Team Equilibrium in Crisis Bargaining Games: The Effects of Collective Choice Rules*

Jeongbin Kim† Thomas R. Palfrey‡ Jeffrey Zeidel§

November 23, 2025

Abstract

Crisis bargaining games model strategic bluffing in bilateral conflicts, akin to poker, where players risk losses but may secure gains if their opponent concedes. Since real-world crises involve collective decision-making within diverse organizations, we apply the *team equilibrium* concept to analyze these games (Kim et al., 2022). In team equilibrium, members share common payoffs with private perturbations and hold rational expectations about opponents. Voting rules shape group decisions under optimal behavior. Our experiment manipulates payoffs, group size, and voting rules, finding that behavior deviates from Perfect Bayesian Equilibrium but closely aligns with team equilibrium predictions about treatment effects on behavior.

JEL Classification: C72, C92, D71, D82
Keywords: Team equilibrium; Crisis bargaining; Signaling; Voting; Collective choice; Experiment

---

*We acknowledge the financial support of the National Science Foundation (SES-1426560 and SES-2343948), the Center for Theoretical and Experimental Social Sciences (CTESS) at Caltech, and the The Ronald and Maxine Linde Institute of Economic and Management Sciences at Caltech. We are grateful to John Duffy, Michael McBride, and the staff of the Experimental Social Science Laboratory (ESSL) at UC Irvine for their support and for granting access to the laboratory and subject pool. We are also grateful to Ryan Oprea and the staff of the Laboratory for the Integration of Theory and Experimentation at UC Santa Barbara for their support and for granting access to the laboratory and subject pool. We thank Michael Gibilisco and seminar audiences at Florida State University, University of Iowa, and University of Southern California for helpful comments and suggestions. This study was approved by the Caltech IRB committee (protocol 20-1053).

†Corresponding Author, Department of Economics, Florida State University. Email: jkim33@fsu.edu.

‡Division of the Humanities and Social Sciences, California Institute of Technology. Email: trp@hss.caltech.edu.

§Center for Behavioral Institutional Design, NYU Abu Dhabi. Email: jrz8904@nyu.edu.

"Treating national governments as if they were centrally coordinated, purposive individuals provides a useful shorthand for understanding problems of policy. But this simplification - like all simplifications - obscures as well as reveals. In particular, it obscures the persistently neglected fact of bureaucracy: the 'maker' of government policy is not one calculating decision maker but is rather a conglomerate of large organizations and political actors."

—Graham Allison, *Essence of Decision, p. 3*

1 Introduction

In one of the most influential political science books of the last century, *Essence of Decision*, Allison (1971) convincingly attacks the unitary actor paradigm of decision-making in international relations and proposes that more realistic approaches should analyze crisis situations by taking account of frictions that arise because the decision makers are teams rather than individuals. Allison's book is a case study of the Cuban Missile Crisis, which occurred in the early years of President John F. Kennedy's administration, when the government of the Soviet Union, led at the time by Nikita Khrushchev, secretly began installing ballistic missiles in Cuba. The book documents how the decision-making team in the Kennedy administration coalesced on a response to this threat. The crux of the argument in the book, illustrated with the case study, is that while it may be plausible to model decisions by individuals as rational, applying the same single-agent rationality model of decision-making to teams cannot be justified.[1] Allison proposes two alternatives to the unitary rational actor model, one of which emphasizes collective choice processes and procedural constraints ("The Organizational Process Model") and the second which emphasizes the heterogeneous idiosyncratic preferences and biases of the individual members of the decision-making organization ("The Governmental Politics Model").

The Allison critique of the unitary rational actor model extends far beyond applications to national security policy and international relations. Similar dynamics arise in many strategic conflicts in economic environments, where decision-making units are not single actors but organizations, i.e., groups of individuals working toward shared goals while also navigating idiosyncratic private objectives and operating collectively under specific organizational rules,

[1] Allison's insight has been shown to be far-reaching, as evidence about differences between team and individual behavior has accumulated from laboratory experiments in economics and psychology over the last several decades in a wide range of strategic games and decision problems. See Charness and Sutter (2012) for a survey of some of this evidence.

procedures and constraints, which shape their decisions. For example, during union and firm negotiations over contract amendments or extensions, union members often vote on strike authorizations, reflecting diverse preferences within the membership. While all members share the same broad objectives, favoring improved working conditions, higher wages, and better benefits, they differ in their willingness to bear the costs and risks of striking. Votes to accept or reject negotiated contracts highlight this internal diversity. Firms, too, are rarely unitary actors; they are typically large corporations, governmental entities, or even coalitions in cases of industry-wide collective bargaining.

To understand inter-group bargaining processes such as these, it is necessary to take the Allison critique seriously and model the diversity of preferences as well as the collective decision-making mechanism by which those preferences are aggregated into a decision within each group. A theoretical framework specifically designed for this purpose, called *team equilibrium* was developed in Kim et al. (2022). In a team equilibrium, group members have rational expectations about opponents' strategies and share common average payoffs, but have idiosyncratic, privately observed payoff perturbations, modeled as additive, mean zero, i.i.d. random disturbances for each team strategy, for each team member. That is, each member of the same team has the same expected payoff for a team action *on average*, plus an unbiased additive disturbance term. The common component of team member expected payoffs captures the fundamental alignment of team member preferences that make them a 'team', a group of individuals with a common interest, while the i.i.d payoff perturbations model the team member heterogeneity that makes collective decision making non-trivial.

The second element of team equilibrium is the voting rule used by teams to select a team decision. In a team equilibrium, given average expected payoffs for each strategy and team member payoff perturbations, each team member votes optimally, given the expected (equilibrium) payoffs of each strategy and their own idiosyncratic payoff disturbances. The voting rule then aggregates votes into one collective decision. Kim et al. (2022) have shown that different collective choice rules can lead to vastly different outcomes. For example, if teams play a prisoner's dilemma but choose strategies using a supermajority rule, the the game essentially becomes transformed into a coordination game.

Thus, team equilibrium allows for heterogeneity of preferences among team members around a common expected preference, and it explicitly models the collective decision making procedure. Understanding the complex interaction of these two characteristics of teams is crucial for understanding the behavior of teams, and how the behavior and outcomes of games played by teams differ from games played by individuals or unitary actors.

This paper investigates a class of simple crisis bargaining games that model the subsequent stages of international crises like the Cuban missile crisis, where, after a threat by one country, a second country decides to acquiesce or escalate the conflict, in anticipation of the responses by the instigator of the crisis. The instigator, in the event that the threatened country fights, makes inferences about the military strength and resolve of the threatened country and decides whether to carry out the threat, resulting in either a military conflict, or back down. This is formally modeled as a very simple bluffing game, taking as given that the threat has already occurred.[2]

The threatened country is the *first mover*, called Player 1, and has private information about their own strength, and decides whether to escalate or acquiesce. The instigator is the *second mover*, called Player 2, and responds to an escalation by either engaging the fight or backing down. If the fight is engaged, then the instigator loses (wins) if the threatened country is strong (weak).[3]

These crisis bargaining games are characterized by four parameters: the probability the threatened country is strong; the *concession payoff*, a, gained by a player if the other player acquiesces or backs down (and an equivalent loss, $-a$, to the other player); the *(sender) risky payoff*, s, (or loss, $-s$) to Player 1 if there is a fight and they win (lose); the *(receiver) risky payoff*, r, (or loss, $-r$) to Player 2 if there is a fight and they win (lose).

Our experiment compares the behavior of 5-person teams operating under two different decision collective choice rules - majority rule and unanimity rule. Team equilibrium makes much different predictions about behavior and outcomes under the two rules. We obtain data for four different payoff variations of crisis bargaining games and also run a parallel series of sessions with 1-person teams, which allows for an evaluation of pure "team" effects. We fix the concession payoff and the probability of a strong Player 1 (0.5) for all four games

[2]The crisis bargaining model also applies directly to the union-firm negotiation example, with the players of the game having threat options such as strikes, lockouts, or violence, and each party responds to the other as the sequence of actions unfolds, with the potential for further escalation. Each party also has the option of acquiescing to the terms offered by the opposing party. There are a variety of other economic applications, for example legal conflicts involving lawsuits and counter-suits, or negotiations over plea agreements in criminal cases.

[3]Such games are sometimes referred to as "simplified poker games". The simple card game in Myerson (1991, Figure 2.1, p. 38) is one example. The class of games also corresponds to the last two stages of the canonical crisis bargaining game form in Fearon (1994b) [Figure 2, p.241]

and only vary the two risky payoffs.

The natural benchmark that guided our choice of four payoff treatments is Perfect Bayesian Equilibrium (PBE), which uniquely pins down the fight probabilities for each player. In these games the strong Player 1 has a dominant strategy to fight; the weak Player 1 bluffs by choosing to fight with a probability strictly between 0 and 1; if Player 1 chooses to fight, then Player 2 responds by backing down with a probability strictly between 0 and 1. Thus, it is a classic bluffing game where PBE mixed strategies are such that the weak Player 1 and Player 2 are both indifferent between fighting and backing down. The payoff variations chosen for the experiment span the four canonical PBE mixing probabilities. In one payoff treatment, both players' PBE fight probabilities are greater than 0.5; in a second payoff treatment, both are less than 0.5; and in the other two payoff treatments, one of the player's PBE fight probabilities is greater than 0.5 and the other's is less than 0.5. This four-payoff design allows clear comparative static predictions of the equilibrium effect of changing payoffs, based on PBE.

All decisions in these games are binary (fight vs. acquiesce for Player 1, fight vs. back down for Player 2), so the voting process with 5-person teams is straightforward. All team members cast simultaneous independent votes when it is their team's turn to make a decision. In the 5-person teams under majority rule, a team's decision is to fight if and only if at least three of the team members vote to fight; with unanimity rule, a team's decision is to fight unless every team member votes *not* to fight.[4] In the 1-person team, the single member's decision is binding, as in a typical 2-person game experiment, without any voting.

Similarly to the PBE, the team equilibrium is defined in terms of the fight probabilities of the teams deciding for each player - strong Player 1 (p_{1S}), weak Player 1 (p_{1W}), Player 2 (p_2) - but with the two additional effects of the random payoff disturbances and the collective choice rule. Formally, a team equilibrium is a solution to the following fixed point problem. Consider any possible team decision probabilities, $p = (p_{1S}, p_{1W}, p_2)$. This in turn implies expected payoffs for fighting or acquiescing for each type of Player 1 and for Player 2. Given these expected payoffs, the distribution of random payoff perturbations implies vote probabilities for each member of each team, which in turn, depending on how the voting rule aggregates votes into team decisions, imply a profile of team decision probabilities, $p' =$

[4]The unanimity rule is asymmetric. The opposite version of unanimity, which we did not study would have the team decision as fight if and only if all members voted to fight.

(p'_{1S}, p'_{1W}, p'_2). We define p to be a team equilibrium if and only if $p' = p$.

The results of the experiment have four main takeaways. First, the collective choice rule matters; outcomes are significantly different for teams operating under majority rule and unanimity rule, and the observed qualitative effects of the voting rule are consistent with team equilibrium. Second, a one-parameter logit specification of the team equilibrium model provides a close fit to the data and explains most of the qualitative patterns of behavior across treatments and games. Third, the data clearly reject the predictions of the PBE model. The *purely random model*, which completely disregards payoff and team effects and predicts that weak Player 1 and Player 2 will each choose IN or OUT with probability 0.5 *across all game variations*, actually fits the data better than PBE. Relatedly, and crucial to the theme of this paper, the team equilibrium model makes specific predictions for all games, voting rules, and player roles about whether the observed IN frequencies should be higher or lower than PBE, and these predictions are borne out in 88% of the cases. Fourth, teams are "more rational" than individuals in the sense that they are more likely to make optimal decisions, given the behavior of the opposing team, consistent with a wide range of other studies of team behavior in games (Charness and Sutter, 2012). On the other hand, outcomes in games played by teams are *not* closer to PBE than outcomes in games played by individuals, as a result of team equilibrium effects.

The rest of the paper is organized as follows. Section 2 explains the connection of our paper with related literature on behavioral game theory and crisis bargaining. Section 3 introduces the class of crisis bargaining games used in the experiment, defines, characterizes, and proves uniqueness of logit team equilibrium under different voting rules, and compares this to PBE. Section 4 describes the experimental design and procedures, provides computational solutions showing the qualitative properties of team equilibrium for all the treatments, and identifies five key research questions that the experiment is designed to answer. Section 5 describes the results of the experiment and the logit estimation of the the team equilibrium model and interprets these results in the light of the five main research questions that motivated the study. Section 6 examines the robustness of the results by analyzing the implications of four alternative models of behavior.

2 Related Literature

The crisis bargaining games we employ in our study follow an extensive line of research using these models to better understand the causes and resolution of international conflicts such as wars, and the role of threats, strategic deterrence, sanctions, alliances, and military interventions. Those studies, which all are in the unitary rational actor paradigm, include a combination of theoretical and empirical analysis. See for example Fearon (1994a), Fearon (1994b), Lewis and Schultz (2003), Signorino (1999), Morrow (1989), Smith (1995), Smith (1999).[5]

The approach of these studies to model state actors as if they are individual decision makers has been widely criticized, as scholars in the field of international relations have argued strongly for more realistic behavioral assumptions. Powell (2017) in particular calls for developing strategic models of international conflict from the bottom up by viewing individuals as the basic building block, but expresses concerns that this presents a difficult challenge theoretically, and so far nobody has successfully developed a tractable approach. Team equilibrium theory offers a tractable modeling framework for such applications, which is the focus of this paper.

The current paper ties in with the recent and expanding body of research that investigates experimentally the behavioral differences between teams and individuals in a wide range of game-theoretic and decision-making environments. See Kocher et al. (2020) and references therein.[6] The series of papers by David Cooper and John Kagel is the most closely related to our study. In Cooper and Kagel (2005), behaviors of individuals (1×1) and two-person teams (2×2) are examined by using the limit pricing game (Milgrom and Roberts, 1982). The team decision-making process unfolds as follows: both members are allotted three minutes to communicate and coordinate their actions. Once their choices align, and there is no alteration in decisions for four consecutive seconds, the team decision becomes binding. In cases where no coordination occurs within the three-minute time frame, one team member is randomly selected to implement their decision as the team's final decision. The authors

[5]The class of crisis bargaining games we study in this paper are similar to the final two-stage subgame of the crisis bargaining games analyzed in Fearon (1994b), Lewis and Schultz (2003), Signorino (1999), and Smith (1999). Signorino (1999) and Lewis and Schultz (2003) compare the properties of PBE and quantal response equilibrium in a slightly different class of crisis bargaining games, which foreshadows the use of payoff disturbances in our team equilibrium framework, except they treat teams as unitary actors.

[6]See also Charness and Sutter (2012) and Kugler et al. (2012) for the early papers in economics that compare teams and individuals.

demonstrate that teams exhibit greater strategic behavior compared to individuals. This is evident in their heightened ability to understand opponent players' incentives and responses, allowing them to adjust their behavior more effectively. As a result, teams are better at transferring their learning to games with different parameters, whereas individuals show no learning transfer between games.

Cooper and Kagel (2009) examine the differences in learning transfer between individuals (1×1) and two-person teams (2×2) in the limit pricing game by employing a similar design. In the experiment, the context is meaningful, for example, with players described as 'firms.' The framing changes as players are asked to choose either quantity or price against their opponent. A similar result is found, where teams consistently choose strategic behavior, while individuals exhibit limited learning across the games.

Although there have been many papers about groups playing games, an environment that corresponds to our design, in which team members vote for an action to determine the team's action, is scarce. One notable exception is Kim and Palfrey (Forthcoming), where behavior of games played by individuals (1×1) is compared to behavior with five-person teams (5×5) in variations of prisoners' dilemma and stag hunt games. Three collective choice procedures are studied: majority rule, majority rule preceded by a poll, and majority rule preceded by chat. They document significant bandwagon effects under the latter two procedures, which tend to generate within-team consensus. Their main results are twofold. First, in prisoners' dilemma games with relatively weaker incentives to defect, teams cooperate more than individuals, which is the opposite of findings from previous prisoners' dilemma experiments with teams. Second, in stag hunt games, teams consistently coordinate more frequently than individuals and, when coordination occurs, are more likely to achieve the payoff dominant outcome.

There is a growing experimental literature in political economy that examines the effect of different voting rules in strategic environments. This includes several studies that compare the effects of majority and unanimity voting rules in legislative bargaining environments (Baron and Ferejohn, 1989). Miller and Vanberg (2013) compare majority and unanimity voting rules in three-person committees. Consistent with theory, they find that the size of coalitions tends to be larger under the unanimity voting rule, while it takes longer for the committees with the unanimity voting rule to reach agreements. Miller and Vanberg (2015) study those two voting rules with committees with different sizes. The finding indicates that while there is no difference in delays between small and large committees under the unanimity voting rules, the majority rule leads to more frequent delays in large than in small committees. In the presence of communication, Agranov and Tergiman (2014) and Agranov and Tergiman (2019) use majority and unanimity voting rules in five-person committees, respectively. While majority voting with communication aligns more closely with theoretical predictions, the unanimity voting rule produces more egalitarian outcomes and less frequent delays in reaching agreements.

