



Group play in games and the role of consent in network formation

Gary Charness^a, Matthew O. Jackson^{b,*}

^a*Department of Economics, University of California Santa Barbara, Santa Barbara, CA 93106-9210, USA*

^b*Department of Economics, Stanford University, Stanford, CA 94305-6072, USA*

Received 20 March 2006; final version received 7 September 2006

Available online 11 December 2006

Abstract

We study games played between groups of players, where a given group decides which strategy it will play through a vote by its members. When groups consist of two voting players, our games can also be interpreted as network-formation games. In experiments on Stag Hunt games, we find a stark contrast between how groups and individuals play, with payoffs playing a primary role in equilibrium selection when individuals play, but the structure of the voting rule playing the primary role when groups play. We develop a new solution concept, *robust-belief equilibrium*, which explains the data that we observe. We provide results showing that this solution concept has application beyond the particular games in our experiments.

© 2006 Elsevier Inc. All rights reserved.

JEL classification: D85; A14; C72; C91; D71; D72

Keywords: Groups; Networks; Game theory; Equilibrium selection; Equilibrium refinement; Majority voting; Group play; Robust-belief equilibrium

1. Introduction

The decisions made by many social and economic organizations or groups involve some sort of consent. While classical models in economics treat firms as if they were individual decision makers, a firm generally comprises many agents who each have some input into the actions of the firm. In most organizations, some committee of members oversees the decisions of the

* Corresponding author.

E-mail addresses: charness@econ.ucsb.edu (G. Charness), jacksonm@stanford.edu (M.O. Jackson)

URLs: <http://www.econ.ucsb.edu/~charness/> (G. Charness), <http://www.stanford.edu/~jacksonm> (M.O. Jackson).

organization and votes on critical decisions. There is typically some specified method of aggregating the agents’ preferences into a group decision, and the nature of this method can have important consequences for the organization’s behavior.

While there is considerable research about how groups behave and how their behavior differs from that of individuals, there is much less research about how groups interact with each other, and in particular, how games are played when the ‘players’ are groups, rather than individuals. This is of obvious importance, since the strategic interaction between competing firms is not between unitary actors, but one between competing groups, and similarly the interaction among countries and other organizations. Although it has long been recognized that treating firms and other organizations as individual decision makers is at best an imperfect approximation, not much is known, theoretically, empirically, or experimentally, that discerns how groups interact with each other in strategic contexts.

In this paper we look explicitly at the play of a game when the ‘players’ are actually groups of players. One of our findings is that the play of games between groups of players is systematically different than the play of games between individuals, even in a situation where all traditional game-theoretic equilibrium concepts do not provide any distinction. This prompts us to develop a new equilibrium concept to explain our observations.

To discuss our approach and analysis, it is easiest to fix ideas directly in terms of the game that we use in the experimental part of our study. Consider the following variation on Rousseau’s classic Stag Hunt game.

		Player 2	
		Stag	Hare
Player 1	Stag	9,9	1,8
	Hare	8,1	8,8

There are two Nash equilibria in pure strategies: one, where both players play Stag, and another where both players play Hare.¹

Let us consider the case where a ‘player’ is actually a ‘group’ or, in particular for our experiments, a couple.² That is, there are two individuals who fill the role of a given player—and both of these individuals receive that player’s payoffs. We now have to be explicit about how the couple determines whether to play Stag or Hare. There are different ways in which a couple could decide on how to play. The three prominent ones are:

- One of the two individuals dictates.
- The couple votes over which strategy to play by one of the following two methods:
 - It takes two votes for Stag to play Stag, and one vote for Hare to play Hare.
 - It takes one vote for Stag to play Stag, and two votes for Hare to play Hare.

Once we specify the method by which each couple makes the decision to play, the overall game is fully specified, so that we can solve for the equilibria of the larger four-player game. It is easy to see that regardless of the way in which the couples make their decisions, there will still be an equilibrium in which both couples play Stag, and another equilibrium in which

¹ The selection between these equilibria (and those in related coordination games) has been the subject of much study. See for instance, Harsanyi and Selten [19], Kandori et al. [24], Young [35], and the literature that followed.

² In our experiments, groups are couples. While our theoretical analysis applies to larger groups, it would be interesting to eventually extend the experimental analysis to larger groups as well.

both couples play Hare. In fact, it is clear that this is a much more general phenomenon: Any pure-strategy equilibrium of a game with ‘individual players’ remains an equilibrium of the game where groups fill the roles of players, regardless of the way in which the groups make decisions.³

Given this observation, one might wonder whether there is anything interesting to learn about groups playing games in situations where each group’s members share the same payoffs. We find that some equilibria are more pertinent when groups play than when individual players do, and that this selection among equilibria depends on the method by which the group makes decisions. We argue that there is a (formal) sense in which one of the two equilibria becomes more ‘focal’. On an intuitive level this is quite easy to see. Suppose that both couples make decisions through a voting system where it takes two votes to play Stag and one vote to play Hare. In this situation, it is ‘easier’ for a couple to decide to play Hare. This in turn can feed into the expectations that one couple has about the other couple’s play, and so forth. This leads to a prediction that the Hare equilibrium is the focal one in cases where it takes two votes to play Stag and one vote to play Hare. Correspondingly, the Stag equilibrium becomes more focal if the voting system is such that it takes two votes to play Hare.

We come back to offer a formal analysis of this in Section 5, but for now we take the view that the focal equilibrium play is the one that only takes one vote of two to undertake.

2. The related literature

There are a modest number of contributions from the economics and game theory literatures on games played between groups.⁴ From the theory side, Duggan [13] defines “group Nash” equilibria and examines the group decision-making processes that allow for the existence of equilibrium when individuals in a group may have different payoff structures.

Experimentally, Cason and Mui [6] consider individual and team choices in a dictator game, with group choices more generous than individual choices. In contrast, Cox [12] finds that groups behave similarly to individuals as senders in an investment game, but send back much less. Blinder and Morgan [3] find that groups were just as quick as individuals to reach decisions and also outperform individuals. Cooper and Kagel [10] study the behavior of individuals and two-player teams in signaling games; teams consistently play more strategically than individuals do. Similarly, Kocher and Sutter [25] find that groups learn more quickly than individuals in a beauty contest (guessing game).

Gary Bornstein and co-authors have a series of experimental papers on group play in social dilemmas, with particular emphasis on inter-group competition on group performance. Perhaps the most closely related paper is Bornstein et al. [4] who examine a variation of Van Huyck’s minimum-effort game, where two groups play the game and only the group with the higher minimum receives a non-zero payoff (in the case of a tie, both groups earn half of this payoff).

³ This is provided the group decision-making system has a full range. Specifying any actions that lead to the candidate strategy leads to a situation where no group member can benefit from a deviation.

⁴ As one would expect, the social-psychology literature has many studies of inter-group relations. Sherif [31] and Tajfel [34] provide extensive surveys of some of the classic theories of how groups interact. While this is a rich literature, we have not found any predictions that would help us select between the multiple equilibria in the games we study.

They find that such a form of inter-group competition leads to higher efforts and improved collective efficiency relative to the usual minimum-effort game.

In all of these games, groups either reach a unanimous decision after some form of communication or adopt the minimum (and therefore unanimous) choice. In contrast, our analysis is one where group members only interact through voting on decisions, and do not even communicate with their partners.⁵ This is clearly an extreme, but one that begins to capture how the procedure by which a group makes a decision matters in strategic interaction. Of course, understanding how communication and procedures interplay is of obvious interest for future research.

Our focus is on how the play of a game depends on the aggregation process of the individual choices in groups, and how this influences the beliefs and behavior of players. The structure of our games bears no resemblance to that in any of the papers mentioned above. The equilibrium analysis that we undertake uses standard game-theoretic (and network-based solutions), as well as one that we develop to explain the data. Our new equilibrium concept, “robust-belief” equilibrium, provides a new refinement of Nash equilibrium play. It builds on tools from the stochastic-stability literature, most notably from Kandori et al. [24] and Young [35].

The new aspect is that voting games, as well as many others, have situations where players are completely indifferent over how they play—for instance in situations where they are not pivotal. In such situations, stochastic stability (and other refinements) has little bite. The insight leading to the concept of robust-belief equilibrium is that players best respond to the process, taking into account the probability that other players might tremble. This breaks the indifference and makes sharp predictions that can actually completely reverse the prediction of stochastic stability (in ways consistent with our experimental observations).

3. Network formation and consent


Our analysis also has a completely different interpretation, and implications for a different area of study. In fact, our initial interest in this project originated from thinking about differences in protocols for network formation and our experimental instructions were written from this perspective, so we would like to explain the connection.

