Strategic Complexity Promotes Egalitarianism in Legislative Bargaining[*]

Marina Agranov[†] S. Nageeb Ali[‡] B. Douglas Bernheim[§]
Thomas R. Palfrey[¶]

July 21, 2025

Abstract

Strategic models of legislative bargaining predict that proposers can extract high shares of economic surplus by identifying and exploiting weak coalition partners. However, strength and weakness can be difficult to assess even with relatively simple bargaining protocols. We evaluate experimentally how strategic complexity affects the ability to identify weak coalition partners, and for the partners themselves to determine whether their positions are weak or strong. We find that, as strategic complexity progressively obscures bargaining strength, proposers migrate to egalitarianism, in significant part because non-proposers begin placing substantial weight on fairness. Greater analytic skill dampens but does not eliminate these patterns.

Keywords: legislative bargaining, egalitarianism, complexity, laboratory experiment

JEL codes: D72, C92, C73

---

[*] We thank 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 thank the Hacker Social Science Experimental Laboratory (SSEL) at Caltech for laboratory facilities and for supporting the software development for the experiment. The paper has benefited from comments by participants at various seminar presentations. Palfrey thanks the Department of Economics at Columbia University for their hospitality during his visit as a Mitchell Professor in Fall 2019. We thank Salvador Candelas, Han Ng, and Jeffrey Zeidel for their excellent research assistance. We gratefully acknowledge the financial support of the National Science Foundation (SES-1426560, SES-1530639, SES-2343948). The experiment was approved by the Caltech IRB (Protocol IR17-0271).

[†] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA

[‡] Department of Economics, Pennsylvania State University, State College, PA

[§] Department of Economics, Stanford University, Palo Alto, CA

[¶] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA

1 Introduction

In legislative negotiations, legislators formulate successful proposals by making concessions in order to build winning coalitions. According to the applicable theoretical literature, two central principles govern this process. First, a skilled negotiator keeps the coalition as small as possible to minimize the number of concessions; second, the negotiator seeks out weak partners to minimize the sizes of those concessions.[1] However, it is also apparent from the pertinent literature that even relatively simple bargaining protocols make it difficult to assess which players are weak. In this paper, our broad objective is to evaluate the degree to which parties to such negotiations identify weak coalition partners and exploit perceived weakness. Our leading hypothesis is that strategic complexity leads people to make and support egalitarian proposals, even in settings where theory predicts that one party should be able to dictate the outcome.

To appreciate the strategic complexity of legislative bargaining, consider the nature of the task facing a proposer who seeks to craft the best deal for herself. The proposer anticipates that another player, if rational, would support a proposal only if its passage makes that player (weakly) better off than its rejection. But what happens following rejection may depend subtly on the dynamics of future play. To identify weak partners, a proposer must therefore understand how the process would subsequently play out. An additional wrinkle is that, even if the proposer can identify players who should, in theory, be willing to acquiesce for a small concession, she may not be entirely confident that weak players see themselves as weak. If weak players incorrectly believe that they are negotiating from positions of strength, proposals that attempt to exploit their objective weakness will fail. In such a scenario, a proposer might be inclined to propose a relatively egalitarian outcome to hedge against the possibility that she or others have incorrectly mapped out the future play.

The following example illustrates some of these subtleties. A three-person legislature consisting of Alice, Bob, and Carol is tasked with dividing one dollar. Bargaining takes place in a series of no more than three rounds. In each round, one legislator makes a proposal, which is then put to a vote. If any proposal receives two or more votes, it is adopted and the process ends. Alice makes the first proposal, Bob makes the second, and Carol makes the third. If all three proposals fail, the parties receive default payoffs (or outside options) $\mathbf{v} \equiv (v_A, v_B, v_C)$. We assume that $v_A + v_B + v_C \leq 1$, and $v_i \geq 0$ for each i. Players do not discount future payoffs.

---

[1] These principles emerge in the seminal work of Baron and Ferejohn (1989) and the subsequent theoretical literature on legislative bargaining.

Which partner is weaker from Alice's perspective, Bob or Carol? One might conjecture that the answer hinges on the default payoffs: a lower value of v_B (respectively v_C) makes Bob (respectively Carol) easier to exploit. And yet, barring a single knife-edge exception, in subgame perfect equilibria, Carol is *always* the weaker partner for Alice, even if v_C is large. In fact, Carol is so weak that she agrees to let Alice claim the entire dollar, which is the unique subgame perfect equilibrium outcome.[2]

To understand this result, suppose that $v_A < v_B$, which implies $v_A < 0.5$. We analyze this game by backward induction. If the first two offers are rejected and the third round is reached, Carol would form a minimal winning coalition with Alice by proposing $(v_A, 0, 1 - v_A)$, given that Bob's default payoff is higher. In the preceding round, Bob (the proposer) and the others expect to receive these continuation payoffs if his proposal is rejected. Because $v_A < (1 - v_A)$, Bob will propose $(v_A, 1 - v_A, 0)$, which will pass with Alice's support. Thus, Carol should expect to obtain zero if the first proposal is rejected and bargaining reaches the second round. Accordingly, in the first round, Alice will recognize and exploit Carol's extreme weakness and form a minimal winning coalition with her by proposing $(1, 0, 0)$.[3] Albeit with different dynamics, the same conclusion follows if $v_A > v_B$.[4] Intuitively, Carol's strong position in the third round induces Bob to exclude her, which in turn makes her extremely vulnerable from Alice's first-round perspective.

One might worry that these theoretical predictions require a high degree of strategic sophistication. For instance, if $v_B < v_C$, Alice might mistakenly infer that Bob is the weaker partner. Alternatively, even if Alice is sophisticated, she might fear that Carol and Bob are naive. If Carol falsely believes that she will benefit in the continuation game because her default payoff is high, she will refuse offers that she ought to accept. Once Alice is uncertain about her own strategic reasoning, or about Bob's and Carol's, she may hedge her bets by treating a coalition partner more generously, or even by making an offer that benefits everyone (i.e., the grand coalition). Thus, even if no player inherently favors egalitarian outcomes, strategic complexity may promote egalitarianism as a hedge against uncertainty concerning the quality of others' or one's own reasoning about the game.

---

[2] As is standard, we consider subgame perfect equilibria in which players do not play weakly dominated actions at the voting stage.

[3] We construe this proposal as a coalition with Carol because, in equilibrium, Carol resolves her indifference in favor of it. There do not exist equilibria in which Carol rejects with positive probability because she has a strict incentive to accept any proposal that offers her $\varepsilon > 0$.

[4] If the third round is reached, Carol will form a minimal winning coalition with Bob by proposing $(0, v_B, 1 - v_B)$. Consequently, if the second round is reached, Bob will form a minimal winning coalition with Alice by proposing $(0, 1, 0)$. Accordingly, in the first round, Alice will form a minimal winning coalition with Carol by proposing $(1, 0, 0)$. The knife-edge exception referenced in the text occurs when $\mathbf{v} = (0.5, 0.5, 0)$. In that case, there is an equilibrium in which Carol's proposal is $(0.5, 0, 0.5)$, Bob's proposal is $(0, 0.5, 0.5)$, and Alice's initial proposal is either $(0.5, 0, 0.5)$ or $(0.5, 0.5, 0)$.

We formalize our conjecture concerning egalitarianism by studying optimal offers in settings where the proposer does not know how others evaluate the continuation game that would follow rejection of her offer. We also assume that strategic complexity causes the proposer to either narrowly frame the problem by focusing only on the current stage game or hold pessimistic beliefs about her own continuation payoff. We show that, as the proposer's beliefs about others' strategic inferences become increasingly diffuse (in the sense of entropy) and their attention to or hopes for the continuation game dwindle, the optimal offer converges to a minimal winning coalition offer with an equal split.

We evaluate these issues, including the aforementioned conjecture, by studying legislative bargaining protocols in a laboratory experiment. We have three main objectives: (1) determine the extent to which people identify and exploit weak coalition partners; (2) ascertain the reasons for failures to do so in settings where they occur; and (3) identify the systematic behavioral patterns (potentially including egalitarianism) that arise in those settings.

With respect to the second objective, we decompose the main theoretical hypothesis into several components. The *exploitative intent hypothesis* holds that proposers wish to extract as much surplus as possible from bargaining partners. The *self-awareness hypothesis* holds that non-proposers understand their own weaknesses or strengths as coalition partners. The *exploitative ability hypothesis* holds that proposers have the ability to exploit weak coalition partners, in terms of both identifying them and having a sufficiently nuanced understanding of continuation play that they can fine-tune their offers for maximal exploitation. If bargaining theory's predictions concerning the exploitation of weakness prove to be inaccurate, the fault could lie in any one of these component hypotheses. In particular, the exploitative intent hypothesis could fail if social motives overwhelm personal interest (e.g., if people are sufficiently fair-minded or altruistic), and the self-awareness and exploitative ability hypotheses could fail if people have limited ability to reason strategically (e.g., if they struggle with backward induction).

In the laboratory, we study three-player bargaining games with majority rule so that the ideal minimum winning coalition includes a single weak partner. Our experimental design involves treatments that differ with respect to the *complexity* of the environment, the *predictability* of bargaining power in future rounds, and the *subject pool*. Because strategic reasoning is more complicated in games with larger numbers of bargaining rounds, we manipulate complexity by varying the length of the game from one round to three. We manipulate the predictability of bargaining power by varying the information subjects have about the identity of future proposers. According to the theoretical analysis of Ali, Bernheim and Fan (2019), greater predictability ought to allow initial proposers to obtain larger shares by weakening potential coalition partners. We vary the subject pool by recruiting participants at two universities: the University of California at Irvine (UCI) and the California Institute of Technology (Caltech). This variation is of interest because strategic reasoning draws on mathematical skills, and Caltech students score higher in mathematics on standardized tests.[5] Finally, all of our analysis focuses on the behavior of subjects who have gained experience with the pertinent protocols.

We begin with an analysis of average payoffs. Across our treatments, the first proposers receive smaller shares than standard theory predicts. The discrepancy is particularly striking for three-round games, where these shares are below one-half of the total prize on average. Contrary to theory, the first proposer's share in three-round games does not vary with the predictability of the proposer recognition order, which ought to affect the identifiability and exploitability of weak partners. Significantly, the first proposer's share declines from one-round to two-round games, and from two-round to three-round games, despite the fact that every multiple-round game can be reduced to a single-round game (with equilibrium payments for the continuation game playing the role of the default payments). This pattern suggests that, in more complex bargaining environments, proposers face greater difficulty identifying and exploiting players who believe they are in weak positions. Moreover, as the bargaining game becomes more complex, outcomes migrate towards egalitarianism. Although social motives offer a contributory explanation for incomplete extraction of surplus by the first proposers, they cannot account for these comparative statics.

The remainder of our analysis examines proposals and voting behavior to determine how participants arrive at these bargaining outcomes.

First, we study the nature of proposals. We classify a proposal as *dictatorial* if all non-proposing members would receive less than five percent, a *minimal winning coalition* (henceforth MWC) if exactly one non-proposing member would receive at least five percent, and a *grand coalition* if each non-proposing member would receive at least five percent. A proposal is *egalitarian* within a coalition type if the difference in shares between coalition partners (including the proposer) is no more than five percent of the prize. In one-round games, Caltech subjects overwhelmingly make MWC offers (86%), always to weak partners. Among those offers, only 9% are egalitarian, and the average proposer share is 77%. Thus, the proposers' behavior in these strategically simple settings shows clear indications of exploitative intent and exploitative ability. In two-round games, the fraction of MWC offers is about the same, but 12% of those offers involve partners who are strong according to theory; these partners appear to be weak based on default payoffs. The average proposer share in MWC offers to weak partners dips slightly to 69%, but egalitarianism remains relatively rare (7%). These patterns point to a small decline in exploitative ability. In three-round games for which theory unambiguously identifies a weak partner, the fraction of MWC offers is 77%, and 25% of those involve theoretically strong partners. Moreover, the average proposer share for MWC offers to weak partners is just over 50%, and 62% of these proposals are egalitarian. All the other proposals in these games are grand coalition offers, two-thirds of which are egalitarian. We interpret these patterns as reflecting a diminished ability to identify weak partners as well as a dramatically reduced inclination to try to exploit them.

---

[5] According to PrepScholar, the average Math SAT scores for UCI and Caltech students are 675 and 795, respectively.

To the extent exploitative intent is universal but exploitative ability is correlated with analytical skill, we might expect to see similar behavior at UCI in one-round games, along with a more immediate and dramatic shift toward egalitarianism as complexity rises. Our findings are generally consistent with that prediction. In one-round games, UCI subjects are significantly less likely to make MWC offers (68%), but nearly all of them are to weak partners (99%); of the MWC offers to weak partners, only 12% are egalitarian, and the average proposer share is 70%. Hence, we confirm exploitative intent, as well as a reasonably high degree of exploitative ability in this simple setting. In two-round games, the fraction of MWC offers increases to 74% but nearly two-thirds of them are to partners who are strong according to theory; such players appear weak because of their smaller default payoffs. Among those offers, the egalitarian fraction rises to 63%, and the average proposer share drops to 53%. UCI subjects also exhibit no ability to identify weak partners in the three-round games for which theory unambiguously identifies a weak partner: with perfect predictability, only 54% of the MWC offers include the weak partner, and with partial predictability, this fraction drops to 31%. Moreover, among MWC offers to weak partners in the perfect-predictability games, egalitarianism runs as high as 90% (with an average proposer payoff of 52%). Among grand coalition offers in games with perfect predictability, egalitarianism rises from 18-19% in one- and two-round games to 75% in three-round games. Thus, we see egalitarianism emerging to a striking degree as a response to strategic complexity.

The next step in our analysis is to examine votes (that is, responses to proposals). We relate voting behavior to three variables: the player's own share, whether the player is a weak or strong partner, and a measure of inequality. As expected, the likelihood of supporting a proposal always increases along with the player's own share.

We evaluate the self-awareness hypothesis by examining the relationship between voting and a player's theoretical bargaining strength. Specifically, we ask whether strong partners are more difficult to recruit as coalition partners. In one-round games, behavior is consistent with this implication. The same is true for two-round games, but the strong-partner effect weakens in the UCI sample. In three-round games, the effect weakens in the Caltech sample and disappears entirely in the UCI sample. Thus, complexity reduces self-awareness of weakness, much as it reduces the ability of proposers to identify and exploit weakness. Interestingly, we find some evidence that the ability to identify weakness declines more rapidly with complexity for proposers than for non-proposers.[6] A possible explanation is that non-proposers are more likely to think ahead to the future proposals they may make.

We provide further evidence concerning the hypothesis that complexity promotes egalitarianism by examining the relationship between voting and fairness. There is no indication that fairness matters to voters in one- and two-round games at Caltech, or in one-round games at UCI, but it does matter to them in three-round games at Caltech, and in two- and three-round games at UCI. These results imply that, to some degree, the observed relationship between the frequency of egalitarian proposals and strategic complexity is a rational response to voters' strategies.

