tradeoff between the size of the neighborhood an the strategic uncertainty which goes along with it. The simplest and most natural extension to the payoff function is to determine the minimum effort only with respect to the neighborhood, but account the payoff only proportionally to the neighborhood size. We therefore introduce the factor $\frac{|N_i(s)|}{n-1}$ into the payoff function:

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

For the case of the maximum neighborhood ($N_i(s) = N \setminus \{i\}$) the payoff function and the degree of strategic uncertainty for the neighborhood game and the baseline game coincide. Moreover it is always beneficial to increase one's neighborhood if the minimum effort does not decrease.

2.3 Theoretical benchmarks

Games involving networks may go along with a huge multiplicity of Nash equilibria. This is also the case in our neighborhood game because we require mutual consent for neighborhood links. For example any fragmentation of the network, where in each component the same effort is played and none of the players wants to interact with players from other components, is a Nash equilibrium. Consequently many actions can be justified as part of an equilibrium strategy by some appropriate beliefs. To get predictive power of our benchmark we apply stochastic stability (Young 1993, 1998) as an equilibrium refinement concept. We chose it over other refinement concepts, because it fits best to our experimental design. Subjects play the game for many rounds and they are able to adjust their

exogenously given fact.

strategy in each round (see next section for details).[9] The analysis yields

Proposition 2.1 *In both, the baseline and the neighborhood game, the unique stochastically stable equilibrium is the complete network with every player playing the lowest effort.*

The proposition summarizes the results of Propositions A.5 and A.7. See the appendix for details and the proofs. Note however, that in both cases the network is complete. While this is due to construction in the baseline game it emerges from the stochastic dynamics in the neighborhood game.

3 Design and hypotheses

Our experiment comprises two treatments; the baseline treatment and the neighborhood treatment. In each treatment subjects played the corresponding game repeatedly for 30 rounds in fixed matching groups of eight. To assure anonymity, subjects did not get to know the identity of the other group members. They referred to themselves as "me" and to the others with the capital letters "A" to "G". For each subject these identifiers remained fixed throughout the experiment. The parameters of the payoff function were set in the same way as in VHBB, i.e. $a = 20$, $b = 10$, and $c = 60$. This yielded an equivalent payoff table (see the payoff table in the instructions in the appendix). When making their decision, subjects had access to the complete history of other group members' previous actions.

3.1 Baseline treatment

In the baseline treatment the full neighborhood was exogenously enforced, i.e. in each round each subject had to interact with all group members and only chose an effort level. The only notable difference to the standard minimum effort game (VHBB, is the additional

[9] Crawford (1995) provides a discussion on the (in-)appropriateness of other equilibrium refinement concept. He also suggests and analyzes a rather general learning model to explain the data of VHBB. This model however requires a one-dimensional strategy space and is therefore not applicable for the neighborhood game.

information of all previous decisions. Moreover, to have the appearance of the screen most similar to the neighborhood treatment (see next section for details) the baseline subjects received a graphical representation of the network.[10]

The major problem, subjects face is the strategic uncertainty. Doubts in the high effort ambitions of other subjects induce playing safe, i.e. low effort. Doubts might be straight or of a more subtle higher order, for example a subject might not doubt another's willingness to coordinate on high efforts, but beliefs that the other is doubtful. As soon as more subjects follow a safe strategy, low-effort outcomes are self-confirming and very difficult to overcome.

This consideration is confirmed by the prevailing evidence from the literature. Subjects do not manage to coordinate on the payoff-dominant equilibrium but rather converge to the lowest effort over time (Devetag and Ortmann 2007). This also holds for extended feedback (VHBB, Engelmann and Normann 2009). Even in the extended feedback condition of Berninghaus and Ehrhart (2001) in which higher than usual effort levels are found, these levels deteriorate over time. The difference in the visual representation of the group might influence subjects' choices but only to a negligible extent. Therefore, we expect to find the usual result:

Hypothesis 3.1 *In the baseline treatment (minimum effort game with fixed exogenously given neighborhoods), subjects behavior will converge to the minimum effort.*

3.2 Neighborhood treatment

In the neighborhood treatment, neighborhoods are endogenous, i.e. in each round subjects do not only choose their effort level but also indicate for all other group members whether they want to interact with them or not. In order to save some time and effort costs, the interaction decisions from the previous round was given as default for the current round.

Before their decisions subjects can browse through the entire history of all other group members' actions. After the decision each subject's neighborhood and corresponding payoff

[10]The representation showed the full network and did not change of course. For a screenshot see the instructions in the appendix.

was individually determined (see Equation 2 in Section 2.2). Thereafter subjects received feedback about the size of their neighborhood, the minimum effort in their neighborhood, and their payoff.

The theoretical benchmarks of both treatments are identical, but from a behavioral aspect there is a difference. Suppose that subjects have an interest in efficient coordination. This is plausible since it does not only satisfy efficiency concerns but also pure self-interest. Suppose further, that nevertheless not all of them play the highest effort due to strategic uncertainty. Unlike in the baseline treatment a subject now faces an alternative way of reacting to low effort providers. Rather than adjusting her own effort downwards, she might exclude them from her neighborhood. Thereby exclusion might serve different motives. One one hand it might be part of a myopic better response[11] or a myopic best response. On the other hand she might also exclude others if this implies costs in order to punish them for their inefficient behavior.

The payoffs are proportional to the size of the subject's neighborhood. An excluded subject forgoes gains. Therefore a possible response is to increase the effort level in consecutive rounds. This gives the excluder an incentive to re-include the excluded subject into her neighborhood. This basic mechanism of efficiency enforcement might lead to maximum effort and a complete network in the group. Weber (2006) shows that starting out from pairs and adding subjects one by one fosters high effort levels up to a certain size of the group. We expect that a similar but endogenous process takes place in our neighborhood treatment. Small subgroups coordinating on high effort levels will form and only admit others to their neighborhood if they elicit high effort. Compared to Weber (2006) this process should even be more robust because deviators can be excluded again and do not have to lead to an effort deterioration. We therefore put forward

Hypothesis 3.2 *In the neighborhood treatment (minimum effort game with endogenously chosen neighborhoods), subjects behavior will converge to the complete network with every subject playing the highest effort.*

[11]With better response we denote a strategy that gives ceteris paribus a higher payoff than the currently played strategy. In contrast to the best response is has not to be optimal.

4 Experimental Procedure

All experimental sessions were conducted at the Behavioral and Experimental Economics Laboratory Maastricht (BEELab) with students from Maastricht University. The vast majority of subjects were from economics and business administration. The others were from law, cultural sciences and from the University College Maastricht. They were recruited through email announcements and announcements on students' intranet. Subjects were not allowed to participate more than once in the experiment. The experiment was computer based and programmed in Z-tree (Fischbacher, 2007). To ensure anonymity, 16 subjects were recruited for every session and were randomly divided into two independent groups of 8 students. Subjects got written instructions which they could study at their own pace (see appendix). Additionally they could ask questions privately.[12] Before the game started subjects had to answer some control questions correctly (see appendix). In total, we conducted 9 sessions: 5 with endogenous neighborhood choice and 4 baseline sessions. Subjects collected their payoffs in points which were paid confidentially in cash immediately after the experiment. The exchange rate was 2 points = € 0.01. On average participants earned € 11.82 in the baseline treatment and € 17.17 in the neighborhood treatment. A typical session in the neighborhood treatment lasted about 80-90 minutes whereas a typical session in the baseline treatment lasted around 50-60 minutes.

To check the robustness of our results we ran additional sessions with large groups of 24 subjects, three sessions with the baseline treatment (baseline XL) and three sessions with the neighborhood treatment (neighborhood XL). The sessions took roughly 30 minutes longer than their small groups counterparts. Baseline XL participants earned € 14.35 on average while neighborhood XL participants got € 20.32 on average.

[12]We did only answer questions about the instructions. No answer was given if it had influenced the subjects expectation or strategy choice.

5 Results

5.1 Effort levels

In the experiment we collected data from 8 groups for the baseline treatment and 10 groups for the neighborhood treatment. Each group forms an independent observation and if not stated differently statistical tests are based on aggregated measures of these independent groups.