In the dynamic public goods environments, Battaglini et al. (2012) compare the effects of majority and unanimity voting rules on investments in the public good. They find that the unanimity voting rule leads to higher investment in the public good than the majority voting rule, which is consistent with the theoretical predictions. Battaglini et al. (2020) examine a dynamic environment that permits lending and borrowing across periods and introduces uncertainty about the value of the future public good. The main finding concerning different voting rules is that the efficiency of public good provision increases with the number of votes required to accept the proposal.

Majority and unanimity rules have also been compared in laboratory experiments in order to understand strategic voting behavior in the information aggregation problem known as the Condorcet jury problem.[7] The first paper to study this application in the laboratory is Guarnaschelli et al. (2000), which compares the effect of voting rules between majority and unanimity, with different group sizes and with/without of straw polls. They find that voters under the unanimity voting rule are more inclined to vote strategically than under the majority voting rule. This effect of unanimity on strategic voting is significantly diminished when the group can communicate via a straw poll, leading to outcomes closer to those observed under majority rule.[8]

3 Crisis Bargaining Games and Team Equilibrium

3.1 Crisis Bargaining Games

We model simple crisis bargaining situations using a sender-receiver signaling games, involving two players, 1 (sender, or first mover) and 2 (receiver, or second mover). A fair coin flip before the game starts determines whether 1 is either Strong or Weak, and this is private

[7]See Palfrey (2013) for a survey of experimental studies of jury voting games with information aggregation.

[8]Goeree and Yariv (2011) find a similar and even stronger effect of preplay communication if a richer communication protocol is used.

information to 1. Player 2 only knows that 1 is strong with probability 0.5. The game takes place in two stages. In the first stage, 1 makes a binary choice between the actions Escalate (IN) and Concede (OUT). If 1 chooses OUT, the game ends with 1 receiving a payoff of $-a$, and 2 receiving a payoff of a, where $a > 0$ is called the *concession payoff*.

If 1 instead chooses IN, the game moves to the next stage and it is 2's turn to choose between the actions Fight (IN) and Concede (OUT). If 2 chooses OUT, then 2 loses the concession payoff of $-a$ and 1 gains a. If 2 chooses IN, conflict ensues and payoffs depend on whether 1 is strong or weak. If 1 is strong, then 1 receives a payoff $a + s$ and player 2 receives a payoff of $-(a + r)$, where $r, s > 0$. If 1 is weak, then 1 receives a payoff of $-(a + s)$ and 2 receives a payoff of $a + r$. We call s the *sender risk* and r the *receiver risk*.

This game models a situation in which two parties, an informed sender and an uninformed responder, can choose - in sequence - between either accepting a certain small loss to the other party, or risking a larger loss (gain) if they are are weaker (stronger) than the other party. The motivating example is one where there is a crisis between two countries. One country (player 1) can threaten to escalate the crisis or suffer the a loss from conceding and letting the crisis be resolved in the second country's favor. If the crisis is escalated, the second country responds by either fighting or conceding.[9]

3.2 Team equilibrium

The experimental design, hypotheses, and analysis of results are guided by the theory of team equilibrium in games (Kim et al., 2022). In a team equilibrium, individual members of each team have rational expectations about the strategies of the other team, and share the same payoffs on average, but have unobserved i.i.d. payoff perturbations. We specify the payoff disturbances as being distributed according to an extreme value distribution with precision λ, so it corresponds to the logit specification of team equilibrium. At each information set,

[9]There are many other applications. For example, a union that is privately informed of the costs to their members of striking and accepting an unattractive contract first decides whether to threaten a strike, and the employer must decide whether to call the bluff or back down upon being threatened. We call such games 'simplified poker', because, in the special case where $r = s$, it corresponds to a poker-like card game, where the first mover is randomly dealt either a high or low card after each player has put in the pot an ante equal to a. The game also corresponds to the "simple card game" in Myerson (1991) where first mover is dealt either a high or low card and can either bet an additional r (IN) or fold (OUT). If she bets, then the second mover can either meet the bet (IN) or fold. If both players bet, then the first (second) player wins the pot if the card is high (low). Simplified poker games have a storied history in the theory of games, with many different variations analyzed by Bellman and Blackwell (1949), Borel and Ville (1938), Morgenstern and von Neumann (1947), Kuhn (1950), and others.

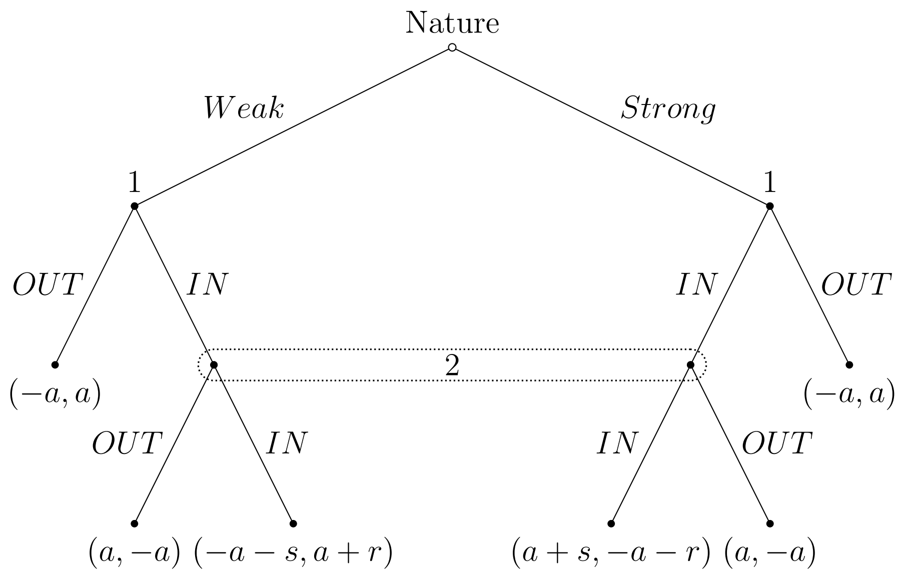

Figure 1: Game Tree of Binary Signaling Game.

each team chooses a strategy either by majority rule voting or unanimity rule. Because of the voting rule and the payoff disturbances, the collective choice rule can result in either more or less aggressive responses to threats, depending on the exact payoff structure in the crisis bargaining game.

3.2.1 Preliminaries

We define and characterize the (logit) team equilibrium for our family of crisis bargaining games as a function of the three payoff parameters, a, r, s, and the responsiveness parameter, λ, assuming equal-sized teams with an odd number, n, of members of each team.

A team equilibrium consists of a profile of behavioral strategies for each team, and a belief that the sender is strong, conditional on the sender choosing IN. Team 1 (the sender) has two information sets and Team 2 (the receiver) has only one information set, so a team equilibrium strategy profile is denoted by a triple, $(p_{1S}^{n*}, p_{1W}^{n*}, p_2^{n*})$, where p_{1S}^{n*} be the equilibrium frequency with which team 1 chooses I (IN) when Strong, p_{1W}^{n*} be the equilibrium frequency with which team 1 chooses I (IN) when Weak, and p_2^{n*} is the equilibrium frequency with which team 2 chooses IN after team 1 chooses IN. The equilibrium belief that the sender is strong, conditional on the sender choosing IN is denoted by μ^{n*}.

A team equilibrium is a fixed point mapping from the set of team mixed strategy profiles into itself, and any fixed point of this mapping is an equilibrium of the team game in the following way. Consider any profile of team (mixed) strategies and the corresponding beliefs of team 2 following an IN choice by team 1, where the beliefs are consistent with Bayes' rule. Any such profile implies expected payoffs for IN and OUT at each information set. Each member of a team whose turn it is to choose at that information set votes either IN or OUT. Because of the payoff disturbances of each voter, the strictly positive probability of each member of the team voting IN at that information set is given by the logit formula, applied using the expected payoffs of IN and OUT. The team decision probabilities then depend on how these votes are aggregated, which is specified by the collective choice rule (either majority rule or unanimity in our experiment). The resulting team decision probabilities define new expected payoffs, which in turn imply a new profile of team (mixed) strategies and beliefs. The profile is a team equilibrium if it is mapped into itself in this way. It is easy to show that a team equilibrium exists for any finite game. In crisis bargaining games, it also turns out that the team equilibrium (like the PBE) is always unique.

3.2.2 Equilibrium conditions for majority rule

We denote by the triple, $(v_{1S}^{n*}, v_{1W}^{n*}, v_2^{n*})$ the vote probabilities for individual members of a strong sender team, a weak sender team, and the responder team. respectively. In a (logit) team equilibrium, these individual member vote probabilities follow best responses. That is:

$$
\begin{aligned}
v_{1S}^{n*} &= \frac{e^{\lambda[p_2^{n*}(a+s)+(1-p_2^{n*})a]}}{e^{\lambda[p_2^{n*}(a+s)+(1-p_2^{n*})a]}+e^{-\lambda a}} \\
v_{1W}^{n*} &= \frac{e^{\lambda[p_2^{n*}(-a-s)+(1-p_2^{n*})a]}}{e^{\lambda[p_2^{n*}(-a-s)+(1-p_2^{n*})a]}+e^{-\lambda a}} \\
v_2^{n*} &= \frac{e^{\lambda[\mu^{n*}(-a-r)+(1-\mu^{n*})(a+r)]}}{e^{\lambda[\mu^{n*}(-a-r)+(1-\mu^{n*})(a+r)]}+e^{-\lambda a}}
\end{aligned}
$$

In case of $n > 1$, the majority decision probability of In is equal to the probability that more than $n/2$ individual voters vote for IN. These probabilities, are given by standard binomial formulas derived from v_{1S}^{n*}, v_{1W}^{n*}, and v_2^{n*}:

$$p_{1S}^{n*} = \sum_{j=\frac{n+1}{2}}^{n} \binom{n}{j} (v_{1S}^{n*})^j (1 - v_{1S}^{n*})^{n-j} \tag{1}$$

$$p_{1W}^{n*} = \sum_{j=\frac{n+1}{2}}^{n} \binom{n}{j} (v_{1W}^{n*})^j (1 - v_{1W}^{n*})^{n-j} \tag{2}$$

$$p_{2}^{n*} = \sum_{j=\frac{n+1}{2}}^{n} \binom{n}{j} (v_{2}^{n*})^j (1 - v_{2}^{n*})^{n-j} \tag{3}$$

In the case of $n = 1$, the individual member vote probabilities and the team majority rule decision probabilities coincide.[10] The fourth equation characterizing the team equilibrium is the Bayesian restriction on player 2's beliefs about player 1's type conditional on player 1 choosing IN:

$$\mu^{n*} = \frac{p_{1S}^{n*}}{p_{1S}^{n*} + p_{1W}^{n*}}.$$

3.2.3 Equilibrium conditions for unanimity rule

The unanimity rule we consider here specifies that the team decision is IN unless every individual member of the team votes for OUT. All of the analysis in the previous section continues to hold, with the exception of the last three equations, which become:

$$p_{1S}^{n*} = \sum_{j=1}^{n} \binom{n}{j} (v_{1S}^{n*})^j (1 - v_{1S}^{n*})^{n-j}$$

$$p_{1W}^{n*} = \sum_{j=1}^{n} \binom{n}{j} (v_{1W}^{n*})^j (1 - v_{1W}^{n*})^{n-j}$$

$$p_{2}^{n*} = \sum_{j=1}^{n} \binom{n}{j} (v_{2}^{n*})^j (1 - v_{2}^{n*})^{n-j}.$$

[10]When $n = 1$, the individual member vote probabilities are presented as follows:

$$p_{1S}^{*} = \frac{e^{\lambda[p_2^*(a+s)+(1-p_2^*)a]}}{e^{\lambda[p_2^*(a+s)+(1-p_2^*)a]} + e^{-\lambda a}} \tag{4}$$

$$p_{1W}^{*} = \frac{e^{\lambda[p_2^*(-a-s)+(1-p_2^*)a]}}{e^{\lambda[p_2^*(-a-s)+(1-p_2^*)a]} + e^{-\lambda a}} \tag{5}$$

$$p_{2}^{*} = \frac{e^{\lambda[\mu^*(-a-r)+(1-\mu^*)(a+r)]}}{e^{\lambda[\mu^*(-a-r)+(1-\mu^*)(a+r)]} + e^{-\lambda a}} \tag{6}$$

3.2.4 Uniqueness of team equilibrium in crisis bargaining games

In this section, we prove that team equilibrium is unique in crisis bargaining games for both majority rule and unanimity rule. We first prove this for the case of $n = 1$, where team equilibrium under either majority rule or unanimity rule are the same. The results for (odd) $n > 1$ follow immediately.

Proposition 1. *Any crisis bargaining game with parameters $a, r, s > 0$ and $n = 1$ has a unique logit team equilibrium for all $\lambda \geq 0$.*

Proof. The result is trivial for $\lambda = 0$, so consider any $\lambda > 0$. Existence is guaranteed by Brouwer's fixed point theorem, so let $p* = (p^*_{1S}, p^*_{1W}, p^*_2)$ be a team equilibrium of the game. Suppose $p^{*\prime} = (p^{*\prime}_{1S}, p^{*\prime}_{1W}, p^{*\prime}_2) \neq p^*$ is another team equilibrium with $p^{*\prime}_{1S} > p^*_{1S}$.[11] Then $p^{*\prime}_2 < p^*_2$ by the following argument. Fixing p_{1S} at any value $\bar{p} \in (0, 1)$, consider player 2's logit reply function to p_{1W}. This must satisfy equation (6), so it is strictly increasing in p_{1W} and exactly equal to 1/2 when $p_{1W} = \frac{\bar{p}r}{2a+r}$. Similarly, consider $1W$'s logit reply function to p_2. This must satisfy equation (5), so it is strictly decreasing in p_2, independent of $p_{1S} = \bar{p}$, and exactly equal to 1/2 when $p_2 = \frac{2a+r}{2a+2r}$. Thus, (p^*_{1W}, p^*_2) is the unique intersection point of these two logit reply functions, when $\bar{p} = p^*_{1S}$. Now consider the logit reply functions for $1W$ and 2 when $\bar{p} = p^{*\prime}_{1S} > p^*_{1S}$. Because $1W$'s logit reply function to p_2 is independent of p_{1S}, it is unchanged. However, 2's logit reply function to p_{1W} shifts, since μ is increasing in p_{1S} for every value of p_{1W}, thereby reducing 2's expected payoff of IN. The logit reply function to p_{1W} for 2 now equals 1/2 when $p_{1W} = \frac{p^{*\prime}_{1S}r}{2a+r} > \frac{p^*_{1S}r}{2a+r}$. Hence, the intersection point of the logit reply functions for $1W$ and 2 when $\bar{p} = p^{*\prime}_{1S}$ has $p^{*\prime}_{1W} > p^*_{1W}$ and $p^{*\prime}_2 < p^*_2$. However, $1S$'s logit reply function in (4) is strictly increasing in p_2 and independent of p_{1W}. This implies that $p^{*\prime}_{1S} < p^*_{1S}$, a contradiction.

□

The uniqueness for $n > 1$ with majority rule and unanimity rule follows immediately, since the team strategy probabilities, $p^n_{1S}, p^n_{1W}, p^n_2$, are strictly increasing in the vote probabilities, $v^n_{1S}, v^n_{1W}, v^n_2$, and the logit reply functions for voting defined by equations (1)-(3) have the same monotonicity properties as equations (4)-(6) for the case of $n = 1$.[12]

[11] The argument that follows is similar if we suppose $p^{*\prime}_{1S} < p^*_{1S}$.

[12] In fact, uniqueness will hold for all collective choice rules that are monotone strictly increasing in the expected payoff of IN.

3.2.5 Perfect Bayesian Equilibrium as the limit of team equilibrium

In any crisis bargaining game there is a unique totally mixed Perfect Bayesian Equilibrium (PBE). The PBE of the crisis bargaining game consists of a 4-tuple, $(p_{1S}^{PBE}, p_{1W}^{PBE}, p_2^{PBE}, \mu^{PBE})$ where p_{1S}^{PBE} is the probability a strong player 1 chooses IN, p_{1W}^{PBE} is the probability a weak player 1 chooses IN, p_2^{PBE} is the probability 2 chooses IN in response to IN, and μ^{PBE} is 2's belief that 1 is strong, conditional on 1 choosing IN, derived from Bayes' rule.