The final portion of our analysis evaluates the optimality of proposals given the voting behavior. We ask whether the migration towards egalitarian outcomes in more complex games is an optimization failure on the part of proposers or approximately optimal given the voting behavior. Our evidence supports the latter possibility: conditional on observed voting behavior, the optimal proposal migrates toward within-coalition egalitarianism as we increase the number of rounds. Consequently, actual offers to MWCs with weak partners are nearly optimal. Egalitarian offers to grand coalitions also become progressively less suboptimal as complexity rises. Accordingly, the nature of voting responses by non-proposing players accounts for much of the observed migration to egalitarianism as complexity rises.

The paper is organized as follows. Section 2 relates our work to the pertinent literature. Section 3 explains the theoretical principles that guide our analysis and provides a conceptual foundation for the hypothesis that greater complexity leads to egalitarianism. Section 4 details the experiment. Section 5 presents our experimental findings. Section 6 concludes. The appendices include tables supporting all bar graphs, additional tables and figures, and sample instructions.

---

[6] Specifically, in two-round games at UCI, non-proposers identify weakness more successfully than proposers.

2 Related Literature

This paper draws on three important but largely separate areas of study: one on coalition formation and legislative bargaining, another on social preferences, and a third concerning the effects of complexity on decision making. The primary contribution of our paper is that it connects these ideas, offering a new rationale for fairness in legislative bargaining.

The literature on legislative bargaining experiments, surveyed in Palfrey (2016) and Agranov (2021), has generally focused on the degree to which proposers form and exploit minimal winning coalitions, whether institutional details such as the use of a closed-rule or open-rule bargaining protocol affects this behavior, and whether the comparative statics highlighted in Baron and Ferejohn (1989) emerge in practice. Because most of the experimental treatments involve symmetrically positioned respondents and do not vary the complexity of the game, the questions at the core of this paper—namely, whether players can identify weak partners and how offers change with varying complexity—have received little attention. In this vein, some work studies three-player coalitional bargaining games with heterogeneous outside options. Diermeier and Gailmard (2006) consider bargaining with a single-round, which resembles a three-player ultimatum game with simple majority rule, and find that the proposer has a slight tendency to target the player with the lowest disagreement value.[7] Miller, Montero and Vanberg (2018) study an environment with an indefinite horizon, and find that proposers are more likely to exclude the player with the higher disagreement payoff. These studies anticipate our finding concerning the validity of the exploitative intent hypothesis by showing that proposers take advantage of obviously weak players.

Several early experiments on two-person offer/counteroffer bargaining with a shrinking pie present results that bear on some of the questions we ask here. Those studies interrogate the conditions under which negotiators act as "fairmen" or "gamesmen" and the extent to which such behavior depends on the length of the bargaining horizon and the rate of pie shrinkage. The results are mixed. In the first such study, Binmore, Shaked and Sutton (1985) compares one- and two-round bargaining games with a shrinking pie. The outcomes they observe in the more complex two-round game are *less* egalitarian than in the original one-stage ultimatum game. First movers in the two-stage game are gamesmen, while first movers in the ultimatum game are fairmen.

---

[7] Knez and Camerer (1995) were the first to study an ultimatum game with three players. Their experiment involved a proposer playing two separate ultimatum games in which the respondents had unequal outside options. The role of the third party was primarily to facilitate social comparisons. Güth and Van Damme (1998) consider a setting in which one of two respondents has veto power, while the other is a dummy player with no ability to influence the outcome. Contrary to the predictions obtained from models with social preferences, they find that the proposer targets the respondent with veto power, and the dummy player receives a very small share.

The Binmore et al. (1985) study triggered a series of follow-ups that involved different variations on the shrinkage rate, the number of periods, and other design features. Güth and Tietz (1988) increased the shrinkage factor and raised the monetary stakes. Their results strongly rejected the game-theoretic solutions in all treatments and suggest that such theoretical failures will inevitably arise in situations where the equilibrium is socially unacceptable. Neelin, Sonnenschein and Spiegel (1988) compared behavior in two-, three-, and five-round bargaining games, with mixed results about the degree to which proposers follow the subgame perfect equilibrium strategy. The findings of Spiegel, Currie, Sonnenschein and Sen (1994) suggest that proposers try to achieve egalitarian outcomes when they are at a bargaining disadvantage. Binmore, Morgan, Shaked and Sutton (1991) identify a potential role for simple rules of thumb and social norms, such as fairness, that bargainers might apply instinctively in some environments. However, the authors are agnostic about the conditions that trigger the application of such rules of thumb and cause proposers to abandon more strategic approaches.

Because these studies focus on bilateral bargaining, negotiators do not need to identify weak partners as in the more complex multi-person bargaining problems that we examine here. Our analysis highlights how, in the latter settings, people are willing to act as “gamesmen” when the bargaining protocol is easy to game, but switch to a hybrid “fairmen/gamesmen” strategy when it is complex, splitting the prize with a single participant while leaving the other with nothing.

Our work is also related to a vast literature on fairness and social preferences; see Fehr and Charness (2024) for an exhaustive survey. In the settings we study, people do not exhibit much intrinsic concern for others’ outcomes. In simple settings, they clearly exhibit exploitative intent, and even in complex settings, most proposals completely exclude one member of the group. Thus, it appears that our subjects are not intrinsically motivated either by fairness or, as in Andreoni and Bernheim (2009), by the desire to appear fair. Nevertheless, egalitarianism *within* coalitions appears to emerge when strategic complexity makes it difficult to understand and anticipate the outcome of subsequent negotiations.

Finally, our work contributes to the broad research agenda on how complexity influences decision making both in general and in strategic settings. Within economics, the study of behavioral responses to complexity traces its roots to Herbert Simon’s seminal work on bounded rationality (Simon, 1956). Over the last few years, interest in understanding how people make complex decisions has exploded (e.g., Bernheim and

Sprenger, 2020; Enke and Graeber, 2023; Oprea, 2024; Ba, Bohren and Imas, 2024; Puri, 2025); see Oprea (2025) for a survey. The inherent complexity of strategic reasoning has inspired the formulation of theories that seek to capture essential aspects of boundedly rational play, such as level-k reasoning (e.g., Nagel, 1995; Stahl and Wilson, 1995; Crawford, Costa-Gomes and Iriberri, 2013), cognitive hierarchies (e.g., Camerer, Ho and Chong, 2004; Alaoui and Penta, 2016), and finite automata (e.g., Neyman, 1985; Abreu and Rubinstein, 1988; Kalai, 1990; Chatterjee and Sabourian, 2009), as well as the search for strategically-simple mechanisms (surveyed in Li 2024). There is also evidence that people favor simple strategies in the infinitely repeated prisoner's dilemmas, and that simplicity can make particular strategies focal (Guan and Oprea, 2024). Our measure of complexity—the number of rounds in a bargaining game—is ordinally equivalent to the number of subgames, and is therefore consistent with the measure used in work on finite automata. Our findings complement the complexity agenda by identifying a new effect of strategic complexity on behavior, namely that it promotes egalitarianism in bargaining.

3 Theoretical Principles

In this section, we describe the theoretical principles that guide our experimental inquiry. Section 3.1 describes equilibrium predictions. Section 3.2 explains why strategic complexity may lead to egalitarianism.

3.1 Equilibrium predictions

Our analysis concerns multi-party finite-round bargaining protocols. We focus on settings with three parties but the key principles discussed in this subsection generalize. In each round $t \in \{1, ..., T\}$, one of the three players, $p_t \in \{A, B, C\}$, proposes a division of a fixed prize which, without loss of generality, we take to be one dollar. We write a proposal as (x^A, x^B, x^C), where $\sum_{i \in \{A,B,C\}} x^i = 1$. All players then vote in favor or against this proposal. A proposal passes if a simple majority votes in favor, in which case negotiations end. If no proposal passes by the end of round T, the parties receive default payoffs (v_T^A, v_T^B, v_T^C), where $\sum_{i \in \{A,B,C\}} v_T^i \leq 1$. We assume each player cares only about their own share of the prize. We also assume no discounting and risk neutrality, but these assumptions are not essential.

The sequence of proposers (or *recognition order*) $P = (p_1, ..., p_T)$ is assigned at random according to some probability distribution μ. The parties learn about the recognition order as negotiations proceed. We consider three informational protocols. For the *Perfect* protocol, the parties learn proposers' identities one period in advance; in other words, p_t is disclosed between rounds $t-2$ and $t-1$, and the parties start out with knowledge of p_1 and p_2. For the *Partial* protocol, they learn the identity of one randomly selected *non-proposer* one period in advance: the parties start out knowing p_1 and that some j will not be the period-2 proposer; $j \neq p_t$ is then disclosed between rounds $t-2$ and $t-1$. For the *None* protocol, the parties have no advance knowledge of proposers' identities. In other words, they start out knowing only p_1, and p_t is disclosed between rounds $t-1$ and t.

This framework encompasses a wide range of possibilities. For example, most theoretical and experimental extensions of Baron and Ferejohn (1989) study the *None* protocol in settings where the selection of each round's proposer is independent and, typically, the probabilities are the same for all three parties. Ali et al. (2019) model the effects of predictability in legislative bargaining, considering a range of informational protocols that include the partial and perfect protocols described above.

As is standard in this literature, we study pure-strategy subgame-perfect equilibria with the following features. First, in each round, strategies prescribe choices that do not depend on the features of previously rejected proposals or the associated profile of votes. Second, when casting votes, non-pivotal voters resolve indifference as if they are pivotal.[8] Conditional on the recognition order P, any such equilibrium induces a vector of expected continuation payoffs, (v_t^A, v_t^B, v_t^C), that the parties receive if they reject the round-t proposal. If $v_t^i < v_t^j$ for $i, j \neq p_t$, we say that i is the *weak round-t partner*, and that j is the *strong round-t partner*.

Some portions of our empirical analysis focus on the match between observed behavior and specific quantitative predictions for individual bargaining games, while other portions test a small collection of less precise predictions that are arguably easier for negotiators to grasp. Next, we describe three theoretical principles that encompass those implications.

Theoretical Result 1*: In equilibrium, proposals always assign the proposer a share of at least 0.5 and at least one other party a share of 0. Strong partners never receive positive shares. All equilibrium proposals pass.*

The logic of Theoretical Result 1 is straightforward and extremely general. A round-t proposal passes as long as $x_t^j \geq v_t^j$ for at least two parties. The best successful proposal is plainly one that secures passage while conceding no more than necessary. The proposer accomplishes this objective by offering $x_t^i = v_t^i$ for $i = \arg\min_{i \neq p_t} v_t^i$, $x_t^j = 0$ for $j \notin \{p_t, i\}$, and $x_t^{p_t} = 1 - v_t^i$ for themselves. Because $v_t^i \leq v_t^j$, we know that $v_t^i \leq 0.5$, so $x_t^{p_t} \geq 0.5$. Furthermore, because $1 - v_t^i \geq v_t^{p_t}$, the best successful proposal delivers a higher share for the proposer than any unsuccessful proposal.

---

[8] Equivalently, they do not make weakly dominated choices at the voting stage.

A common interpretation of this first result is that proposers form minimum winning coalitions (MWCs) with weak partners, to the extent they exist. If $v_t^i = 0$, the proposer offers nothing to the other parties, which somewhat obscures the MWC interpretation. In the associated equilibrium, party i resolves their indifference in favor of the proposal, thereby joining the MWC. Were i to resolve her indifference differently, the proposer would offer i a small token share, thereby securing almost all of the prize.

Theoretical Result 2: *Suppose $p_{t+1} \neq p_t$. When the identity of the next proposer is known one round in advance (as is the case under the Perfect protocol), p_{t+1} is never the weak round-t partner for the proposer p_t.*

Theoretical Result 2 is also general and intuitive. From Theoretical Result 1, we know that p_{t+1} will obtain a payoff of at least 0.5 if round t is reached. Because the identity of p_{t+1} is known throughout round t, it follows that $v_{t+1}^p \geq 0.5$. But $v_t^i + v_t^{p_{t+1}} \leq 1$ for $i \notin \{p_t, p_{t+1}\}$, so $v_t^i \leq v_t^{p_{t+1}}$, which means p_{t+1} is not a weak partner in round t.

An important and subtle implication of this result is that the identity of the weak partner can be completely unrelated to the default payoffs (v_T^A, v_T^B, v_T^C), and solely determined by the recognition order. We conjecture that, in practice, failures to properly identify weak partners may be related to this distinction.

Theoretical Result 3: *Predictability of the recognition order strengthens the proposer's bargaining position by weakening at least one of her potential partners. Perfect predictability of the next proposer enables her to secure the entire prize.*[9]

Ali, Bernheim and Fan (2019) make the second point with considerable generality. Here we illustrate it for the special case in which (i) the selection of each round's proposer is independent and based on a distribution that assigns equal probability to all three parties, and (ii) all players' default payoffs are zero. We focus on symmetric equilibria in which, whenever a proposer is indifferent between two potential coalition partners, she chooses each with equal probability.[10]

---

[9] For the illustration in the introduction, we made the stronger assumption that the *entire* recognition order is known in advance.

[10] In an infinite-round setting, Baron and Ferejohn (1989) show that this property follows from stationarity.

First imagine that the recognition order is entirely unpredictable (the *None* protocol). In that case, the problem's symmetry and our focus on symmetric equilibria ensures that $(v_t^A, v_t^B, v_t^C) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for all $t < T$. It follows that, for such t, the round-t proposal assigns $\frac{2}{3}$ to the proposer and $\frac{1}{3}$ to a randomly selected partner.

Next imagine that the identity of each proposer is known one round in advance (the *Perfect* protocol). With default payoffs of $(0, 0, 0)$, the proposer in the final round claims the entire prize. The proposer in the penultimate round can therefore always identify a potential partner who will receive zero in the continuation—specifically, a party who will not make the last offer. Consequently, the penultimate proposer also claims the entire prize. Through a repeated application of the same argument, we see that the proposer claims the entire prize in every round with the support of a party who does not propose next. Comparing this outcome to the one that emerges under the *None* protocol, we see that predictability of the recognition order augments proposer power, effectively turning each proposer into a dictator.

According to Ali, Bernheim and Fan (2019) (Theorem 1), the preceding conclusion holds even when information about the next proposer is imperfect, as long as it always rules out more than half of the parties. An ability to rule out only one party (the *Partial* protocol) does not satisfy this requirement. Consequently, theory predicts that the proposer should fare substantially better with the *Perfect* protocol than with the *Partial* protocol.