We begin the analysis with the effort choices over time. Figures 1(a) and 1(b) show the cumulative distribution of effort levels for the baseline treatment and the neighborhood treatment, respectively. In the first round, we observe only little differences. Regarding

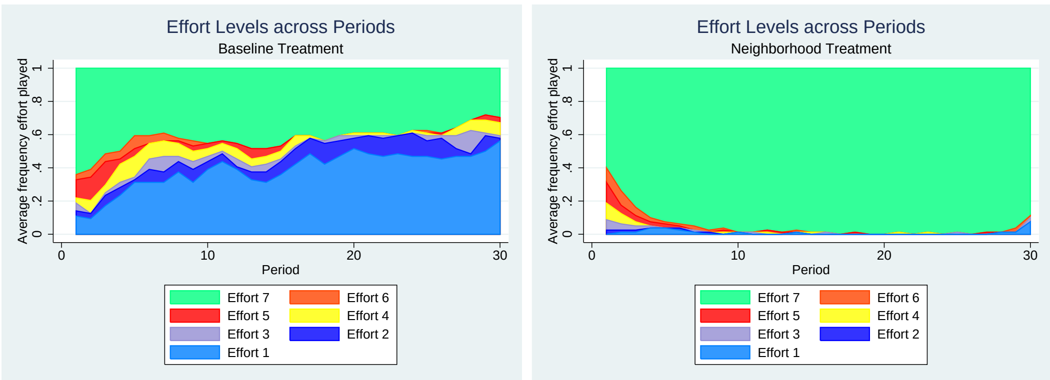

(a) Baseline Treatment (b) Neighborhood Treatment

Figure 1: Cumulative Distribution of Efforts across Rounds

the effort level as cardinal and applying a MWU-test on individual first round effort levels across treatments yields no significance (MWU, n=144, p=0.8919, two-sided).[13] However, behavior follows completely different trends in the consecutive rounds. In the baseline treatment the lowest effort becomes more frequent and ends being most frequentin the last round, more often played than all other effort levels together. In the neighborhood treatment, in contrast, the highest effort level becomes very quickly dominant while the frequencies of the other levels dwarf to a negligible size. To test this difference in trends we calculate each group's Spearman rank correlation coefficient between round and effort

[13]The distributions of effort levels are not identical though (Fisher exact test, n=144, p=0.027). There is more mass on the extreme effort levels in the baseline treatment.

level. The coefficients in the neighborhood treatment are (marginally) significantly more positive than those of the baseline treatment (MWU, n=18, p=0.0756, two-sided).[14]

A similarly clear picture yields the look at the average effort levels and the average minimum effort levels reported in Figure 2. Both curves for the neighborhood treatment

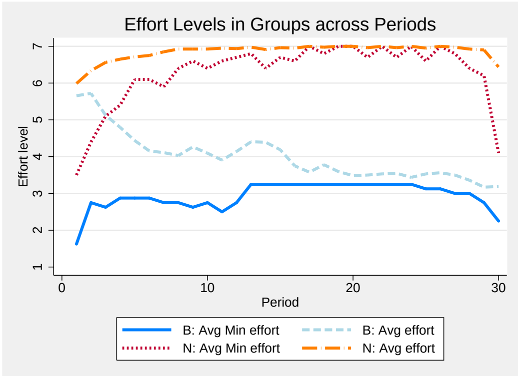

Figure 2: Average and average minimum efforts across rounds. Average minimum refers to the minimum effort levels played in the group averaged across groups.

are clearly above their counterparts for the baseline treatment. To test this we average effort levels and minimum effort levels across rounds for each group. The differences between treatments are significant (MWU on groups' average effort levels over all rounds, n=18, p=0.001, two-sided; MWU on groups' minimum effort levels averaged over all rounds, n=18, p=0.0033, two-sided).

In the baseline treatment, three out of eight groups manage to coordinate on the highest effort during rounds 13 to 24. The average effort of the other groups is approaching 1. Most of this convergence is caused by subjects who adapt their behavior to the minimum effort played in the groups. Moreover some subjects even overshoot, i.e. they play a lower level than the minimum level of the previous round. This behavior causes a drop in the minimum effort level and shows up in the data within the first 11 and the last 6 rounds. Towards the end we observe an end game effect in two of the three efficient groups (see Section 5.5).

The baseline treatment results are in line with previous findings although they are rather on the Pareto-efficient end of the spectrum from the literature (compare for example

[14]Stronger results can be achieved with a Kolmogorov-Smirnov test (KS, n=18, p=0.012).

Devetag and Ortmann 2007). Possible reasons could be the relatively small group size, the high number of rounds, complete feedback, cultural differences, and lower social distance due to the graphical representation of the network on the decision screen. All factors increase the chance of Pareto-efficient coordination and except for the last factor some (weak) supporting evidence exists. See the introduction for details.

In the neighborhood treatment all but one group converge within the first ten rounds to a complete network with each player playing the maximum effort level. One group remains unstable and does not converge to any common effort level. Nevertheless, the average level in this group is above 6 from round 10 onwards. Also here we observe an end game effect in six of the ten groups which drags down the overall average (see Section 5.5).

5.2 Neighborhood formation

Next to the development of the effort levels we are interested in the neighborhood formation process. Depending on the players' interaction choices I_i and I_j, four different situations may result. Besides mutual interaction ($j \in I_i$ and $i \in I_j$) and mutual exclusion ($j \notin I_i$ and $i \notin I_j$) player i might be the sole proposer of an interaction ($j \in I_i$ and $i \notin I_j$) or the one to reject j's interaction proposal ($j \notin I_i$ and $i \in I_j$). Figure 3 shows the frequencies of these situations over time.[15] We see that the groups are almost completely interlinked from the

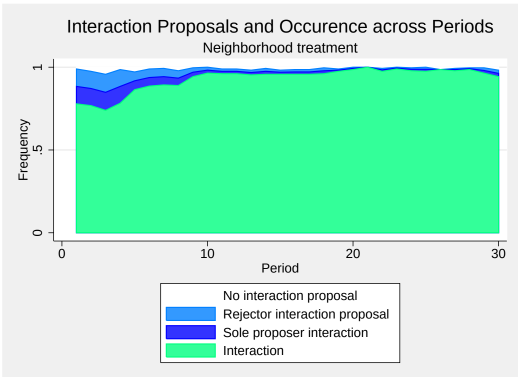

Figure 3: Interactions Proposals across Rounds

[15]Note that the frequencies to be a sole interaction proposer and a rejector of a proposal have to be equal by construction.

beginning. Moreover ther is a clear increase of the interaction frequency over time. This is confirmed by a Wilcoxon signed rank test applied to groups' Spearman rank correlation coefficients between frequency of mutual interaction and round number (Wilcoxon, n=10, p=0.043, two-sided). An interesting fact to mention is that the frequency of interactions that do not take place after round 10 is caused by one group.[16] The remaining groups approach the complete network within this time.

5.3 Exclusion

The frequencies of highest effort and interaction density seem to be in strong accordance. This suggests a causal relationship between both and calls for a more thorough investigation. In the introduction we already discussed that exclusion of others from the neighborhood extends a player's strategic possibilities to react to low effort providers. Exclusion could either be used to myopically best respond or to discipline others. In the following we investigate this to more detail. We look at the development of dyadic relationships among all subject pairs i and j for a two period time interval ($t-1$ and t). In a first step we analyze the consequences which arise for low effort providers; do they face a higher risk of being excluded by others?

Therefore we categorize the dyadic relations by the behavior of subject i in round $t-1$ into three distinct classes. The first class includes all cases where i provided at least as much effort as j ($e_i \geq e_j$). The second class includes cases where i provided less effort than j but still more than the minimum in j's neighborhood ($e_i < e_j$ but $e_i > \min_{k \in N_j}\{e_k\}$). Eventually the third class includes all cases where i's effort was not only lower than j's but also the minimum in j's neighborhood ($e_i < e_j$ and $e_i = \min_{k \in N_j}\{e_k\}$). We finally check whether in round t subject j excluded i from her neighborhood.

The data in Table 1 (rounds $t-1$ and t) shows that playing at least the effort of the other subject goes along with a very small risk of being excluded (0.6%). In contrast, the risk is quite high when playing a lower effort level than the other subject (25.0%) and it further increases if this effort level is the minimum of the other subject's neighborhood (38.8%). To

[16]This is the same group which also showed unstable behavior with respect to the effort choices.

test these differences we calculate the exclusion rates for each independent group separately and apply a Wilcoxon signed rank test. All three pairwise comparisons are significant (Wilcoxon signed rank test, two-sided: 0.6%<25.0%, n=8, p=0.039; 25.0%<38.8%, n=8, p=0.023; 0.6%<38.8%, n=10, p=0.002).