The strong player 1 has a dominant strategy to choose IN, so $p_{1S}^{PBE} = 1$. Equilibrium requires that weak player 1 mix by choosing IN with probability p_{1W}^{PBE} such that 2 is indifferent between responding IN or OUT. Similarly, player 2 mixes to make weak player 1 indifferent between IN and OUT. It is easy to see that the PBE is pinned down by three equations:

$$\begin{aligned} p_2^{PBE}(-a-s) + (1-p_2^{PBE})a &= -a \\ \mu^{PBE}(-a-r) + (1-\mu^{PBE})(a+r) &= -a \\ \mu^{PBE} &= \frac{1}{1+p_{1W}^{PBE}} \end{aligned}$$

and hence the PBE strategies are: $p_{1S}^{PBE} = 1$; $p_{1W}^{PBE} = \frac{r}{2a+r}$; $p_2^{PBE} = \frac{2a}{2a+s}$.

A natural question to ask is whether, holding fixed the number of members on each team, the unique team equilibrium converges to the PBE in crisis bargaining games as the logit response parameter, λ, diverges to ∞. Here we show that this convergence result also holds for both the majority and unanimity voting rules.[13]

Proposition 2. *Fix n. For both the majority voting rule and unanimity voting rule, the team equilibrium in crisis bargaining games converges to the PBE as $\lambda \to \infty$. That is:*

$$\lim_{\lambda\to\infty} (p_{1S}^{n*}(\lambda), p_{1W}^{n*}(\lambda), p_2^{n*}(\lambda)) = (1, \frac{r}{2a+r}, \frac{2a}{2a+s})$$

Proof. We prove the result for unanimity rule. The proof for majority rule is similar.

It is easy to see that $p_{1S}^{n*}(\lambda) \to 1$ since IN is strictly dominant for strong player 1. It is also clear that if $p_{1W}^{n*}(\lambda) \to \frac{r}{2a+r}$ then $\mu^{n*}(\lambda) \to \frac{r}{2a+r}$. Hence, we only need to show that $(p_{1W}^{n*}(\lambda), p_2^{n*}(\lambda)) \to (\frac{r}{2a+r}, \frac{2a}{2a+s})$. Suppose to the contrary that $p_{1W}^{n*}(\lambda) \to p > \frac{r}{2a+r}$. Then $\mu^{n*}(\lambda) \to \mu < \frac{2a+r}{2a+2r}$ so IN is strictly better than OUT for player 2. Hence, $v_2^{n*}(\lambda) \to 1$,

[13] In fact, the convergence result holds for general threshold voting rules that require at least m out of n voters to vote for IN in order for IN to be the team decision, for $m = 1, 2, ..., n$. Moreover, the threshold does not have to be the same for both teams.

which in turn implies that $p_2^{n*}(\lambda) = \sum_{j=1}^{n} \binom{n}{j}(v_2^{n*}(\lambda))^j(1 - v_2^{n*}(\lambda))^{n-j} \to 1$. But $p_2^{n*}(\lambda) \to 1$ implies that OUT is a strict best response for weak player 1, so $v_{1W}^{n*}(\lambda) \to 0$, implying that $p_{1W}^{n*} = \sum_{j=1}^{n} \binom{n}{j}(v_{1W}^{n*})^j(1 - v_{1W}^{n*})^{n-j} \to 0$, a contradiction. A similar logic reaches a contradiction, if we suppose that $p_{1W}^{n*}(\lambda) \to p < \frac{r}{2a+r}$ or if we suppose that $p_2^{n*}(\lambda) \to p \neq \frac{2a}{2a+s}$.
□

Given this convergence, the PBE of the game serves as a natural benchmark for comparing team and individual behavior. It motivates the hypothesis that multi-member teams will exhibit behavior more closely aligned with PBE predictions than individuals.[14] However, since these games lack a general monotonicity property in convergence to PBE, it is theoretically possible—particularly for small or intermediate values of λ, which are common in experiments—that the team equilibrium could deviate further from the PBE in multi-member settings.

4 Experimental Design and Research Questions

4.1 Games and Treatments

We implement a 3×4 treatment design that varies both the team structure and the strategic environment. The three team treatments differ in the number of players and the decision rule used within each team. In the Individual treatment, individual participants compete in a 1v1 version of the game. In the Majority treatment, teams of five players make decisions via majority voting. In the Unanimity treatment, teams of five make decisions via unanimity voting, with the default action set to IN if consensus is not reached.

The four games differ in their payoff parameters, leading to distinct equilibrium IN probabilities for the weak first mover and the second mover. These probabilities vary not only in magnitude but also in their relative ordering across players. Crucially, the set of games is constructed to encompass all four possible orderings of these probabilities, ensuring comprehensive coverage of strategically diverse environments. Table 1 presents the specific payoff parameters, a, s, r, and the corresponding PBE mixed strategies for each game. For example, in the LH game (Game 1), the equilibrium IN probability is below 0.5 for the weak first mover and above 0.5 for the second mover.

[14]This hypothesis is also supported by evidence on team behavior in a wide range of games, as documented in the survey by Charness and Sutter (2012).

Table 1: Game Payoffs and PBE Mixed Strategies

	Concession payoff (a)	Sender risk (s)	Receiver risk (r)	p_{1W}^{PBE}	p_2^{PBE}
Game 1 (LH)	10	5	5	$\frac{1}{5}$	$\frac{4}{5}$
Game 2 (LL)	4	16	4	$\frac{1}{3}$	$\frac{1}{3}$
Game 3 (HL)	3	15	15	$\frac{5}{7}$	$\frac{2}{7}$
Game 4 (HH)	4	4	16	$\frac{2}{3}$	$\frac{2}{3}$

This design allows us to test how team decision-making responds to variations in strategic incentives, including cases where one or both players are predicted to play IN with high or low probability. Each game thus serves as a distinct testbed for comparing individual and group behavior relative to PBE benchmarks.

4.2 Team Equilibrium Predictions

4.2.1 Convergence and Non-monotonicity

In this subsection, we illustrate the computation of the team equilibrium for Game 1 (LH) across the three team treatments and two roles. We also provide an explanation of the equilibrium's comparative statics for this game as a representative example. As established in Proposition 2, team equilibrium frequencies of the action IN always converge to the PBE frequencies as $\frac{1}{\lambda}$, the variance of team members' payoff disturbances, approaches zero. In this case, the PBE mixing probabilities are $p_{1W}^{PBE} = 1/5$ and $p_2^{PBE} = 4/5$. This example highlights how convergence to the PBE can be non-monotonic: for small values of λ, the logit team equilibrium choice probabilities initially move further away from the PBE as λ increases. Moreover, the team decision rule (majority versus unanimity) can systematically influence the logit team equilibrium by biasing team choices toward one of the actions.

The left panel of Figure 2 shows how the Game 1 (LH) team equilibrium IN frequency for weak first movers, p_{1W}^* varies with λ, for each of the three team treatments. The right panel shows the same for second movers, p_2^*. Each curve traces out the logit team equilibrium for values of λ between 0 and 2, on the x-axis.[15] We will first discuss the comparative statics of the team equilibrium model with respect to λ for majority rule teams and individuals,

[15]Appendix 2 contains analogous figures for Games 2-4.

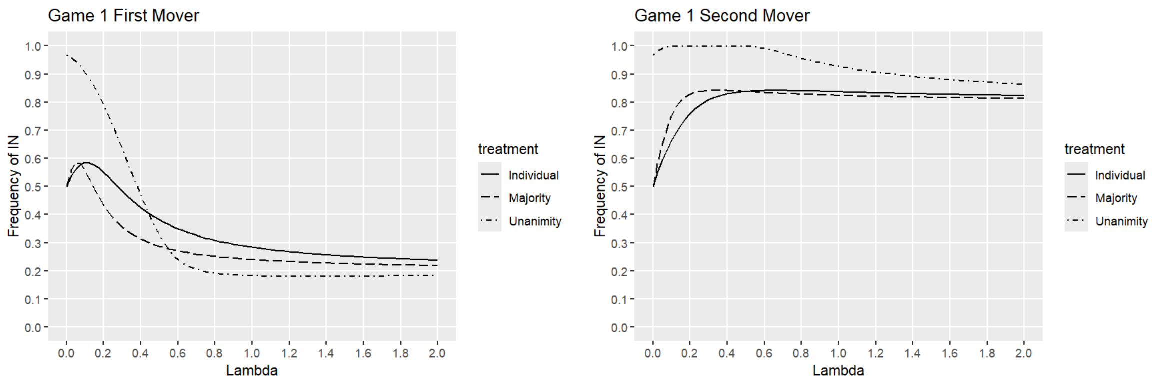

Figure 2: Game 1 Team Equilibrium

then the equilibrium effects of the majority rule treatment, and finally discuss the unanimity treatment.

The figure shows how team equilibrium always converges to the game 1 PBE of 1/5 for weak first movers and 4/5 for second movers as λ goes to infinity, irrespective of the choice rule. At $\lambda = 0$, team members are completely unresponsive to payoffs and mix uniformly in their voting behavior between IN and OUT, resulting in a team frequency of 1/2 for majority rule and individuals.[16]

The non-monotonicity of convergence to PBE for Weak first movers arises in Game 1 (LH) because, for values of λ close to 0, the expected payoff of IN is approximately $\frac{1}{2}a + \frac{1}{2}(-a - s) = -\frac{1}{2}s$, which is strictly greater than the OUT payoff of $-a$ since $s = 5$ and $a = 10$. Team members become more responsive to expected payoff differences as λ increases, and so in this game, they begin to vote for IN at higher rates as λ increases from 0. As λ increases, the second mover behavior changes as well, increasing from $\frac{1}{2}$ until crossing the PBE level of 0.8. At that point, the expected payoffs of IN and OUT become equal for Weak first movers by the definition of equilibrium, and so p^*_{1W} exactly equals $\frac{1}{2}$.

Understanding the comparative statics of p^*_2 with respect to λ is more subtle, as one needs to consider the effects on 2's incentives resulting from changes in a 50/50 *mixture* of Strong and Weak type first movers choosing IN. The more Weak types go IN relative to Strong types, the greater the incentive for second movers to also go IN. Because strong and weak first mover team members behave identically at $\lambda = 0$, a Bayesian second mover does not gain any information from the fact that the first mover team chose IN. The probability of

[16]Not displayed in the graph is the team equilibrium IN frequency for Strong first movers, which also starts out at 1/2 for majority rule and individuals, and very quickly converges to 1.

strength conditional on choosing IN is the same as the prior. Consequently, the second mover expected payoff from going IN against a first mover with $\lambda = 0$ is $\frac{1}{2}(a + r) + \frac{1}{2}(-a - r) = 0$, which is greater than the payoff of $-a$ from choosing OUT. For infinitesimal increases in λ we can ignore the effect of changes in first mover team behavior on expected payoffs, and conclude that second mover team members become more likely to vote IN.

In Game 1 (LH), under the majority rule and individual decision treatments, weak first movers and second movers both converge to the PBE from above. This occurs because, under these choice rules, teams are more likely in any team equilibrium to select the action with the higher expected payoff. Whenever a team equilibrium frequency crosses that team's PBE, the corresponding opponent's equilibrium frequency must cross $\frac{1}{2}$ at that value of λ, since the PBE mixing probability equalizes the opponent's expected payoffs. In this game, the second mover's PBE frequency is above $\frac{1}{2}$. Therefore, for large values of λ, IN must yield a higher expected payoff that OUT for 2, so weak first movers must choose IN too frequently relative to the PBE. Similarly, since the first mover's PBE frequency is below $\frac{1}{2}$, second mover teams also choose IN too often in equilibrium. As a result, both roles select IN more frequently than prescribed by the PBE when λ is large.

In the unanimity treatment, second mover teams begin above the PBE when $\lambda = 0$ because the voting rule biases team behavior toward choosing IN. Specifically, at $\lambda = 0$, all team members mix uniformly between voting for IN and OUT, resulting in a team-level IN frequency of 31/32, which exceeds the second mover PBE level in all of our games. As a result, OUT becomes the best response for weak first movers, and the probability that any team member votes for IN must decline as λ increases. As noted earlier, second mover team members gain no information from observing that $\lambda = 0$ first movers choose IN, which leads second mover team equilibrium IN frequencies to initially increase with λ at low values. As shown in the right panel of Figure 2, second mover teams choose IN with near certainty at low λ, and only begin to noticeably decrease toward the PBE once first movers are already close to the PBE level of 1/5.

Under the unanimity rule and for large values of λ, team members may need to vote for IN with very low probabilities in order for the team-level frequency of IN to approach the PBE, even when the PBE exceeds $\frac{1}{2}$. Because team members are more likely to vote for the action with the higher expected payoff, OUT must yield a higher expected payoff than IN in this regime for all of our games. As a result, weak first mover teams under the unanimity rule must converge to the PBE from below, while second mover teams must converge from above.

4.2.2 Equilibrium effect

We now explain the divergence between equilibrium behavior under the majority rule and individual decision-making. In the team equilibrium model, voters are assumed to vote sincerely for the team action they prefer, based on their individual expected payoff, which includes a random disturbance. These disturbances are assumed to follow the same distribution regardless of team size or the decision rule used. Consequently, for a given value of λ and fixed expected payoffs for IN and OUT, the probability that a team member votes for IN is identical to the probability that an individual decision maker chooses IN unilaterally. Under majority rule, however, these individual choice probabilities are amplified through a Condorcet jury-type mechanism. When each team member votes for IN with a probability less than (greater than) $\frac{1}{2}$, the probability that IN wins the team vote is pushed closer to 0 (or 1). We refer to this as the reinforcement effect of majority rule voting.

In Figure 2, near $\lambda = 0$ the dashed line representing the majority rule team equilibrium increases faster than the solid line representing the individual team equilibrium, due to this reinforcement effect. In addition to this effect, there is an 'equilibrium effect' created by the influence of majority rule voting on the opponent team behavior. Since second mover majority rule teams reach the PBE level of IN frequencies at lower values of λ than individuals, the majority rule Weak first mover IN frequency must begin to decrease at lower values of λ than do individuals. This means that majority rule teams vote for IN more frequently than individuals go IN for some values of λ, and less frequently for higher values of λ. A similar effect can be seen in the team equilibrium of the second movers, however the separation between treatments is less pronounced for second movers than for first movers for this particular set of payoffs.

For sufficiently large values of λ, majority rule teams are closer to PBE than individuals. In this way our theory captures the typical result of the prior literature that teams are closer to equilibrium than individual decision makers. For these crisis bargaining games, this result only holds if subjects have sufficiently small payoff disturbances and use a neutral voting rule such as majority rule voting. With larger payoff disturbances, in other words, with more heterogeneity between team members, or non-neutral voting rules, team equilibrium effects may drive teams either further or closer to equilibrium depending on payoff parameters. These observations lead naturally to the research questions we present in the next section.

4.3 Research Questions

Since team equilibrium serves as a benchmark model for the games used in this paper, the following first question is of primary importance:

Question 1: How well (or poorly) does team equilibrium account for the variation of behavior across the treatments?

We now turn to the role of collective choice rules in shaping behavior. In previous experiments, the most commonly employed collective decision-making procedure has been a consensus rule, where team members reach a joint decision through deliberation, typically via face-to-face discussion or online chat. Given this relatively uniform approach in the literature, our comparison of different collective choice rules, specifically majority and unanimity, in strategic interactions represents a novel contribution. The team equilibrium framework predicts systematic effects of the voting rule on team choices and outcomes, which is directly related to Question 2. The significant behavioral differences observed under the two choice rules highlight the limitations of the unitary actor approach and emphasize the importance of explicitly modeling internal team decision-making processes.

Question 2: Does the collective choice rule matter for team behavior in crisis bargaining games and, if so, how?

A common claim in the literature, supported by previous experimental findings, is that teams behave more "rationally" than individuals. In a decision-theoretic context, this claim can be evaluated straightforwardly by assessing whether teams are more likely to make optimal choices. However, in strategic settings such as our crisis bargaining games, the question becomes more complex because the optimality of a team's decision is endogenous and depends on the strategy chosen by the opposing team.

To address this, we focus on responsiveness to differences in empirical expected payoffs, given the empirical strategies of the opponent. According to the team equilibrium model, a weak type's expected utility either increases or decreases monotonically with their own probability of choosing IN, holding the opponent's behavior constant. Thus, by examining whether teams or individuals are more responsive to these differences in expected payoffs when choosing IN or OUT, we can address the following question.

Question 3: Are teams more rational than individuals in crisis bargaining games?

A common corollary of the claim that teams exhibit superior rationality is that, in many applications, teams are expected to produce outcomes closer to the Nash equilibrium than individuals. However, the predictions of team equilibrium suggest that this is not always the case, due to the potential non-monotonicity in convergence to the PBE. This leads to the following question, which helps clarify which model is more appropriate for explaining behavior in teams.

Question 4: Are the outcomes of crisis bargaining games played between teams closer to PBE than the same games played between individuals?

There are other alternative models to team equilibrium that could potentially explain the data as well or better than team equilibrium. The most obvious candidate is PBE, which makes clear unique predictions that vary with the payoff parameters of the crisis bargaining game, but do not vary with the collective choice rule. Because PBE is nested in team equilibrium, it is easy to test whether or not PBE can be rejected. We also consider the logit quantal response equilibrium (QRE) as an alternative to team equilibrium. Finally, all of these models assume risk neutrality, but crisis bargaining games involve players choosing between a same alternative (OUT) and a risky alternative (IN), so we also examine whether risk aversion is useful in as a possible explanation of the data. This leads to the final question we address:

Question 5: Do other models of behavior explain the variation of behavior across the treatments, perhaps even better than team equilibrium?