To refine this prediction, we solve for the equilibrium of a three-round bargaining game under the *Partial* protocol. The round-3 proposer always obtains the entire prize. There are two possibilities for the round-2 proposer. If she learns she will not be the final round-3 proposer (a case we henceforth label *Excl*), then her two potential coalition partners each have an expected continuation payoff of $\frac{1}{2}$. Therefore, her optimal strategy is to offer $\frac{1}{2}$ to a randomly selected partner, leaving a share of $\frac{1}{2}$ for herself, in which case each non-proposer receives an expected payoff of $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$. If she learns she might be the next proposer (a case we henceforth label *Incl*), then her best option is to claim the entire prize by forming a coalition with the non-proposer who will not be the round-3 proposer. Consequently, the expected payoff for the round-2 proposer is $\frac{1}{3} \times \frac{1}{2} + \frac{2}{3} = \frac{5}{6}$, and the expected payoff for each round-2 non-proposer is $\frac{1}{12}$. Likewise, there are two possibilities for the round-1 proposer. If she learns she will not be the next proposer (the *Excl* case), then her potential partners each have an expected continuation share of $\frac{1}{2} \times \frac{5}{6} + \frac{1}{2} \times \frac{1}{12} = \frac{11}{24}$. Her best option is to claim a share of $\frac{13}{24}$ and offer $\frac{11}{24}$ to a randomly selected partner. If she learns she might be the next proposer (the *Incl* case), then one of her potential partners has an expected continuation fraction of $\frac{1}{12}$, which means she can claim $\frac{11}{12}$. Overall, her expected payoff is $\frac{1}{3} \times \frac{13}{24} + \frac{2}{3} \times \frac{11}{12} = \frac{55}{72}$, which substantially exceeds $\frac{2}{3}$.

The preceding reasoning demonstrates that, even though a small degree of recognition-order predictability does not render the first proposer a dictator, she still obtains a higher expected payoff with the *Partial* protocol than with the *None* protocol. The higher payoff materializes only when she learns she might be the second proposer (the *Incl* case); in fact, she is worse off relative to the *None* protocol when she learns she will not be the second proposer (the *Excl* case), but she gains on average. In contrast, with the *Perfect* protocol, her payoff does not depend on whether she learns she definitely will be the next proposer (which serves as the *Incl* case for the *Perfect* protocol), or definitely will not be the next proposer (which serves as the *Excl* case for the *Perfect* protocol). This difference is attributable to the fact that she can identify a weak partner even in the *Excl* case for the *Perfect* protocol (i.e., the other non-proposer for the next round), but not with the *None* protocol.

3.2 Why Complexity May Lead To Egalitarianism

The strategic reasoning underlying equilibrium predictions becomes increasingly complex as the number of stages and subcases proliferate and may be too subtle for typical negotiators to grasp. Even savvy negotiators who appreciate the logic of equilibrium may lack confidence that others share their strategic sophistication. For instance, they may worry that potential partners are excessively optimistic or pessimistic about their own continuation payoffs. Consequently, a key issue that we investigate experimentally is whether variations in the complexity of the negotiating environment affect the accuracy of the theoretical predictions, leading systematically to alternative behavioral patterns. In particular, if complexity increases a proposer's uncertainty about her ability to exploit coalition partners and makes her more concerned about her vulnerability to exploitation by others, offering more egalitarian proposals may be a tempting safe alternative. In that case, changes to the environment that add complexity could encourage more egalitarian behavior.

We formalize this point using a simple framework in which a proposer is uncertain about the offers other players would support. Imagine the current-round proposer (Player A) believes that player $i \in \{B, C\}$ has an acceptance threshold τ_i such that i votes for the proposal iff x_i is at least τ_i.[11] One interpretation of τ_i is that it reflects player

---

[11] Thus, players resolve indifference in favor of supporting the proposal. In addition, each player's acceptance decision depends only on her own share.

i's perceived continuation value. In complex settings, the proposer (player A) may be highly uncertain about τ_B and τ_C because she recognizes that each player's perceptions of continuation values may stray far from equilibrium predictions. We model this uncertainty by assuming that the proposer's beliefs about others' thresholds correspond to some joint CDF that allows for arbitrary correlations. The proposer also believes that she will receive an expected payoff of $d < 1$ in the continuation game if her offer is rejected. Finally, we assume the proposer is risk neutral.

We are interested in behavior near a boundary case in which the game's complexity has two effects on the proposer. First, it causes the proposer to hold diffuse beliefs reflecting deep uncertainty about the other players' thresholds. Second, it leads the proposer to either (i) narrowly frame the problem by focusing on the stage game in isolation, or (ii) hold pessimistic beliefs about the continuation game. We model diffuse beliefs as involving a bivariate CDF H for which the marginal distributions of each τ_i $(i = B, C)$, written H_i, are uniform on $[0, \overline{\tau}]$, where $\overline{\tau} \geq \frac{1}{2}$. The assumption concerning the magnitude of $\overline{\tau}$ allows for the possibility that a coalition partner might insist on equal division within an MWC but does not necessarily require a belief that the partner might insist on a share close to unity. In addition to having the intuitive interpretation that all possibilities are equally likely, the uniform distribution reflects "maximal uncertainty" because it achieves the highest entropy among all continuous distributions on its support. Consequently, other studies have also used it to represent minimally informed beliefs.[12] Notably, this boundary case allows for correlation, positive or negative, between the two thresholds. The following assumption places a modest restriction on this correlation:

Assumption 1. *For every $x > 0$, $H(x, x) > 0$.*

Assumption 1 states that for every strictly positive value x, there is a strictly positive probability that players B and C would both support the offer $(1 - 2x, x, x)$. Notably, this assumption rules out perfect negative correlation between players' thresholds. It is always satisfied when H has full support on $[0, \overline{\tau}] \times [0, \overline{\tau}]$.

We examine sequences $(H_n, d_n) \to (H, d)$ where, for the purpose of defining convergence, we endow the space of CDFs with the weak topology. In stating our proposition, we follow the convention that player B is the individual to whom the proposer offers the weakly larger share $(x_B \geq x_C)$. The following proposition shows that optimal offers in the neighborhood of (H, d) involve minimum winning coalitions.

---

[12] For example, in the literature on levels of strategic thinking, the beliefs held by higher level players about a level-0 player's strategy are often assumed to be uniform. See, for example, Camerer et al. (2004) and Crawford et al. (2013).

Proposition 1. *As $(H_n, d_n) \to (H, d)$, the optimal offers converge to* $\left(\frac{1+d}{2}, \frac{1-d}{2}, 0\right)$.

We are especially interested in the possibility that, in response to the problem's complexity, the proposer either narrowly frames the problem by focusing exclusively on the stage game, or makes highly pessimistic assumptions about the continuation game. Such possibilities correspond to the boundary case of $d = 0$. Proposition 1 implies that as we approach $(H, 0)$, the optimal offer converges to complete egalitarianism within an MWC, i.e., $(1/2, 1/2, 0)$. Even if the proposer is not completely pessimistic about her continuation payoff d, her optimal offer may still involve a reasonably equal split. For instance, if she expects a share of 20% in the continuation game, and she has diffuse beliefs about others' acceptance thresholds, she will propose a 60-40 split in an MWC.

While Proposition 1 is intuitive, its proof requires attention to various technical details.[13] We outline the key steps here, relegating the complete argument to Appendix A. One key step is to show that, for the limit case (H, d), the optimal offer takes the form described in the proposition.[14] In another key step, we show that, for arbitrary beliefs, (H, d), the proposer's payoff function is upper semicontinuous in the offer, which guarantees existence of an optimal offer. Further, we show that, in a neighborhood of (H, d), the optimal offer correspondence is upper hemicontinuous in the proposer's beliefs about other players' thresholds and her own expected payoff in the continuation game following a rejection.[15] Therefore, when beliefs are close to the boundary case of $d = 0$, egalitarian offers targeting minimum winning coalitions are approximately optimal.

The preceding analysis shows how deep uncertainty about others' thresholds coupled with narrow framing or pessimism about the continuation game can give rise to egalitarian MWC offers. Given that we also observe grand coalition (GC) offers in our experiment, one may wonder why proposers also find those offers attractive. We highlight three potentially relevant considerations.

First, strategic uncertainty can substantially reduce the difference in expected payoffs from MWC and GC offers. For example, while the difference between the proposer's payoffs in the egalitarian MWC and GC allocations is 0.17 (0.5 versus 0.33), the difference in expected payoffs from the corresponding offers with uniform-independent priors, $\overline{\tau} = 1$, and $d = 0$ is only 0.065 (0.25 versus 0.185), partly because the GC offer provides a hedge. A preference for fairness is, therefore, more likely to outweigh the expected monetary gain, especially if that preference pertains to the offer itself rather than to the implemented allocation—for example, because acting fairly has deontological force as in Andreoni, Aydin, Barton, Bernheim and Naecker (2020) or signaling value as in Andreoni and Bernheim (2009).

---

[13] We benefited from discussing this argument with Andreas Kleiner.

[14] For this step, the restriction to three-player bargaining games plays a role. With more than three players, one can show that the optimal MWC offer is egalitarian when $d = 0$ if players' thresholds are distributed independently. However, offers that seek to form larger coalitions may do better, at least for large $\overline{\tau}$. For example, with five players, $\overline{\tau} = 1$, and $d = 0$, an offer with equal shares among three players (including the proposer) yields an expected payoff of 0.0370, while offering shares of 0.22 to three non-proposers yields an expected payoff of 0.0420.

[15] We establish this result using a variant of the Theorem of the Maximum that holds for upper semicontinuous payoff functions. Berge's Theorem does not apply because the proposer's payoff function is discontinuous.

Second, proposers may believe that the fairness of an outcome elevates the probability of acceptance. In other words, there may be a consensus that equal division is reasonable in settings that feel mostly random due to their complexity. Notably, this consideration can also reinforce the attraction of egalitarianism within MWCs.

Third, despite understanding that one potential partner is likely weaker than the other, a proposer may not be sure which is which. Thus, she may believe that τ_i and τ_j are negatively correlated. This possibility increases the value of the hedge embedded in a GC offer. To illustrate, suppose a proposer with $d = 0$ believes τ_i is uniformly distributed on $[0, 1]$, and that $\tau_j = 1 - \tau_i$. In that case, when offering s_i and s_j, the proposer's expected payoff is $(1 - s_i - s_j)(s_i + s_j)$. Because the two first-order conditions are redundant, the optima consist of the locus $s_i = \frac{1}{2} - s_j$. Consequently, offering $\frac{1}{2}$ to one party and nothing to the other is still optimal, but so is offering $\frac{1}{4}$ to each party. Notably, the latter option minimizes the Gini coefficient among all proposals yielding the maximized payoff. Consequently, if the proposer has a tiny preference for fairness or considers such offers slightly more likely to be accepted, a Grand Coalition offer emerges as the most attractive alternative. Furthermore, the associated difference between the expected payoff from an optimal proposal and an egalitarian GC proposal falls to 0.028 (0.250 versus 0.222), making the latter option even more attractive for those who care about fairness even to a small degree.

Taken as a whole, the analysis of this section underscores an important conceptual point: with strategic uncertainty, egalitarianism among MWCs, GCs, or both may become increasingly attractive in complex settings.

4 Experimental Design

The experiment examines behavior in three-person finite-round bargaining games under majority rule.[16] In each game, participants negotiate the division of a fixed prize, denominated in "points," which are convertible to dollars at a known rate.[17] The games we consider differ with respect to the number of rounds and the informational protocol.[18]

For three-round negotiations, we examine all informational protocols mentioned in the preceding section (*Perfect*, *Partial*, and *None*). As in our examples, we set all the default payments equal to zero.

For the two-round negotiations, all parties know the identities of both proposers at the outset. We set default shares to 0.05 for the first-round and second-round proposers (who always differ), and to 0.90 for the non-proposer. Our objective is to determine how the players would behave in a three-round bargaining problem—one with default payments of zero—if they thought through the third-round outcome by performing one round of backward induction (approximately) correctly. By comparing behavior between these two-round games and the *Perfect* three-round games, we can determine the extent to which departures from theoretical predictions in the latter result from difficulties associated with thinking forward to third-round outcomes.

For the one-round negotiations, there is no distinction between the three informational protocols. We set default shares to 0.05 for the first-round proposer and one of the non-proposers, and to 0.90 for the other non-proposer. Our objective is to determine how the players would behave in a three-round *Perfect* bargaining problem if they thought through the second-round and third-round outcomes by performing two rounds of backward induction (approximately) correctly. By comparing behavior between the one-round and three-round games, we can determine the extent to which departures from theoretical predictions in the latter result from difficulties associated with thinking forward to both second-round and third-round outcomes.

Notably, in the one-round bargaining problems, it is reasonably obvious that the weak potential partner is the one with the smaller default share (0.05). In contrast, in the two-round bargaining problems, the weak potential partner is the one with the *larger* default share (0.90). The first proposer must reason out that strength and weakness are consequences of the recognition order (specifically, who proposes second) rather than the default payments.

---

[16] We also ran some small-scale pilots involving indefinite-horizon games with discounting and communication among committee members. The data from these treatments are available upon request.

[17] For example, in some treatments the participants divide a prize of 240 points. In that case, a share of 0.90 corresponds to 216 points.

[18] There are some other inessential differences in details across the treatments, including the total number of points and the conversion rate from points to dollars.

Table 1 summarizes the equilibrium offers for each experimental treatment. We explained the predictions for three-round bargaining games in the previous section. The logic behind the implications listed for the one-round and two-round bargaining games is straightforward. We attach the *Perfect* label to the one-round and two-round games to underscore their conceptual connection to the three-round *Perfect* games.

Table 1: Equilibrium Offers in Bargaining Games

Game	Round 1	Round 2	Round 3
1-round, *Perfect*	0.95^*	NA	NA
2-round, *Perfect*	1.0^*	0.95^*	NA
3-round, *Perfect*	1.0^* (*Incl*) 1.0^* (*Excl*)	1.0^* (*Incl*) 1.0^* (*Excl*)	1.0
3-round, *Partial*	0.92^* (*Incl*) 0.54 (*Excl*)	1.0^* (*Incl*) 0.50 (*Excl*)	1.0
3-round, *None*	0.67	0.67	1.0

Notes: The table shows the share of the prize the proposer seeks to keep for herself. The proposer offers the residual share to a single partner. * indicates that the proposer targets a weak partner.

We conducted twelve sessions at UCI and nine at Caltech, each drawing on the standard student population of subjects. Each session consisted of either 15 or 20 repetitions of the same bargaining game. We refer to each bargaining game as a *treatment*, and to each repetition as a *match*. Table 2 summarizes the numbers of matches, sessions, and subjects for each treatment at UCI and Caltech.

The number of subjects for each session was a multiple of 3, and ranged from 9 to 24. For the first match of each session, we randomly divided the subjects into 3-member groups, which we call *committees*. In each subsequent match, we reshuffled the groupings randomly so that the composition of the committees changed unpredictably over the course of a session.