Table 1: Reactions to being excluded

Round	Action								
$t-1$	i's effort $e_i \geq e_j$			i's effort $e_i < e_j$ but $e_i > \min_{k \in N_j}\{e_k\}$			i's effort $e_i < e_j$ and $e_i = \min_{k \in N_j}\{e_k\}$		
t	j excluded i excl. rate: 0.6% cases: 84/14738			j excluded i excl. rate: 25.0% cases: 26/104			j excluded i excl. rate: 38.8% cases: 100/258		
$t+1$	i's reaction			i's reaction			i's reaction		
		$j \in I_i$	$j \notin I_i$		$j \in I_i$	$j \notin I_i$		$j \in I_i$	$j \notin I_i$
	$e_i \uparrow$	11.8% (9)	2.6% (2)	$e_i \uparrow$	69.2% (18)	11.5% (3)	$e_i \uparrow$	61.7% (58)	9.6% (9)
	$e_i =$	68.4% (52)	14.4% (11)	$e_i =$	3.8% (1)	3.8% (1)	$e_i =$	19.1% (18)	0.0% (0)
	$e_i \downarrow$	1.3% (1)	1.3% (1)	$e_i \downarrow$	11.5% (3)	0.0% (0)	$e_i \downarrow$	9.6% (9)	0.0% (0)

There are two possible motivations for exclusion. Players may act in their self-interest and play a best or at least a better response[17] to the behavior of the others in the previous round. However, players may also perceive low efforts as an intentional or stupid negative act and want to educate the others by excluding them. These motivations are not mutually exclusive, but the self-interest explanation only holds if the exclusion of the other player is not costly.

[17] With better response we denote a weaker concept than the best response. Rather than being payoff maximizing a better response only ensures a higher (or at least equal) payoff compared to the previous round ceteris paribus. A best response is always also a better response.

It turns out that 45% of all exclusions can be explained by a best response and 24% by a better response. In contrast 31% of all exclusions can only be assigned to reciprocal behavior. This means that in a majority of cases (24%+31%) players incur short term costs or forgo short term gains in order to exclude low effort providers.

The question remains whether excluded players change their behavior or not. Only by keeping their interaction wish and increasing their effort level they give the excluding player an incentive to re-establish the link. To investigate this issue in detail we extend the analysis of the dyadic relationships by another period $t+1$. We investigate the behavior of those subjects i who have been excluded in round t. Subjects may react in two dimensions. Firstly, they may keep the interaction wish to j ($j \in I_i$) or break the link on their part ($j \notin I_i$). Secondly, they may stay (e_i=), go up ($e_i \uparrow$) or down ($e_i \downarrow$) with their effort level.

Table 1 (round $t+1$) reports the results.[18] Roughly 86% of the cases excluded subjects keep their interaction wish in round $t+1$ (81.6%, 84.6%, and 90.6% for the respective classes). With respect to the effort level subjects make their decisions dependent on how much effort they provided in round $t-1$. If they were at least equal to the excluding subject, they most often stick to their effort level in round $t+1$ (82.9%). If they were lower than the excluding subject, they most often increase their effort level in round $t+1$ (80.8% and 71.3% for the respective classes).

5.4 Welfare

Failure of coordination can create huge potential welfare losses. We have seen that the possibility of exclusion promotes coordination to higher effort levels. However, exclusion comes with some costs and we do not know yet whether or not the gains from higher effort levels outweigh these costs. We therefore calculate a group's welfare level by the summation of payoffs of all group members. Additionally we calculate two benchmark levels: firstly the optimal level (every group member plays maximum effort and interacts with all other group members) and secondly a fictitious welfare level which had been obtained in the

[18] Note that the sum of cases in round $t+1$ might be lower than in round t because for $t=30$ no further round exists.

neighborhood treatment if subjects stuck to their effort level but were fully interlinked with all group members. 4 depicts the average welfare level over time for each treatment as well as the corresponding benchmarks. Both average welfare levels increase over time.

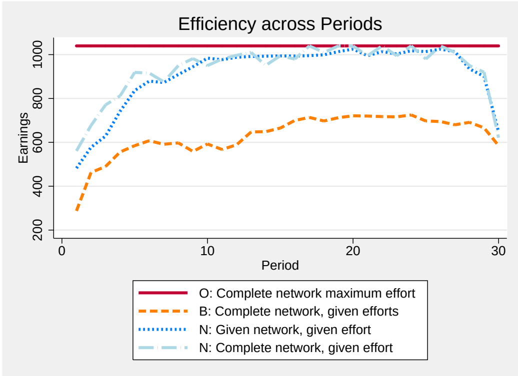

Figure 4: Welfare Levels across Rounds

However the average welfare level of the neighborhood treatment significantly outreaches that of the baseline treatment (MWU on groups average welfare, n=18, p=0.0129, two-sided).

The difference between the fictitious benchmark and the welfare level of the neighborhood treatment provides us with a measure of the net costs of exclusion. As can be seen it is small in the beginning and turns negligible from round 10 onwards. In some particular rounds (e.g. round 14) exclusion is even efficiency enhancing.

5.5 End game effect

As already briefly discussed in Section 5.1 we see an end game effect in many groups. This end game effect happens in two out of three efficient groups in the baseline treatment. In one group it happened in the last round only in the other it started in round 25 already.

In the neighborhood treatment an end game effect happened in six out of ten groups. In five out of the six groups, the effect is caused by *one* player deviating from the highest effort, in the other group this was caused by four players. In four of the six groups, the end game effect occurred in the last round only, in two groups it started earlier (in rounds 28 and 29 respectively).

Strictly speaking lowering the effort in the last rounds could be part of an overall equilibrium strategy. From a myopic perspective, however, it seems rather implausible. Why should one leave a “functioning” and efficient equilibrium? In the following we discuss three possible explanations and report the corresponding evidence in the data.

A first possible reason is revenge. Subjects who played a high effort in the beginning rounds and either by exclusion or by good example dragged others into the Pareto-efficient equilibrium earn a lower total payoff compared to their peers who started off with low effort. These efficiency minded subjects might use the last rounds to equalize total payoffs by playing a low effort.

Another reason only applies for the neighborhood treatment. Some subjects might lack the insight that high effort levels are beneficial for all. By exclusion they feel forced into a high effort equilibrium. Towards the end of the experiment, facing a diminishing threat of exclusion they fell back into their preferred behavior.

Neither pattern of behavior was frequent in the data. Therefore we have to come back to the rather vague explanation that for various reasons trust vanishes and some subjects start to doubt that others will continue to play the highest effort. As already mentioned in the introductions this doubts may be of a very subtle higher order type.

6 Large groups

Our results, in particular those from the neighborhood treatment, are strikingly clear, but an important question which we can not answer from our data is that of the robustness of the efficiency enforcing mechanism. While there are several dimensions in which robustness can be checked, we want to focus on group size, because of two reasons: Firstly, group size has a strong impact to strategic uncertainty. Adding another player decreases the chance of efficient coordination in a multiplicative way. Secondly, also the existing literature acknowledges group size as an important factor (Van Huyck, Battalio and Beil, 1990 and Van Huyck Battalio and Rankin, 2007).

We conducted six additional sessions, three with the neighborhood treatment and three with the baseline treatment. Each session consisted of one large group of 24 subjects.[19] In the following we will call them XL baseline treatment and XL neighborhood treatment. Besides seeing whether our results are robust with respect to size, we were also interested in whether the group gets fragmented into separate neighborhoods, for example one coordinating on high and the other on low effort. In the following, we will briefly report the results.

By and large we get the same treatment differences than with the small groups. Figure 5 shows the distribution of effort levels, which develops in a similar manner as in the smaller groups. In contrast to the small group treatments find already a clear difference in the effort levels in the first round. Neighborhood XL subjects show a significantly higher effort in round 1 than baseline XL subjects (MWU, n=144, p= 0.0238, two-sided).