4.4 Procedures

We conducted a total of 19 sessions with subjects recruited either from the Experimental Social Science Laboratory (ESSL) at UC Irvine, or from UC Santa Barbara. In each session, one team treatment was implemented. The individual treatment was conducted in 8 sessions, the majority treatment in 6 and the unanimity rule treatment in 5.

In each session subjects participated in 10 rounds of each of the 4 games for a total of 40 rounds. In half of the sessions, the order of games was 1234, and in the other half of sessions, the reverse order, 4321, was used. Subjects were randomly re-matched with new team members and a new opponent team between every round. Subject roles and the team treatment was held fixed for all rounds in a session. The subject interface and software for the experiment was programmed in oTree (Chen et al., 2016).

At the beginning of each session the experimenter read the instructions aloud, including specific payoff information for the first game of the session, and displayed examples of the subject interface on a projector in front of the room.[17] After these instructions were finished, a short comprehension quiz was completed by the subjects. Subjects were required to provide correct answers to all questions before moving on to the first round of the experiment. After the 10 rounds of the first game were finished, the experimenter announced the new payoffs for the second game. Subjects then participated in 10 rounds of the second game, and so forth until all 40 rounds of the experiment were completed.

In each round of the experiment, the game was conducted sequentially as illustrated in Figure 1. First movers were initially informed of the result of a virtual coin flip, which determined whether the first mover was strong (Heads) or weak (Tails).

In the 1×1 session where each team was composed of a single individual, first movers were then prompted to choose an action, either IN or OUT. Second movers then observed their paired first mover's action choice, but not the result of the coin flip. If the first mover chose IN then second movers were prompted to make a choice. If the first mover chose OUT, the game ended without the second mover making a choice. After all decisions were made, feedback about the outcome was given to all subjects and second movers learned the state of the world.

In the 5×5 session, first movers were then prompted to *vote* for an action, either IN or OUT, with the first mover team's decision decided by the voting rule, either majority or unanimity, depending on the session. Second movers then observed their paired first mover team's decision, but not the result of the coin flip. If the first mover team chose IN then second movers were prompted to vote for either IN or OUT, with the second mover team's decision decided by the voting rule, either majority or unanimity, depending on the session. If the first mover team chose OUT, the game ended without the second movers making a choice. After all decisions were made, feedback about the outcome was given to all subjects.

[17] A sample of the instructions for a reverse-order session, and the corresponding Powerpoint screens, are provided in the supplemental online appendix.

Second movers learned the state of the world, and in team treatments subjects were told the vote totals for both teams.

After the conclusion of all 40 rounds, one round from each game (total of 4 rounds) was randomly chosen for each subject to determine the payments. Sessions lasted on average 90 minutes, including instructional time. Subject earnings averaged $32.1, which included a fixed payment of $7 for showing up on time and a completion payment of $5 for completing all rounds of the experiment. Subjects were paid for their decisions in an artificial currency, points, where each point had a value of $0.20.[18]

5 Results

5.1 Descriptive summary

5.1.1 Behavior of strong first movers

Strong first movers face a strategically trivial decision task, since choosing IN is a strictly dominant strategy for strong first movers. This is reflected in our data. In the Majority and Unanimity treatments, strong first-mover teams select IN 100% of the time, and in the Individual treatment IN is chosen in more than 96.7% of cases. Because behavior is essentially degenerate for strong first movers, our analysis focuses on weak first movers and second movers, where strategic uncertainty is central.

5.1.2 Behavior of weak first movers and second movers

The strategic calculus is nontrivial for both the weak first movers and the second movers. First movers find it optimal to "bluff" by choosing IN if they expect second movers are sufficiently likely to choose OUT, and that threshold level of expectation varies across the

[18]When the initial $1-1$ sessions were conducted, in-person experiments were not feasible due to the pandemic, so these sessions were conducted online through Zoom video conferencing. Online sessions are clearly identified in Table 3 in Appendix 1. In these sessions, occasionally a subject would become disconnected during the experiment, so it was necessary to address this issue by substituting a 'robot' stand-in player to replace a dropout player. If a player became disconnected during the session, the session was briefly paused to allow the experimenter to make a public announcement informing all remaining subjects. A robot player then took the place of the dropout, and chose IN with 50% probability and OUT with 50% probability at every opportunity, for the remainder of the session. Subjects were subsequently advised in any future round when they were paired with the robot, via a private announcement. It was announced publicly that rounds played against the robot player would not be selected for subjects' payoffs. The data from any such unincentivized games were discarded. In session 10 (the only online 5×5 session), a subject disconnection caused the experimental software to crash immediately following Round 20, so that session did not include any data from two of the games. See Table 3 in Appendix 1.

four games we study. Similarly, it is only optimal for second movers to choose OUT if they think they believe weak first movers are choosing IN with sufficiently low probability, and the threshold level of second mover beliefs varies across the four games as well.

The Perfect Bayesian Equilibrium (PBE) IN frequencies provide an obvious natural benchmark To compare the behavior of weak first movers and second movers across the four games. This comparison was the basis for choosing the four games in the design of the experiment, such that the PBE probability of choosing IN varied systematically across the games and between first and second movers.

Unfortunately, PBE has little predictive power across the games and treatments. Figure 3 displays bar plots of the observed IN vote frequencies and IN team decision frequencies for weak first movers (left panel) and second movers (right panel), with the PBE predictions superimposed as dashed lines. If behavior were consistent with PBE and there were no effect of team size or voting rule, and all of the bars should align with the dashed lines. Instead, the data reveal systematic and significant deviations from PBE and variation across the treatments within games.

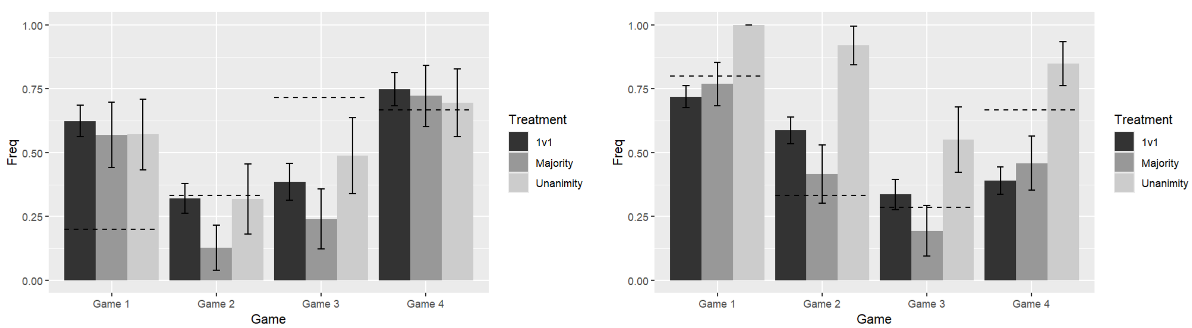

(a) Weak First Mover Team Decisions

(b) Second Mover Team Decisions

Figure 3: Empirical IN frequencies compared to PBE predictions.

As shown in the left panel of Figure 3, weak first movers' IN decision frequencies often diverge sharply from PBE predictions. In the Individual treatment, choices differ significantly from PBE in two of the four games at the 1% level.[19] In the Majority and Unanimity treatments, deviations are significant in three and two games, respectively.

The right panel of Figure 3 shows that second movers also deviate from PBE across treatments. In the Individual treatment, IN frequencies differ significantly from PBE in

[19]Unless otherwise noted, significance is based on binomial tests.

three of the four games, and in the Unanimity treatment the deviations are significant in all four. Majority-rule teams show smaller but still systematic departures, with significant deviations in Game 4.

Overall, the evidence demonstrates that PBE does not organize the data well. In both roles and across treatments, observed play frequently overshoots or undershoots the PBE mixing probabilities, and the direction of deviation varies with the collective choice rule. This lack of systematic alignment underscores the need for an alternative model—such as team equilibrium—to account for behavior.

5.2 How well does Team Equilibrium account for variations of behavior across treatments?

While the PBE framework fails to account for the systematic deviations in our data, the team equilibrium (TE) model provides a remarkably successful explanation. The core idea of TE is that team members evaluate expected payoffs with small idiosyncratic payoff disturbances and cast votes accordingly, with collective choice rules aggregating these votes into team decisions.

As a first step, in Section 5.2.1 we fit the data to the (logit) team equilibrium model, estimating a single logit responsiveness parameter, λ, to the entire data set and compare the predicted values of IN team decision frequencies for each of the 24 *game* (1-4) – *treatment* (ind/maj/unan) – *role* (first, second) combinations to the observed decision frequencies. Thus the team equilibrium model must simultaneously capture variation across different payoff structures, team sizes, voting rules and player roles—a highly demanding test of the model.

The second part of this section evaluates whether the directional predictions of team equilibrium deviations from PBE are borne out in the data. As seen in Figure 3, the observed team equilibrium theoretical decision frequencies differ quite dramatically from the PBE IN frequencies, and these differences vary systematically across the four games, two player roles, and three choice rule treatments. In some cases the team equilibrium predicts higher IN frequencies than PBE and in other cases it predicts lower IN frequencies than PBE. Section 5.2.2 documents the extent to which the data reflect these patterns.

5.2.1 Estimating the logit team equilibrium model

The estimation is conducted using the individual choice data for weak first movers and second movers.[20] The estimated individual vote frequencies imply predicted team equilibrium decision frequencies which we then compare to the empirical decision frequencies. The results are striking. Figure 4 plots observed frequencies of IN decisions against TE predicted values for all 24 game–treatment–role combinations. If TE organized the data perfectly, the points should all lie along the 45-degree line. This is almost exactly what happens: the fitted relationship is extremely tight (intercept = 0.04, slope = 0.85, R^2 = 0.76). Given the heterogeneity of our design (four distinct payoff environments, three decision rules, and two roles), such alignment is exceptional. TE delivers a coherent account of the data with a single free parameter.

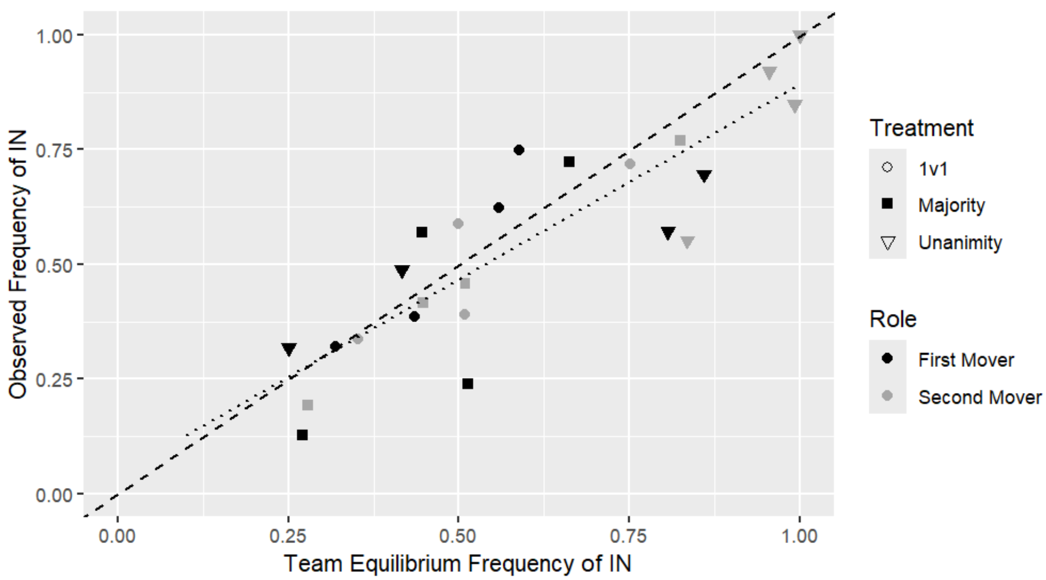

Figure 4: Observed versus predicted frequencies of IN under the fitted team equilibrium model. Each point corresponds to one game × treatment × role.

The contrast with PBE is stark. Figure 5 plots observed frequencies of IN against PBE predictions. The fitted regression line has an intercept of 0.36 and a slope of only 0.36, which is not statistically different from zero, with an R^2 of just 0.12. In short, there is only a weak and noisy relationship between PBE and observed play. A breakdown by role underscores

[20] Inclusion of strong first mover data leads to almost identical results, since, except for implausibly low values of λ, the logit equilibrium IN frequencies are very close to 1, as is the data. See the appendix for the estimation results including the strong first mover observations.

this point: for second movers the slope rises modestly to 0.57 ($R^2 = 0.26$), while for first movers it falls to 0.12 ($R^2 = 0.02$), and neither differs significantly from zero.

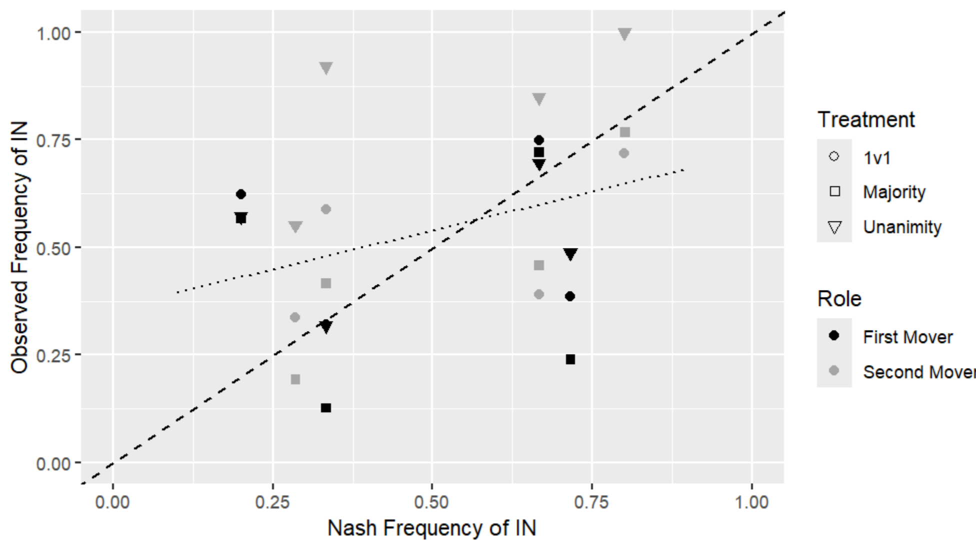

Figure 5: Observed versus PBE-predicted frequencies of IN. Each point corresponds to one game × treatment × role.

This side-by-side comparison delivers a clear message: team equilibrium dramatically outperforms PBE as an explanatory framework for crisis bargaining behavior. With just one parameter, TE accounts for the broad patterns of play, the systematic treatment effects, and the sign of departures from PBE. In the subsections that follow, we show that this success is not accidental but systematically reinforced by the ways in which collective choice rules shape outcomes and by the greater payoff responsiveness of teams relative to individuals.

5.2.2 Qualitative predictions of team equilibrium departures from PBE

Equally important, TE captures not only aggregate frequencies but also the comparative statics across treatments. Because there are so many predictions about the treatment, game, and role effects, a useful lens for focusing and organizing these comparative static predictions is in terms of the TE-predicted qualitative departures from PBE.

How do the observed deviations from the PBE IN frequencies compare with the differences between the TE predictions and the PBE frequencies; in particular, do the observed positive (negative) deviations from PBE coincide with cases where the TE-predicted deviations are also positive (negative)? The data follow exactly this pattern. Across all games, roles, and voting rules, 21 out of 24 TE predictions about whether observed behavior should lie above or below the PBE benchmark are correct.

Furthermore, in 12 out of 24 cases, the ordering between team equilibrium and PBE is *invariant* with respect to the payoff-noise parameter λ. That is, for all admissible λ, the team equilibrium mixing probability is strictly above the PBE value or strictly below it. In 100% of these 12 cases, the sign of the observed deviation from PBE coincides with this "global" TE prediction. When the team equilibrium lies above PBE for all values of λ, observed frequencies exceed the PBE prediction; when it lies below, observed frequencies fall short of the PBE prediction. This alignment is meaningful because it does not rely on structural estimation at all; it reflects robust, parameter-free comparative statics.