At the beginning of each match, we randomly assigned each of the three committee members an ID number (1, 2, or 3), which they kept for the duration of the match. We randomly assigned new ID numbers in subsequent matches. Committee member 1 always made the first proposal. For the 2-round treatment, we told subjects at the outset of the match that committee member 2 would be the second proposer. For the 3-round *Perfect* treatments, we publicly announced the ID number of the next-round proposer if the current proposal failed. For the 3-round *Partial* treatments, we publicly announced the ID numbers of two players who were eligible to be the next-round proposer if the

Table 2: Treatments and Sessions

	Label	# Sessions	# Matches	# Subjects
UCI				
1-round, *Perfect*	1-Perfect	3	15	57
2-round, *Perfect*	2-Perfect	3	15	57
3-round, *Perfect*	3-Perfect	2	20	42
3-round, *Partial*	3-Partial	2	20	48
3-round, *None*	3-None	2	20	48
CALTECH				
1-round, *Perfect*	1-Perfect	3	15	33
2-round, *Perfect*	2-Perfect	3	15	33
3-round, *Perfect*	3-Perfect	3	20	39

Notes: The 1-round and 2-round treatments were conducted within the same session at both locations (Part I and Part II of the same experiment, respectively).

current proposal failed. For the 3-round *None* treatments, we provided no information about the identity of the next-round proposer.

5 Results

In the analyses that follow, we focus on the second half of all treatments, i.e., the last eight matches in the 1-Perfect and 2-Perfect treatments and the last 10 matches in the 3-None, 3-Partial, and 3-Perfect treatments. We impose this filter because our interest is in behavior after subjects have had the opportunity to experience the game and converge to a stable behavior. We refer to these matches as *experienced*.[19] Statistical tests are performed using regression analyses, in which we cluster standard errors at the session level to account for the interdependencies that come from re-matching subjects between bargaining games. To compare average outcomes between treatments, we regress the variable of interest on a constant and an indicator for one of the two treatments. The p-values we report are those associated with the estimated coefficient on the dummy for the treatment variable.

---

[19] Analyses for all matches, both the first and second halves of each treatment, appear in Appendix D. The behavior is largely similar.

5.1 Bargaining Outcomes

We begin by describing, at a coarse level, the observed bargaining outcomes. We first examine the extent to which the first proposer successfully appropriates surplus (Section 5.1.1) and then turn to the distribution of payoffs (Section 5.1.2).

5.1.1 How Successful Is The First Proposer?

Figure 1 depicts the realized share of resources appropriated by the first proposer in each treatment, focusing on first-round proposals that were accepted.[20] For comparability, we organize the treatments by location. number of rounds, and bargaining protocol.

Figure 1: Realized Proposer Share of Accepted First-round Offers

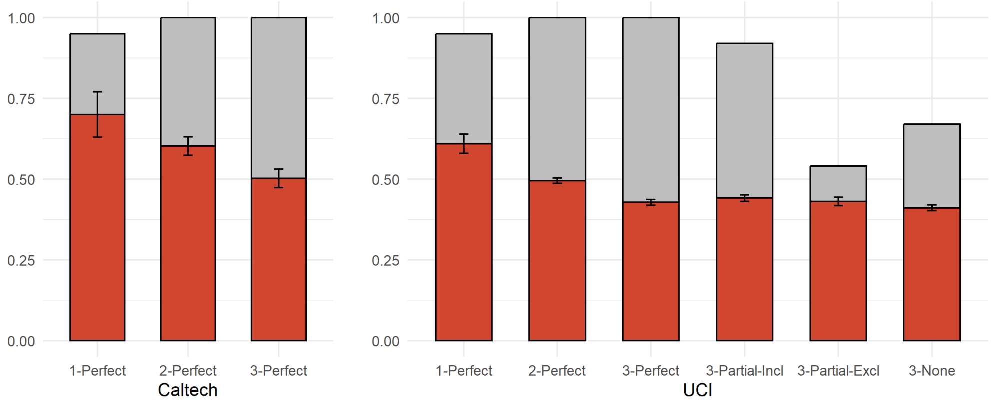

Notes: The total height of each bar represents the predicted equilibrium share for the first proposer while the darker portion is the empirical average appropriated by first proposers in the experiment. The error bars mark the 95% confidence intervals computed based on robust standard errors (clustered at the session level). The empirical averages and standard errors used in this figure are reported in Table 6 of Appendix B.

Three features of this figure are especially salient. First, standard equilibrium theory predicts much higher first-proposer shares than we observe. For example, theory implies that the first proposer ought to obtain all or nearly all of the surplus in the treatments with perfectly predictable recognition orders. However, at both locations, the average share is always well below 75%, and often much lower. While it is always higher at Caltech than at UCI, the gap between theory and observation is large at both locations. We explore the reasons for the difference between outcomes at UCI and Caltech in subsequent sections.

---

[20] We limit the sample to matches in which the first offer is accepted because including payoffs for matches in which the first offer is rejected confounds the continuation values in multi-round games with the exogenously specified default values for one-round games. That said, the qualitative patterns noted below are similar when we include final payoffs for all matches; see Appendix C.3.

Second, standard equilibrium theory fails to anticipate a striking comparative static: for Perfect games, the first proposer's share declines significantly as the potential number of bargaining rounds rises at both Caltech and UCI. It turns out that this pattern is symptomatic of a more general phenomenon involving inequality of the payoff distribution across committee members. We turn to that topic in the next subsection.

Even if proposers do not fully understand that, in equilibrium, their bargaining positions are very strong, they might nevertheless recognize that they should be able to obtain a share of at least 50% (Theoretical Result 1 in Section 3). Consistent with this principle, the average first proposer's share is always at least 50% in all Caltech treatments, although it does fall to that level in 3-Perfect treatment. However, behavior at UCI is inconsistent with this principle: the first proposer's average share falls below 50% in all four three-round treatments as well as the 2-Perfect treatment. It nevertheless remains above 40%, so proposers are able to exploit some advantage on average in all our treatments.

Third, we find no support for the theoretical implication that greater predictability of the recognition order should strengthen the first proposer's bargaining position and enable them to obtain a larger share (Theoretical Result 3). No such pattern is visible across the last four bars of Figure 1: the first proposer's share hovers between 41% and 44%, and the variations do not align with theoretical predictions.

5.1.2 Payoff Inequality Across Committee Members

We measure the inequality of payoff distributions across the three committee members using the Gini coefficient. Formally, for a division (x_1, x_2, x_3) of a dollar with shares written in ascending order, the Gini coefficient is given by the formula:

$$G(x_1, x_2, x_3) \equiv \frac{|x_2 - x_1| + |x_3 - x_1| + |x_3 - x_2|}{3(x_1 + x_2 + x_3)} = \frac{2}{3}(x_3 - x_1).$$

Higher values indicate greater inequality. Specifically, the Gini coefficient is 0 for the completely egalitarian prize distribution, $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Equality within a minimal winning coalition, $(0, \frac{1}{2}, \frac{1}{2})$, results in a Gini coefficient of $\frac{1}{3}$. The dictatorial outcome $(0, 0, 1)$ yields a Gini coefficient of $\frac{2}{3}$, the theoretical maximum for this environment.

Figure 2, which displays the cumulative distribution functions (CDFs) of Gini coefficients separately for 1-Perfect, 2-Perfect, and 3-Perfect bargaining games at Caltech

Figure 2: CDFs of Gini Coefficients of Final Payoffs of First-round Accepted Offers

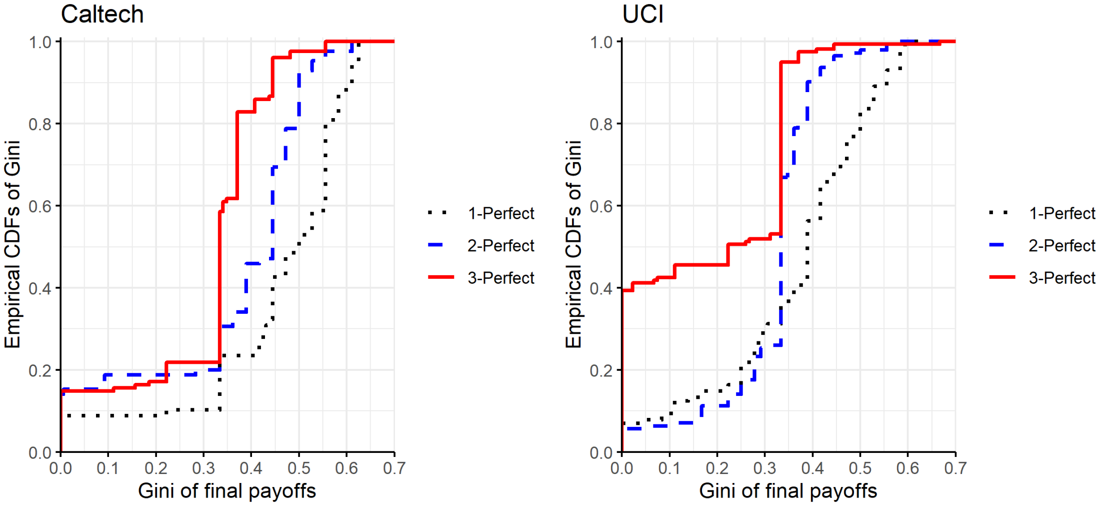

Notes: We present the cumulative distribution functions of Gini coefficients computed based on final payoffs of accepted first-round offers in treatments 1-Perfect, 2-Perfect, and 3-Perfect in both locations. Figure 11 in Appendix C.3 displays qualitatively similar inequality distributions using all (experienced) matches, including matches where the first-round offer failed to pass.

and UCI, allows us to compare the payoff inequality among committee members across different treatments.[21] The horizon of the game significantly impacts payoff inequality: longer-horizon games generally yield a more equitable distribution of resources among committee members, as indicated by lower GINI coefficients.[22]

Strikingly, the CDFs above exhibit some clear mass points (vertical jumps), and the concentrations at these points are plainly related to the length of the game. Looking first at Caltech, equal division within MWCs (i.e., a Gini coefficient of $\frac{1}{3}$) accounts for between from 11% to 13% of allocations in one and two-round negotiations versus 37% of allocations in three-round negotiations. Qualitatively, we observe a similar (albeit less dramatic) pattern for equal division within grand coalitions (i.e., a Gini coefficient of 0), which accounts for 9% of allocations in one-round problems, 14% in two-round problems, and 15% in three-round problems. At UCI, equal division within MWCs (i.e., a Gini coefficient of $\frac{1}{3}$) accounts for only 5% of allocations in one-round games versus 41%-42% in two-round and three-round games. Furthermore, equal division within grand coalitions (i.e., a Gini coefficient of 0) accounts for only 6% to 7% in one and two round games versus 39% for the three-round games. Altogether, at Caltech, equal splits within a coalition (either a MWC or a grand coalition) account for 52% of allocations in three-round games versus 22% in one-round games; at UCI, the corresponding figures are 81% and 13%. Accordingly, moving from one round to three, egalitarianism more than doubles at Caltech and increases by a factor of more than six at UCI. In other words, we see a dramatic migration towards egalitarian splits in games with longer horizon.

---

[21] For the same reason as in the previous section, we restrict attention to accepted first-round offers.

[22] The Kolmogorov-Smirnov statistics for equality of distributions are, for Caltech, 0.4 in 1-Perfect versus 2-Perfect, 0.62 in 1-Perfect versus 3-Perfect, and 0.49 in 2-Perfect versus 3-Perfect; for UCI, 0.39 in 1-Perfect versus 2-Perfect, 0.58 in 1-Perfect versus 3-Perfect, and 0.39 in 2-Perfect versus 3-Perfect; all of these differences are highly significant with p-values below 0.1%.

The number of rounds in a bargaining game provides a natural proxy for the game's strategic complexity. A longer game requires a subject to think through far more contingencies and to form more elaborate beliefs about how others will reason and behave. Viewed from this perspective, our results offer clear evidence that increased complexity pushes subjects towards more egalitarian allocations. However, these results pertain only to the final outcomes of negotiations. In the next section, we provide more detailed analyses of first-round proposals, voters' responses to these proposals, and the optimality of proposals in light of voters' behavior.

5.2 How Do Groups Arrive At These Bargaining Outcomes?

5.2.1 First-Round Proposals

In this section, we analyze all first-round proposals, including those that do not pass. This analysis uncovers the extent to which proposers target and exploit weak coalition partners, and how that exploitation depends on the complexity of the game.

We categorize proposals using the following taxonomy. A *coalition partner* is a non-proposing player to whom the offer assigns at least 5% of the budget. A *minimal winning coalition* (MWC) proposal is one that includes exactly one coalition partner, a *grand coalition* proposal is one that includes two coalition partners, and a *dictatorial* proposal is one that includes no coalition partners.[23] We classify an MWC offer as targeting a weak (respectively, strong) coalition partner if, according to theory, the partner has a lower (respectively higher) continuation payoff than the excluded player. If theory implies that both non-proposing players have the same continuation payoff, we classify the coalition

---

[23] Others have employed similar criteria for categorizing coalitions; see, for example, Frechette, Kagel and Morelli (2005) and Agranov and Tergiman (2014). Our results are not sensitive to reasonable alternatives to the 5% threshold. There is some ambiguity as to how we should classify proposals that offer both non-proposing partners less than 5%. If either or both non-proposing players have a continuation payoff of zero, then one could construe the proposal that assigns the entire prize to the proposer as an MWC or grand coalition offer (assuming that either or both non-proposing players resolve indifference favorably). In practice, dictatorial proposals are rare, accounting for fewer than 3% of proposals in any of our treatments.

Figure 3: Coalition types in first-round proposals

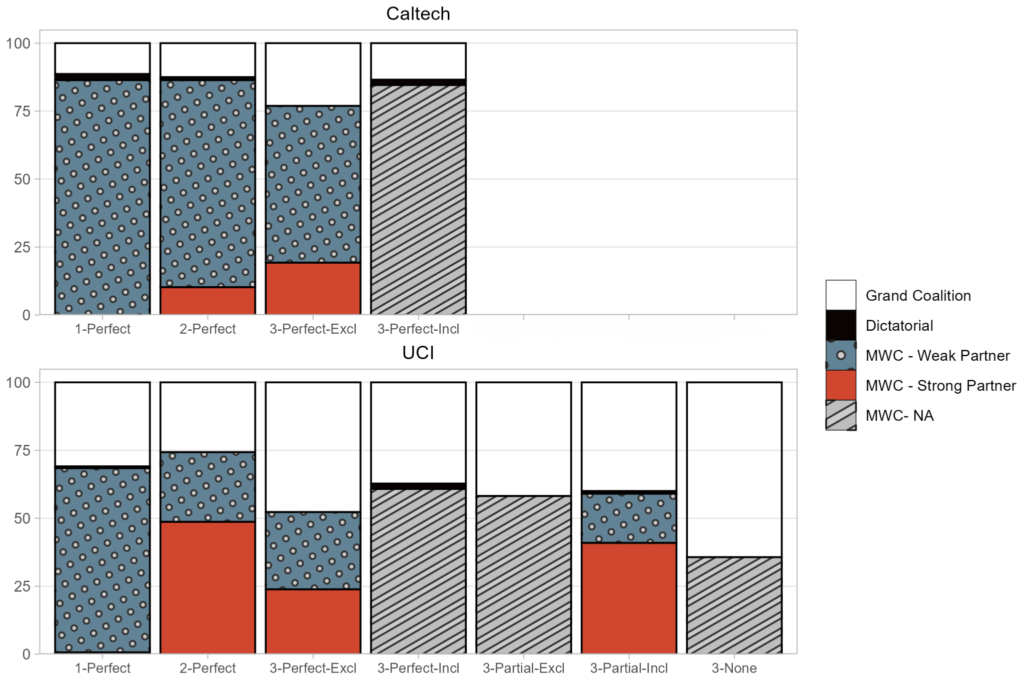

Notes: The empirical frequencies used in this figure are reported in Table 8 of Appendix B.

as MWC-NA.