Like in the small group treatments we find opposite trends with respect to the efforts development. In the baseline XL treatment, the absolute majority of subjects converge towards playing the lowest possible effort level, whereas in the neighborhood XL treatment, subjects converge towards the highest effort (MWU, n=6, p=0.05, one-sided). Figure 6

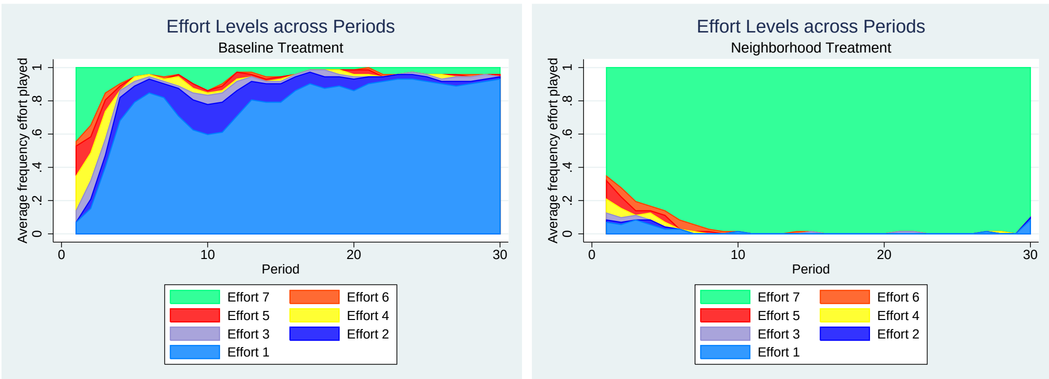

(a) Baseline Treatment (b) Neighborhood Treatment

Figure 5: Distribution of Efforts in Groups across Rounds

shows the average effort level and the average minimum effort level. Despite the small number of independent observations, the differences are significant (MWU on groups average effort level, n=6, p=0.05, one-sided; MWU on average minimum effort level, n=6, p=

[19]To the best of our knowledge this has been the largest groups ever being tested with the minimum effort game.

0.05, one-sided).[20]

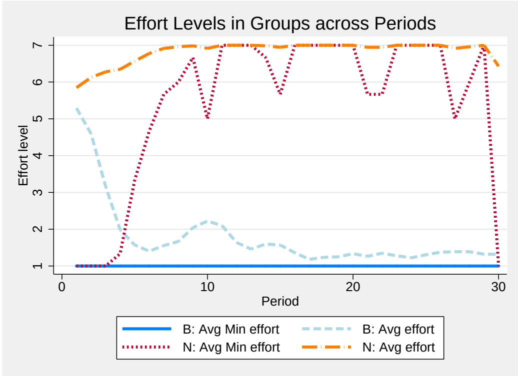

Figure 6: Average and Minimum Efforts across Rounds

In each of the neighborhood XL sessions we observe an end game effect in the last round. It was triggered by one, two and four subjects, respectively. A notable observation is that all downward movements of the average as well as minimum effort, except for the last round, are caused by *one* single group where *one* single individual played low effort levels recurrently. The other two groups reach the highest effort after ten rounds.

A similar picture yields the analysis of the link structure (see Figure 7). In the very

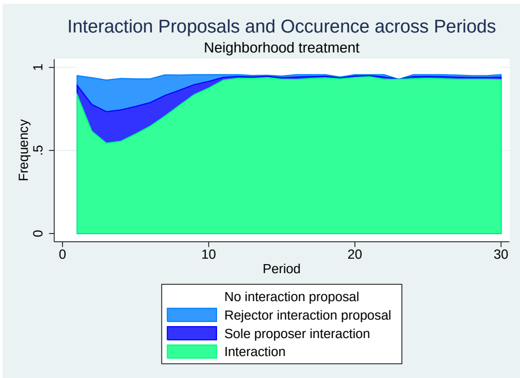

Figure 7: Interactions across Rounds

beginning the group forms an almost complete network, but quickly the interaction frequency drops down to a bit more than 0.5, before recovering to almost full interaction by

[20]The tests are one-sided because given our results from the small groups we had a clear hypothesis for the direction of the difference.

round 11. From then on all cases of non-interaction involve the one single subject who repeatedly deviated from the highest possible effort.

As with the small groups we find that in the neighborhood XL treatment exclusion is triggered by playing the minimum effort within the neighborhood. In each situation where a subject played an effort below the maximum, this person was excluded by at least one other person. In the most extreme case, a person who played a lower effort level got excluded completely.

Being excluded made subjects realize that they should increase their effort level to the maximum. After this increase, they were included again by all other subjects. This does not hold however for the particular subject who repeatedly decreased the effort level. Other subjects became more and more reluctant to interact with this person. On average this person only received five interaction proposals per round, whereas the remaining subjects receive 21 proposals per round on average.

7 Conclusions

With the real live examples of the minimum effort game we pointed out a major difference between situations where players are bound to their neighborhood and others where players may freely choose with whom to interact. We extended the minimum effort game by endogenous neighborhood choice and conducted an experiment with two treatments, a baseline treatment where the neighborhood was fixed and a neighborhood treatment where subjects could endogenously determine their neighborhood.

The results show that endogenous neighborhood choice solves the coordination problems and ensures efficient outcomes. The basic mechanism is quite simple. In standard minimum effort games subjects do not have the means to force others into the payoff dominant equilibrium. Consequently, as soon as there exists subjects unwilling to exert the highest effort or doubting that others will exert high efforts, efficient coordination will fail. The only way to ensure reasonable payoffs in such a situation is to also provide low effort. By this self-confirming process the outcome gets mired into the risk dominant equilibrium.

Endogenous neighborhood choice gives subjects the possibility to exclude others with unwanted behavior. It turns out that exclusion is used by high-effort players against low-effort players although it is sometimes costly. Moreover exclusion works, i.e. low-effort players quickly realize that they have to increase effort in order to be included again by other subjects. This behavior forced all subjects of the neighborhood treatments into the efficient equilibrium.[21]

With respect to our real world examples our data indicate that the worrying results of VHBB and successors do not apply for the type of situations where the minimum effort game is played in networks and individual players are able to choose their interaction partners. Regulatory means seem not be necessary for this class of games. However we do not want to stretch our results too far. Our neighborhood treatments give efficient coordination a good shot. Our subjects played the game for many rounds; they had full information about others choices (effort levels as well as neighborhood choices); and eventually the costs of exclusion were rather moderate. Hence, we consider our experiment to provide two benchmarks that show what is possible with endogenous networks. The region between the benchmarks is largely unexplored. It is not clear where the line of border between efficient and inefficient coordination runs. In our view a possible and interesting avenue to proceed is to check the robustness of the results under varying conditions. Our results are remarkably robust with respect to group size. Other dimensions like feedback information have to be tested still. For example coordination on the payoff dominant equilibrium has a harder stand if subjects only receive information about those players' actions whom they interacted with. A fragmentation of the group into different sub-groups becomes more likely then, because linking up to "unknown" players involves risk. A further avenue is to vary the characteristics of links, e.g. add linking costs or allow for one-sided linking. While linking costs might lead to an increase of efficiency in the minimum effort game due to a forward induction argument, one-sided linking completely offsets the exclusion power and will most likely lead to outcomes as in the standard minimum effort game experiments. Further interesting variations may include non-linear exclusion costs or a limited amount of links. Once these and other robustness checks have been done, more

[21] Only one subject in the neighborhood XL treatment repeatedly fell back on playing low effort levels.

detailed statements can be made about real world situations.

References

1. Berninghaus, Siegfried. K., and Karl-Martin Erhart. 2001. "Coordination and Information: Recent Experimental Evidence." *Economics Letters*, 73(3):345.351.

2. Brandts, Jordi, and David J. Cooper. 2006. "A Change Would Do You Good: An Experimental Study on How to Overcome Coordination Failure in Organizations". *American Economic Review*, 96(3): 669-693.

3. Bryant, John.1993. "A Simple Rational- Expectations Keynes Type Model." *Quarterly Journal of Economics*, 98(3):525-528.

4. Devetag, Giovanna, and Andreas Ortmann. 2007. "When and Why? A Critical Survey on Coordination Failure in the Laboratory." *Experimental Economics*, 10(2): 171-178.

5. Fischbacher, Urs.2007. "Z-Tree: Zurich Toolbox for Ready-made Economic Experiments."*Experimental Economics*, 10(3): 331-344.

6. Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil.1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review*, 80(1):23448.

7. Van Huyck, John B., Raymond C. Battalio, and Frederick W.Rankin. 2007. "Evidence on Learning in Coordination Games." *Experimental Economics*, 10(3): 205220.