In many social and economic networks it takes the consent of both parties to maintain a link or tie. This includes such varied applications as friendships, political alliances, trading relationships, partnerships, etc. In some other contexts a tie can be formed unilaterally—one web site can link to another or one author can cite another, without having mutual consent. How does this difference in the consent needed to form a link affect the network that emerges?


Network formation maps directly to the analysis of the Stag Hunt as follows. The choice of play by a couple is whether or not to form a link. The ‘voting rule’ then translates into whether or not one needs the consent of both individuals to form a link or whether just one individual can

⁵ After the original version of our paper, we became aware of related independent work by Elbittar et al. [15]. They examine the role of voting by 3-player groups on ultimatum-game proposals, varying the decision rule for acceptance of an offer. They find no effect of the voting rule on individual responder behavior, although individual proposer behavior anticipates the effect of the voting rule on the likelihood of acceptance, by becoming more aggressive as it takes more votes to eject a proposal. In contrast, we find the voting rules have strong effects on individual choices, and we also observe more complex strategic differences.


form a link. For instance, consider starting at the following network, with corresponding payoffs to each one of the players. From here players can choose whether or not to add the missing links.



Players A and D might be thought of as the row group in the Stag Hunt game, and players B and C might be thought of as the column group in the Stag Hunt game. Playing Hare corresponds to adding a link, and playing Stag corresponds to not adding a link. Thus, if A and D add a link between themselves, while B and C do not, then we end up with the following network and payoffs:



If both groups add a link, then we end up with the following network and payoffs:



If mutual consent is required to add a link, this is equivalent to having both players have to vote Hare in order to play Hare. If either player in a group can add a link unilaterally, then this

corresponds to a single vote being enough to play Hare. Thus the two variations on this network game where we vary the consent needed to form a link correspond exactly to the Stag Hunt game played between two groups where the voting rule is varied.⁶

The reasoning that corresponds to that of what is ‘focal’ when groups play Stag Hunt is as follows. In the game where unilateral consent is needed to form a link, it is ‘easier’ to form links, and one should expect them to form. If it is more likely that the other pair will add their link, then it is a best response to add a link. In contrast, when it takes mutual consent to form a link, then it can be more difficult for either group to form a link and thus the network where no links are formed becomes more focal.

In terms of predictions from the network-formation literature, we again find that there are multiple equilibria.⁷ For instance, if mutual consent is needed to form a link then we can use the concept of pairwise stability from Jackson and Wolinsky [23] to develop predictions of which links will form. This asks that no two players wish to add a link, and no single player wishes to delete a link. Here, (subject to the constraint that the initial four links cannot be altered) both the network where neither link is formed and the network where both links are formed are pairwise stable. In the case where single consent is needed, then the game can be viewed as a variation on a link-formation game proposed by Myerson [28] (see also [1]), and Nash equilibrium can be used as a solution concept. Again, there are two (pure-strategy) equilibrium network configurations: one where neither link is added and the other where both links are added. In order to select among the equilibria, we need to employ some coalitional refinement. In the case of mutual consent, strong stability (see [21]) will single out the networks that lead to the payoffs of nine, and similarly strong Nash equilibrium will do the same in the case of unilateral link formation.

There are also notions of farsighted network formation, such as those of Page et al. [29] and Dutta et al. [14]. In the context of the networks we have examined, they would select equilibria where Stag is played.⁸ This is not the behavior observed here, although perhaps because our setting is sufficiently different in structure so as not to favor such forward-looking behavior.⁹

4. Experimental description and hypotheses

4.1. The experiment

We conducted a series of experiments in six separate sessions at the University of California at Santa Barbara. There were 16 participants in each session, with average earnings of about \$15

⁶ Note that there are other variations of the above network formation game that we can consider. In addition to the two variations that we have already described, we can also reverse the payoff structure. That is, we could have payoffs of 8 when neither link is added, payoffs of 9 when both links are added, and, when exactly one link is added, payoffs of 1 for the group adding a link and payoffs of 8 for the group not forming a link. While these are seemingly different variations on the games, they are again completely strategically equivalent to the others. We ran these variations to check that the framing was not important, and indeed found the results to be indistinguishable across framing, as we detail below.


⁷ See Jackson [20] for a survey of models of network formation.

⁸ Dutta et al.’s [14] definition is in the context of an infinite game of network formation, and would require a sufficiently high discount factor. Their analysis would not have any specific prediction in the context of our game as it is outside of their setting. Nevertheless, the reasoning that they are attempting to capture is consistent with some of the attempts of players in the game to play Stag, with the hope that others will see this and follow suit. The analysis of Page et al. [29] is independent of a particular protocol, as it builds on Chwe’s [9] largest consistent set.

⁹ See Pantz and Ziegelmeyer [30] for a more detailed study of myopic versus farsighted behavior in network formation.


(including a \$5 show-up payment) for a one-hour session. The experiments were framed in terms of the formation of links. The complete instructions are presented in Appendix D. Our experiment was programmed using the z-Tree software [16].

Participants were sorted randomly into groups, typically with four people in a group. We imposed the initial link structure of a ‘square’, where each person was linked to two adjacent parties, with no diagonal links. In three of our sessions, this was the initial link structure:



The issue of choice was whether to add a diagonal link, as we described in the previous section. If neither diagonal link was created, each person received nine experimental units. If both diagonal links were created, each person received a payoff of eight. If only one diagonal link was created, the people at the nodes of this link received eight, while the people not connected by a diagonal link received only one.

In the other three sessions, we started with lower default payoffs and also altered the way that payoffs were determined by the links added.



In these sessions, if no diagonal links were added, then each person received a payoff of eight experimental units. If both diagonal links were added, then each person received a payoff of nine. If only one diagonal link was created, the people at the nodes of this link received only one, while the people not connected by a diagonal link received eight.

The contrast between the two different starting payoff structures is simply one of framing, as we shall discuss more fully below.

There were 60 periods in a session, comprised of four 15-period segments with different decision rules for whether a diagonal link was created. These segments differed with respect to the decision rules for creating a diagonal link. The people in each group were matched with each other for

the 15 periods in a segment; after each segment, participants were randomly re-matched. Each person was involved once in each of these three different cases.¹⁰

- (1) *Mutual consent*: People are in 4-person groups. Consider the two people diagonally across from each other: a link is added if and only if *both of these people* wish to add it.
- (2) *Unilateral consent*: People are in 4-person groups. Consider the people diagonally across from each other: a link is added if and only if *at least one of these people* wishes to add it.
- (3) *Dictators*: People are in 4-person groups. However, only two of these people will make choices in the game (say players A and C in the above diagrams), and they can unilaterally decide whether or not to add their corresponding links. Each one of these people is paired with a silent partner (say players B and D in the above diagrams) who is inactive in the game but whose payoff depends on the play of the game.

Participants always knew which case applied before making a decision. Once the decisions for the period were reached, each person learned his or her payoff, the choice of his or her partner (diagonal counterpart), and whether or not the other group (off-diagonals) added a link. Participants also knew that they would learn this information before making subsequent choices.

4.2. Hypotheses

The reason for running variations where the default payoffs differ was to simply check for framing effects. That is, consider two variations: one where the starting payoffs are 8's and it takes mutual consent to add a link, and another where the starting payoffs are 9's and a link can be added unilaterally. These two situations are strategically equivalent: they differ only in their framing. In either situation if a single player wants to go for the payoff of 8's (the "hare" play) that player can guarantee it, while it takes two players to try for the payoff of 9 (the "stag" play). In principle, play should be the same across strategically equivalent games.

As noted above, solution concepts such as pairwise stability and Nash equilibrium result in multiple equilibria in our games, and thus leave much open in terms of how play might proceed. We can look to ideas such as risk and payoff dominance to try to select among the equilibria. The notion of risk dominance is consistent with both groups choosing Hare and the notion of payoff dominance is consistent with both players choosing Stag.¹¹ We wish to remain agnostic regarding the relative merits of payoff dominance and risk dominance, but note that play should be the same across all treatments if either notion were a consistent underlying force. We find nothing else in conventional theory to guide us. Thus, we formulate null hypotheses based on the view that behavior *will not vary* across our experimental games.

¹⁰ Subjects also played 15 periods in a fourth type of game, which were standard two-player games (so, without groups). There are differences in play between the "dictator" games and the two-player games. We report on those in a separate paper, Charness and Jackson [8]. Overall, we find fairly strong support for the notion of *responsibility-alleviation* [7]. If being responsible for another player's payoff has no effect on how individuals act, then the play in these two treatments should be exactly the same. However, we find that the behavior of about one-third of the population is sensitive to the issue of being responsible for another person's welfare. In about 90% of these cases, the decider takes on less risk when he or she is the agent for another party than when acting only for him or herself. However, this lessened risk comes at a social cost, as the average payoffs are reduced substantially when the deciding agent also represents another party.