Figure 3 shows distributions of proposals across these categories for each treatment. The top panel depicts results for Caltech subjects. In one-round games, the vast majority of proposals (86%) seek to form MWCs, and all of them target weak partners. These findings confirm that Caltech subjects have exploitative intent; they also show that these subjects have exploitative ability in the simplest bargaining contexts. The same fraction of proposals seek to form MWCs in two round games, and 88% of those target weak partners. In other words, proposers try to behave much as they do in one-round games, but occasionally choose the wrong partner. Behavior changes more significantly in three-round games: 77% of proposals seek to form MWCs when the first-round proposer is not the second-round proposer (3-Perfect-Excl); when the first-round proposer remains the second-round proposer (3-Perfect-Incl), the figure is 85%. Equilibrium predictions identify one non-proposing player to be weaker than the other only in the first setting, and 75% of those proposals target the weak partner. Thus, we see a greater shift toward grand coalition offers in the three-round game alongside a continued decline in exploitative ability.

We find different patterns at UCI: in relatively simple games, players are much more inclined to make offers that either target grand coalitions, or that mistakenly target strong partners in MWCs. In one-round games, 68% of the proposals seek to form MWCs, and 99% of those target weak coalition partners. Consequently, at least half of the UCI subjects have exploitative intent, and they also have exploitative ability in the simplest bargaining contexts. In two-round games, the fraction of MWC offers increases to 74%, indicating even greater exploitative intent. However, in contrast to Caltech subjects, UCI subjects often fail to realize that the player with the higher default payoff is in the weaker bargaining position because they do not propose next. Consequently, contrary to Theoretical Result 2, almost two-thirds of MWC offers mistakenly target the strong partner. Despite the differences between Caltech and UCI for one-round and two-round games, the patterns of offers exhibit greater similarity for the more complex three-round games. At UCI, 52% of proposals seek to form MWCs when the first proposer is not the second proposer (3-Perfect-Excl). These proposals target weak and strong partners with roughly equal frequencies, indicating an inability to discern theoretical bargaining strength. When the first proposer is also the second proposer (3-Perfect-Incl), the frequency of MWC offers is slightly higher. The share of proposals that seek to form MWCs is similar for the 3-Partial-Excl and 3-Partial-Incl treatments but falls to 36% for the 3-None treatment. UCI subjects also tend to select the stronger player, rather than the weaker one, as a coalition partner in 3-Partial-Incl.

Table 3 summarizes the information presented in Figure 3 in a slightly different way, in order to address directly the exploitative ability hypothesis (as well as Theoretical Result 2). For each treatment, column (1) shows the difference between the frequency with which proposers make MWC offers targeting (theoretically) strong players and the frequency with which they make MWC offers targeting (theoretically) weak players. Lower (more negative) numbers indicate greater exploitative ability, zero implies no ability to distinguish between weak and strong players, and numbers greater than 0 reflect a tendency to systematically confuse weak and strong players. At Caltech, subjects exhibit exploitative ability in all treatments, but this ability declines substantially as the environment becomes more complex. At UCI, subjects exhibit exploitative ability only in one-round bargaining games. They commit systematic errors with respect to discerning potential partners' bargaining strengths in two-round games and some three-round games. Column (2) explores the same issue using a measure of discernment that does not depend on the prevalence of MWC offers: the probability of choosing strong partner conditional on making an MWC offer. The pattern is similar.

Next, we ask how the tendency to make egalitarian offers evolves as the bargaining environment becomes more complex. Recall that we classify an offer as *egalitarian* if,

Table 3: Ability to Identify Weak Partners

	P(Strong) - P(Weak) (1)	P(Strong | MWC) (2)
Caltech, 1-Perfect	-0.864^{***}	0.000
	(0.025)	(0.000)
Caltech, 2-Perfect	-0.659^{***}	0.118
	(0.221)	(0.091)
Caltech, 3-Perfect-Excl	-0.385^{***}	0.250^{***}
	(0.140)	(0.067)
UCI, 1-Perfect	-0.671^{***}	0.010
	(0.064)	(0.008)
UCI, 2-Perfect	0.230	0.655^{***}
	(0.215)	(0.142)
UCI, 3-Partial-Incl	0.229^{***}	0.694^{***}
	(0.034)	(0.027)
UCI, 3-Perfect-Excl	-0.046	0.456^{***}
	(0.046)	(0.041)
Note:		*p<0.1; **p<0.05; ***p<0.01

Notes: P(strong) (resp. P(weak)) denotes the overall frequency with which first proposers make MWC offers targeting strong (resp. weak) partners. Standard errors are clusterered by session. We omit the treatment 3-Perfect-Incl as both non-proposing players would be weak partners.

within a coalition, no two players' shares differ by no more than 5% of the prize. For each treatment, Figure 4 shows the distribution of offers across three categories: MWC egalitarian, grand coalition egalitarian, and non-egalitarian. The patterns are striking. At Caltech, the total frequency of egalitarian offers rises from 15% in the one-round treatment to 25% in the two-round treatment, to 60% in the three-round treatment when the first and second proposers differ (3-Perfect-Excl).[24] The same qualitative pattern emerges at UCI but the magnitudes are much larger: the total frequency of egalitarian offers rises from 14% in the one-round treatment to 51% in the two-round treatment to 83% in the 3-Perfect-Excl treatment.[25] Thus, egalitarianism tends to emerge progressively

---

[24] The frequency of egalitarianism is lower (32%) in three-round settings when the first proposer enjoys the advantage of proposing twice in a row (3-Perfect-Incl).

[25] Again, the egalitarian share is a bit lower (71%) when the first-proposer enjoys the advantage of proposing twice in a row (3-Perfect-Incl). The egalitarian share is generally 60% or above for all three-round games in UCI.

Figure 4: The Frequency of Egalitarian Offers

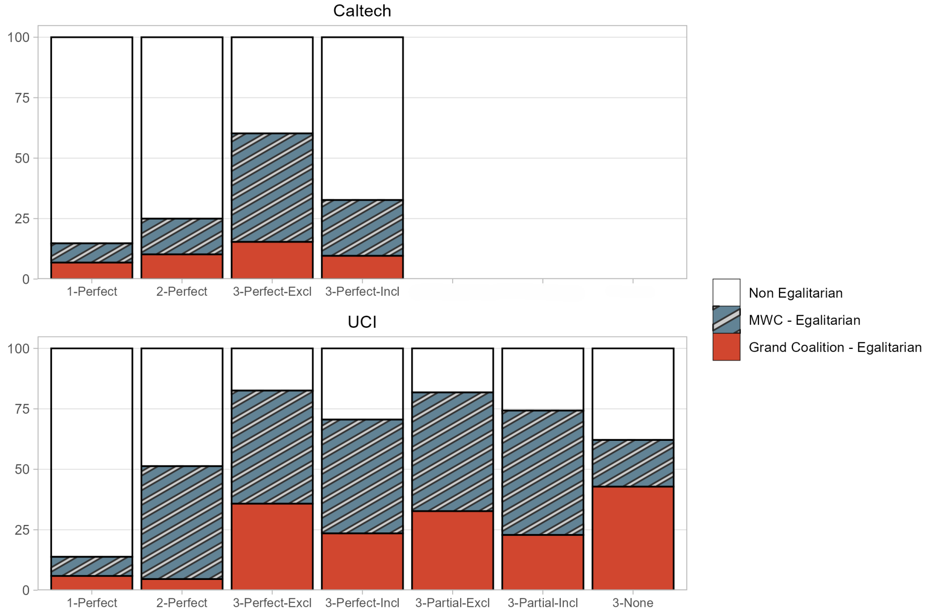

Notes: The empirical frequencies used in this figure are reported in Table 9 of Appendix B.

as complexity increases.

Notably, we observe the same migration to egalitarianism even among subjects who conform most closely to theory, in the sense that they make MWC offers to weak players. For each pertinent treatment (i.e., those for which non-proposing players are asymmetrically positioned), Figure 5 displays the frequency of egalitarianism among MWC offers to weak players, as well as the average share the first proposers hope to keep for themselves. At Caltech, fewer than 10% of these offers are egalitarian in either one-round or two-round games. Even so, adding the second round leads the first proposers to behave more conservatively, in the sense that the proposer's share declines. With the addition of a third round, the egalitarian-offer share jumps up to 62% and the average proposed share for the first proposer is barely more than 50%. Once again the pattern is qualitatively similar at UCI but the quantitative changes are much larger. While 12% of offers are egalitarian in one-round games, the share jumps to roughly one-third in two-round games and to roughly 90% (resp. 84%) in three-round games when the first and second proposers differ (resp. are the same). The proposed share for the first proposer declines from 70% in the one-round game, to 56% in the two-round game, to barely more than 50% in the three-round games. Accordingly, even for the most theory-compliant participants, we see a striking migration toward egalitarianism as the game becomes more

Figure 5: Frequency of Egalitarianism Among MWC Offers to Weak Partners

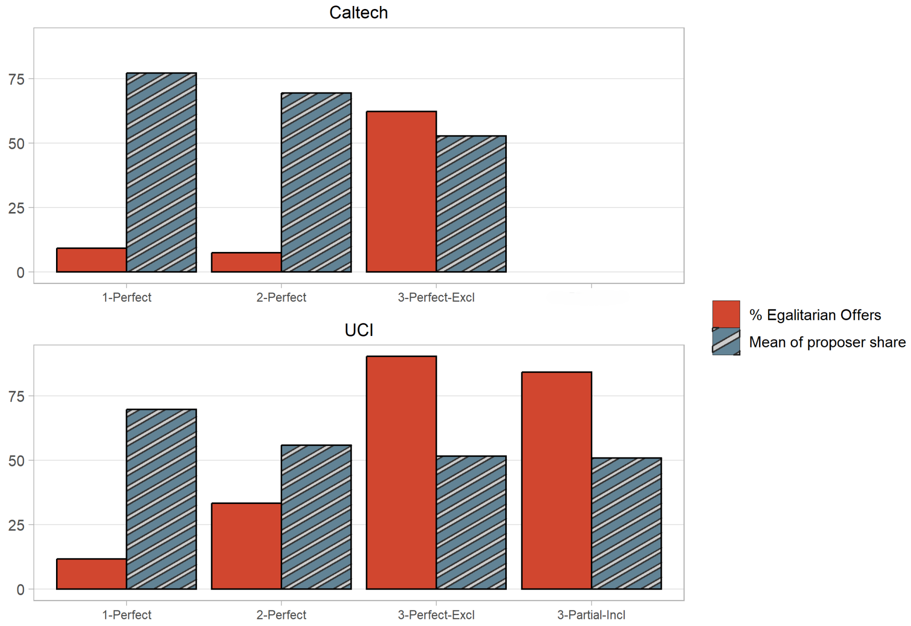

Notes: The empirical frequencies used in this figure are reported in Table 10 of Appendix B.

complex.

5.2.2 Votes And Coalitions

Having completed our analysis of the first proposers' offers, we next examine how others vote on these offers. Figure 6 displays the raw data for the one-, two-, and three-round *Perfect* treatments. In each case, we plot first-round offers on a simplex. The numbers along the right face label the level grid lines for the first proposer's share, which reaches a maximum of 100% at the top vertex. The numbers along the bottom face label the level grid lines for the strong partner's share, which reaches a maximum of 100% at the bottom-right vertex. The numbers along the left face label the level grid lines for the weak proposer's share, which reaches a maximum of 100% at the bottom-left vertex. In each case, a player's share is constant along any of the grid lines emanating from the face that shows their share. The figure distinguishes between accepted proposals (shown as black circles) and rejected proposals (gray circles). The size of each circle reflects the number of observations at that point.

For Caltech, we see similar patterns for the one-round and two-round treatments. In both cases, there is a cluster of MWC proposals that target the weak partner by offering

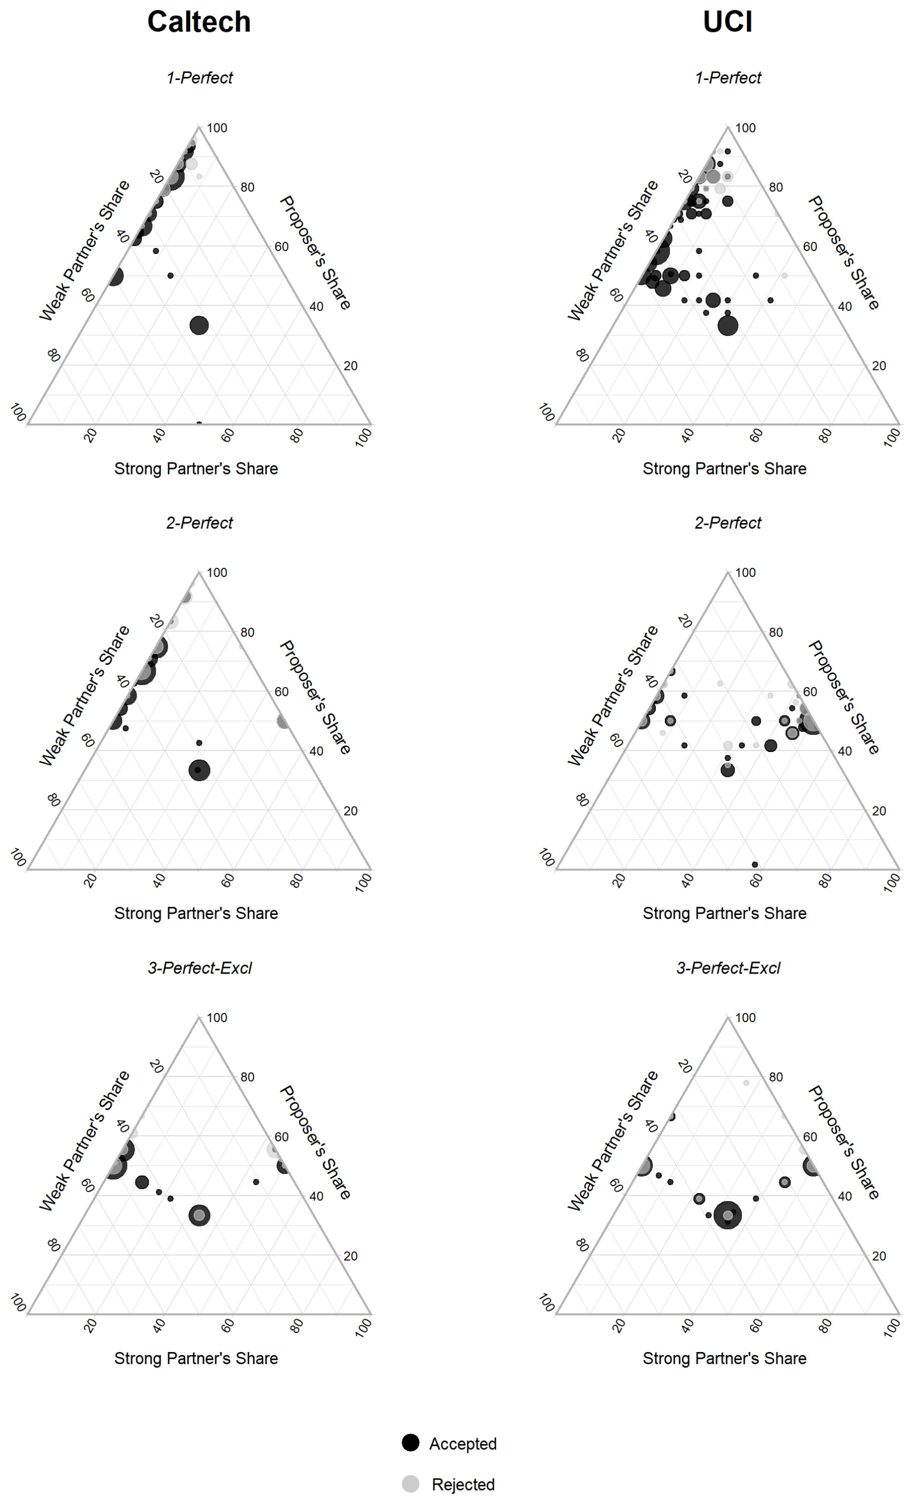

Figure 6: Distributions of Offers and Acceptances

Notes: For the 3-Perfect treatment, we show results for cases in which the first and second proposers differ (3-Perfect-Excl). Figure 8 in Appendix C.1 shows similar plots for 3-Perfect-Incl, as well as the other three-round treatments at UCI with less than perfect predictability.