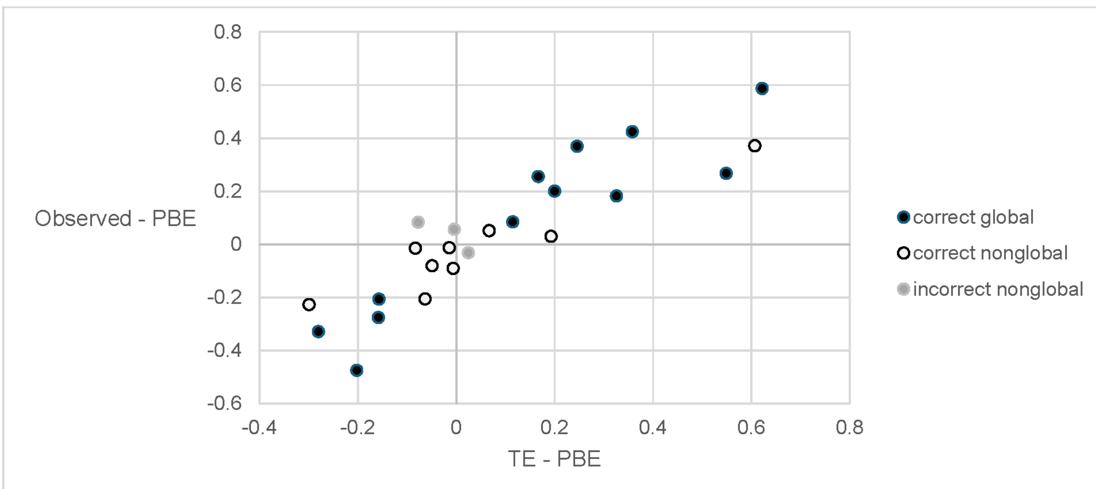

Figure 6: Observed deviation from PBE IN frequencies of IN (horizontal axis) plotted against the predicted difference between TE and PBE IN frequencies. Each point corresponds to one game × treatment × role.

In 9 of the remaining 12 cases, the ordering between the TE-predicted IN frequencies obtained from the structural estimation of team equilibrium and PBE is correctly mirrored in the observed IN frequencies. The only three exceptions occur for Player 1 in Game 4 (1x1 and Majority) and Player 2 in Game 1 (Majority). Both the TE-predicted and the observed deviations from PBE IN frequencies are very small (less than 0.09 in all three cases). In such cases the model does not imply a significant sign prediction.

A final observation relates to the magnitude of the departures from PBE and how these magnitudes line up closely with the TE-predicted deviations. The scatter diagram of all 24 cases in Figure 6 shows the strong positive correlation between the observed departures from PBE and the TE-predicted deviations.[21] The solid dark circles are the 12 cases where TE makes a global prediction about the direction of deviations from PBE. Three cases where the direction of observed deviations does not match the TE prediction are indicated in grey, and are all located very close to (0,0). The remaining 9 cases are displayed as open circles. While our data clearly reject the notion that teams are systematically closer to PBE than individuals, they support a stronger claim: team equilibrium organizes the sign of deviations from PBE in a consistent and theoretically grounded manner. When the equilibrium benchmark is misspecified, greater rationality does not move teams closer to PBE. It moves behavior toward the comparative statics captured by team equilibrium.

5.3 Do Collective Choice Rules Matter for Behavior and Outcomes?

We now turn to the role of collective choice rules. Whereas most previous experiments have studied teams under a consensus rule based on (face-to-face) discussion, our design directly compares majority and unanimity voting in a strategic signaling environment. This allows us to isolate how the voting rule alone affects team behavior.

Figure 3 illustrates how collective choice rules systematically shape team choice behavior. For weak first movers, the vote frequencies of IN are consistently lower under unanimity than under majority in all four games ($p < 0.01$), and as shown in Figure 3, these lower voting rates translate into significantly lower team choice frequencies in games 2 and 3 ($p = 0.027$ and $p = 0.017$). Unanimity therefore makes weak first movers less willing to choose IN, because a single IN vote is sufficient for the team to select IN, which makes cautious members pivotal and amplifies their reluctance to support that action. In effect, unanimity magnifies the influence of the most cautious team members and generates a structural bias toward restraint on the sender side.

For second movers, the pattern is reversed. Both vote frequencies and team choice frequencies are consistently higher under unanimity than under majority, with differences strongly significant across all games ($p < 0.01$). Here the unanimity rule biases second-mover choices toward IN unless all members coordinate on OUT, so the willingness of even a single member to call a bluff is amplified into a team decision. In this role, unanimity magnifies the influence of the most aggressive members, pushing second-mover teams toward more

[21] See also Table 4 in Appendix 2.

frequent confrontation.

Taken together, these figures show that unanimity dampens IN choices by weak first movers while encouraging IN choices by second movers. This sharp, role-dependent divergence underscores that collective choice rules are not neutral procedural details: they fundamentally reshape strategic incentives and drive systematic differences in behavior. PBE, which predicts identical frequencies across voting rules, cannot account for these effects. By contrast, team equilibrium incorporates precisely these mechanisms and therefore provides a natural explanation for the observed patterns.

5.4 Are Teams More Rational Than Individuals?

The explanatory power of team equilibrium rests on the idea that teams are more sensitive to payoff differences than individuals. In this sense, rationality means adjusting choices more strongly in the direction favored by expected payoffs. If teams indeed react more sharply to payoff incentives, this would provide a natural explanation for why the team equilibrium model captures the data so successfully. We therefore test whether teams behave more rationally than individuals in our crisis bargaining games.

In equilibrium, weak first movers and second movers should respond monotonically to their opponents' strategies: if the opponent plays IN too often relative to equilibrium, the best response is to play OUT with probability one, and vice versa. Only in the knife-edge case where the opponent mixes exactly at the equilibrium probability are players indifferent. This logic suggests a straightforward test: if teams are more rational, their choice frequencies should shift more sharply with changes in expected payoff differences.

To implement this test, we compute the expected payoff difference between IN and OUT in each game and treatment, given observed opponent play, and use this difference as the explanatory variable in a logit regression of team choices. Table 2 reports the results. Across all specifications, the coefficient on payoff difference is positive and highly significant, showing that teams are more likely to choose IN when the expected payoff from IN is relatively high. The first two columns confirm that this responsiveness is robust across roles: first movers and second movers react similarly to payoff incentives.

Table 2: Team Choice Logit Regressions

	Dependent variable:		
	Team Choice of IN		
	(1)	(2)	(3)
Constant	0.156***	0.143***	0.060
	(0.036)	(0.049)	(0.043)
Majority			−0.309***
			(0.104)
Unanimity			1.439***
			(0.159)
First Mover		0.0005	
		(0.076)	
Pay Diff	0.134***	0.146***	0.141***
	(0.009)	(0.012)	(0.011)
Majority × Pay Diff			0.064**
			(0.029)
Unanimity × Pay Diff			0.070**
			(0.029)
First Mover × Pay Diff		−0.027	
		(0.019)	
Observations	3,285	3,285	3,285
Log Likelihood	−2,139.164	−2,138.076	−2,070.334
Akaike Inf. Crit.	4,282.328	4,284.151	4,152.668

Note: *p<0.1; **p<0.05; ***p<0.01

The third column adds interactions with the collective choice rules and reveals important treatment effects. Majority rule teams are significantly more likely to choose OUT than individuals when payoffs are balanced (the negative constant), yet they are also significantly more responsive to payoff differences. Unanimity teams, in contrast, exhibit a strong bias toward IN at indifference, consistent with the structural tilt of the rule. However, they too respond more strongly to payoff differences than individuals. Thus, both majority and unanimity teams are systematically more sensitive to incentives than individual decision-makers, even though their baseline tendencies differ.

Figure 7 illustrates this result for majority rule teams compared to individuals. Each

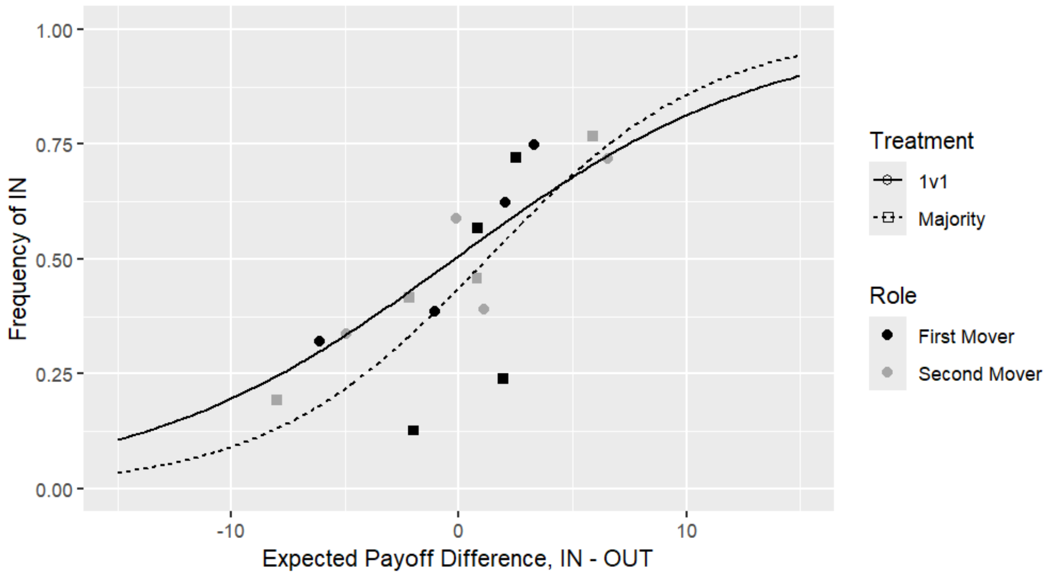

Figure 7: Majority Voting Logistic Regression

dot corresponds to observed behavior under the Individual treatment, and each square to behavior under the Majority treatment. The x-axis measures the expected payoff difference between IN and OUT, and the y-axis the observed frequency of IN. If teams were playing according to PBE, all points would cluster around zero payoff difference. If teams were perfect best responders, all points would lie at 0 or 1 depending on the sign of the payoff difference. In practice, we observe upward-sloping response functions. The dashed line for majority rule teams is steeper than the solid line for individuals, indicating that majority voting dampens random play and makes teams more responsive to incentives.

Figure 8 shows the same comparison for unanimity teams. Here the in-built bias of the unanimity rule is visible in the very high probability of choosing IN at indifference. Yet the slope of the response function is again steeper than that for individuals. In particular, when IN is the best response (as for second movers facing overly aggressive weak first movers), unanimity teams are especially effective at selecting IN.

These findings make clear that teams are not merely averaging individual biases. Voting rules increase their sensitivity to payoff incentives, leading to behavior that is closer to best responses. This enhanced rationality is a central reason why team equilibrium provides such a strong account of our data, and it raises the natural next question of whether this greater rationality also brings teams closer to equilibrium predictions than individuals.

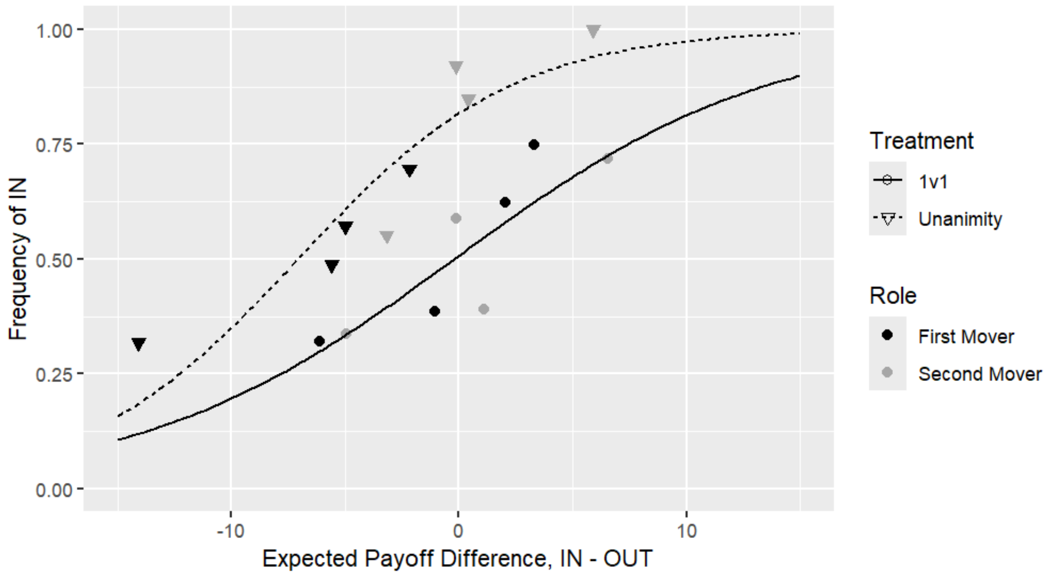

Figure 8: Unanimity Voting Logistic Regression

5.5 Are Teams Closer to PBE than Individuals?

Although PBE predictions largely fail to explain behavior, the literature has often suggested that teams may nevertheless behave more closely in line with equilibrium play than individuals. The logic is that if teams are more rational and payoff-responsive, their choices should deviate less from PBE. We therefore examine whether teams in our experiment are in fact closer to PBE than individuals.

The evidence from aggregate team choice frequencies, shown in figure 3, provides little support for this idea. In several cases, such as second movers under unanimity, team behavior actually deviates further from PBE than individual behavior. In others, deviations are of similar magnitude. Overall, there is no consistent evidence that teams converge toward PBE relative to individuals.

While our data clearly reject the notion that teams are systematically closer to PBE than individuals, the results reported in Section 5.2.2 showed that they support a much stronger claim: team equilibrium organizes the sign of deviations from PBE in a consistent and theoretically grounded manner. When the equilibrium benchmark is misspecified, information aggregation by voting does not move teams closer to PBE. Rather, it moves behavior in the direction of the comparative statics captured by team equilibrium.

6 Other Models

As a robustness check on our main findings, we compare the explanatory power of team equilibrium to two natural reference points: a no-information null model that assumes purely random play, and an omniscient perfect-fit model that matches the data exactly. These bounds allow us to assess where the performance of team equilibrium lies on the spectrum from random noise to perfect prediction. We then briefly consider alternative structural models such as QRE and extensions of TE with risk aversion.

6.1 Two Benchmark Models

(1) No Information Benchmark (NI). Suppose one tries to predict aggregate behavior in our four crisis bargaining games without any knowledge of the payoff structure. A natural benchmark is to assume that weak first movers and second movers play IN and OUT with equal probability. This *no information* model yields a log likelihood of −4770.2. Strikingly, this random-guessing benchmark fits the data better than the PBE model (log likelihood = −4817.8), underscoring the failure of PBE. By contrast, the team equilibrium model produces a much higher likelihood of −4357.2, representing a highly significant improvement even relative to NI.[22]

(2) Perfect Fit Benchmark (PF). At the opposite extreme, one can consider an *omniscient* model that perfectly reproduces the observed frequencies in all games, roles, and team treatments. This *perfect fit* benchmark, which uses 24 free parameters, achieves a log likelihood of −4163.1. No model can improve on this fit, since it simply restates the data.

Relative Fit. To place team equilibrium on this spectrum, we compute a pseudo-R^2 measure analogous to McFadden's index:

$$Fit_{TE} = \frac{0.087}{0.127} = 0.685.$$

This measure is normalized so that NI has a value of 0 and PF has a value of 1. Team equilibrium, with a single free parameter, explains nearly 70% of the variation that could possibly be explained by any model, despite the heterogeneity of payoffs, roles, and collective choice rules. This result confirms that TE captures the essential systematic patterns in the data while remaining extremely parsimonious.

[22] The NI model is nested in the TE model, corresponding to $\lambda = 0$.

6.2 Quantal Response Equilibrium

As pointed out in Section 3.2, the team equilibrium with $n = 1$ is equivalent to the logit quantal response equilibrium at every value of λ. However, the case of $n = 1$ is special, and the team equilibrium with team sizes $n > 1$ is different from a quantal response equilibrium. To analyze a quantal response equilibrium when $n > 1$, one looks at the $2n$ person voting game (10-person in our experiment, with 5 voters for each team). In a quantal response equilibrium the i.i.d. payoff disturbances are applied for each voter to the expected payoffs of *voting* for either IN or OUT, rather than the expected payoff of the *team decision* to choose IN or OUT. In a 5×5 team game, the quantal response equilibrium is an equilibrium of this larger voting game. In contrast to the uniqueness of Perfect Bayesian (and team) equilibrium for the $n = 1$ special case, there are typically many Nash equilibria in multi-player voting games, under either majority rule or unanimity, and there can be multiple logit quantal response equilibria as well. In fact, there are bifurcations in the graph of the logit quantal response equilibrium correspondence that arise at values of λ where additional logit equilibria are picked up.

We are able to compute the unique continuous selection of the logit quantal response equilibrium correspondence that converges to the Perfect Bayesian equilibrium of the baseline game for the 5×5 majority games. [23] This allows us to use maximum likelihood estimation to compare the fit of the logit quantal response equilibrium model to the team equilibrium model in our experiment, using all of our pooled data except the 5×5 unanimity games. The team equilibrium model fits significantly better than the logit quantal response equilibrium model ($\chi^2 - statistic > 20$).