¹¹ We need to be a bit careful here, as with four players the ideas apply slightly differently than with just two. Nevertheless, the risk-dominant play is still for every player to play Hare, regardless of the voting rule, and payoff dominance has both groups playing Stag, regardless of the voting rule.

Table 1
Individual votes by game

Game	Stag	Hare
1 vote for stag	1047 (72.7%)	393 (27.3%)
2 votes for stag	456 (31.7%)	984 (68.3%)

Hypothesis 1 (*No effect of payoff structure*). There will be no difference between the tendency to vote Stag when it takes two votes to go Stag and the tendency to vote Hare when it takes two votes to go Hare.

Hypothesis 2 (*No effect of voting rule*). There will be no difference in voting for Stag across the ‘two-votes-to-go-stag’ and the ‘one-vote-to-go-stag’ treatments, all else held constant.

The above hypotheses are not nested nor are they mutually exclusive, so there are many alternative hypotheses to various sets of the above. Rather than list all of these possible alternatives, let us mention a few.

There are two different views as to how payoffs might affect play, depending on whether one looks to payoff or risk dominance. Strong stability and strong Nash equilibrium give clear predictions in Stag, while risk dominance gives a clear prediction in favor of play resulting in Hare. We thus have two alternative hypotheses.

Hypothesis 1a. The tendency to vote for Stag when it takes two votes to go Stag will be stronger than the tendency to vote for Hare when it takes two votes to go Hare.

Hypothesis 1b. The tendency to vote for Hare when it takes two votes to go Hare will be stronger than the tendency to vote for Stag when it takes two votes to go Stag.

If we adopt the view that links will form if and only if formation requires only one vote, we have an alternative hypothesis:

Hypothesis 2a. There will be more votes for the Stag when it takes one vote to go stag than when it takes two, and similarly for Hare.

This is the critical alternative hypothesis that the method by which groups make decisions affects play.

5. Analysis of the experimental results

5.1. Descriptive statistics

We observe considerable variation in play, according to the type of game. Table 1 displays the pattern of play in the four different types of games and already gives us a strong look at the hypotheses, and makes it pretty clear what the conclusions will be.

While we proceed shortly to careful statistical tests of the hypotheses, let us first discuss what these aggregate statistics suggest. The largest difference in play can be found by comparing the

Table 2

Group outcomes: number of observations (percentages)—by category

Game	Both groups stag	Mis-coordination	Both groups hare
1 vote for stag	310 (86.1%)	12 (3.9%)	38 (10.5%)
2 votes for stag	46 (12.8%)	25 (6.9%)	289 (80.3%)

four-player games. People are much more likely to vote for Stag when a group choice of Stag requires only one vote out of two than when both votes are required. Although ex post this seems quite intuitive, we had not anticipated it ex ante and, as we shall see, this result is not captured in standard equilibrium notions.

Hypothesis 1, that the payoff structure does not matter, will be accepted—as the voting for Hare or Stag almost exactly reverses itself based on the voting rule and not on the label of Hare or Stag (and hence not on the payoffs). Hypothesis 2, that the voting rule does not matter, will be rejected in favor of the alternative that voting goes in favor of the strategy requiring only one vote—that is, the group decision-making rule is an (the) important factor in determining play and it goes in favor of the decision that requires the least votes or consent to achieve.

From the standpoint of network formation and social-welfare analysis, it is also relevant to consider group outcomes in addition to individual voting behavior. Table 2 shows these group outcomes across different voting rules:

The patterns seen in Table 1 are also apparent in Table 2, and are more striking. This is a consequence of the voting rule: For instance, suppose that players were simply flipping coins in order to decide whether or not to vote to form a link. In that case, more links would be formed under unilateral consent than mutual consent, and not because their play was influenced by the structure of the consent rule, but simply because of the consent rule.

Aggregating across four-voter games, the (9,9) outcome happens 86.1% of the time when the consent rule mandates that the group go Stag if there is at least one vote to do so, compared to 12.8% of the time when both votes must be Stag. We also observe that the unattractive mis-coordination (1,8) outcome is relatively rare, eventuating in only 37 of 720 cases across all games, or 5.1% of the time; in the few cases where there is coordination failure in initial periods, we often see rapid movement to successful coordination.

We can also consider the choice tendencies of each individual in each game. The complete profiles across all games each person played are shown in Appendix C, and we come back to discuss the determinants of individual plays below.

5.2. Hypothesis tests

As we have repeated interaction among players, we must be careful in performing statistical tests of our hypotheses. While there is no obvious reason why behavior should unravel over time, one could be concerned that there is nevertheless a repeated-game flavor to the data. One method for eliminating such concerns is to consider only the last period in the 15-period segments; these data might also be seen as reflecting more ‘settled’ behavior. Further, since the behavior of the individuals within a group may well be correlated, we consider group performance, as measured by the number of votes for links in this last period. As this severely limits the number of observations and the power of the tests, in some cases we also consider a less conservative test involving the behavior of each individual in each segment, yielding more observations at a cost of some loss of independence.

Table 3
Stag vote totals for groups in last period of segment

Voting rule	Number of groups with				
	0 stag votes	1 stag vote	2 stag votes	3 stag votes	4 stag votes
2 votes to go Stag	13	5	2	0	4
1 vote to go Stag	3	0	2	8	11

Table 4
Link choices by game

Game	No link	Link
<i>Start at 9's</i>		
Mutual consent	512 (71.1%)	208 (28.9%)
Unilateral consent	191 (26.5%)	529 (73.5%)
<i>Start at 8's</i>		
Mutual consent	455 (63.2%)	265 (37.8%)
Unilateral consent	185 (25.7%)	535 (74.3%)

The group-level data we use for most of our hypothesis testing are shown in Table 3:

Hypothesis 1 (*No effect of payoff structure*). Table 4 summarizes the link choices by four-person game, illustrating the effect of the voting rule.

The choice of adding a link when starting at 9's under mutual consent is the same as not adding a link when starting at 8's under unilateral consent, and similarly for the other case. Another way to see this is to break play down player by player, as in the following figures that track the frequency of individual stag play across strategically equivalent variations (Figs. 1 and 2):

Statistical tests confirm the visual pattern observed in these figures, that there is little difference due to framing.¹² Given that the play across strategically equivalent games is statistically indistinguishable, we generally pool the data in what follows (although on some tests we also disaggregate to be especially cautious).

We cannot reject Hypothesis 1, as the Wilcoxon test on group outcomes in the last period of the game gives $Z = 0.24$, very slightly in favor of Hypothesis 1b, but far short of statistical significance ($p = 0.810$, two-tailed test). If we further increase (quadruple) the sample size by

¹² The most careful tests use group choices made in the last-period of a segment, to avoid both interactions and repeated-game effects. The Wilcoxon–Mann–Whitney ranksum test (see [32]) finds no significant difference in rates of stag play across either the mutual-consent game starting at 9's and the unilateral-consent game starting at 8's ($Z = 0.61$, $p = 0.271$, one-tailed test) or across the mutual-consent game starting with 8's and the unilateral-consent game starting with 9's ($Z = 0.12$, $p = 0.452$, one-tailed test). If we instead consider individual last-period data, these comparisons are 13 people of 48 versus 12 people of 48, and 10 people of 48 versus 15 people of 48. Using the test of proportions (see [18]), the test statistics for these differences are $Z = 0.23$ and $Z = -1.16$, respectively, with $p = 0.409$ and $p = 0.877$ for the one-tailed tests.




Fig. 1. Stage play if one vote needed to go Stag.




Fig. 2. Stage play if two vote needed to go Stag.

considering individual voting in the last period of each 15-period segment, there is no difference at all: 25 of 96 individuals voted to go Stag when going Stag took two votes and 25 of 96 individuals voted to go Hare when going Hare took two votes.