shares between zero and one-half, and a cluster of egalitarian offers targeting the grand coalition. Rejections are concentrated among the most selfish MWC offers (close to the top vertex). In the two-round treatment, we see some egalitarian MWC offers to the strong partner, who rejects them. Moving from the two-round to the three-round games, we see the migration to egalitarian offers, both MWC and grand coalition, noted previously. We also see the emergence of "hybrid egalitarian" offers which allocate less than one-third to one of the non-proposers and splits the remainder equally between the proposer and the other player. These are the allocations in the figure that lie between the centroid (equal-split) and the midpoint of either the left or right face of the simplex. Rejections are still more concentrated among most selfish offers, particularly for those that target strong partners, but the pattern is less striking than in the one-round and two-round games.

For UCI, the pattern of offers and responses in one-round games is broadly similar to the one observed at Caltech, although we see more variation in offers, as well as more hybrid egalitarian offers, which are not observed at Caltech in either one- or two-round settings. Rejections are once again concentrated among the most selfish offers, but also among those that target strong partners. For two-round games, the patterns at UCI and Caltech are much less similar. For UCI, we see a high concentration of egalitarian (or nearly egalitarian) MWC offers to strong partners, a smaller number of egalitarian (or nearly egalitarian) MWC offers to weak partners, as well as a smattering of egalitarian or hybrid egalitarian grand coalition offers. Committees reject the most selfish offers more frequently, but it is not obvious that the strong partners remain less accommodating than the weak partners. For three-round games, the UCI and Caltech patterns are once again broadly similar, in that we see concentrations of egalitarian MWC offers, egalitarian grand coalition offers, and hybrid egalitarian offers. Egalitarian grand coalition offers pass with the highest frequency, and committees systematically reject the most selfish proposals. The striking symmetry of the three-round UCI figure suggests that neither proposers nor their potential partners fully understand how to differentiate weak and strong coalition partners.

Table 4 presents estimates of logit models relating the binary voting outcome (yes or no) for a non-proposing players to a constant and three variables: (1) an indicator specifying whether the player is weak or strong according to theory ("Strong Partner"), (2) the share proposed for that player ("Own Share"), and (3) the inequality of the offer ("Gini Coefficient").[26] As expected, in every treatment, the likelihood of voting in favor

---

[26] We experimented with more complicated models, including ones with MWC indicators, quadratic terms, and alternative measures of inequality, but found that they did not meaningfully improve the fit or

Table 4: Determinants of Proposal Acceptance

	Caltech: 1-Perf (1)	Caltech: 2-Perf (2)	Caltech: 3-Perf-Excl (3)	UCI: 1-Perf (4)	UCI: 2-Perf (5)	UCI: 3-Part-Incl (6)	UCI: 3-Perf-Excl (7)
Strong Partner	-2.472^{***}	-2.800^{***}	-1.358^{***}	-2.183^{***}	-1.028^{***}	-0.074	-0.195
	(0.647)	(0.707)	(0.501)	(0.641)	(0.382)	(0.386)	(0.403)
Own Share	8.271^{***}	9.211^{***}	10.851^{***}	12.561^{***}	9.522^{***}	9.224^{***}	15.478^{***}
	(2.968)	(1.883)	(3.434)	(2.662)	(1.241)	(1.963)	(4.560)
Gini Coefficient	-0.040	-1.628	-4.519^{**}	0.514	-1.626	-3.542^{***}	-8.610^{***}
	(2.022)	(1.293)	(1.903)	(1.884)	(1.348)	(1.363)	(2.800)
Constant	-0.420	-0.895	-2.428^{*}	-1.596	-2.435^{***}	-2.239^{***}	-3.707^{**}
	(1.506)	(0.717)	(1.345)	(1.328)	(0.515)	(0.622)	(1.546)
Pseudo R2	0.48	0.51	0.46	0.58	0.36	0.35	0.46
Observations	176	176	156	304	304	210	218

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Notes: This table reports estimates of logit models relating the binary voting outcomes (yes or no) for individual non-proposing players to an indicator specifying whether the player is weak or strong according to theory, the share proposed for that player, and a measure of inequality (the Gini coefficient). Standard errors are clustered by session.

of a proposal is increasing in the player's own share. Moreover, the magnitude of this effect is similar across treatments and locations (Caltech versus UCI).

The coefficients of "Strong Partner" speak to the self-awareness hypothesis. At Caltech, theoretically strong partners are much more likely to reject offers than theoretically weak partners in one-round games. This finding indicates a high degree of awareness concerning their own bargaining strength. For two-round games, the corresponding coefficient indicates a similar degree of self-awareness. For three-round games, the magnitude of the coefficient attenuates substantially, indicating a lower level of self-awareness. However, even in three-round games, the coefficient remains negative and statistically significant. Accordingly, Caltech subjects do have some ability to discern their bargaining strengths even in the most complex settings we consider. At UCI, we similarly find that theoretically strong partners are much more likely to reject offers than theoretically weak partners in one-round games. However, in contrast to the Caltech pattern, this tendency is substantially attenuated in two-round games. Surprisingly, UCI voters retain some ability to distinguish their own strength or weakness in the 2-Perfect game even though, as shown above, proposers target strong players, apparently believing them to be weak. Possibly voters feel greater motivation to think through their own prospects than proposers. Significantly, in three-round games, there is no relationship between voting responses and theoretical bargaining strength. These findings imply that UCI

---

change the results.

subjects reliably self-identify as weak or strong only in the simplest bargaining settings. Their ability to do so declines rapidly with the game's strategic complexity, and vanishes entirely in the most complex settings we consider.

The pattern of coefficients for the "Gini Coefficient" shows that, as subjects become less able to determine their own bargaining strength, they place progressively greater weight on fairness. Strikingly, fairness has no impact on voting responses at either Caltech or UCI in one-round games, where players understand their bargaining strengths. Accordingly, there is no indication that they care about fairness intrinsically in the settings we consider. In two-round games, the coefficients of our inequality measure turn negative (meaning that greater inequality reduces the odds of passage), but the effect lacks statistical significance in both locations. In three-round games, wherein players demonstrably struggle to determine their own bargaining strengths, the corresponding coefficients become more negative and highly statistically significant, especially at UCI. Thus, the tendency to place substantial weight on equity appears to emerge in this setting as a response to strategic complexity (which obscures which players are in weak or strong positions), rather than as an expression of innate preference.

5.2.3 Optimality Of Proposals

Next, we ask whether proposals are approximately optimal given the voting behavior documented in the previous subsection. In other words, do actual proposals differ from equilibrium proposals because proposers fail to optimize correctly, or because they respond rationally to empirical voting patterns that differ from theoretical predictions?

To evaluate optimization failures, we determine the proposals that would maximize the proposer's expected payoff given the observed mapping from proposals to votes. To make this calculation, we assume that the continuation outcome in the event of rejection does not depend on the terms of the rejected proposal. This assumption is consistent with the implications of backward induction, assuming players vote as if they are pivotal (a common premise) and that indifference is insufficiently prevalent to support multiple equilibria. It also has empirical support, in that continuation outcomes are not systematically related to rejected offers. Under this assumption, for each treatment, we can estimate the first proposer's expected payoff conditional on rejection of her offer by calculating the average payoff for first proposers when voters reject their offers (the *mean rejection payoff*, henceforth MRP). We list these average payoffs in Appendix C.4.

We use the logit regression models shown in Table 4, which describe non-proposing players' voting responses to proposals, along with the MRP, to determine the mapping from proposals to expected first-proposer payoffs. Specifically, we use the regression models to compute the probability a proposal passes as a function of its terms. We then multiply the proposer's share by this probability and add the product of the MRP and the estimated rejection probability. Figure 7 uses contour heat maps to exhibit the first proposers' expected payoff mappings for the one-, two-, and three-round *Perfect* treatments.[27] In each case, we associate offers with points on a simplex, exactly as in Figure 6.

In all cases, the offer that maximizes the proposer's expected payoff is an MWC offer. Interestingly, it sometimes delivers an expected payoff below 50% of the prize, contrary to Theoretical Result 1, because observed voting patterns are less favorable to the proposer than equilibrium voting behavior. At Caltech, the ideal proposal targets the weak player and seeks to retain 76% of the prize in one-round games, 69% of the prize in two-round games, and 53% in three-round games (for which there is a weak partner). At UCI, the ideal proposal targets the weak player in one- and two-round games, and seeks to retain 73% of the prize in the first instance and 59% in the second. For three-round games, the ideal proposal at UCI seeks to form an MWC and retains only 51% of the prize, but whether it targets the weak player or the strong player is nearly inconsequential. Thus, in both locations, being egalitarian within an MWC becomes increasingly attractive from a purely selfish perspective as strategic complexity increases. Accordingly, the nature of voting responses by non-proposing players accounts for much of the observed migration to egalitarianism.

Next, we evaluate the quality of first proposers' offers by comparing them to optimized proposals. We perform two versions of this calculation. For the first, we calculate the first proposer's expected payoff based on the average offer for each coalition type and then divide by the optimized expected payoff; to gauge overall performance, we compute the average of these measures weighted by the frequencies of the coalition types. We depict this measure in Section 5.2.3. For the second, we calculate the average ratio between the first proposer's actual payoff and the optimized expected payoff, both in the aggregate and for each coalition type chosen by the first proposer. We provide results based on this measure in Appendix C.4. The relative merits of these two approaches depend on how one thinks about the observed variation in offers within treatments. If the variation mainly reflects observational noise, the first approach is more relevant, while if it arises mainly from true heterogeneity with respect to preferences or strategic thinking, the second approach is more appropriate.

---

[27] For the 3-Perfect treatment, we show results for cases in which the first and second proposers differ (3-Perfect-Excl). In Appendix C, we show similar plots for 3-Perfect-Incl, as well as the other three-round treatments at UCI.

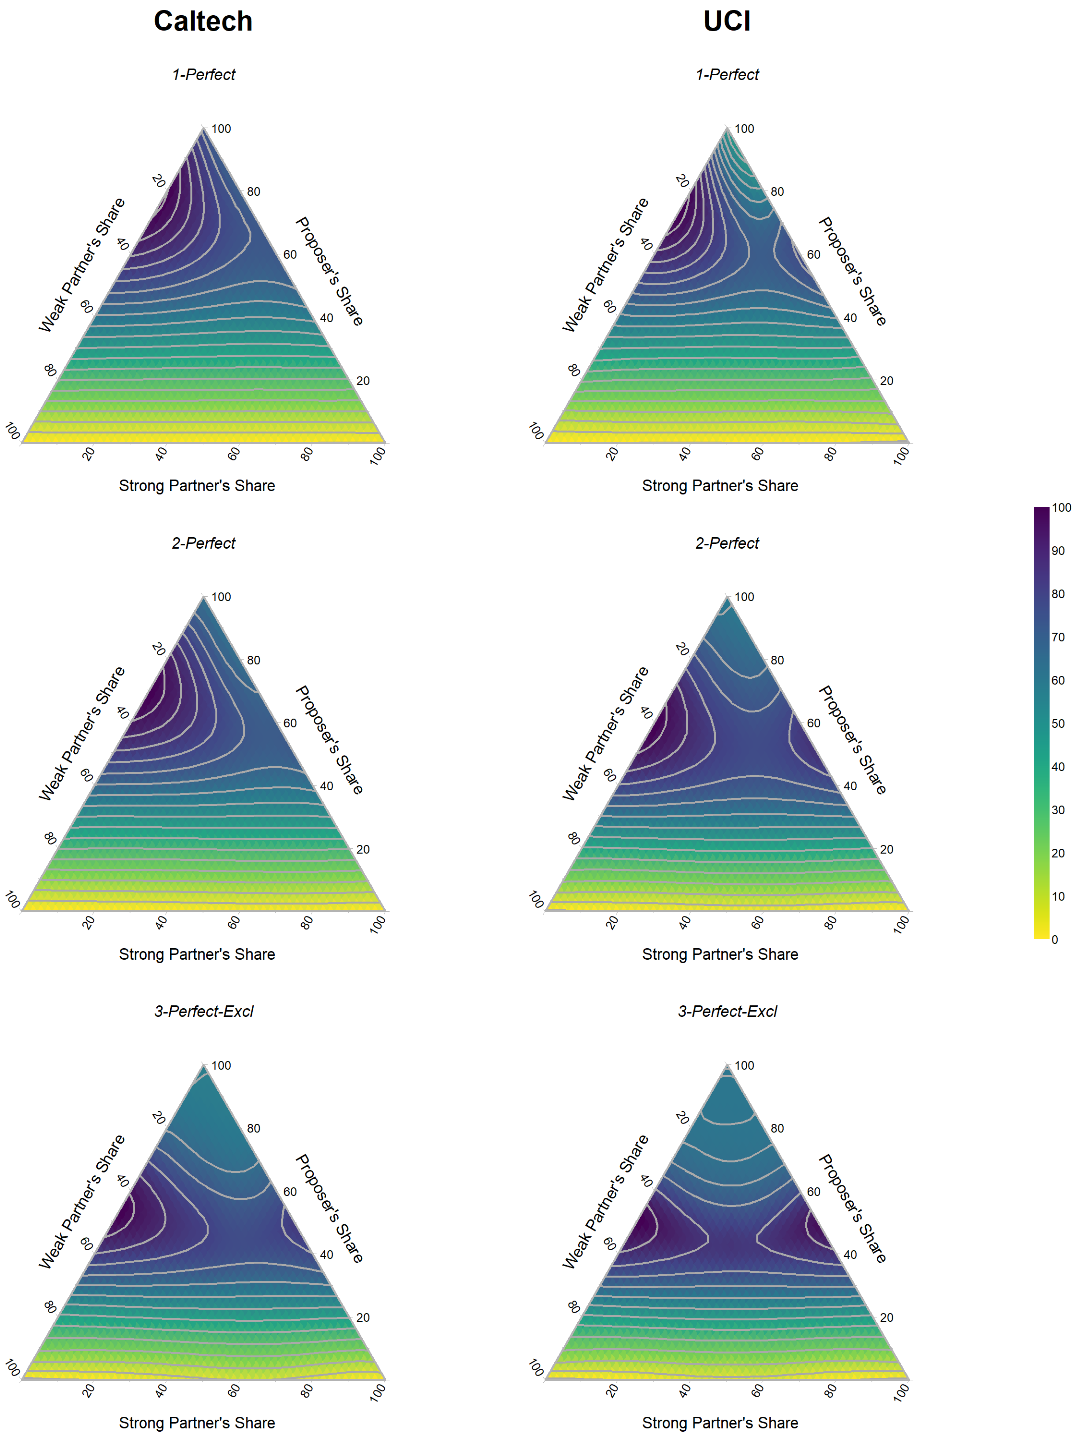

Figure 7: *Mappings from First Proposals to Expected Payoffs*

Notes: For the 3-Perfect treatment, we show results for cases in which the first and second proposers differ (3-Perfect-Excl). Figure 9 in Appendix C.2 shows similar plots for 3-Perfect-Incl, as well as the other three-round treatments at UCI with less than perfect predictability.