8. Sandler, Todd. 1998. "Global and regional public goods: a prognosis for collective action." *Fiscal Studies*, 19(3):221-247.

9. Weber, Roberto A. 2006. "Managing Growth to Achieve Efficient Coordination in Large Groups." *American Economic Review*, 96(1): 114126.

10. Young, H. Peyton. 1993. "The Evolution of Conventions." *Econometrica*, 61: 57-84.

11. Young, H. Peyton. 1998. *Individual Strategy and Social Structure*, Princeton University Press.

A Theoretical benchmarks

We show that for our parameter setting the fully interlinked network with every player playing the lowest effort is the unique stochastically stable equilibrium for the minimum effort game with neighborhood choice. Thereafter we show that this also holds for the restricted case of the minimum effort game without neighborhood choice.

A.1 The game

The analysis rests on a one-shot game where each player simultaneously chooses an effort level and the set of other players whom he or she wants to interact with. The basic elements of the game are:

- $N = \{1, 2, 3, ..., n\}$ is a finite set of players.
- $E = \{1, 2, ..., m\}$ is a set of effort levels.
- $s_i = (e_i, I_i)$ is the strategy of player i with $e_i \in E$ is the player's chosen effort level and $I_i \subseteq N \setminus \{i\}$ is the set of players with whom i wants to interact with.
- $s = (s_1, ..., s_n) = ((e_j, I_j))_{j \in N}$ is a strategy profile. Later we will interpret it as a state in a Markovian chain. With $s^{\tilde{e}}$ we denote the strategy profile where every player wants to interact with every other player and all play the same effort $\tilde{e}$, i.e. $s^{\tilde{e}} = ((\tilde{e}, N \setminus \{j\}))_{j \in N}$.
- $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n) = ((e_j, I_j))_{j \in N \setminus \{i\}}$ is the vector of strategies of all players except i.
- Given a strategy profile s, players i and j are linked if both want to interact with each other, i.e. $j \in I_i$ and $i \in I_j$. If an interaction wish of i with j is not reciprocated, i.e. $j \in I_i$ but $i \notin I_j$ then j is called a dangling link of i.

- Given a strategy profile s, the neighborhood of a player i is the set of all players i is linked to, i.e. $N_i(s) = \{j | j \in I_i \wedge i \in I_j\}$. With $|N_i(s)|$ we denote the cardinality of $N_i(s)$, in other words the size of the neighborhood.

- For a given s players i's payoff is
$$\pi_i(s) = \frac{|N_i(s)|}{n-1} \left[a \left(\min_{j \in N_i(s) \cup \{i\}} \{e_j\} \right) - be_i + c \right]$$
with $a > b > 0$ and $c > 0$.

- Our specific experimental settings are $m = 7$, $a = 20$, $b = 10$ and $c = 60$ for $n = 8$ and $n = 24$, respectively.

A.2 The process

We follow a variant of the approach of Young (1993,1998). Assume discrete and successive time rounds ($t = 1, 2, ..$). Each round t every player may play a myopic best response to the other players' strategies of round $t-1$ with some positive probability σ. This construction constitutes a Markov chain with the state space being the set of all strategy profiles and the transition probabilities depending on σ and the payoffs π. We will show that only states with the complete network and everybody playing the same effort are absorbing states and that no other recurrent class of states exist. Thereafter we take the parametrization of our experimental design and calculate the number of errors which are necessary to get from one absorbing state to the other. We use the method introduced by Young (1993,1998) to show that the state with the complete network and the minimum effort played by all players is the only stochastically stable one.

A.3 Some useful observations

To start, we state some almost trivial propositions. We will use them later. The corresponding proofs are straightforward.

Proposition A.1 *Given an outcome $s = ((e_i, I_i), s_{-i})$. Then the players i's payoff is independent of dangling links, i.e. $\pi_i((e_i, I_i), s_{-i}) = \pi_i((e_i, I'_i), s_{-i})$ with $I'_i = \{j | j \in I_i \wedge i \in I_j\}$.*

Proof: Dangling links are not decisive for the neighborhood of a player, i.e. $N_i((e_i, I_i), s_{-i}) = N_i((e_i, I_i'), s_{-i})$. Hence neither the removal nor the addition of dangling links change payoffs. As a corollary we get that any best response remains a best response if dangling links are removed or added. ■

Proposition A.2 *For $N_i(s) \neq \emptyset$ the condition $e_i = \min_{j \in N_i(s)}\{e_j\}$ is necessary for $s_i = (e_i, I_i)$ being a best response to $s_{-i} = ((e_j, I_j))_{j \in N \setminus \{i\}}$.*

Proof: Assume $N_i(s) \neq \emptyset$ and that the condition does not hold, then i can improve his payoff by adjusting e_i to $\min_{j \in N_i(s)} e_j$ while holding I_i constant. ■

A.4 Absorbing states and recurrent classes

Proposition A.3 *A state $s^{\tilde{e}}$ with $s_i^{\tilde{e}} = (\tilde{e}, N \setminus \{i\})$ for each player $i \in N$ and some $\tilde{e} \in E$ is an absorbing state.*

Proof: We have to show that $s_i = (\tilde{e}, N \setminus \{i\})$ is the only best response to $s_{-i} = (s_j)_{j \in N \setminus \{i\}}$ with $s_j = (\tilde{e}, N \setminus \{j\})$.

Removing players from I_i strictly decreases $|N_i(s)|$ and therefore strictly decreases π_i. The marginal payoff change with respect to e_i is $-a + b$ (< 0) if i decreases e_i and $-b$ (< 0) if i increases e_i. Any combination of changes in effort and interaction also leads to negative payoff consequences. It follows that s_i is the only best response to s_{-i}. This holds for each player. ■

Each of the absorbing states forms a recurrent class. In the following we will show that no other recurrent class exists.

Proposition A.4 *No other recurrent class than those defined by the absorbing states in proposition A.3 exists.*

Proof: We introduce a hierarchy of sets of states into the state space. Let

S be the set of all possible states.

S' be the subset of S for which in any state s' players who are linked play the same effort level and no dangling link exists. This means that the network induced by s' consists of one or more components with all players within a component are playing the same effort level.

S'' be the subset of S' where in any state s'' the link relation is transitive, i.e. if i is linked to j and j to k then also i is linked to k. This means that the network induced by s'' consists of one or more fully interlinked components without dangling links and with all players within a component are playing the same effort level.

S^a be the set of the absorbing states.

It is obvious that $S \supset S' \supset S'' \supset S^a$. The last inclusion follows from Proposition A.3, an absorbing state is characterized by the complete network with every player playing the same effort.

The proof comprises three steps. We will show that for each state in the supersets S, S', and S'' a finite path of best responses into the subsets S', S'', and S^a exists which the players follow with positive probability.

Step $s \in S \to s' \in S'$

Assume state $s \in S$. The following algorithm generates a finite sequence of best responses that occurs with a positive probability and leads to a state $s' \in S'$. Assume that players update their strategies one by one.

1. Sort the players into a two lists A and B.[22]

   List A contains players that do not have any dangling link in their strategy and whose effort level is at most the minimum effort level of their neighborhood.

   List B contains the other players.

[22] We chose to have lists rather than sets because list B must enable an order for processing the elements (first in first out principle).

2. Take the first player i from list B and calculate all best responses to s_{-i}. There exists a best response s_i without dangling links and $e_i = \min_{j \in N_i(s)}\{e_j\}$ (see propositions A.1 and A.2). Update player i's strategy to such a best response and put him at the end of list A. The update of s_i may cause some players in A to violate the A-conditions. Put them at the end of list B.

3. If there are players left in B then continue with step 2.

The generated sequence is finite because for each application of step 2 there is exactly one player moving from B to A. A move of a player i from A to B can only happen if a neighbor j moves from B to A and breaks up a link or decreases his effort level such that $e_j < e_i$. In the first case the link between i and j will be irreversibly deleted.[23] In the second case i updates by keeping the link and lowering e_i or by breaking the link. Since E is finite, lowering e_i can only happen a finite number of times before the link has to be broken. Because the breaking of a link is irreversible and there is only a finite number of links to be broken, there can only be a finite number of moves from A to B.