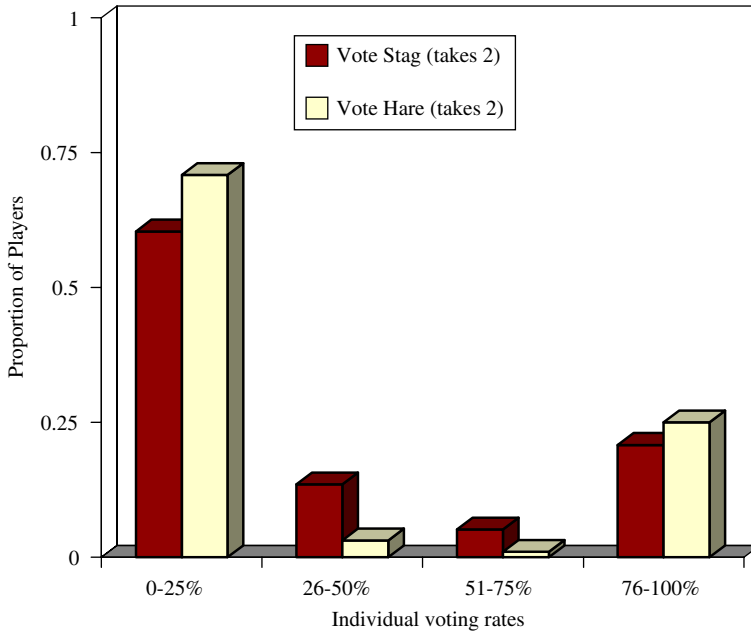


Fig. 3. Voting behavior.

Fig. 3 examines whether or not the payoff structure affects individual voting behavior, by comparing the rate of voting to go Stag when this requires two votes to the rate of voting to go Hare when this requires two votes.

Overall, there is no clear visual difference across these conditions, although there are slightly more cases of one voting an intermediate proportion of the time for Stag than for Hare.

Hypothesis 2 (*No effect of voting rule*). If we combine the mutual plays and unilateral plays across starting payoffs, we see the difference between consent/voting structures most vividly: Fig. 4 illustrates the central finding of our experiments: Play is largely dictated by the consent/voting role, essentially independently of the payoff structure.

Quite to the contrary of the null hypothesis, play is essentially completely reversed by the voting rule. For a conservative statistical test, we use the Wilcoxon–Mann–Whitney ranksum test (see [32]) on the last-period group-level data for four-voter games. Consistent with the visual evidence, there are strong differences according to the voting rule. The Wilcoxon test gives $Z = 3.56$; thus, we can easily reject Hypothesis 1 in favor of Hypothesis 1a, as the corresponding significance level is $p < 0.001$.¹³

¹³ We can also break things down to the specific games, not pooling across strategically equivalent treatments. Comparing the start-at-9's games, the test gives $Z = 3.09$, while the test gives $Z = 2.02$ for the start-at-8's games, with $p = 0.002$ and $p = 0.042$, respectively, using two-tailed tests.

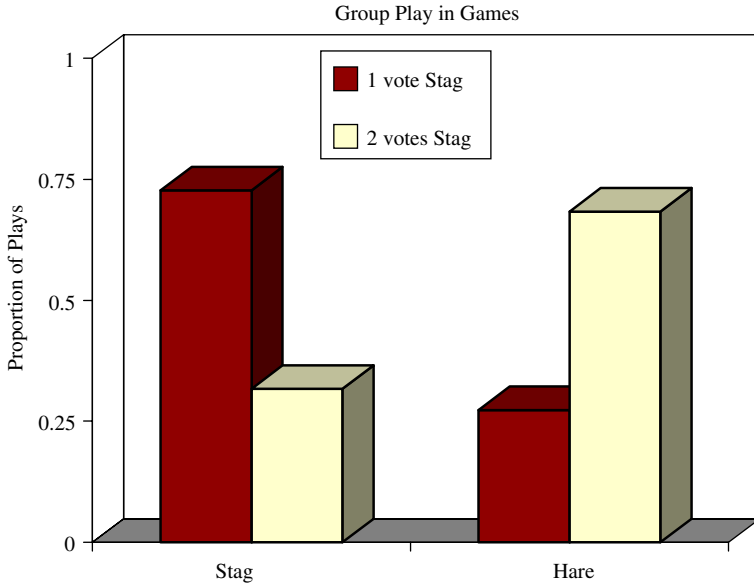


Fig. 4. Individual votes.

Table 5
Rates of risky play over time, by game

Game	Rate of risky play (Stag) in		
	First five periods	Middle five periods	Last five periods
1 vote to go stag	.7125	.7271	.7417
2 votes to go stag	.4445	.2792	.2583

5.3. Changes over time and determinants

Does the tendency to play the risky strategy (Stag) change appreciably over time? If so, can we isolate any determinants of this change? Table 5 shows the average rate of risky play in the first five periods of the segment, the middle five periods of the segment, and the last five periods of the segment for each of our games:

While risky play increases slightly over time when only one vote for Stag is sufficient to go Stag, it decreases substantially when two Stag votes are required. This latter result suggests an initial attraction for payoff dominance that meets with poor results.

We can also examine individual choice tendencies and how common it was for an individual to change his or her strategy over the course of a 15-period segment, and in which direction (from Stag to Hare, or vice versa). This is shown in Table 6.

We see that individual behavior mirrors the patterns observed in Tables 2 and 4, and switching behavior (changes from one’s choice in the first period of a segment to the last period of a segment) mirrors the pattern seen in Table 5. Overall, people changed their choice in 58 of the 192 possible

Table 6
Individual choice tendencies, by game

Game	Stag	Mix	Hare	Stag⇒Hare	Hare⇒Stag
1 vote to go stag	71	1	24	5	12
2 votes to go stag	19	12	65	36	5

“Stag” or “Hare” means the individual made this choice at least 10 times (out of 15). “Mix” means that an individual chose each action at least six times (out of 15). Switching (rightmost columns) reflects a difference between play in the first and last periods of a segment.

Table 7
Determinants for risky play

Independent variables	Dependent variable	
	Vote for Stag (1)	Vote for Stag (2)
Two votes to go Stag	−1.732*** [0.078]	−1.755*** [0.079]
One bad outcome	−1.158*** [0.151]	−1.194*** [0.150]
Two bad outcomes	−2.635*** [0.245]	−2.754*** [0.244]
Three or more bad outcomes	—	−0.815*** [0.213]
Rho	0.710*** [0.034]	0.726*** [0.036]
Constant	0.468*** [0.101]	0.503*** [0.116]
# Observations	2880	2880
Log-likelihood	−1002.4	−995.2

Standard errors are in brackets.

***Indicates significance at 0.1%.

cases, or 30.2% of the time. We see that changes go predominantly in the direction of the voting rule, as 12 of the 17 people who switch change from Hare to Stag when Stag is focal, and 36 of the 41 people who switch change from Stag to Hare when Hare is focal. A simple binomial test rejects the hypothesis that 48 of 58 choices going in one direction reflect random changes ($Z = 4.99$, $p < 0.00001$).¹⁴

There is also some support for the notion that people have some initial attraction to payoff dominance, as there is an inference about initial attraction from the fact that only 12 people switch from Hare to Stag in the first row of Table 6, while 36 people switch from Stag to Hare in the second row; the binomial test rejects the hypothesis that these data reflect randomness ($Z = 3.46$, $p < 0.001$).

¹⁴ Given the grouping of subjects, there might be (small amounts of) correlation so that the observations might not satisfy the appropriate independence condition to make the binomial test fully valid. We feel these tests are still illuminating enough to present, especially in conjunction with the group level tests above and the regression analysis below.

Why do people switch their behavior over the course of a segment? One would naturally suspect that this decision is correlated with having previously experienced bad (payoff = 1) outcomes during the segment. In fact, random-effects probit regressions (with the individual as the random effect) strongly confirm this conjecture:

In these games, people are much more likely to play safe (hare) when it takes two votes for the group to make the risky (stag) choice. A player who has experienced one bad outcome is substantially less likely to vote for the risky play (stag). The effect of a second bad outcome is much stronger, as the coefficient is even larger than that for the treatment effect. However, a person who persists in voting for the risky play (stag) after two bad outcomes doesn't seem to care much about bad outcomes, as we see in specification (2) that the coefficient for the effect of three or more bad outcomes is actually smaller than the coefficient for having had only one bad outcome. Rho (the random-effects term) is highly significant in both regressions, so that individual variation is seen to be an important factor (Table 7).

6. A foundation for the selection of equilibrium

Our experiments exhibit patterns and regularities in the way that groups play these games. In particular, while the games in question have multiple strict Nash equilibria, there is a tendency to focus in on specific ones: those 'favored' by the voting rule, as evidenced by the sound rejection of Hypothesis 1. This is exhibited even when the payoffs are altered, and so the selection cannot be attributed to equilibrium refinements such as risk dominance or payoff dominance. None of the other standard equilibrium refinements, such as trembling hand perfection or stochastic stability, predict the selection that the participants are consistently making either.¹⁵ Nevertheless, there is a clear logic behind the behavior of the subjects, which suggests that there should be an equilibrium concept to capture it. In this section, we use the intuition and observed play in these experiments to develop a new equilibrium concept. It builds on tools from perfection and stochastic stability, but in a way to capture the logic that seems to be underlying the behavior in our experiments.