Location	Treatment	Optimal Payoff	MWC Weak	MWC Strong	MWC NA	Grand Coalition	Dictatorial	Agg
Caltech	1-Perfect	64	100	–	–	58	80	95
Caltech	2-Perfect	60	100	72	–	59	67	92
Caltech	3-Perfect-Excl	46	100	82	–	75	–	91
Caltech	3-Perfect-Incl	55	–	–	87	78	80	85
UCI	1-Perfect	65	98	71	–	73	54	90
UCI	2-Perfect	49	99	87	–	75	–	87
UCI	3-Perfect-Excl	44	100	98	–	80	–	90
UCI	3-Perfect-Incl	47	–	–	81	79	–	79
UCI	3-None	45	–	–	76	77	–	77
UCI	3-Partial-Excl	41	–	–	90	84	–	88
UCI	3-Partial-Incl	42	100	99	–	79	41	91

Table 5: Proposer Optimization Rates

Notes: This table displays the optimization rates of the first proposer using our first measure, which captures how well the average offer within a coalition type performs relative to the empirically optimal payoffs, reporting this ratio as a percentage. The final "Agg" column reflects an average of the performance across coalition types weighted by the frequency of each coalition type.

Several lessons emerge from this analysis. First, proposers perform reasonably well: across all protocols and locations, they achieve 77% to 95% of optimized payoffs according to the first measure. Table 5 provides a detailed breakdown of these optimization rates.[28] To put these ranges into perspective, in settings where the first proposer is also not the second proposer, our calculations imply that first proposers would achieve between 8% and 61% of the optimized payoff by making a rejected proposal, between 40% and 80% if they made dictatorial offers, and between 50% and 80% if they offered an equal split within a grand coalition.[29] Relative to these measures, the performance of proposers indicates a meaningful degree of sophistication.

Second, Caltech subjects perform slightly better than UCI subjects: averaged across the four treatments, the Caltech subjects achieved 91% of the optimized payoff according to the first measure and 88% according to the second; for UCI subjects, the corresponding figures are 88% and 84% respectively. This comparison is consistent with other evidence of greater sophistication on the part of Caltech subjects discussed elsewhere in this paper.

Third, we also see no clear relationship between complexity and performance. In other words, these measures of performance do not decline systematically as we move from one-round to two-round to three-round games. Apparently, the tendency for voting players to evaluate proposals based on simple egalitarian criteria, along with proposers' intuitive understanding of those criteria, offsets the greater strategic complexity of the environment.

---

[28] The results using the second measure are nearly identical, with optimized payoff rates ranging from 80% to 92%. See Table 12 in Appendix C.4 for details.

[29] See Tables 13 and 14 in Appendix C.4.

Finally, according to the first measure, the average offer that seeks to form an MWC with a weak player always comes close to achieving expected-payoff maximization. This finding is of particular interest because, in discussing Figure 5, we noted the sharp progression toward egalitarianism among MWC offers targeting weak players that coincides with increases in strategic complexity. As it turns out, despite being contrary to equilibrium implications, this pattern is nearly optimal for the proposers.[30]

6 Conclusion

In this study, we examined legislative bargaining protocols experimentally to determine (1) the extent to which people identify and exploit weak coalition partners, (2) the reasons for failures to do so in settings where they occur, and (3) the systematic behavioral patterns that arise in those settings. Our experiment focuses on the role of strategic complexity by varying the number of bargaining rounds from one to three in a setting with asymmetrically positioned players, which makes strength and weakness progressively more difficult to ascertain. We employed two distinct subject pools, one of which (Caltech) displays greater quantitative aptitude and skill than the other (UCI). Our primary finding is that, regardless of analytic sophistication, players gravitate toward egalitarianism as strategic complexity progressively obscures relative negotiating strengths. Analytic sophistication only affects the speed with which this migration occurs. More specifically, we reach the following seven main conclusions.

First, in all settings, the first proposers receive far smaller shares than predicted by standard bargaining theory, and this discrepancy increases with the game's degree of strategic complexity. For three-round games, theoretical predictions concerning comparative statics (involving variations that ought to affect the identity and exploitability of the weak partner) fail entirely.

Second, as the number of rounds increases, there is a dramatic increase in the frequency with which the selected allocations involve equal division, either within minimum winning coalitions or grand coalitions.

---

[30] We noted in Section 5.2.1 that in two and three round games, first proposers at UCI frequently target the strong partner rather than the weak one, which is sub-optimal given the predicted voting behavior. That said, according to either measure, these offers do not reduce their payoffs substantially: in the two-round games, such proposals accrue 86%-87% of the optimized payoff, and more than 97% in the three-round game.

Third, results for relatively simple (one- and two- round) settings confirm that subjects nevertheless have exploitative intent, in that they make MWC offers with high frequency, and seek to extract relatively high shares.

Fourth, as strategic complexity increases, exploitative ability declines: proposers have greater difficulty identifying weak partners, and increasingly adhere to within-coalition egalitarianism. Even the non-egalitarian offers exhibit a hybrid form of egalitarianism—i.e., a small share for one player and an equal split between the other two. A comparison of results for Caltech and UCI suggests that analytic sophistication dampens but does not eliminate these patterns.

Fifth, increasing strategic complexity reduces self-awareness of bargaining power among non-proposers. In one-round games, strong partners are significantly less accommodating than weak players. This differential attenuates substantially as strategic complexity increases at Caltech, and vanishes entirely at UCI.

Sixth, in the simplest settings, non-proposing players place no weight on fairness when casting their votes, but fairness begins to play a much larger role in their evaluations once strategic complexity obscures their own bargaining power.

Seventh, the migration toward egalitarian that accompanies rising complexity (our second main result) is, in substantial part, a rational response to non-proposers' increasing emphasis on fairness (our sixth main result). Most notably, egalitarian offers to MWCs become nearly optimal.

Taken together, our findings suggest a new rationale for fairness in legislative bargaining: when it is difficult to discern the strength of others' bargaining positions, a proposer may offer a fair split (within a coalition type) to hedge against her own strategic uncertainty or that of others. The proposer's motivation is fundamentally strategic—she offers a fair split not from a desire for fairness or to appear fair, but because such a division is likely to pass, even if she cannot identify weak players or weak players mistakenly believe they are strong. An important direction for future research is to formalize this strategic theory, building on the framework in Section 3.2, and to better understand its implications for organizations and politics.

References

Abreu, Dilip and Ariel Rubinstein, "The structure of Nash equilibria in repeated games with finite automata," *Econometrica*, 1988, *56*, 1259–1282.

Agranov, Marina, "Legislative Bargaining Experiments," in Emin Karagozoglu and Kyle Hyndman, eds., *Bargaining: Current Research and Future Directions*, Palgrave Macmillan, 2021.

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

Alaoui, Larbi and Antonio Penta, "Endogenous depth of reasoning," *The Review of Economic Studies*, 2016, *83* (4), 1297–1333.

Ali, S. Nageeb, B. Douglas Bernheim, and Xiaochen Fan, "Predictability and power in legislative bargaining," *The Review of Economic Studies*, 2019, *86* (2), 500–525.

Andreoni, James and B. Douglas Bernheim, "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects," *Econometrica*, 2009, *77* (5), 1607–1636.

__ **, Deniz Aydin, Blake Barton, B. Douglas Bernheim, and Jeffrey Naecker**, "When Fair Isn't Fair: Understanding Choice Reversals Involving Social Preferences," *Journal of Political Economy*, 2020, *128* (5), 1673–1711.

Ba, Cuimin, J. Aislinn Bohren, and Alex Imas, "Over- and Underreaction to Information," 2024. Working Paper.

Baron, David P. and John A. Ferejohn, "Bargaining in legislatures," *American Political Science Review*, 1989, *83* (4), 1181–1206.

Bernheim, B Douglas and Charles Sprenger, "On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting," *Econometrica*, 2020, *88* (4), 1363–1409.

Billingsley, Patrick, *Probability and Measure*, third edition ed., John Wiley & Sons, 1995.

__ , *Convergence of Probability Measures*, second edition ed., John Wiley & Sons, 1999.

Binmore, Ken, Avner Shaked, and John Sutton, "Testing Noncooperative Bargaining Theory: A Preliminary Study," *American Economic Review*, 1985, *75* (5), 1178–1180.

__ **, Peter Morgan, Avner Shaked, and John Sutton**, "Do People Exploit Their Bargaining Power? An Experimental Study," *Games and Economic Behavior*, 1991, *3*, 295–322.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Chatterjee, Kalyan and Hamid Sabourian, "Game Theory and Strategic Complexity," in R Meyers, ed., *Encyclopedia of Complexity and Systems Science*, Springer, 2009.

Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, March 2013, *51* (1), 5–62.

Diermeier, Daniel and Sean Gailmard, "Self-interest, inequality, and entitlement in majoritarian decision-making," *Quarterly Journal of Political Science*, 2006, *1* (4), 327–350.

Enke, Benjamin and Thomas Graeber, "Cognitive uncertainty," *The Quarterly Journal of Economics*, 2023, *138* (4), 2021–2067.

Fehr, Ernst and Gary Charness, "Social Preferences: Fundamental Characteristics and Economic Consequences," *Journal of Economic Literature*, 2024.

Frechette, Guillaume, John H Kagel, and Massimo Morelli, "Nominal bargaining power, selection protocol, and discounting in legislative bargaining," *Journal of Public Economics*, 2005, *89* (8), 1497–1517.

Guan, Menglong and Ryan Oprea, “Three Faces of Complexity in Strategic Choice,” 2024. Working Paper.

Güth, Werner and Eric Van Damme, “Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study,” *Journal of mathematical Psychology*, 1998, *42* (2-3), 227–247.

_ and Reinhard Tietz, “Ultimatum Bargaining for a Shrinking Cake: An Experimental Analysis,” in Reinhard Tietz, Wolf Albers, and Reinhard Selten, eds., *Bounded Rational Behavior in Experimental Games and Markets*, Springer-Verlag, 1988, pp. 111–128.

Kalai, Ehud, “Bounded Rationality and Strategic Complexity in Repeated Games,” in Tatsuro Ichiishi, Abraham Neyman, and Yair Tauman, eds., *Game Theory and Applications*, Academic Press, 1990.

Knez, Marc J and Colin F Camerer, “Outside options and social comparison in three-player ultimatum game experiments,” *Games and Economic Behavior*, 1995, *10* (1), 65–94.

Li, Shengwu, “Designing Simple Mechanisms,” *Journal of Economic Perspectives*, 2024, *38* (4), 175–192.

Miller, Luis, Maria Montero, and Christoph Vanberg, “Legislative bargaining with heterogeneous disagreement values: Theory and experiments,” *Games and Economic Behavior*, 2018, *107*, 60–92.

Nagel, Rosemarie, “Unraveling in guessing games: An experimental study,” *The American economic review*, 1995, *85* (5), 1313–1326.

Neelin, Janet, Hugo Sonnenschein, and Matthew Spiegel, “A Further Test of Noncooperative Bargaining Theory: Comment,” *American Economic Review*, 1988, *78* (4), 824–836.

Neyman, Abraham, “Bounded complexity justifies cooperation in the finitely repeated prisoners’ dilemma,” *Economics letters*, 1985, *19* (3), 227–229.

Oprea, Ryan, “Decisions under risk are decisions under complexity,” *American Economic Review*, 2024, *114* (12), 3789–3811.

_ , “Complexity and Its Measurement,” in XXX, ed., *Handbook of XXXX*, Vol. 1, XXX, 2025.

Palfrey, Thomas R, “Experiments in political economy,” in John Kagel and Alvin E. Roth, eds., *The Handbook of Experimental Economics*, Vol. 2, Princeton, N.J.: Princeton University Press, 2016, pp. 347–434.

Puri, Indira, “Simplicity and risk,” *The Journal of Finance*, 2025, *80* (2), 1029–1080.

Santambrogio, Filippo, *A Course in the Calculus of Variations: Optimization, Regularity, and Modeling*, Springer, 2023.

Simon, Herbert A, “Rational choice and the structure of the environment.,” *Psychological review*, 1956, *63* (2), 129.

Spiegel, Matthew, Janet Currie, Hugo Sonnenschein, and Arunava Sen, “Understanding When Agents are Fairmen or Gamesmen,” *Games and Economic Behavior*, 1994, *7*, 104–115.

Stahl, Dale O. and Paul W. Wilson, “On players’ models of other players: Theory and experimental evidence,” *Games and Economic Behavior*, 1995, *10* (1), 218–254.

A Appendix: Proof of Proposition 1

We begin with notation used in the argument below. Let G denote a generic bivariate CDF on $\mathbb{R}^2$ and d denote a generic expected payoff for the proposer in the continuation game following rejection of her offer. We restrict attention to distributions G such that $G(x,x) > 0$ for some $x < (1-d)/2$; we later show that this case is relevant for our analysis as our limit belief satisfies this property (Assumption 1) and we consider a sequence of beliefs that converge to this limit belief (in the weak topology).