6.3 Risk Aversion

In our crisis bargaining games, OUT is a safe option yielding a guaranteed payoff, while IN is a risky option that produces a binary lottery. Risk preferences are therefore potentially relevant for strategic play. The team equilibrium framework accommodates this naturally by embedding constant relative risk aversion (CRRA) preferences. Each payoff x is mapped to utility $u(x) = \frac{x^{1-\rho}}{1-\rho}$, with ρ denoting the coefficient of relative risk aversion, and logit responses are then applied to expected utilities rather than expected payoffs. The fixed-

[23] This was not possible for the 5×5 unanimity voting games to multiple equilibria and bifurcations.

point logic of team equilibrium remains unchanged.

We estimate two variants. First, we consider PBE with homogeneous CRRA preferences but no payoff disturbances ($\lambda = \infty$). The resulting estimate implies negative risk aversion, i.e. risk-seeking behavior, which is implausible. Moreover, this PBE with risk aversion model fits only marginally better than risk-neutral PBE and still worse than the no-information benchmark.

Second, we estimate a team equilibrium model with both CRRA preferences and payoff disturbances. The fitted coefficient is $\rho = 0.485$, close to quadratic utility and consistent with prior experimental evidence (e.g. Goeree et al., 2002, 2003). The fitted responsiveness parameter λ is somewhat higher, reflecting the rescaling of payoffs due to concave utility, but the qualitative fit is essentially unchanged. Figure 9 shows that the scatter plot of observed versus predicted frequencies is nearly identical to the risk-neutral case (compare with Figure 4). The OLS fit yields an intercept of 0.09, a slope of 0.82, and an R^2 of 0.77, virtually the same as before.

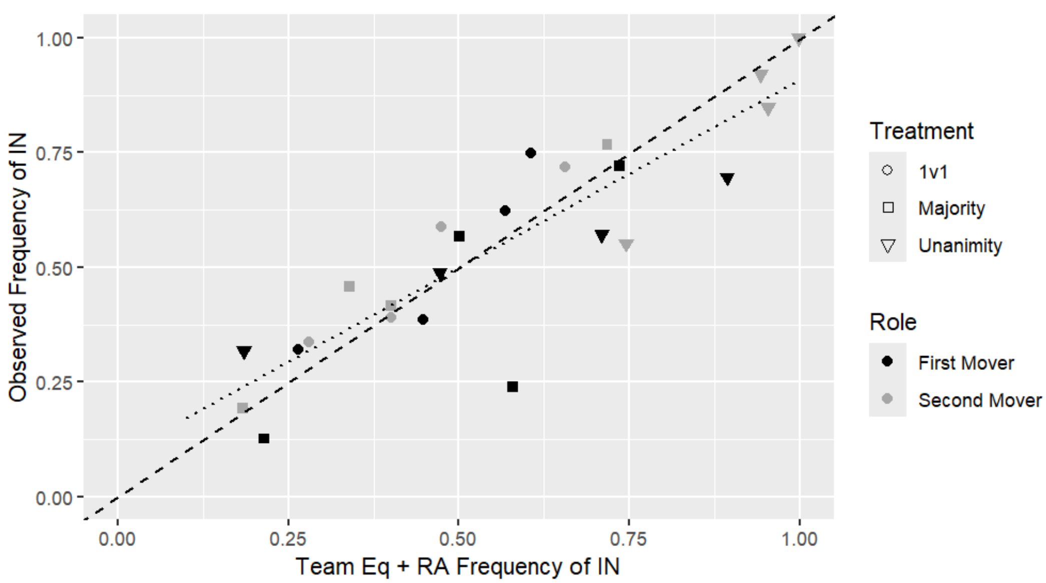

Figure 9: Team Equilibrium Fit with Risk Aversion

Allowing for risk aversion therefore produces plausible parameter estimates and a slight quantitative improvement in fit, but does not change the qualitative implications of team equilibrium. The central explanatory power of the model comes from heterogeneity in payoff disturbances, not from risk preferences. The effect of risk aversion is second-order.

7 Conclusion

The application of a unitary rational actor model to analyze crisis decision-making in international conflict situations has been sharply criticized, on account of the diversity of interests in group decision-making and organizational rules that constrain how these diverse interests are aggregated into a group decision. Crisis bargaining games capture three essential features of these conflict environments - asymmetric information, sequential timing, and strategic calculation - and thus provide a valuable paradigm for analyzing these situations.

In this paper, we apply team equilibrium to these games, which is a general framework that extends the standard single-actor equilibrium analysis of games by incorporating both diversity of team members' interests and organizational rules that constrain the group decision-making process. For this class of games, team equilibrium implies sharp predictable systematic differences from the standard PBE model of behavior, in terms of both group decisions and bargaining outcomes.

The experiment was designed to test the full range of these team equilibrium predictions, using four different payoff variations from a simple class of crisis bargaining games. The four games were carefully selected because they represent four starkly different patterns bluffing and bluff-calling behavior by first and second movers, respectively: in game 2, PBE predicts that both players should choose IN less than half the time; in game 4, both players should choose IN more than half the time; in game 1, only player 2 chooses IN more than half the time; and in game 3 only the first mover chooses IN more than half the time.

We summarize the findings of the experiment in terms of the answers to the five main research questions.

How well (or poorly) does team equilibrium account for the variation of behavior across treatments? A one-parameter logit specification of the team equilibrium model closely tracks the data and captures essentially all the salient effects of the payoff variations across the four games, as well as the effects of the group size and collective choice procedure. The regression of IN frequencies predicted by the logit model on observed IN frequencies produces an $R^2 = .76$ with a slope of 0.85; the qualitative predictions about the deviations of IN frequencies (higher or lower) relative Nash equilibrium were observed in 21 out of 24 cases.

Does the collective choice rule matter for team behavior and if so how? Yes, the collective choice rule has strong effects, as predicted by team equilibrium. Unanimity rule induces first mover teams to choose IN less frequently than majority rule, and the opposite is true for second movers.

Are teams more rational than individuals? Yes, consistent with team equilibrium theory, based on the a logit regression, teams have a stronger positive response of choosing IN to payoff differences between IN and OUT than individuals, a comparison that holds true for both collective choice rules. Unanimity rule biases teams toward choosing IN.

Are the outcomes played between teams closer to PBE than the same games played between individuals? There is no systematic difference between teams and individuals in terms of how close the outcomes are to Nash equilibrium.

Do other models of behavior explain the variation of behavior across the treatments, perhaps even better than team equilibrium? Team equilibrium fits the data significantly better than quantal response equilibrium. A two-parameter estimation of the logit team equilibrium model that includes a risk aversion parameter fits the data slightly better ($R^2 = .77$ vs $R^2 = .76$) but makes identical qualitative predictions to the one-parameter model that assumes risk neutrality.

This study has implications for understanding strategic decision-making in conflict situations and we conclude with three remarks about this. First, the experiment demonstrates that group decision-making in crisis bargaining situations cannot be adequately modeled by standard game-theoretic approaches that rely on unitary rational actors. By integrating the diversity of interests and the impact of collective choice rules, the team equilibrium framework offers a more nuanced and accurate description of group behavior in strategic contexts.

Second, the findings emphasize the critical role of institutional structures in shaping collective decisions. Variations in group size and voting rules systematically influence behavior and outcomes, highlighting the importance of organizational design in strategic interactions. For policymakers and institutions managing crises, this points to the need to carefully consider how decision-making processes and rules impact group behavior, particularly in high-stakes strategic situations.

Third, the robustness of the team equilibrium model across different payoff structures and decision-making rules opens promising avenues for further research. Future work could explore how factors such as communication within teams, the role of leadership, and heterogeneity in team members' payoffs shape outcomes. Extending this framework to more complex strategic interactions could offer deeper insights into how groups navigate uncertainty and strategic conflict.

References

Agranov, Marina and Chloe Tergiman, "Communication in multilateral bargaining," *Journal of Public Economics*, 2014, *118*, 75–85.

_ **and** _ , "Communication in bargaining games with unanimity," *Experimental Economics*, 2019, *22*, 350–368.

Allison, Graham T, *Essence of decision: Explaining the Cuban Missile Crisis*, Little, Brown: Boston, 1971.

Baron, David P and John A Ferejohn, "Bargaining in legislatures," *American political science review*, 1989, *83* (4), 1181–1206.

Battaglini, Marco, Salvatore Nunnari, and Thomas R Palfrey, "Legislative bargaining and the dynamics of public investment," *American Political science review*, 2012, *106* (2), 407–429.

_ **,** _ **, and** _ , "The political economy of public debt: a laboratory study," *Journal of the European Economic Association*, 2020, *18* (4), 1969–2012.

Bellman, Richard and David Blackwell, "Some Two Person Games Involving Bluffing," *Proceedings of the National Academy of Sciences*, 1949, *35* (10), 600–605.

Borel, Emile and Jean Ville, *Applications de la Theorie des Probabilites aux Jeux de Hasard*, Vol. 4, Gauthier-Villars: Paris, 1938.

Charness, Gary and Matthias Sutter, "Groups make better self-interested decisions," *Journal of economic perspectives*, 2012, *26* (3), 157–176.

Chen, Daniel L., Martin Schonger, and Chris Wickens, "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Cooper, David J and John H Kagel, "Are two heads better than one? Team versus individual play in signaling games," *American Economic Review*, 2005, *95* (3), 477–509.

_ **and** _ , "The role of context and team play in cross-game learning," *Journal of the European Economic Association*, 2009, *7* (5), 1101–1139.

Fearon, James D., "Domestic Political Audiences and the Escalation of International Disputes," *American Political Science Review*, 1994, *88* (3), 577–592.

_ , "Signaling Versus the Balance of Power and Interests: An Empirical Test of a Crisis Bargaining Model," *Journal of Conflict Resolution*, 1994, *38* (2), 236–269.

Goeree, Jacob K and Leeat Yariv, "An Experimental Study of Collective Deliberation," *Econometrica*, 2011, *79* (3), 893–921.

_ , Charles A Holt, and Thomas R Palfrey, "Quantal Response Equilibrium and Overbidding in First Price Auctions," *Journal of Economic Theory*, 2002, *104* (1), 247–272.

_ , _ , and _ , "Risk Averse Behavior in Generalized Matching Pennies Game," *Games and economic behavior*, 2003, *45* (1), 97–113.

Guarnaschelli, Serena, Richard D McKelvey, and Thomas R Palfrey, "An experimental study of jury decision rules," *American Political science review*, 2000, *94* (2), 407–423.

Kim, Jeongbin and Thomas R Palfrey, "An Experimental Study of Prisoners' Dilemma and Stag Hunt Games Played by Teams of Players," *Games and Economic Behavior*, Forthcoming.

_ , _ , and Jeffrey R Zeidel, "Games played by teams of players," *American Economic Journal: Microeconomics*, 2022, *14* (4), 122–157.

Kocher, Martin G, Matthias Praxmarer, and Matthias Sutter, "Team decision-making," *Handbook of labor, human resources and population economics*, 2020, pp. 1–25.

Kugler, Tamar, Edgar E Kausel, and Martin G Kocher, "Are groups more rational than individuals? A review of interactive decision making in groups," *Wiley interdisciplinary reviews: Cognitive science*, 2012, *3* (4), 471–482.

Kuhn, H. W., "Simplified Two-Person Poker," *Contributions to the Theory of Games: Volume 1*, 1950, pp. 97–103.

Lewis, Jeffrey B. and Kenneth A. Schultz, "Revealing Preferences: Empirical Estimation of a Crisis Bargaining Game with Incomplete Information," *Political Analysis*, 2003, *11*, 345–367.

Milgrom, Paul and John Roberts, "Limit pricing and entry under incomplete information: An equilibrium analysis," *Econometrica: Journal of the Econometric Society*, 1982, pp. 443–459.

Miller, Luis and Christoph Vanberg, "Decision costs in legislative bargaining: an experimental analysis," *Public Choice*, 2013, *155*, 373–394.

_ and _ , "Group size and decision rules in legislative bargaining," *European Journal of Political Economy*, 2015, *37*, 288–302.

Morgenstern, Oskar and John von Neumann, *Theory of Games and Economic Behavior*, Princeton University Press, 1947.

Morrow, James D, "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining," *American Journal of Political Science*, 1989, *33* (4), 941–972.

Myerson, Roger B., "Game Theory: Analysis of Conflict," *The president and fellows of Harvard College, USA*, 1991, *66.*

Palfrey, Thomas R, "Experiments in political economy," *Handbook of experimental economics*, 2013, *2*, 347–434.

Powell, Robert, "Research Bets and Behavioral IR," *International Organization*, 2017, *71* (Supplement), S265–S277.

Signorino, Curtis S., "Strategic Interaction and the Statistical Analysis of International Conflict," *American Political Science Review*, 1999, *93* (2), 279–297.

Smith, Alastair, "Alliance Formation and War," *International Studies Quarterly*, 1995, *39* (December), 405–425.

_ , "Testing Theories of Strategic Choice: The Example of Crisis Escalation," *American Journal of Political Science*, 1999, *43* (4), 1254–1283.

Appendices

Appendix 1: Experimental Design Details

Table 3: Summary of Sessions

SESSION#	TREATMENT	#SUBJECTS	LOCATION
1	1×1	18	Online UCI
2	1×1	20	Online UCI
3	1×1	10	Online UCI
4	1×1	14	Online UCI
5	1×1	14	Online UCI
6	1×1	12	Online UCI
7	1×1	18	UCI
8	1×1	8	UCI
9	5×5 Majority	20	UCI
10	5×5 Majority	20	Online UCI
11	5×5 Majority	20	UCSB
12	5×5 Majority	20	UCSB
13	5×5 Majority	20	UCSB
14	5×5 Majority	20	UCSB
15	5×5 Unanimity	20	UCI
16	5×5 Unanimity	20	UCI
17	5×5 Unanimity	20	UCI
18	5×5 Unanimity	20	UCSB
19	5×5 Unanimity	20	UCSB

Figure 10: Game Payoffs

(a) Payoffs for Game 1:

First Mover Team's Choice of Action	OUT	10,35

If HEADs was selected by coin flip

		Second Mover Team's Choice of Action	
		OUT	IN
First Mover Team's Choice of Action	IN	30,15	35,10

If TAILs was selected by coin flip

		Second Mover Team's Choice of Action	
		OUT	IN
First Mover Team's Choice of Action	IN	30,15	5,40

(b) Payoffs for Game 2:

First Mover's Choice of Action	OUT	20,32

If HEADs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	44,20

If TAILs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	4,36

(c) Payoffs for Game 3:

First Mover's Choice of Action	OUT	21,27

If HEADs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	27,21	42,6

If TAILs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	27,21	6,42

(d) Payoffs for Game 4:

First Mover's Choice of Action	OUT	20,32

If HEADs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	32,8

If TAILs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	16,48

Note: These payoffs were constructed from the payoffs displayed in Table 1 by adding a constant to all payoffs in a game, which does not affect the PBE or the team equilibrium in any of the games. This was done to ensure that all payoffs were positive. Specifically, in Game 1, the added constant was 20 for the first mover and 25 for the second mover. The corresponding pair of added constants for Games 2-4 were (24, 28), (24, 24), and (24, 28), respectively.