We begin by explaining the underlying ideas; to do so we consider an abstract two-by-two game that has two pure-strategy equilibria. For this discussion, consider a coordination game where, as will become clear, the symmetry of payoffs is irrelevant to the general development of the solution concept.

		Player 2 Role (Group 2)	
		A	B
Player 1 Role (Group 1)	A	a, a	c, d
	B	d, c	b, b

Let $a > d$ and $b > c$, so that the game has two pure-strategy equilibria, A,A and B,B. Assume that $(a-d) > (b-c)$, so that A,A is the risk-dominant equilibrium in the sense of Harsanyi and

¹⁵ Even the application of the theory of global games (see [5]) does not work here. With the right specifications of uncertainty, this sort of game can be analyzed using the techniques of Morris et al. [27]. However, depending on the structure of uncertainty that is introduced, it seems that the selection is made either in favor of the risk dominant, or in favor of the payoff dominant payoffs (viewed as a two-player game); and that the selection is *not* being made relative to the voting rule. There are other behavioral based solution concepts, such as the "Lk" theory of Stahl and Wilson [33] (see also, e.g., [11]) that can be examined. Again, the details can affect which equilibrium is selected, but the theory does not make a unique prediction that aligns with what we observed here.

Selten (1988).¹⁶ The game in our experiments fits into this setting. Let each player role be filled by some group of at least two individuals. Describe a group's method of selecting a strategy by a quota rule q , where q in $\{1, \dots, n\}$ denotes the minimum number of votes for strategy A needed before A is played. Thus, if q or more members of the group vote to play A, then A is played, while if fewer than q members of the group vote to play A, then B is played.

Let a state be a list of each individual's play—so it is a vector of length $2n$ where entries are either A or B. Given a state, each player has a best response. That is, considering a conjecture that other players will play as described in a state, each player has a best response. A pure-strategy Nash equilibrium is then just a state where the state's prescription of each player's vote is a best response given the anticipation that the state correctly describes the voting behavior of the players.

There are many pure-strategy equilibria of such a game. Any profile of votes where both groups pick A is an equilibrium; any profile of votes where both groups pick B is an equilibrium; and any profile of votes where one group picks A, the other group picks B, and no voter is pivotal, is also an equilibrium. In fact, all of these states are undominated Nash equilibria.¹⁷ Moreover, all voting A is a trembling-hand-perfect equilibrium, as is all voting B.

A critical aspect of this game is that at any profile where a voter is not pivotal, the voter is completely indifferent as to their action. It is this aspect that makes standard refinements inadequate here.

The behavior in our experiments, however, picks out the equilibrium that is 'favored' by the voting rule. That is, if q is less than or equal to $n/2$, then all voting A is observed, while if q is greater than $n/2$ then all players voting B is observed.¹⁸ On one level we might think of this as a 'focal' effect, in that the bias in the voting rule focuses attention on one of the equilibria. Yet, there are many different things that might be used to 'focus' the attention of the players. For instance, why isn't it that the highest payoff equilibrium becomes the 'focal' one? The challenge for an equilibrium concept is to capture the focal nature of the voting rule, and to do it in a robust way that is not ad hoc to the games in question.

Let us go through some heuristic reasoning that suggests an equilibrium concept. Consider a player forecasting the play of the others. If the player starts agnostically, and simply assumes random voting by the players in the other group, then the forecast would be that the other players would end up playing the strategy favored by the voting rule. While this is rather simplistic reasoning that does not account for payoffs, we can take it further to consider further levels of reasoning. Imagine now that a player starts with any forecast of other players' behavior based on the payoffs and the structure of the game. The player can still think about how sensitive this forecast is to the play of the others. For instance, in our games where $n = 2$ and the quota is 1 to play A, then the player realizes that if either of the players in the other group plays A then it

¹⁶ This equilibrium has the property that each player is choosing a strategy that is also a best response to the other player mixing 50/50. Thus, A is the strategy that is a best response to the largest set of beliefs over possible plays of the opponent, and so the risk-dominant equilibrium is the pure-strategy equilibrium with a larger basin of attraction than the other pure-strategy equilibrium. Specifically, playing A is a player's best response if the probability that his opponent plays A is greater than or equal to $(b-c)/(a-d+b-c) < \frac{1}{2}$. Thus if $a = 3$, $b = 1$, and $c = d = 0$, then playing A is a best response if the belief is that the probability is $\frac{1}{4}$ or more that a player's opponent plays A.

¹⁷ Some of these can be ruled out by 'trembling' arguments, but we still end up with no real selection in outcomes.

¹⁸ We have only examined particular parameter values for the game and number of players, so this is clearly an extrapolation of what is occurring in these games.

is in her interest to play A. Moreover, it is enough for her to believe that one of the other players reasons this way to induce one of the other players to play A. Or, it is enough for her to believe that one of the other players might believe that she is going through this logic, etc. This suggests looking at basins of attraction of equilibria to try to capture some of this reasoning.

The primary equilibrium selection notion that builds on basins of attraction is stochastic stability.¹⁹ Can we bring stochastic stability off the shelf and apply it here? There are two problems. First, stochastic stability is generally used in evolutionary analyses where agents either play games against randomly selected opponents repeatedly over time, or else are playing strategies simultaneously against many other players. This is easily overcome, as it is merely a matter of interpretation.²⁰ That is, the machinery of stochastic stability still captures ideas underlying basins of attraction and robustness of equilibria even when the game is only played once by a given set of players. The second problem is the more crippling one: stochastic stability does not make much of a selection here. Let us discuss this in the context of the game we used in our experiments, where the voting rule is that it takes one vote to go Stag and two votes to go Hare.

Consider movement among states, and start at a state where all players vote to go Stag. In this case, every player has two best responses: either voting Stag or Hare. So, completely consistently with best-response behavior, we can move to a state where all players play Hare. However, if all players are voting Hare, then the unique best response of each player is to vote Hare. So, it takes at least one ‘tremble’ or ‘error’ on the part of some player to switch to a state where at least some players play Stag. Since it took no trembles to get us to move from the equilibrium where all players vote Stag to one where all players vote Hare, but at least one tremble to move in the reverse direction, the unique stochastically stable state turns out to be where all players vote Hare!²¹ However, this is exactly the opposite of what we observe in the experiments.

What’s missing from this reasoning? Suppose we start from a state where all players vote Stag, and we also know that there is some unpredictability (real or imaginary) in players’ votes so that they might vote Hare with a small probability. That is, players are fully aware of the noise in the process and take that into account when making their decisions. Now, choosing Hare makes it more likely that my group will cross the threshold and fail to coordinate, as I realize that one tremble by my partner will lead to coordination failure (whereas, by playing Stag it will require two errors for coordination failure). Straightforward calculations show that now Hare is the unique best response.²² Now, it takes at least two errors or trembles to get from voting Stag to voting Hare, and only one to go back. Using this reasoning, players being somewhat unsure of the other players’ play leads to a unique selection from the game. The key difference between this and the usual definition of stochastic stability is that players are not best-responding to the state, but rather

¹⁹ See Foster and Young [17], Kandori et al. [24], and Young [35] and the literature that follows.

²⁰ See Jackson and Watts [22] for a model of stochastic stability in network formation and some discussion of possible interpretations.

²¹ We could also consider a stochastic process where players’ strategy choices are modified one at a time, according to a best response with errors. In that case, both all voting Hare and all voting Stag are stochastically stable. Regardless of the process, stochastic stability fails to select the equilibrium observed in our experiments.

²² There is a probability of $(1 - \varepsilon)^3$ that all others vote stag, a probability on the order of ε that one other player votes Hare, and a probability on the order of ε^2 or less that two or more other players vote Hare. This means voting Stag leads to a probability on the order of ε^2 for coordination failure, while voting Hare leads to a probability on the order of ε for coordination failure. For small ε the expected payoff is strictly higher from voting Stag.

to the state with the understanding of the full process including errors.²³ This seems like a subtle difference, but it is a critical one not only in changing the predictions of the solution concept, but also in capturing the reasoning that might underlie what we observe in the experiments.

Now let us move to a formal definition of the solution concept. It uses ideas from stochastic stability to capture basins of attraction and robustness, and yet has players best-responding to the actual process to capture the importance of their beliefs (rather than viewing players as random mutations in an evolutionary soup).