Step $s' \in S' \to s'' \in S''$

Assume state $s' \in S'$, i.e. N is subdivided into several disjoint components $C_1, C_2, \ldots, C_k$. Each component is (not necessarily fully) linked and is free of dangling links. Each player within a component is playing the same effort level. Consider a component C_l in which effort level $\tilde{e}$ is played. Consider further a player $i \in C_l$ who plays strategy $s_i = (\tilde{e}, I_i)$ with $I_i \subseteq C_l$. Then strategy $s_i = (\tilde{e}, C_l \setminus \{i\})$ is a best response to s_{-i} because adding dangling links do not change payoffs (see Proposition A.1). There is a positive probability that each player is playing this best response and that this happens in each component. This means that we reach a state $s'' \in S''$ where components are fully connected.

[23]Note that only players moving from B to A may change their strategy. Player j who is now in A can only reestablish the link to i, after he has been moved back to B. Since he will be put at the end of the list, i will have deleted the dangling link before j's next turn.

Step $s'' \in S'' \rightarrow s^a \in S^a$

Assume state $s'' \in S''$. Assume two disjoint components $\underline{C}, \overline{C} \subset N$ and that each components is fully connected. Each player in $\underline{C}$ and $\overline{C}$ plays effort $\underline{e}$ and $\overline{e}$, respectively. Without loss of generality assume that $\overline{e} \geq \underline{e}$.

By the condition

$$\frac{|\underline{C}| + |\overline{C}| - 1}{n-1}\left(a\underline{e} - b\underline{e} + c\right) > \frac{|\overline{C}| - 1}{n-1}\left(a\overline{e} - b\overline{e} + c\right) \tag{3}$$

we distinguish two cases:

If condition (3) holds, then a player $i \in \overline{C}$ prefers to play $\underline{e}$ and to link up to all players in both components rather than to stay with component $\overline{C}$. Because of Proposition A.1 it happens with positive probability that all players in $\underline{C}$ will create a dangling link to all players in $\overline{C}$. Once this happens it becomes a best response for each player from $\overline{C}$ to connect to all players from $\underline{C}$ and switch to effort level $\underline{e}$.

If the converse of condition (3) holds, then

$$\frac{|\underline{C}| + |\overline{C}| - 1}{n-1}\left(a\underline{e} - b\underline{e} + c\right) < \frac{|\overline{C}|}{n-1}\left(a\overline{e} - b\overline{e} + c\right) \tag{4}$$

follows. This means that player $i \in \underline{C}$ prefers to join $\overline{C}$, to switch the effort level to $\overline{e}$ and to break up all links to his neighbors from $\underline{C}$, rather than being connected to both components and to play an effort level of $\underline{e}$. Because of Proposition A.1 it will happen with positive probability that all players in $\overline{C}$ will offer to establish a link with all players in $\underline{C}$. Once this happens it becomes a best response for all players from $\underline{C}$ to connect to all players from $\overline{C}$, switch the effort level to $\overline{e}$ and break the links with their old neighborhood. The resulting component is not completely connected but connects completely with positive probability and within finitely many steps (for the proof see step $s' \in S' \rightarrow s'' \in S''$).

The previous paragraphs show that regardless of the result of condition 3 there is a positive probability that two components merge to one fully interlinked component $\underline{C} \cup \overline{C}$ where each player plays the same effort level (either $\underline{e}$ or $\overline{e}$). A repeated application of this part of the proof shows that with finitely many steps and positive probability we reach a single fully interlinked component with all players playing the same effort level; i.e. we reach a state $s^a \in S^a$. ■

A.5 Stochastically stable equilibria

So far we showed that the best reply dynamic almost surely converges to one of the absorbing states, making them the only candidates for stochastically stable equilibria. We continue by determining the number of mistakes that is necessary to reach from one absorbing state to a basin of attraction of another absorbing state. First, note that mistakes in the form of breaking links only, cannot lead to a different absorbing state unless players are excluded completely from their neighborhood. This requires at least $n-1$ mistakes. Fewer mistakes are necessary if players change their effort level as well. Note further that a number of uncoordinated deviations does never create a higher temptation for other players to change their strategies than the same number of coordinated deviations to another effort level. Therefore it is sufficient to look at the cases where with a subgroup's joint deviation to another effort level we can reach the basin of attraction of another absorbing state.

Consider the system to be in an absorbing state s^a with N fully interconnected and $e_i = \tilde{e}$ for all players i. Assume that $d_\downarrow$ players deviate to $\underline{e} < \tilde{e}$. The following condition must hold to make other players switching to $\underline{e}$ instead of breaking the links to the deviators:

$$a\underline{e} - b\underline{e} + c \geq \frac{n-1-d_\downarrow}{n-1}\left(a\tilde{e} - b\tilde{e} + c\right).$$

This resolves to

$$d_\downarrow \geq \frac{(n-1)(\tilde{e}-\underline{e})}{\tilde{e}+\frac{c}{a-b}}.$$

Consider $d_\uparrow$ players who deviate to $\overline{e} > \tilde{e}$. The following condition must hold to make players switching to $\overline{e}$ and cutting all the links to the players of the lower effort level, instead of not changing the strategy:

$$\frac{d_\uparrow}{n-1}\left(a\overline{e} - b\overline{e} + c\right) \geq a\tilde{e} - b\tilde{e} + c.$$

This resolves to

$$d_\uparrow \geq \frac{(n-1)\left(\overline{e}+\frac{c}{a-b}\right)}{\tilde{e}+\frac{c}{a-b}}.$$

In the remainder we focus on our parameter settings of the experiment. In short this means: $a = 20, b = 10, c = 60, E = \{1, 2, \ldots, 7\}$, and $N = \{1, 2, \ldots, 8\}$ or $N =$

$\{1, 2, \ldots, 24\}$, respectively.[24] The following tables show the resistances between the absorbing states, i.e. the number of deviations needed to move from an absorbing state (row entry) to a basin of attraction of another absorbing state (column entry).

For the case of 8 players we get:

$r(s^i, s^j)$	s^1	s^2	s^3	s^4	s^5	s^6	s^7
s^1	0	7	6	5	5	5	4
s^2	1	0	7	6	6	5	5
s^3	2	1	0	7	6	6	5
s^4	3	2	1	0	7	6	6
s^5	3	2	2	1	0	7	6
s^6	3	3	2	2	1	0	7
s^7	4	3	3	2	2	1	0

For the case of 24 players we get:

$r(s^i, s^j)$	s^1	s^2	s^3	s^4	s^5	s^6	s^7
s^1	0	21	18	17	15	14	13
s^2	3	0	21	19	17	16	15
s^3	6	3	0	21	19	18	16
s^4	7	5	3	0	21	20	18
s^5	9	7	5	3	0	22	20
s^6	10	8	6	4	2	0	22
s^7	11	9	8	6	4	2	0

Proposition A.5 *For both, the 8-player case and the 24-player case s^1 is the only stochastically stable equilibrium of the minimum effort game with endogenous neighborhood choice.*

[24] There are more general results achievable. For example all absorbing states may be stochastically stable for some appropriate parameter settings. These results are however not important for our central goal, the calculation of the theoretical benchmark for our experimental settings.

Proof: We have to show that s^1 is the state with the minimum stochastic potential. This means that among all rooted trees, those with the minimum total resistance (sum of all resistances among the edges) have the root s^1.

Consider a rooted tree T^m with root $s^m \neq s^1$ and total resistance $r(T^m)$. Then we can construct a new rooted tree T^{m-1} with root s^{m-1} by connecting s^m to s^{m-1} and removing the edge from s^{m-1} to state s^k on the path to s^m. The new tree has total resistance

$$r(T^{m-1}) = r(T^m) + r(s^m, s^{m-1}) - r(s^{m-1}, s^k) \leq r(T^m).$$

The last inequality holds because $r(s^m, s^{m-1}) \leq r(s^{m-1}, s^k)$ as can be verified from the tables.

We can iterate this process till we construct T^1 with root s^1 and total resistance

$$r(T^1) = r(T^2) + r(s^2, s^1) - r(s^1, s^k) < r(T^2).$$

The last inequality holds because $r(s^2, s^1) < r(s^1, s^k)$ as can be verified from the tables.