Appendix 2: Observed and Theoretical Decision Frequencies

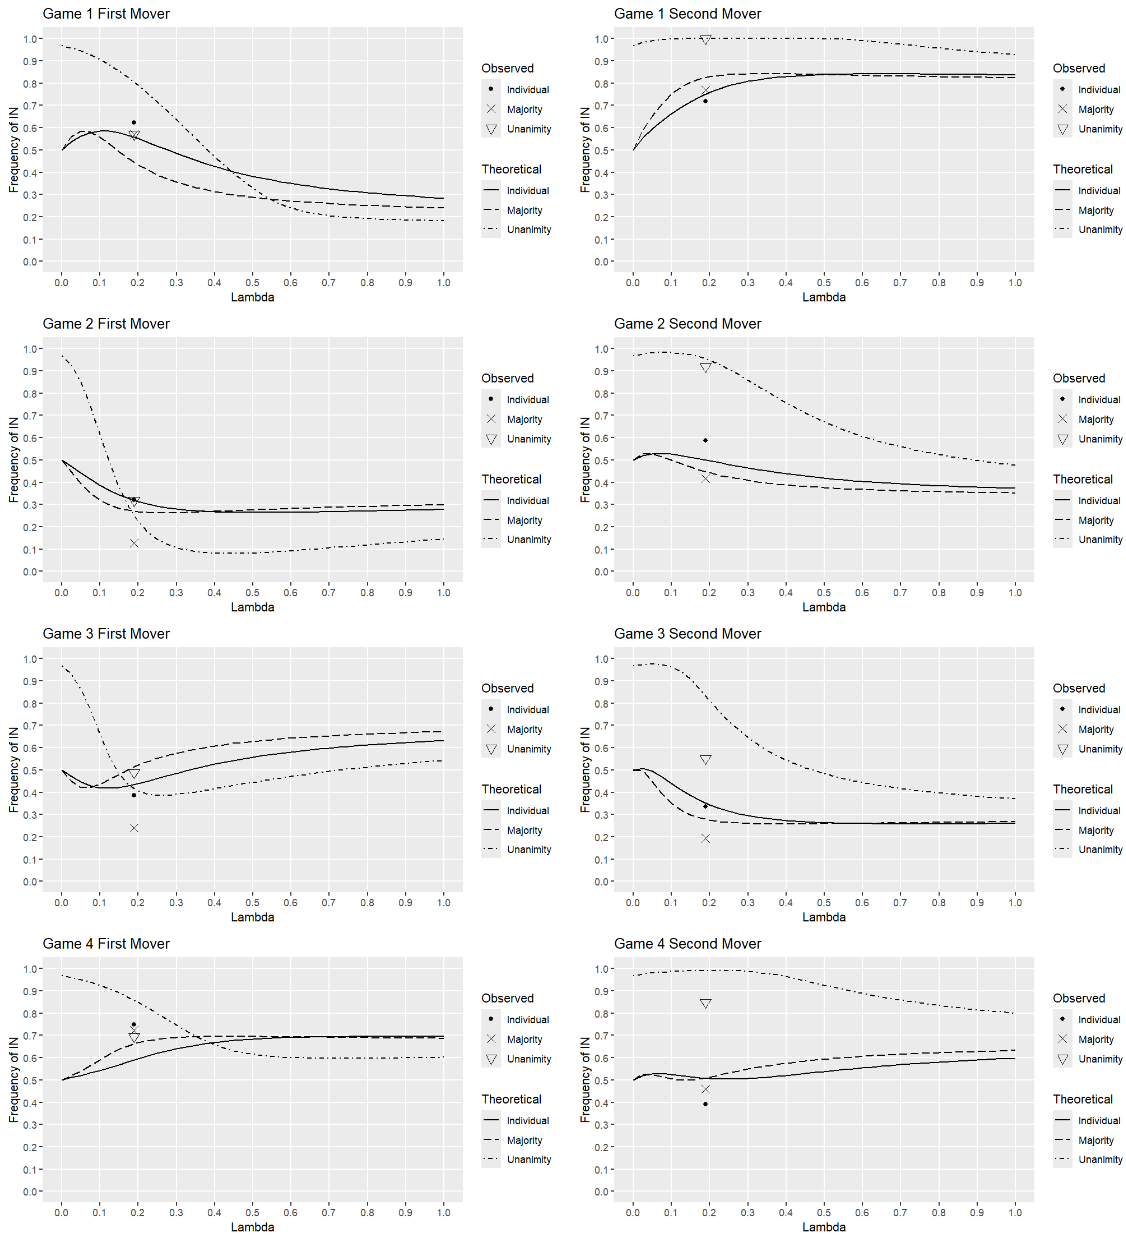

Each panel of the above figure displays the team equilibrium for one game and one player role. Each line traces out the team equilibrium for one team treatment, for all λ values from 0 to 1. The points are plotted on the x-axis at the estimated value of λ, and on the y-axis at the observed team decision frequency.

Table 4: Supplemental information for Figure 6

			IN Frequency						
Game	Rule	Player	Obs	TE	Nash	TE Prediction	Correct?	Obs−Nash	TE−Nash
1	1v1	First	0.624	0.558	0.200	Obs>NE*	YES	0.424	0.358
1	1v1	Second	0.719	0.751	0.800	Obs<NE	YES	−0.081	−0.049
1	Maj	First	0.569	0.446	0.200	Obs>NE*	YES	0.369	0.246
1	Maj	Second	0.768	0.825	0.800	Obs>NE		−0.032	0.025
1	Unan	First	0.571	0.807	0.200	Obs>NE	YES	0.371	0.607
1	Unan	Second	1.000	1.000	0.800	Obs>NE*	YES	0.200	0.200
2	1v1	First	0.320	0.319	0.333	Obs<NE	YES	−0.013	−0.014
2	1v1	Second	0.588	0.500	0.333	Obs>NE*	YES	0.255	0.167
2	Maj	First	0.127	0.270	0.333	Obs<NE	YES	−0.206	−0.063
2	Maj	Second	0.417	0.448	0.333	Obs>NE*	YES	0.084	0.115
2	Unan	First	0.318	0.250	0.333	Obs<NE	YES	−0.015	−0.083
2	Unan	Second	0.920	0.955	0.333	Obs>NE*	YES	0.587	0.622
3	1v1	First	0.386	0.435	0.715	Obs<NE*	YES	−0.329	−0.280
3	1v1	Second	0.336	0.352	0.285	Obs>NE	YES	0.051	0.067
3	Maj	First	0.240	0.513	0.715	Obs<NE*	YES	−0.475	−0.202
3	Maj	Second	0.194	0.279	0.285	Obs<NE	YES	−0.091	−0.006
3	Unan	First	0.488	0.416	0.715	Obs<NE	YES	−0.227	−0.299
3	Unan	Second	0.552	0.834	0.285	Obs>NE*	YES	0.267	0.549
4	1v1	First	0.749	0.589	0.666	Obs<NE		0.083	−0.077
4	1v1	Second	0.390	0.508	0.666	Obs<NE*	YES	−0.276	−0.158
4	Maj	First	0.722	0.662	0.666	Obs<NE		0.056	−0.004
4	Maj	Second	0.459	0.509	0.666	Obs<NE*	YES	−0.207	−0.157
4	Unan	First	0.696	0.859	0.666	Obs>NE	YES	0.030	0.193
4	Unan	Second	0.848	0.992	0.666	Obs>NE*	YES	0.182	0.326

*Global prediction applies for all values of λ

Online Appendix: Instructions and screenshots

Sample instructions

These are the instructions used for the 1-1 in-person experiments with the reverse game order (4321). Modifications for the 5-5 majority rule treatment is noted in double brackets at the appropriate places. Instructions for the unanimity rule treatment were nearly identical to the majority rule instructions, with the exception of the voting rule explanation. The instructions were handed out to all subjects so they could follow along while the experimenter read the instructions aloud. They could also refer back to the instructions during the experiment if they wished.

For the team experiments, these instructions were slightly modified to explain that each subject's action choice was a *vote* of IN or OUT and all final team decisions was determined by majority [unanimity] rule. The instructions for online sessions included an explanation of the procedures that would be followed in case any subject became disconnected. The online sessions were conducted on Zoom and also used a utility software, Experimentalist, which streamlined the login, connection, and payoff protocols.

General instructions

Thank you for coming. You are about to participate in an experiment on decision-making. Your earnings will depend partly on your decisions, partly on the decisions of others, and partly on chance. This experiment requires your undivided attention. Please refrain from other activities for the duration of the experiment.

The entire session will take place through computer terminals, and all interaction between participants will take place through the computers. Please do not attempt to communicate in any way with other participants during the experiment.

Some of your decisions will be randomly selected for payment. Your earnings are denominated in points, and each point has a value of $0.20. In other words, every 100 points generates $20 in earnings for you. In addition to your earnings from decisions, you will receive a show-up fee of $7, and a completion fee of $5. At the end of the experiment, your earnings will be rounded up to the nearest dollar amount. All your earnings will be paid by cash.

Main task

Overview of the experiment:

This experiment consists of 4 matches, and each match has 10 rounds. We will now review the instructions for match 1. After the conclusion of each match, instructions for the next match will be read.

Type: First mover and Second mover

At the beginning of the experiment, you will be randomly assigned to one of two types: First mover or Second mover. Once you become a first mover or a second mover, your type will remain fixed throughout the entire experiment. As will be explained later, in every round, first movers make a decision first, and then second movers follow.

Matching and decision

At the beginning of each round, you will be randomly paired with another subject of the opposite type. If you are a first mover, you will be paired with a second mover, and vice versa. In each round, when it is your turn to move, you will make a decision to choose an action (IN or OUT).

[[*Team Assignment*

At the beginning of each round, first movers will be randomly sorted into first mover teams, and second movers will be randomly sorted into second mover teams. Each team has 5 members with the same type. In each round, your 5-member team will be randomly paired with another 5-member team of the opposite type. If your team is a first mover team, your team will be paired with a second mover team, and vice versa.]]

[[*Team Decision*

In each round, when it is your team's turn to move, your team will make a collective decision to choose an action (IN or OUT). You will be asked to vote for an action in each round, and your team decision will be determined by majority rule. For instance, if there are 2 member of your team who votes for action In and 3 members who vote for action OUT, then your team's collective decision will be action OUT.]]

Payoff table

Payoff table

		Second Mover's Choice of Action	
First Mover's Choice of Action	IN	OUT	IN
		1000, 220	40, 505

On the decision-making screen, you will see three tables that show the payoffs in points from each combination of choices. The payoffs in the example table above are only for illustration, they are not the actual payoffs used in the experiment. The rows always correspond to the actions the first mover [[first mover team]] can choose, and the columns correspond to the actions the second mover [[second mover team]] can choose. The first entry in each cell represents the first mover's [[first mover team members']] earnings, while the second entry represents the earnings of the second mover [[second mover team members]]. For instance (1) if the first mover [[team]] chooses action IN and the second mover [[team]] chooses action OUT, the first mover receives [[each of the first mover team members receives] 1000 points, while the second mover [[each of the second mover team members]] receives 220 points. (2) if the first mover [[team]] chooses action IN and the second mover [[team]] chooses action IN, the first mover receives [[each of the first mover team members receives]] 40 points, while the second mover receives [[each of the second mover team members receives]] 505 points.

Coin flip: HEADs or TAILs

There will be three different payoff tables on your screen, one upper table and two lower tables. At the beginning of each round, the computer flips a virtual, fair coin which determines either HEADs or TAILs. The probability that HEADs (or TAILs) is selected is 50%. A separate virtual, fair coin is flipped for each round, for each pairing. The result of each coin flip is independent of every other coin flip.

If the result of the coin flip is HEADs, then, among the two lower tables, the table on the left side of the screen shows the true payoffs to each combination of actions. If the result is TAILs, then the table on the right shows the true payoffs. As will be explained later, the upper payoff table is not related to the coin flip result. Only the first mover [[first mover team members]] knows [[know]] whether HEADs or TAILs has been selected before he chooses an action [[they vote]]. The second mover [[second mover team members]] only learns [[learn]] this information at the end of the round, after all decisions have been made. This will be further explained in detail below.

Decision-making procedures

There are three stages of decision-making procedures-(1) action choice of the first mover [[voting of the first mover team members]], (2) action choice of the second mover [[voting of the second mover team members]], and (3) feedback about decisions and payoffs. We will now review each of these stages by looking at screenshots of the actual experimental user interface. All payoffs in these screenshots are the real payoffs used in match 1.

(The screenshots were handed out to subjects and they followed along as the experimenter read the following script. Copies of the screenshots are reproduced at the end of this appendix.)

(1) Action choice of the first mover First mover: [Slide 1] On the first screen, the first mover [[first mover team members]] will see three tables that show the payoffs in points from each combination of actions and a random computerized coinflip. The top table, which has one row and two cells, shows the payoffs resulting from the first mover [[first mover team]] choosing action OUT. If the first mover [[first mover team]] chooses action OUT, the round immediately ends and subjects receive the payoffs shown in this table. That is, regardless of the result of the coin flip, the first mover [[first mover team members]] receives 10 points, while the second mover [[second mover team members]] receives 35 points. The bottom two tables show the payoffs if the first mover [[first mover team members]] chooses action IN, for every combination of computerized coin flip and action of the second mover. The first mover [[first mover team members]] can choose [[vote for]] an action by clicking a row in the payoff tables. To finalize a choice [[vote]], the Submit button highlighted in red should be clicked.

Second mover: [Slide 2] The second mover [[second mover team members]] will wait for the first mover [[first mover team members]] to choose [[vote for]] an action. On this screen, the computer's random selection of HEADs or TAILs will not be revealed to the second mover [[second mover team members]].

(2) Action choice of the second mover. First mover: [Slide 3] This screen appears only if the first mover [[first mover team]] chooses IN. The first mover [[first mover team]] is asked to wait for the second mover [[second mover team]] to choose an action in the two payoff tables below. Choosing [[voting for]] an action, the second mover [[second mover team members]] does [[do]] not know whether HEADs or TAILs was randomly selected.

Second mover: [Slide 4] This screen appears only if the first mover [[first mover team]] chooses IN. The second mover is asked to choose an action in the two payoff tables below. Choosing [[voting for]] an action, the second mover [[second mover team members]] does not know whether HEADs or TAILs was randomly selected. The second mover [[second mover team members]] can choose an action by clicking a column in the payoff tables. To finalize a choice, the Submit button highlighted in red should be clicked.

(3) Receiving feedback about the payoffs

First mover: [Slide 5] This is the feedback screen for the case in which OUT was chosen by the first mover [[first mover team]]. The first mover [[first movers]] and the second mover [[second movers]] will receive feedback about their payoffs and the computer's random choice of HEADs or TAILs.

[Slide 6] This is the feedback screen for the case in which IN was chosen by the first mover. The action of each is highlighted, and the payoffs you will receive are highlighted in pink. The computer's random choice of HEADs or TAILs will also be revealed to the second mover.

Second mover: [Slide 7] This is the feedback screen for the case in which OUT was chosen by the first mover [[first mover team]]. The first mover[[s]] and the second mover[[s]] will receive feedback about their payoffs and the computer's random choice of HEADs or TAILs.

[Slide 8] This is the feedback screen for the case in which IN was chosen by the first mover [[first mover team]]. The computer's random choice of HEADs or TAILs will also be revealed to the second mover [[second mover team]]. For the randomly selected payoff table, the action of each mover is highlighted, and the payoffs you will receive are highlighted in purple. On the feedback screen, all subjects should click the Confirm button to advance the page.

Payoffs

At the end of the experiment, one round from each Match (4 rounds in total) will be randomly selected for determining the payment of each subject. The rounds chosen may be different for different subjects. You will be paid the amount of points you earned in the rounds randomly selected for you.

Quiz

For your comprehension, before Match 1 begins, you will be asked to solve quiz problems. You can participate in the experiment only if you enter correct answers for all problems.

Summary of instructions

1. At the beginning of the experiment, you will be randomly assigned to one of two types: First mover or Second mover. Once you become a first mover or a second mover, your type will remain fixed throughout the entire experiment.

2. Each of the 4 matches has 10 rounds.

3. At the beginning of each round, you will be randomly paired with a subject of the other type. In each round, when it is your turn, your will choose an action. [[At the beginning of each round, you will be randomly sorted into a team of subjects with the same type, and your team will be randomly paired with another team of the other type. In each round, when it is your team's turn to move, your team will make a collective decision to choose an action.]]

4. If the first mover [[team]] chooses OUT, the round is over and the second mover does [[team member do]] not make a choice. The payoffs do not depend on the result of the coin flip.

5. If the first mover [[team]] chooses In, the second mover [[team members vote for]] chooses an action in the two payoff tables below. After all members of both teams have voted in their turn, the outcome of the coin flip of HEADs or TAILs and each team's choice will determine the payoffs.

First Mover's Choice of Action	OUT	20,32

If HEADs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	32,8

If TAILs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	16,48

Figure 11: Match 1 payoffs

[Subjects answer comprehension questions, after which the first round of Match 1 begins. After 10 rounds, the experiment is paused and the experimenter reads the next script to announce the change of payoffs for Match 2.]

Match 2

We have reached the end of match 1 and match 2 will now begin. The rules and procedures used in this match are exactly the same as in the first match. However, the payoffs are now different. The new payoff tables are pasted at the bottom of this page, below the instructions summary.

Summary of instructions

1. At the beginning of the experiment, you will be randomly assigned to one of two types: First mover or Second mover. Once you become a first mover or a second mover, your type will remain fixed throughout the entire experiment.

2. Each of the 4 matches has 10 rounds.

3. At the beginning of each round, you will be randomly paired with a subject of the other type. In each round, when it is your turn, your will choose an action. [[At the beginning of each round, you will be randomly sorted into a team of subjects with the same type, and your team will be randomly paired with another team of the other type. In each round, when it is your team's turn to move, your team will make a collective decision to choose an action.]]

4. If the first mover [[team]] chooses OUT, the round is over and the second mover does [[team member do]] not make a choice. The payoffs do not depend on the result of the coin flip.

5. If the first mover [[team]] chooses In, the second mover [[team members vote for]] chooses an action in the two payoff tables below. After all members of both teams have voted in their turn, the outcome of the coin flip of HEADs or TAILs and each team's choice will determine the payoffs.

First Mover's Choice of Action	OUT	21,27

If HEADs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	27,21	42,6

If TAILs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	21,27	6,42

Figure 12: Match 2 payoffs

[After 10 rounds, the experiment is paused and the experimenter reads the following script to announce the change of payoffs for Match 3.]

Match 3

We have reached the end of match 2 and match 3 will now begin. The rules and procedures used in this match are exactly the same as in the previous match. However, the payoffs are now different. The new payoff tables are pasted at the bottom of this page, below the instructions summary.

Summary of instructions

1. At the beginning of the experiment, you will be randomly assigned to one of two types: First mover or Second mover. Once you become a first mover or a second mover, your type will remain fixed throughout the entire experiment.