We define *robust-belief equilibrium* for an arbitrary game as follows:

Consider an arbitrary finite game $G = (N, S, u)$, where (with an abuse of notation $N = \{1, \dots, N\}$) is a set of players, $S = S_1 \times \dots \times S_N$ is the finite set of pure strategies available to players, and $u = (u_1, \dots, u_N)$ is the profile of von Neumann–Morgenstern utility functions, $u_i : S_i \rightarrow R$. Given a game G and an $\varepsilon > 0$, let $G(\varepsilon)$ be the following variation on the game. Each player trembles from their prescribed strategy with probability ε , $1 > \varepsilon > 0$. In particular, if a player chooses a given strategy then with probability $1 - \varepsilon$ the player's strategy is played and with probability ε the player's strategy is chosen to be any action with equal probability.²⁴

Given a game variant $G(\varepsilon)$, define a Markov process as follows. A state is a profile of strategies s in S . For each state, let players pick a best response to that state *in the game* $G(\varepsilon)$ —that is forecasting that the other players will try to play according to s , but may tremble. This defines a function β from S to S . There may be multiple best responses, so there are may be multiple functions β . For small enough ε , some β 's remain a best-response selection for all smaller ε . Let the set of such β 's that remain best-response selections for small ε 's be denoted $B(G)$. Any selection of best responses β in $B(G)$, together with the tremble probabilities defines a transition matrix across states. We end up with a Markov process $M(G, \varepsilon, \beta)$ that is irreducible and aperiodic. Thus, $M(G, \varepsilon, \beta)$ has a unique steady-state distribution.

Given a game $G = (N, S, u)$, a *robust-belief state* is any state that has positive probability under the limit of the steady-state distribution of $M(G, \varepsilon, \beta)$ as ε goes to 0 for some best-response selection β in $B(G)$. If s is a state that is both a Nash equilibrium and a robust-belief state, then it is said to be a *robust-belief equilibrium*.²⁵

Given a game $G = (N, S, u)$, a *strongly robust-belief state* is any state that has positive probability under the limit of the steady-state distribution of $M(G, \varepsilon, \beta)$ as ε goes to 0 for *all* best response selections β in $B(G)$. If s is a state that is both a Nash equilibrium and a strongly robust-belief state, then it is said to be a *strongly robust-belief equilibrium*.

²³ Thus, this has the flavor of Quantal Response Equilibrium of McKelvey and Palfrey [26] in that players are best responding to the actual process including errors. Note, however, that there generally are multiple Quantal Response equilibria of this game, even for very small error sizes, as all intending to vote Stag, and all intending to vote Hare are QRE for small error probabilities. Therefore, while QRE is a useful econometric tool, it does not provide any selection here. It is critical to our approach to look at limits and relative probabilities, as from the stochastic stability approach, at the same time as having players best responding to the actual process.

²⁴ One can modify our definitions to allow for different ε 's for different players, and to have non-uniform probabilities on the actions that are chosen in a 'tremble'. We note however, that just as with stochastic stability, the relative size of trembles can be important: changing from a situation where trembles are all of the same relative size, to situations where some are of a smaller or larger order can lead to different conclusions, along the same lines as shown by Bergin and Lipman [2] for stochastic stability.

²⁵ In some games (such as matching pennies) the process will cycle among a set of states, none of which are equilibria. In that case, these states might be thought of as robust and 'likely' states, but are not equilibria in and of themselves. This means that while robust-belief states will always exist, robust-belief equilibria might not, as this is a pure-strategy concept.

Let us now return to the coordination game mentioned above.

Theorem 1. *There is a unique robust-belief state in the group-play coordination game where both groups vote using the quota q : it is all players voting A if q is less than or equal to $n/2$, and all players voting B otherwise. Moreover, this state is a strongly robust-belief equilibrium.*

An important observation about Theorem 1 is that the conclusions are completely determined by the voting rule and independent of the levels of the payoffs (a,b,c,d—subject to the inequalities defining the coordination structure). Indeed this is consistent with what we observe in the experiments.

Theorem 1 is actually a corollary to a more general theorem covering a wide class of two-group games with an arbitrary (finite) number of pure strategies for each group, and potentially different sized groups.

Theorem 2. *Consider a two-group finite game, with corresponding group sizes n_1 and n_2 , such that there is a strict Nash equilibrium in the game with corresponding pure strategies s_1 and s_2 . If for each group i , playing any strategy other than s_i requires that more than $n_i/2$ members of the group vote for that strategy (and otherwise s_i is played), then there is a unique robust-belief state in the group-play game: all players in group i vote for s_i . Moreover this is a strongly robust-belief equilibrium.*

The formal proof appears in the Appendix, but we present a brief outline here. For small enough ε , any best-response correspondence is such that if the other group j is electing to play s_j , then any player in group i should vote to play s_i . This comes from the fact that s_i is the unique best response to s_j (given the strictness of the equilibrium) and agents are anticipating that there might be trembles by other players and acting accordingly. This means that from any state where the groups are choosing (s_1, s_2) , all voters will best-respond by voting for these strategies. It is possible that there are also other states, which are also (strict) Nash equilibria, and where all players best respond by voting for those states, or some states where the two groups fail to coordinate, and end up selecting strategies that are not equilibria, and perhaps cycling among several such strategy combinations. In order to determine what the robust-belief states are, we need to see how hard it is to move from (each) one of these situations to another. The key to the theorem is that it takes more ‘trembles’ in order to leave the situation where all voters are voting to play (s_1, s_2) , than it does to leave other states, and to get back to the situation where all players vote for (s_1, s_2) . This is entirely determined by the voting rule and is independent of the specifics of the payoffs (other than the strictness of the equilibrium).

While Theorem 2 covers a broad class of games, it does not cover games where one group’s voting rule favors one equilibrium while the other group’s rule favors another equilibrium. Indeed, it is easy to see (simply by symmetry arguments) that such games will not always have a unique pure-strategy selection. Battle of the sexes, where one group’s rule favors one equilibrium and the other group’s rule favors the other equilibrium is such an example. Nor does it cover more complicated voting rules, where some players have veto rights over certain strategies, or other such variations on voting that one can imagine.

We emphasize that the concept of robust-belief equilibrium should be useful more generally, especially as a refinement in games where players’ best response correspondences are multi-valued in some situations. Best responding to the noisy play can tie down play uniquely.

7. Conclusion

We have explored behavior in a Stag Hunt that can be interpreted as either a study of how the voting rule that a group uses for decision making impacts the play of a game or as endogenous link-formation in networks. We find a stark contrast between how groups play and how individuals play: in the group games, the voting rule is the primary determinant of play, independently of the payoff structure—when hunting Stag requires unanimous consent (two votes) and Hare only requires one vote, then most subjects vote to hunt Stag; while when the voting rule is reversed so that hunting Stag requires only one vote and Hare requires two votes, then most subjects vote to play Hare. Thus, even when group members all have the same payoff structure, we see different behavior in terms of how groups play games against other groups compared to how individuals play games against other individuals. As existing game-theoretic solution concepts allow for multiple equilibria and do not distinguish between these cases, we develop a new solution concept, *robust-belief equilibrium*, which offers a unique prediction that is consistent with the data.

As our experiments were limited to a specific game, it would be interesting to see to what extent group decision making processes are critical determinants of play in other sorts of games, and more generally how this varies with the structure of the game, the number of players in a group, the type of group decision-making procedure (which one might even endogenize), etc. It would also be interesting to further investigate both the theoretical properties and other applications of robust-belief equilibria, which should be a useful tool far beyond our setting.

Acknowledgments

Financial support under NSF Grant SES-0316493 is gratefully acknowledged, as is support from the Lee Center for Advanced Networking, and the Center for Advanced Studies in the Behavioral Sciences. We thank the associate editor and two anonymous referees for comments on an earlier draft.

Appendix A

Proof of Theorem 2. A theorem from Young [35] is instrumental in the proof of Theorem 2. Before stating Young's theorem, the following definitions from Young [35] are needed.

Consider a stationary Markov process on a finite state space X with transition matrix P .

A set of *mutations* of P is a range $(0, a]$ and a stationary Markov process on X with transition matrix $P(\varepsilon)$ for each ε in $(0, a]$, such that: (i) $P(\varepsilon)$ is aperiodic and irreducible for each ε in $(0, a]$; (ii) $P(\varepsilon) \rightarrow P$, and (iii) $P(\varepsilon)_{xy} > 0$ implies that there exists $r \geq 0$ such that $0 < \lim_{\varepsilon \rightarrow 0} \varepsilon^{-r} P(\varepsilon)_{xy} < \infty$.