Hence, for any rooted tree T^m we find a chain of rooted trees such that $r(T^1) < r(T^2) \leq \cdots \leq r(T^m)$. Therefore the root of the tree with the minimum total resistance must be s^1.
■

A.6 The case without neighborhood choice

For the case without neighborhood choice I_i is restricted to $I_i = N \setminus \{i\}$. A strategy is therefore $s_i = e_i$. Furthermore $N_i(s) = N \setminus \{i\}$ and hence $|N_i(s)| = n - 1$. This reduces the payoff function to

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

Proposition A.6 *A state s with $s_i = \tilde{e}$ for each player $i \in N$ and some $\tilde{e} \in E$ is an absorbing state in the minimum effort game without neighborhood choice and no other recurrent class exists.*

Proof: This follows immediately from the restriction $I_i = N \setminus \{i\}$ and Proposition A.2.■

The resistance table for the case of 8 players is:

$r(s^i, s^j)$	s^1	s^2	s^3	s^4	s^5	s^6	s^7
s^1	0	7	7	7	7	7	7
s^2	1	0	7	7	7	7	7
s^3	1	1	0	7	7	7	7
s^4	1	1	1	0	7	7	7
s^5	1	1	1	1	0	7	7
s^6	1	1	1	1	1	0	7
s^7	1	1	1	1	1	1	0

For the case of 24 players we get:

$r(s^i, s^j)$	s^1	s^2	s^3	s^4	s^5	s^6	s^7
s^1	0	23	23	23	23	23	23
s^2	1	0	23	23	23	23	23
s^3	1	1	0	23	23	23	23
s^4	1	1	1	0	23	23	23
s^5	1	1	1	1	0	23	23
s^6	1	1	1	1	1	0	23
s^7	1	1	1	1	1	1	0

Proposition A.7 *For both, the 8-player case and the 24-player case s^1 is the only stochastically stable equilibrium of the minimum effort game without endogenous neighborhood choice.*

The proof is identical to the proof of Proposition A.5.

B Instructions

In the following we present an integrated version of the instructions for the neighborhood treatment. The instructions for the baseline treatment were the same except for paragraphs marked with [**N**] and [*]. Paragraphs with [**N**] were only given in the neighborhood treatment. Paragraphs with [*] were given in both treatments but appropriately reformulated for the baseline treatment. A complete set of instructions is available from the authors.

Instructions

Introduction

Welcome to this decision-making experiment. In this experiment you can earn money. How much you earn depends on your decisions and the decisions of other participants. During the experiment your earnings will be counted in points. At the end of the experiment you get your earned points paid out privately and confidentially in cash, according to the exchange rate:

2 points = 1 eurocent.

It is important that you have a good understanding of the rules in the experiment. Therefore, please read these instructions carefully. In order to check that the instructions are clear to you, you will be asked a few questions at the end of the instructions. The experiment will start only after everybody has correctly answered the questions. At the end of the experiment you will be asked to fill in a short questionnaire. Thereafter you will receive your earnings.

During the whole experiment, you are not allowed to communicate with other participants in any other way than specified in the instructions.

If you have a question, please raise your hand. We will then come to you and answer your question in private.

Explanation Experiment

During the experiment every participant is in a **group of eight**, that is in a group with seven other participants. The group you are in will not change during the experiment. You will not receive any information about the identity of the persons in your group, neither during the experiment, nor after the experiment. Other participants will also not receive any information about your identity.

Each person in your group is indicated by a letter. You will receive the name "**me**". The other seven persons in your group will be indicated by the letters **A, B, C, D, E, F** and **G**. The same letter always refers to the same person.

The experiment consists of 30 rounds. In each round you can earn points. Your total earnings in this experiment is the sum of your earnings in each of the 30 rounds.

[*] In each round, you - and each other person in your group - will have to make two decisions which will influence your earnings. You have to make a **decision** called **"With whom would you like to interact?"** and a **decision** called **"Which number do you choose?"** Your decisions and the decisions of the other participants in your group will influence your earnings (as well as the earnings of the other group members). These decisions are explained in detail below.

Note: During the whole experiment, during all 30 rounds the other participants in your group will stay the same persons

Decisions (in one single round)

[N] Decision: "With whom would you like to interact?"

[**N**] You have to decide with whom you would like to interact. You can **propose** an interaction to any other person in your group and you can make as many proposals as you want. (You can also decide not to make any proposal.) Your interaction proposals - together with the proposals of the other persons in your group - determine with whom you actually interact (in the respective round) as explained below:

- [**N**] You will interact with a person to whom you proposed to interact only if the other person also proposed to interact with you. That is, **mutual consent** is needed for an interaction to actually take place.
- [**N**] You will not interact with another person if either only you or only the other person proposed to interact.
- [**N**] You will not interact with another person if neither of you proposed to interact with each other.

[**N**] For convenience we will call those persons in your group with whom you **interact: your neighbors**. Your neighbors are therefore those persons to whom you proposed to interact and who at the same time also proposed to interact with you.

Decision: "Which number do you choose?"

In each round, each person in your group has to choose one number from 1 to 7; i.e. either 1, 2, 3, 4, 5, 6, or 7. Your earned points in each round depend on

1. your **own choice of number**
2. [*] the **smallest number** chosen by your neighbors ***and*** yourself

3. **[N]** the number of neighbors you have

[N] Note: You can not choose different numbers for different neighbors. You can, however, choose different numbers as well as different persons to interact with in the different rounds.

Here is the payoff table:

		Smallest number chosen by your group including yourself						
		7	6	5	4	3	2	1
	7	130	110	90	70	50	30	10
	6	-	120	100	80	60	40	20
Your	5	-	-	110	90	70	50	30
chosen	4	-	-	-	100	80	60	40
number	3	-	-	-	-	90	70	50
	2	-	-	-	-	-	80	60
	1	-	-	-	-	-	-	70

[*] Since one's choice can be a number from 1 to 7, the smallest number can range from 1 to 7. Your payoff is determined by the cell in the row of "your chosen number" and the column of the "smallest number chosen by your neighbors and yourself". An example is given below.

[*] In the table there are cells with "-". This indicates that such a combination of "your chosen number" and the "smallest number chosen by your neighbors and yourself" is not possible. For example, if "your chosen number" is 4, the smallest number chosen by your neighbors and yourself cannot be 7, 6, or 5.

[N] Your earned points in a round will be the payoff as given in the table multiplied by

$$\frac{\text{number of neighbors}}{7}.$$

[N] For each person in your group with whom you do not interact (i.e. all persons who are not your neighbors) you earn 0 points. For example, if you have no neighbors in a round, then you earn 0 points in that round.

Examples:

[*] Suppose you have four neighbors. You chose 3 and the smallest number chosen by your neighbors and yourself was 3, you earn 4/7 * 90 = 51 3/7 points.

[*] Suppose you have three neighbors. You chose 5 and the smallest number chosen by your neighbors and yourself was 3, you earn 3/7 * 70 = 30 points.

[*] Suppose you have four neighbors. You chose 5 and the smallest number chosen by your neighbors and yourself was 4, you earn 4/7 * 90 = 51 3/7 points.

[*] Suppose you have three neighbors. You chose 7 and the smallest number chosen by your neighbors and yourself was 4, you earn 3/7 * 70 = 30 points.

Information about Computer Screen (in one single round)

You now get information about the computer screen that you will see during the experiment. You received a print-out of the computer screen **(Example screen 1)** from us. Take this print-out in front of you. The screen consists of two windows: **History** and **Decision**.

- **History:** This window holds information about past round(s). At the beginning of a new round you will automatically receive information in this window about decisions made in the previous round, (In the example, this is round 2; see upper part of the screen). In the window there are 8 circles, named **me, A, B, C, D, E, F** and **G**. **Me** always refers to you. The letters refer to the other seven persons in your group.
  - [N] A **thick complete** line between two persons (letters or 'me') indicates that they both proposed to interact with each other, that is they were neighbors and, hence, did actually interact with each other. (See, e.g., the line between **me** and **D** on the example screen).

- [**N**] A **thin incomplete** line between two persons indicates that only one of them proposed to interact. That is, they were not neighbors and, hence, did not interact with each other. Such a line starts on the side of the person that proposed to interact, and stops just before the circle of the person that did not want to interact. (See, e.g., the line between **me** and **F** on the example screen: **me** proposed to interact with **F**, but **F** did not propose to interact with **me**.)
- [**N**] No line between two persons indicates that neither of them proposed to interact. That is, they were not neighbors and hence, did not interact with each other.
- Next to the letters you see numbers between 1 and 7. These numbers indicate the chosen numbers of the persons in your group. The number next to letter A shows the chosen number of A. The number next to letter B shows the chosen number of B and so forth. (For example in screen 1, persons **A** and **G** have chosen number 5, while the persons **me** and **E** have chosen number 7.)
- At the bottom of this window you find two buttons called **Previous Round** and **Next Round**. You can use these buttons to look at the decisions in all previous rounds. The button **Most Recent Round** brings you back to the last round played.
- Your **earnings** (in points) in the corresponding round can be found just above the graph next to **Round Earnings**.