2. Each of the 4 matches has 10 rounds.

3. At the beginning of each round, you will be randomly paired with a subject of the other type. In each round, when it is your turn, your will choose an action. [[At the beginning of each round, you will be randomly sorted into a team of subjects with the same type, and your team will be randomly paired with another team of the other type. In each round, when it is your team's turn to move, your team will make a collective decision to choose an action.]]

4. If the first mover [[team]] chooses OUT, the round is over and the second mover does [[team member do]] not make a choice. The payoffs do not depend on the result of the coin flip.

5. If the first mover [[team]] chooses In, the second mover [[team members vote for]] chooses an action in the two payoff tables below. After all members of both teams have voted in their turn, the outcome of the coin flip of HEADs or TAILs and each team's choice will determine the payoffs.

First Mover's Choice of Action	OUT	20,32

If HEADs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	44,20

If TAILs was selected by coin flip

		Second Mover's Choice of Action	
		OUT	IN
First Mover's Choice of Action	IN	28,24	4,36

Figure 13: Match 3 payoffs

[After 10 rounds, the experiment is paused and the experimenter reads the following script to announce the change of payoffs for Match 4.]

Match 4

We have reached the end of match 3 and match 4 will now begin. The rules and procedures used in this match are exactly the same as in the previous match. However, the payoffs are now different. The new payoff tables are pasted at the bottom of this page, below the instructions summary.

Summary of instructions

1. At the beginning of the experiment, you will be randomly assigned to one of two types: First mover or Second mover. Once you become a first mover or a second mover, your type will remain fixed throughout the entire experiment.

2. Each of the 4 matches has 10 rounds.

3. At the beginning of each round, you will be randomly paired with a subject of the other type. In each round, when it is your turn, your will choose an action. [[At the beginning of each round, you will be randomly sorted into a team of subjects with the same type, and your team will be randomly paired with another team of the other type. In each round, when it is your team's turn to move, your team will make a collective decision to choose an action.]]

4. If the first mover [[team]] chooses OUT, the round is over and the second mover does [[team member do]] not make a choice. The payoffs do not depend on the result of the coin flip.

5. If the first mover [[team]] chooses In, the second mover [[team members vote for]] chooses an action in the two payoff tables below. After all members of both teams have voted in their turn, the outcome of the coin flip of HEADs or TAILs and each team's choice will determine the payoffs.

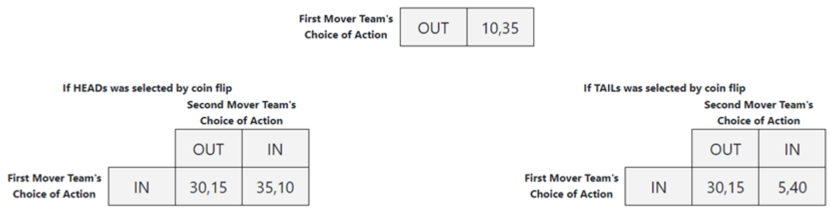

Figure 14: Match 4 payoffs

End of Experiment (Read after Match 4 ends): We have reached the end of match 4. Thank you for your participation. Please take your time to review your randomly payoff on this screen. When you are ready, please click the confirm button and enter your Venmo username on the next screen for payment. For security reasons, please close your browser after entering your Venmo username. This concludes the experiment, thank you for your participation.

Screenshot slides

These are the slides referred to in the instructions and distributed as handouts to subjects for the 1-1 sessions. The slides for the 5-5 sessions were similar, except for including the voting outcomes of each team in the results screen and minor wording changes to be consistent with teams instead of individuals.

Figure 15: Slide 1

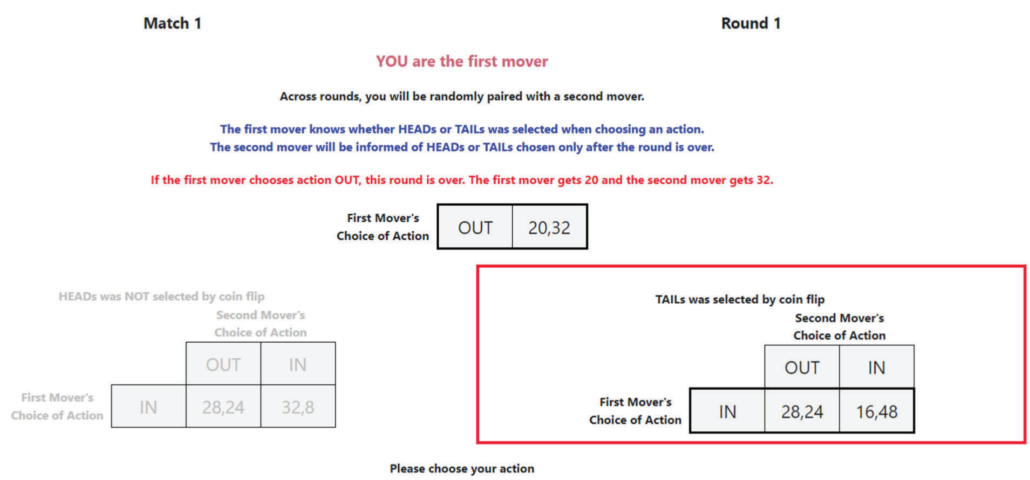

Figure 16: Slide 2

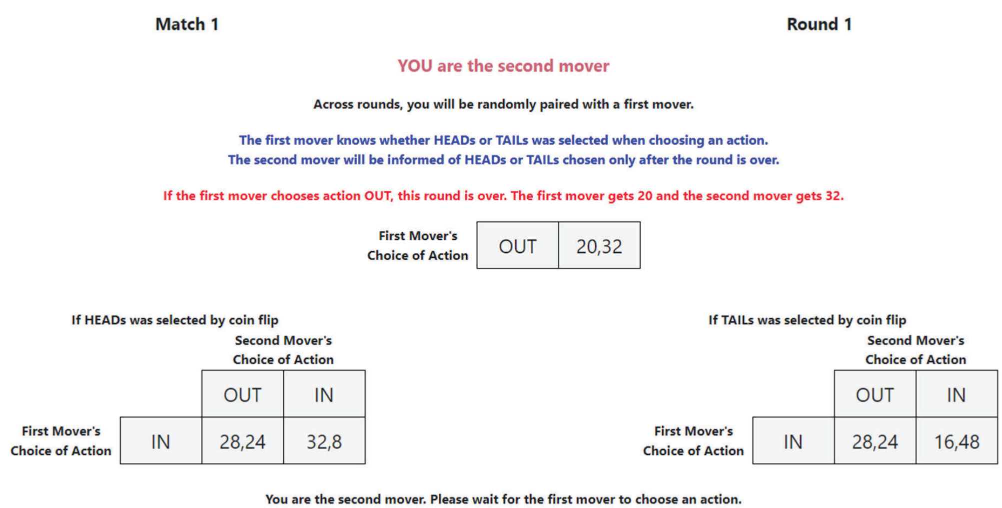

Figure 17: Slide 3

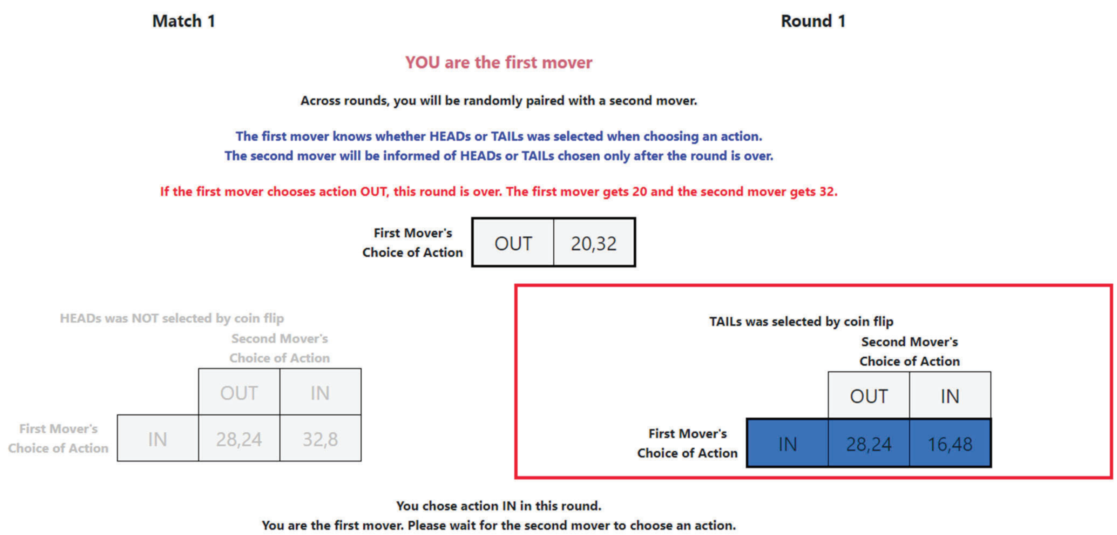

Figure 18: Slide 4

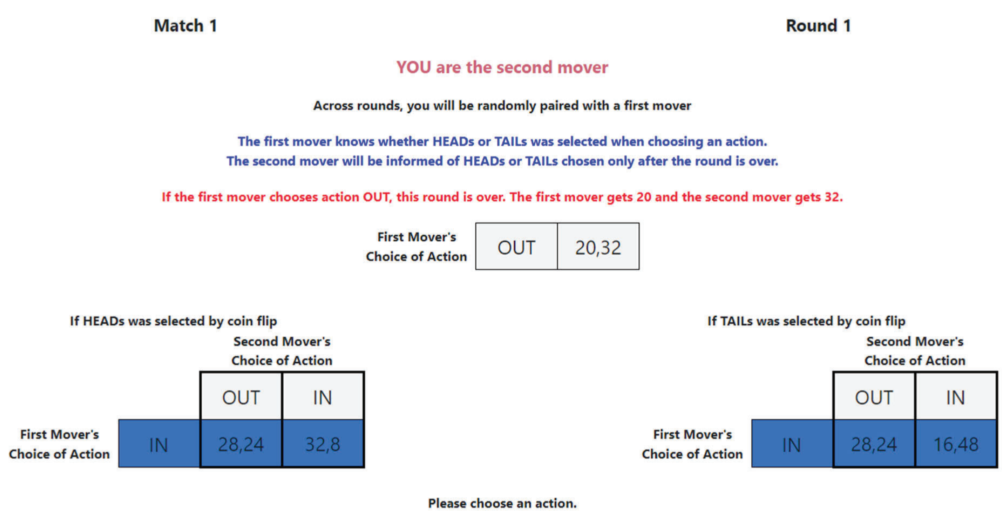

Figure 19: Slide 5

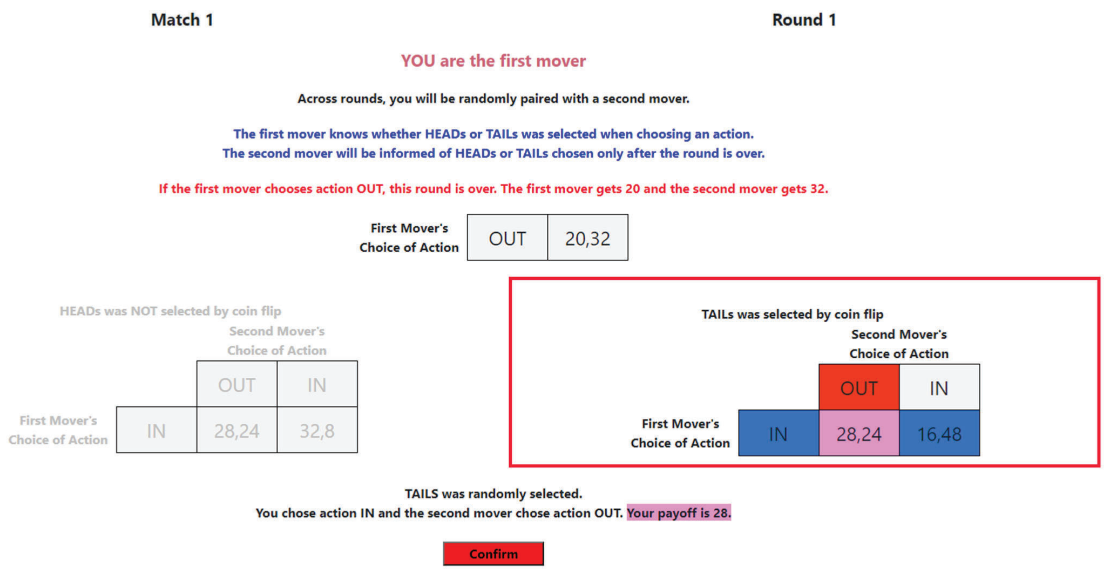

Figure 20: Slide 6

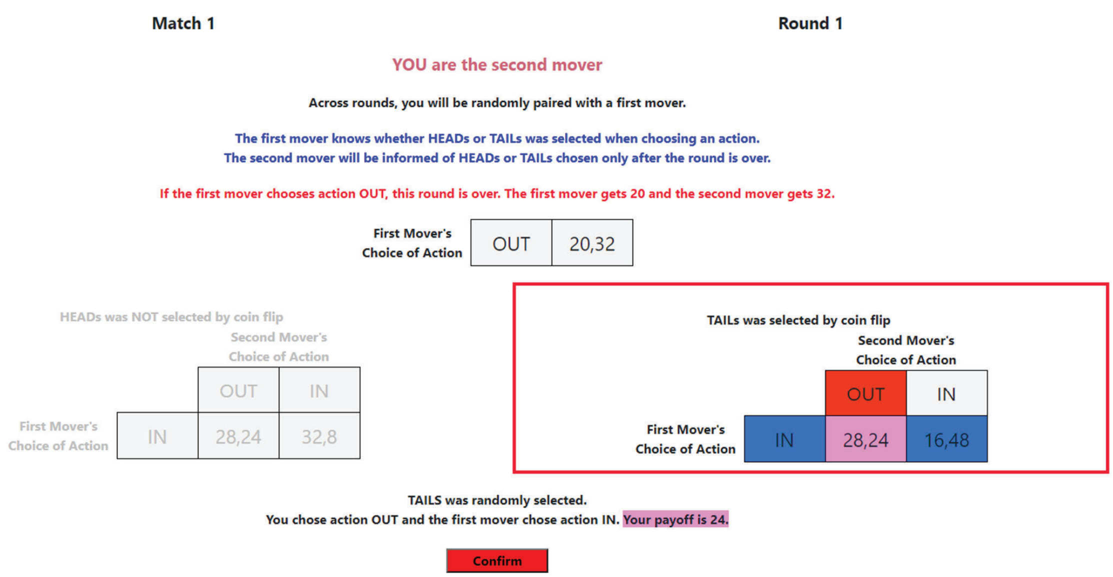

Figure 21: Slide 7

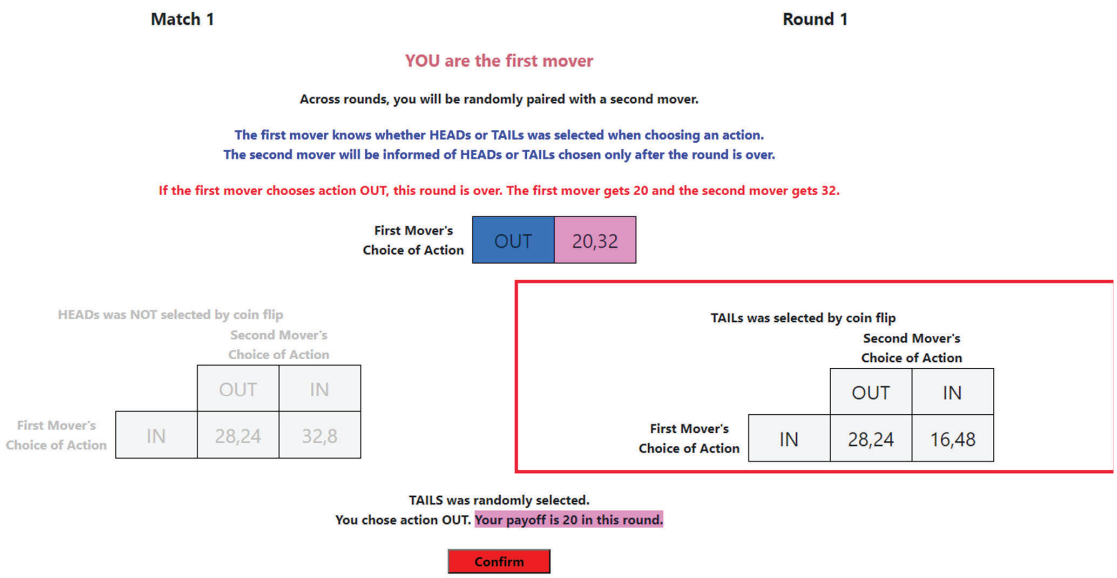

Figure 22: Slide 8