The number r in (iii) above is the *resistance* of the transition from state x to y . There is a path from x to z of *zero resistance* if there is a sequence of states starting with x and ending with z such that the transition from each state to the next state in the sequence is of zero resistance. Note that from (ii) and (iii), this implies that if there is a path from x to z of zero resistance, then the n th order transition probability associated with P of x to z is positive for some n .

The *recurrent communication classes* of P , denoted X_1, \dots, X_J , are disjoint subsets of states such that: (i) from each state there exists a path of zero resistance leading to a state in at least one recurrent communication class; (ii) any two states in the same recurrent communication class are connected by a path of zero resistance (in both directions), and (iii) for any recurrent

communication class X_j and states x in X_j and y not in X_j such that $P(\varepsilon)_{xy} > 0$, the resistance of the transition from x to y is positive.

For two communication classes X_i and X_j , since each $P(\varepsilon)$ is irreducible, it follows that there is a sequence of states x_1, \dots, x_K with x_1 in X_i and x_K in X_j such that the resistance of transition from x_k to x_{k+1} for each k from 1 to $K - 1$ is defined by (iii) and finite. Denote this by $r(x_k, x_{k+1})$. Let the resistance of transition from X_i to X_j be the minimum over all such sequences of $\sum_{k=1}^{K-1} r(x_k, x_{k+1})$, and denote it $r(X_i, X_j)$.

Given a recurrent communication class X_i , an *i-tree* is a directed graph with a vertex for each communication class and a unique directed path leading from each class j ($\neq i$) to i . The *stochastic potential* of a recurrent communication class X_j is then defined by finding an *i-tree* that minimizes the summed resistance over directed edges, and setting the stochastic potential equal to that summed resistance.

Given any state x , an *x-tree* is a directed graph with a vertex for each state and a unique directed path leading from each state y ($\neq x$) to x . The resistance of x is then defined by finding an *x-tree* that minimizes the summed resistance over directed edges. \square

The following theorem is a combination of Theorem 4 and Lemmas 1 and 2 in Young:

Theorem (Young [35]). *Let P be the transition matrix associated with a stationary Markov process on a finite state space with a set of mutations $\{P(\varepsilon)\}$ and with corresponding (unique) stationary distributions $\{m(\varepsilon)\}$. Then $m(\varepsilon)$ converges to a stationary distribution m of P , and a state x has $m_x > 0$ if and only if x is in a recurrent communication class of P which has a minimal stochastic potential. This is equivalent to x having minimum resistance.*

Now let us return to the details of Theorem 2. Let us describe what must be true of any β in $B(G)$. Consider a player in Group i . We argue that if we are at a state where the players in the other Group j choose s_j , then regardless of the voting behavior of the other players in Group i , the best response for the player anticipating s in $G(\varepsilon)$ for small enough ε is to vote s_i . This is clear if the player were determining the action of his or her group. As the trembles of the group are independent of those of the other group, for small enough ε the player prefers that his or her group choose s_i . Given that it takes at least a strict majority to choose any strategy other than s_i then a player by voting for a strategy other than s_i could only be pivotal between that strategy and s_i . If the player's vote is pivotal (which occurs with positive probability under the trembles), then the player strictly prefers to vote for s_i .

From this, we conclude that if we are in some state where at least one of the groups votes for the corresponding strategy from (s_1, s_2) , then the best responses to that either cycle, with at each point at least one of the groups voting for the corresponding strategy from (s_1, s_2) and which group it is switching in each period and never reaching (s_1, s_2) ; or else the best responses always include at least one of the groups voting for the corresponding strategy from (s_1, s_2) and eventually reaching (s_1, s_2) .

This implies that the recurrent communication classes are of three sorts:

- (1) the class consists of a single state where both of the groups votes for the corresponding strategy from (s_1, s_2) ,
- (2) the class consists of a cycle of states, corresponding to cycles in the best responses, and in each state one of the groups votes for the corresponding strategy from (s_1, s_2) and which group it is continually switches along the cycle,

- (3) the class consists of states such that neither group ever votes for the corresponding strategy from (s_1, s_2) .

Now consider the associated Markov process with small ε . We establish the result by showing that for any tree with minimum resistance for some class other than all players voting for (s_1, s_2) , there is a tree with lower resistance for the class where all players vote for (s_1, s_2) .

So start with a tree that has minimum stochastic potential for some recurrent communication class. If that class is of type (2) above, then alter the tree as follows. Point from that class to the state where all players vote for (s_1, s_2) , and erase the link that was leaving the state where all players vote for (s_1, s_2) . The new link requires no more than $n_1/2$ trembles (as there is a state where group 2 votes for s_2 and group 1 does not, and we need at most half the group 1 players to tremble to voting for s_1) and the old link required more than $n_1/2$ trembles. Thus the new tree has lower resistance.

If that class is of type (3) above, they construct a new tree as follows. First, suppose that there is some class of type (2). Then direct an edge from the original class to the class of type (2). Then point from the class of type (2) to the class (1). Delete the old link out of the class of type (2) and the link out of class (1). Each of the new links requires at most $n_1/2$ trembles (having group 1 change to voting for s_1), while each of the old links required at least that many, and the link out of (1) required more. Thus the new tree has lower resistance. If there is no class of type (2), then instead of pointing from the class of type (3) to the class of type (2), point directly to the class (1). Since there is no class of type (2), then by the definition of (2) it must be that any cycle that includes states where at least one group votes for the corresponding (s_1, s_2) has a path of zero resistance that gets to class (1). This means that the resistance of pointing from a class of type (3) to (1) is no more than $n_1/2$ trembles, while the old link out of (1) required at least that many. The conclusion then follows.

Appendix B. Aggregated rates for adding links, by individual and game

Details of the games starting with 9's and 8's are given in Tables B1 and B2.

Table B1
Games starting with 9's

Participant #	Mutual	Unilateral
1	0.00	0.53
2	1.00	0.67
3	0.00	0.87
4	0.00	0.60
5	0.00	0.93
6	0.00	0.87
7	0.00	1.00
8	0.13	0.53
9	0.00	0.87
10	0.00	0.87
11	0.00	0.73
12	0.87	0.87
13	0.07	0.93
14	0.07	1.00
15	0.00	1.00

Table B1. *Continued*

Participant #	Mutual	Unilateral
16	0.47	0.53
33	1.00	1.00
34	0.00	0.87
35	0.20	0.27
36	0.00	1.00
37	0.00	1.00
38	1.00	0.93
39	0.27	0.40
40	1.00	1.00
41	0.00	0.53
42	0.27	1.00
43	0.00	0.67
44	0.87	0.00
45	0.87	0.93
46	0.20	0.93
47	0.87	0.53
48	1.00	0.93
65	0.00	0.00
66	0.00	0.93
67	0.00	0.40
68	0.00	0.00
69	0.87	1.00
70	1.00	1.00
71	0.00	0.87
72	0.87	1.00
73	0.00	0.80
74	0.00	0.00
75	1.00	1.00
76	0.00	0.00
77	0.00	0.93
78	0.00	0.67
79	0.00	1.00
80	0.00	0.87

Table B2

Games starting with 8's

Participant #	Mutual	Unilateral
17	1.00	1.00
18	1.00	1.00
19	1.00	1.00
20	1.00	1.00
21	0.00	0.07
22	0.00	0.07
23	0.00	0.00
24	0.00	0.00
25	0.93	1.00
26	0.00	0.93
27	0.07	1.00

Table B2. *Continued*

Participant #	Mutual	Unilateral
28	0.20	1.00
29	0.00	0.00
30	0.60	0.87
31	0.47	1.00
32	0.27	1.00
49	0.93	1.00
50	0.00	1.00
51	0.87	1.00
52	0.00	1.00
53	0.13	1.00
54	0.13	1.00
55	0.20	1.00
56	0.07	0.00
57	0.20	1.00
58	0.13	1.00
59	0.13	1.00
60	0.93	1.00
61	0.00	1.00
62	0.00	0.93
63	0.00	0.93
64	0.07	0.93
81	1.00	0.93
82	0.53	1.00
83	1.00	1.00
84	0.33	0.00
85	0.07	0.07
86	0.00	0.00
87	1.00	1.00
88	0.00	1.00
89	0.00	0.93
90	0.93	0.00
91	1.00	1.00
92	0.00	0.07
93	0.07	0.73
94	0.40	1.00
95	1.00	1.00
96	0.00	1.00

Appendix C

Regarding the mutual and unilateral game-pair, 39 people out of 96 (40.6%) played about the same in both games. We see that in games starting with 9's, many more people did not add links under mutual consent and added links under unilateral consent than the other way around; the reverse is true in the games starting with 8's. Overall, for the 57 people whose choice behavior differed across the mutual- and unilateral-consent treatments, 54 of the differences (94.7%) favor being an individual being more aggressive when going Stag (trying for 9) requires only one vote of the two (mapping to mutual consent when starting with 9's or to unilateral consent when starting with 8's) (Tables C1 and C2).