• **Decision:** In this window you see which round you are in and here you have to make your decisions.

1. [**N**] **With whom would you like to interact?** Below this question you see the seven letters which refer to the seven other participants in your group. You can propose to interact with another person in your group by clicking the button "yes" (the first button), that is the first button to the right of that person's letter. If you do not want to interact with a person or if you want to remove a proposal to interact, you activate the button "no", that is the second button to the right of that person's letter. **Note:** At the beginning of each new

round the buttons (i.e. proposals) you have chosen in the previous round will be activated. In each new round you can change your choices, i.e., proposals (not) to interact in the way described above.

2. **Which number do you choose?** In the small window next to **My Number** you type in the number you want to choose.

[*] When you are satisfied with all your decisions (that is, with **both** the proposals (not) to interact and your chosen number), you have to confirm these decisions by clicking on the button **"Ok"**.

[*] **Note:** After each round you will receive information about all the decisions made (that is, all interaction proposals made and the number choices) by all persons in your group. All other persons in your group will also receive information about all your decisions. This is the end of the instructions. You will now have to answer a few questions to make sure that you understood the instructions properly. If you have any questions please raise your hand. After you have answered all questions please raise your hand. We will then come to you to check your answers. The experiment will begin only after everybody has correctly answered all questions.
If you are ready please remain seated quietly.

C Instructions Large Networks

The instructions for the robustness check with large groups were similar to the small group instructions, except that subjects were called "**me, N1, N2, ..., N23**" rather than "**me, A, B, ..., G**", the calculation examples have been modified, and the layout of the history window has been slightly modified.

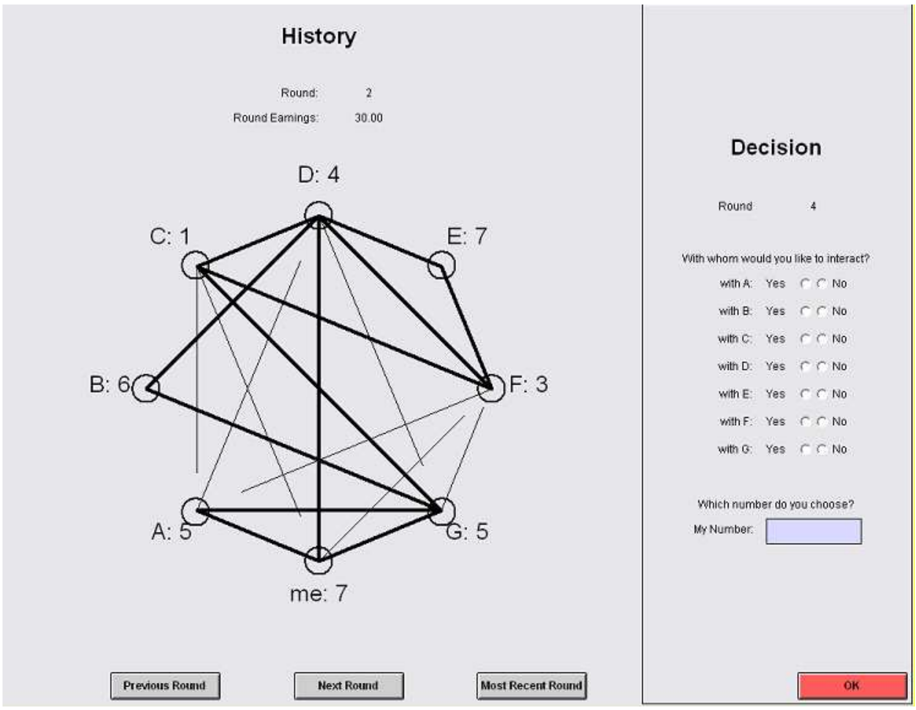

Figure 8: [**B**] Example screen 1

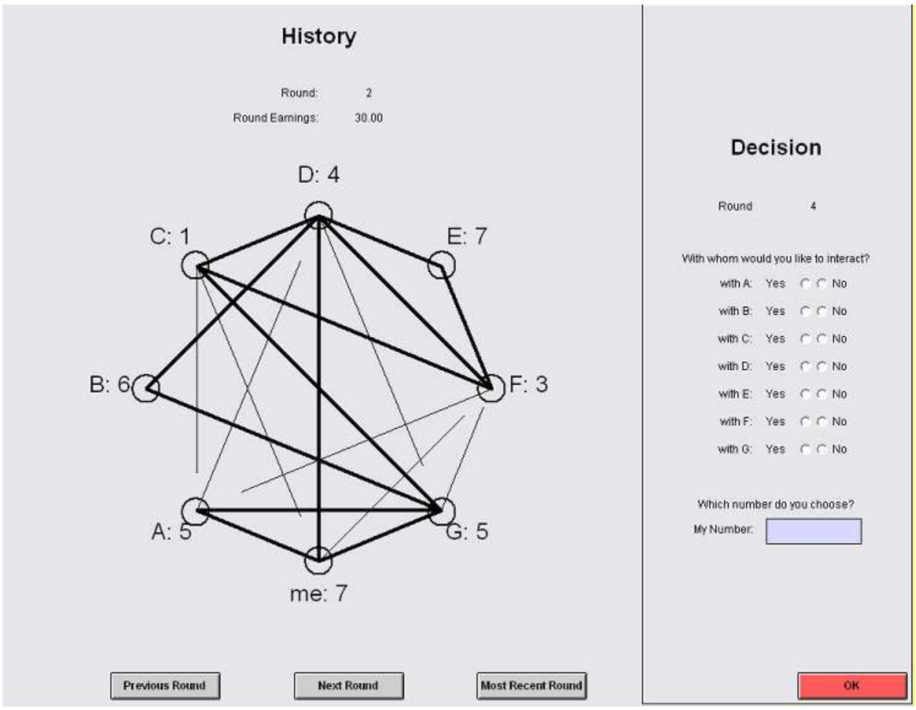

Figure 9: [**N**] Example screen 1

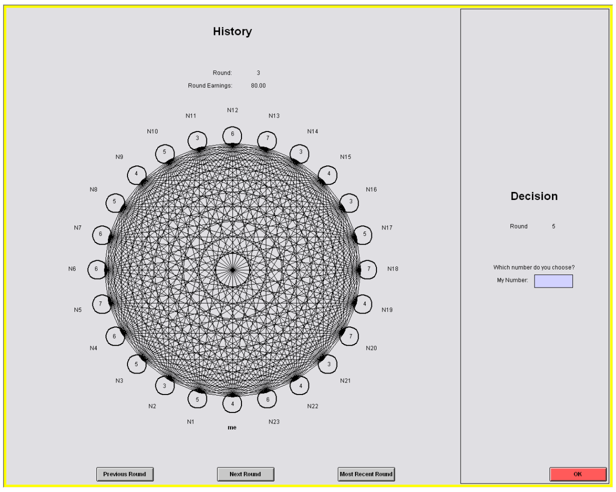

Figure 10: [**B**] Example screen 1

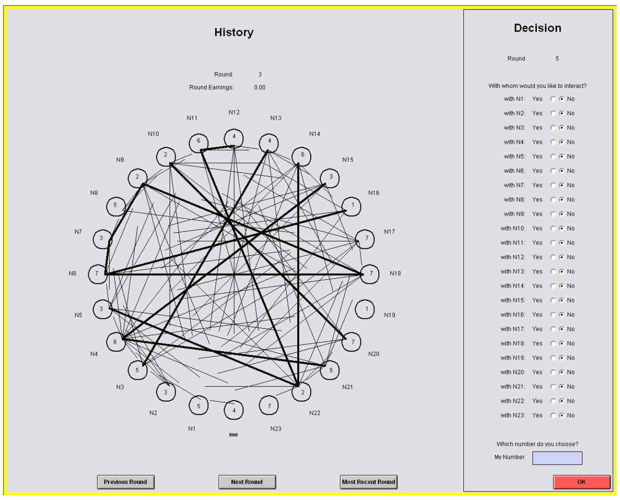

Figure 11: [**N**] Example screen 1