Table C1
Individual profiles in games

Game pair	NN	NM	NL	MN	MM	ML	LN	LM	LL
<i>Start with 9's</i>									
Mutual & unilateral	5	6	23	0	1	0	1	1	11
<i>Start with 8's</i>									
Mutual & unilateral	10	0	21	0	0	4	1	0	12

In this table, “N” means voted for No Link, “M” means Mix, and “L” means voted for Link.

Table C2
Within-subject individual profiles across games

Profile—start with 9's	# of occurrences
A, A	11
A, M	1
A, D	1
M, M	1
D, A	23
D, M	6
D, D	5
All	48
Profile—start with 8's	# of occurrences
A, A	12
A, D	1
M, A	4
D, A	21
D, D	10
All	48


The order of the games in both profiles is mutual, unilateral. “A” means an individual chose to add a link 10 or more times (of 15 in the respective segment), “D” means he or she chose to not add a link 10 or more times, and “M” means he or she chose both actions at least six times.

Appendix D. Experimental instructions (start at 9's)

Welcome to our experiment. You will receive \$5 for showing up, in addition to your earnings from the session.

There will be a total of 60 periods in the session. You will be paired with a group of people (and your position in the group will stay the same) for 15 periods, and then your pairing will change for the next 15 periods, etc.

Payoffs for each person are determined by the links that exist in the *network* below at the end of a period. Your group begins with the following links:




If this network is in place at the end of the period, then each of A–D would receive 9 units, as is indicated by the bold number in parentheses below each letter.


Some or all of the people in the network will simultaneously indicate on their computer their choice concerning whether or not to add a link to the person diagonally opposite. There are four different cases, and you will make choices in each case for 15 periods during the session:

- (1) People are in 4-person groups. Consider the people diagonally across: a link is added if *at least one of these two people* wishes to add it.
- (2) People are in 4-person groups. Consider the people diagonally across: a link is added if and only if *both of these two people* wish to add it.
- (3) People are in 4-person groups; however, only two of these people (not diagonally across from each other) choose whether to add links. Each person’s decision is implemented and two silent participants in the group receive payoffs according to these choices.
- (4) People are in 2-person groups, so that each person controls 2 vertices. If a person chooses to add a link, it is added and he or she receives the payoffs from one vertex.

You will always know which case applies to your decision before you make your decision.

There are four possible networks that could result from the process of adding or not adding links. The first is the network pictured above, where there are no changes. The other three possibilities arise if at least one link is added. These networks, and the corresponding payoffs to the participants are as follows:





The diagrams show that if exactly one diagonal link is included in the final network, the people connected by the link receive 8 units and the people not connected by the diagonal link receive 1 unit.

If both diagonal links are included in the final network, then everyone in the network receives 8 units.

We will randomly choose one period from each of the 15-period blocks for payment, so that only four periods will actually count towards monetary payoffs. These periods will be chosen at the end of the session. We will add up your payoffs from these four periods and convert them to actual dollars at the rate of \$0.30 for each unit.

At the end of the experiment, we will pay each participant individually and privately.

We encourage you to ask questions about the instructions by raising your hand.

Thank you again for your participation in our research.

References

- [1] V. Bala, S. Goyal, A non-cooperative model of network formation, *Econometrica* 68 (2000) 1181–1229.
- [2] J. Bergin, B. Lipman, Evolution with state-dependent mutations, *Econometrica* 64 (1996) 943–956.
- [3] A. Blinder, J. Morgan, Are two heads better than one? An experimental analysis of group versus individual decision making, *J. Money, Credit, Banking* 37 (5) (2005) 789–811.
- [4] G. Bornstein, U. Gneezy, R. Nagel, The effect of intergroup competition on group coordination: an experimental study, *Games Econ. Behav.* 41 (2002) 1–25.
- [5] H. Carlsson, E. van Damme, Global games and equilibrium selection, *Econometrica* 61 (1993) 989–1018.
- [6] T. Cason, V.-L. Mui, A laboratory study of group polarisation in the team dictator game, *Econ. J.* 107 (1997) 1465–1483.
- [7] G. Charness, Responsibility and effort in an experimental labor market, *J. Econ. Behav. Org.* 42 (2000) 375–384.
- [8] G. Charness, M. Jackson, The role of responsibility in group play, Mimeo.
- [9] M. Chwe, Farsighted coalitional stability, *J. Econ. Theory* 63 (1994) 299–325.
- [10] D. Cooper, J. Kagel, Are two heads better than one?, Team versus individual play in signaling games, Mimeo, 2003.
- [11] M. Costa-Gomez, V. Crawford, B. Broseta, Cognition and behavior in normal—form games: an experimental study, *Econometrica* 69 (5) (2001) 1193–1235.
- [12] J. Cox, Trust, reciprocity, and other-regarding preferences: groups vs. individuals and males vs. females, in: R. Zwick, A. Rapoport (Eds.), *Advances in Experimental Business Research*, Kluwer Academic Publishers, Dordrecht, Netherlands, 2002, pp. 331–350.
- [13] J. Duggan, Non-cooperative games among groups, Mimeo, 2001.
- [14] B. Dutta, S. Ghosal, D. Ray, Farsighted network formation, Mimeo, University of Warwick and NYU, 2003.
- [15] A. Elbittar, A. Gomberg, L. Sour, Group decision making and Voting in Ultimatum Bargaining: An Experimental Study, *Microeconomics* 0511005, Econ WPA, 2005.

- [16] U. Fischbacher, *z-Tree—Zurich Toolbox for Readymade Economic Experiments—Experimenter’s Manual*, Institute for Empirical Research in Economics, University of Zurich, Switzerland, 1999.
- [17] D. Foster, H.P. Young, Stochastic evolutionary game dynamics, *Theoret. Popul. Biol.* 38 (1990) 229–232.
- [18] D. Glasnapp, J. Poggio, *Essentials of Statistical Analysis for the Behavioral Sciences*, Merrill Publishing Co., Columbus, OH, 1985.
- [19] J. Harsanyi, R. Selten, *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge, MA, 1988.
- [20] M.O. Jackson, A survey of models of network formation: stability and efficiency, in: G. Demange, M. Wooders (Eds.), *Group Formation in Economics; Networks, Clubs and Coalitions*, Cambridge University Press, Cambridge, MA, 2004(Chapter 1).
- [21] M.O. Jackson, A. van den Nouweland, Strongly stable networks, *Games Econ. Behav.* 51 (2005) 420–444.
- [22] M.O. Jackson, A. Watts, The evolution of social and economic networks, *J. Econ. Theory* 106 (2002) 265–295.
- [23] M.O. Jackson, A. Wolinsky, A strategic model of social and economic networks, *J. Econ. Theory* 71 (1996) 44–74.
- [24] M. Kandori, G. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* 61 (1993) 29–56.
- [25] M. Kocher, M. Sutter, The decision maker matters; individual versus group behaviour in experimental beauty-contest games, *Econ. J.* 115 (2005) 200–233.
- [26] R. McKelvey, T.R. Palfrey, Quantal response equilibrium, *Games Econ. Behav.* 10 (1995) 6–38.
- [27] S. Morris, D.M. Frankel, A. Pauzner, Equilibrium selection in global games with strategic complementarities, Cowles Foundation Working Paper No. 1336, 2001.
- [28] R. Myerson, *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA, 1990.
- [29] F. Page, M. Wooders, S. Kamat, Networks and farsighted stability, *J. Econ. Theory* 120 (2005) 257–269.
- [30] K. Pantz, A. Zeigelmeyer, An experimental study of network formation, Mimeo, Max Planck Institute, 2003.
- [31] M. Sherif, *In Common Predicament: Social Psychology of Intergroup Conflict and Cooperation*, Houghton-Mifflin, Boston, MA, 1966.
- [32] S. Siegel, N. Castellan Jr., *Nonparametric Statistics for the Behavioral Sciences*, McGraw-Hill, Boston, MA, 1988.
- [33] D. Stahl, P. Wilson, On players’ models of other players: theory and experimental evidence, *Games Econ. Behav.* 10 (1995) 218–254.
- [34] H. Tajfel, *Social Identity and Intergroup Relations*, Cambridge University Press, Cambridge, MA, 1982.
- [35] H.P. Young, The evolution of conventions, *Econometrica* 61 (1993) 57–84.