Let $s \equiv (s_A, s_B, s_C)$ denote a generic division of the dollar and S be the set of all divisions. The probability that offer $s = (s_A, s_B, s_C)$ is accepted by at least one non-proposing player is

$$\Lambda(s_B, s_C, G) \equiv \int_{\tau_B \times \tau_C} \mathbb{1}_{\max\{s_B - \tau_B, s_C - \tau_C\} \geq 0} \, dG.$$

This term is the probability of the event $[(-\infty, s_B] \times \mathbb{R}] \cup [\mathbb{R} \times (-\infty, s_C]]$. We denote the proposer's expected payoff from offer s by $\pi(s, G, d) \equiv d + (s_A - d)\Lambda(s_B, s_C, G)$.

We first argue that the proposer's optimal offer gives her at least her disagreement payoff d.

Lemma 1. *For every s in which $s_A \leq d$, there exists s' in which $s'_A > d$ such that $\pi(s', G, d) > \pi(s, G, d)$.*

Proof. Observe that $\pi(s, G, d) \leq d$. Consider an offer s' in which $s'_B = s'_C = x \in (0, \frac{1-d}{2})$ and $G(x,x) > 0$. Such an x exists given the restriction in the opening paragraph. By construction, $s'_A > d$. Moreover, $\Lambda(x, x, G) \geq G(x,x) > 0$, where the first inequality follows from $(-\infty, s_B] \times (-\infty, s_C]$ being a subset of $[(-\infty, s_B] \times \mathbb{R}] \cup [\mathbb{R} \times (-\infty, s_C]]$. Therefore, $\pi(s', G, d) > d$. $\square$

We now consider the limit case (H, d), where H is a joint distribution on $[0, \overline{\tau}] \times [0, \overline{\tau}]$ where $\overline{\tau} \geq 1/2$ that has uniform marginals and satisfies Assumption 1. We use H_{τ_i} to denote the marginal CDF for τ_i.

Lemma 2. *Under distribution H, every optimal offer is an MWC offer in which the highest share offered to a non-proposing player, s_B, is no more than $\overline{\tau}$.*

Proof. Lemma 1 establishes that every offer s in which $s_A \leq d$ is suboptimal. Therefore, it suffices to consider only offers s in which $s_A > d$.

We show that every offer s in which $s_B + s_C > \overline{\tau}$ is suboptimal.[31] Consider an alternative offer s' in which $s'_A = 1 - \overline{\tau}$, $s_B = \overline{\tau}$, and $s_C = 0$. Observe that $s'_A > s_A > d$ and $\Lambda(s_B, s_C, H) \leq 1 = \Lambda(s'_B, s'_C, H)$. Therefore, $\pi(s', H, d) > \pi(s, H, d)$.

We now consider an offer s in which $s_B + s_C \leq \overline{\tau}$ and s_B, s_C are each strictly positive. Let s' be the offer in which $s'_A = s_A$, $s'_B = s_B + s_C$, and $s'_C = 0$. Observe that

$$\begin{aligned}
\Lambda(s_B, s_C) &= H_{\tau_B}(s_B) + H_{\tau_C}(s_C) - H(s_B, s_C) \\
&= \frac{s_B}{\overline{\tau}} + \frac{s_C}{\overline{\tau}} - H(s_B, s_C) \\
&\geq \frac{s_B}{\overline{\tau}} + \frac{s_C}{\overline{\tau}} - H(\min\{s_B, s_C\}, \min\{s_B, s_C\}) \\
&> \frac{s_B}{\overline{\tau}} + \frac{s_C}{\overline{\tau}} \\
&= \Lambda(s'_B, s'_C),
\end{aligned}$$

where the first equality uses the inclusion-exclusion principle; the second equality uses the fact that H has uniform marginals on $[0, \overline{\tau}]$; the first inequality uses the fact that H is non-decreasing in each argument; the second inequality uses the fact that $\min\{s_B, s_C\} > 0$ and therefore Assumption 1 implies that $H(\min\{s_B, s_C\}, \min\{s_B, s_C\}) > 0$, and the final equality follows by construction. As $s'_A = s_A > d$, we reach the conclusion that $\pi(s', H, d) > \pi(s, H, d)$. □

We now characterize the optimal MWC offer in the limit case (H, d).

Lemma 3. *The optimal offer in the bargaining problem (H, d) is an MWC offer of the form $\left(\frac{1+d}{2}, \frac{1-d}{2}, 0\right)$.*

Proof. Lemma 2 establishes that every optimal offer is an MWC offer in which $s_B \leq \overline{\tau}$. Therefore,

$$\begin{aligned}
\arg\max_{s \in S} \pi(s, H, d) &= \arg\max_{x \in [1-d, \overline{\tau}]} \{(1 - x - d)\Lambda(x, 0, H) + d\} \\
&= \arg\max_{x \in [1-d, \overline{\tau}]} (1 - x - d)(x / \overline{\tau}).
\end{aligned}$$

The optimum to this strictly concave problem is characterized by its first-order condition, which delivers the optimal $x = (1 - d)/2$. □

Having considered the limit case, we establish some results for an arbitrary bivariate CDF G such that $G(x, x) > 0$ for some $x < (1 - d)/2$. In light of Lemma 1, an offer optimizes $\pi(\cdot, G, d)$ if and only if it optimizes the modified payoff function $\hat{\pi}(s, G, d) \equiv d + \min\{s_A - d, 0\}\Lambda(s, G)$. We argue that $\hat{\pi}$ is upper semicontinuous in s.

---

[31] This case is relevant only if $d + \overline{\tau} < 1$.

Lemma 4. *The payoff function $\hat{\pi}(s, G, d)$ is upper semicontinuous in s.*

Proof. Recall that the product of two nonnegative upper semicontinuous functions is upper semicontinuous. Observe that $\min\{s_A - d, 0\}$ is nonnegative and continuous in s, and $\Lambda(s_B, s_C, G)$ is, by definition, nonnegative. We establish that $\Lambda(s_B, s_C, G)$ is upper semicontinuous in s by first showing that Λ is right-continuous in its first two arguments; to minimize notation, we suppress G in the argument below.

We establish right-continuity in s_B as a symmetric argument applies for s_C. Consider a sequence of sets A_n where $A_n \equiv \left[(-\infty, s_B + \frac{1}{n}] \times \mathbb{R}\right] \cup [\mathbb{R} \times (-\infty, s_C]]$. Observe that $A_1 \supseteq A_2 \supseteq \ldots$, and $\lim_{n\to\infty} A_n = [(-\infty, s_B] \times \mathbb{R}] \cup [\mathbb{R} \times (-\infty, s_C]]$. Therefore, by Theorem 2.1(ii) of Billingsley (1995), $\Lambda(s_B, s_C) = \lim_{n\to\infty} \Lambda\left(s_B + \frac{1}{n}, s_C\right)$, which establishes that Λ is right-continuous in its first argument.

We now prove that Λ is upper semicontinuous in (s_B, s_C). Fix $\varepsilon > 0$. By right continuity of Λ in its first argument, there exists $\delta_B > 0$ such that

$$\Lambda(s_B + \delta_B, s_C) - \Lambda(s_B, s_C) < \varepsilon/2.$$

Analogously, there exists $\delta_C > 0$ such that

$$\Lambda(s_B + \delta_B, s_C + \delta_C) - \Lambda(s_B + \delta_B, s_C) < \varepsilon/2.$$

Therefore, $\Lambda(s_B + \delta_B, s_C + \delta_C) - \Lambda(s_B, s_C) < \varepsilon$. Because Λ is non-decreasing in each argument, we obtain that for any $(s'_B, s'_C) \leq (s_B + \delta_B, s_C + \delta_C)$, $\Lambda(s'_B, s'_C) < \Lambda(s_B, s_C) + \varepsilon$. Thus, there exists a δ-ball around (s_B, s_C) such that for all (s'_B, s'_C) in that δ-ball, $|\Lambda(s'_B, s'_C) - \Lambda(s_B, s_C)| < \varepsilon$. □

Lemma 5. *There exists an offer that maximizes $\pi(\cdot, G, d)$.*

Proof. Because $\hat{\pi}(\cdot, G, d)$ is upper semicontinuous and the unit simplex S is compact, there exists an offer s that maximizes $\hat{\pi}(\cdot, G, d)$ over S. By Lemma 1, any such offer also maximizes $\pi(\cdot, G, d)$. □

Finally, we complete the argument by establishing that the optimal solution close to this limit case approaches the optimal offer of the limit case. For a bargaining problem (G, d), let $s^*(G, d) \equiv \arg\max_{s\in S} \hat{\pi}(s, G, d)$ be the correspondence of optimal proposals and $\pi^*(G, d) \equiv \max_{s\in \overline{S}} \hat{\pi}(s, G, d)$ be its value.

Lemma 6. *The value function $\pi^*(\cdot)$ is continuous at (H, d) and the correspondence of optimal proposals $s^*(\cdot)$ is upper hemicontinuous at (H, d).*

Proof. Consider a sequence (H_n, d_n) where $H_n \to_W H$ and $d_n \to d$.

Step 1: Fix a small $\varepsilon > 0$. Because $d_n \to d$, there exists N_d such that for all $n \geq N_d$, $d_n < d + \varepsilon$. Assumption 1 assures that for every $x \in \left(0, \frac{1-d-\varepsilon}{2}\right)$, $H(x, x) > 0$. Therefore, the open set $O \equiv (-\infty, \frac{1-d-\varepsilon}{2}) \times (-\infty, \frac{1-d-\varepsilon}{2})$ is of strictly positive H-measure. The Portmanteau Theorem (Billingsley, 1999, Theorem 2.1) then implies that there exists N_H such that for every $n \geq N_H$, $Pr_{H_n}(O)$ is strictly positive. Now consider $n \geq \max\{N_d, N_H\}$. For any such bargaining problem (H_n, d_n), there exists $x < \frac{1-d-\varepsilon}{2} < \frac{1-d_n}{2}$ such that $H_n(x, x) > 0$. Therefore, the assumption that there exists $x < \frac{1-d_n}{2}$ such that $G(x, x) > 0$ applies to all $G = H_n$ for $n \geq \max\{N_d, N_H\}$. We therefore restrict attention to such distributions and invoke Lemmas 1, 4 and 5 going forward.

Step 2: We show that for every $s \in S$, and for every sequence $(s^n)_{n=1,2,\ldots}$ where $s^n \to s$, $\limsup_n \hat{\pi}(s^n, H_n, d_n) \leq \hat{\pi}(s, H, d)$. To establish this claim, observe that

$$\begin{aligned}
&\limsup_n \hat{\pi}(s^n, H_n) - \hat{\pi}(s, H) \\
\leq &\limsup_n [\hat{\pi}(s^n, H_n) - \hat{\pi}(s, H_n)] + \limsup_n [\hat{\pi}(s^n, H_n, d_n) - \hat{\pi}(s^n, H_n, d)] \\
&+ \limsup_n [\hat{\pi}(s, H_n) - \hat{\pi}(s, H)].
\end{aligned}$$

The first term on the RHS is weakly negative because $\hat{\pi}(\cdot, H_n)$ is upper semicontinuous by Lemma 4; the second term is zero because $\hat{\pi}$ is continuous in its third argument, and the final term on the RHS is weakly negative by the Portmanteau Theorem because $\hat{\pi}$ is upper semicontinuous and bounded from above.

Step 3: We show that for every $\varepsilon > 0$, for every $s \in S$, there exists a sequence $(s^n)_{n=1,2,\ldots}$ such that $\liminf_n \pi(s^n, H_n, d_n) \geq \pi(s, H, d) - \varepsilon$.

For s in which $s_A \leq d$, we obtain that $\hat{\pi}(s, H, d) = d$. This step then holds trivially as it suffices to take any s^n in which $s_A^n = 0$.

Now consider $s_A > d$. Let $s^\varepsilon = (s_A - 2\varepsilon, s_B + \varepsilon, s_C + \varepsilon)$ for $\varepsilon \in (0, \frac{s_A - d}{2})$. Also let

$$\Lambda^o(s_B, s_C, H) \equiv \int_{\tau_B \times \tau_C} \mathbb{1}_{\max\{s_B - \tau_B, s_C - \tau_C\} > 0} \, dH,$$

denote the probability of the open set $[(-\infty, s_B) \times \mathbb{R}] \cup [\mathbb{R} \times (-\infty, s_C)]$ under measure

H. Because $\Lambda^o\left(s_B+\varepsilon, s_C+\varepsilon, H^n\right) \leq \Lambda\left(s_B+\varepsilon, s_C+\varepsilon, H^n\right)$,

$$\begin{aligned}
\hat{\pi}(s^{\varepsilon}, H_n, d_n) &= d_n + \max\{s_A-\varepsilon-d_n, 0\}\Lambda\left(s_B+\varepsilon, s_C+\varepsilon, H_n\right) \\
&\geq d_n + \max\{s_A-\varepsilon-d_n, 0\}\Lambda^o\left(s_B+\varepsilon, s_C+\varepsilon, H_n\right).
\end{aligned}$$

Given that $H_n \rightarrow_W H$ and $[(-\infty, s_B) \times \mathbb{R}] \cup [\mathbb{R} \times (-\infty, s_C)]$ is an open set, the Portmanteau Theorem implies that

$$\liminf_n \Lambda^o(s_B+\varepsilon, s_C+\varepsilon, H^n) \geq \Lambda^o(s_B+\varepsilon, s_3+\varepsilon, H). \tag{1}$$

Therefore,,

$$\begin{aligned}
\liminf_n \hat{\pi}(s^{\varepsilon}, H_n, d_n) &\geq d + \max\{s_A-\varepsilon-d, 0\}\Lambda^o\left(s_B+\varepsilon, s_C+\varepsilon, H\right) \\
&\geq d + \max\{s_A-\varepsilon-d, 0\}\Lambda(s_B, s_C, H) \\
&\geq \hat{\pi}(s, H) - \varepsilon,
\end{aligned}$$

where the first inequality follows from $d_n \rightarrow d$ and (1), the second inequality follows from $\Lambda^o\left(s_B+\varepsilon, s_C+\varepsilon, H\right) \geq \Lambda(s_B, s_C, H)$, and the final inequality follows from $\Lambda(s_B, s_C, H) \leq 1$.

Step 4: The desired conclusion then follows from (Santambrogio, 2023, Propositions 7.4-7.5): $\hat{\pi}(\cdot, H_n, d_n)$ then Γ-converges to $\hat{\pi}(\cdot, H, d)$, which then implies that $\pi^*(H_n, d_n)$ converges to $\pi^*(H, d)$ and the correspondence $s^*(\cdot)$ is upper hemicontinuous at (H, d). □

Online Appendices

Our Online Appendices are organized as follows:

- Appendix B shows the tables supporting the bar graphs.
- Appendix C shows additional tables and figures.
- Appendix D shows exhibits for all matches.
- Appendix E contains some sample instructions.

B Tables Supporting Bar Graphs

Table 6: Table For Figure 1

Location	Treatment	Mean Payoff (%)	Std
Caltech	1-Perfect	70.00	7.07
Caltech	2-Perfect	60.23	2.87
Caltech	3-Perfect	50.19	2.85
UCI	1-Perfect	60.94	2.94
UCI	2-Perfect	49.54	0.83
UCI	3-Perfect	42.82	0.92
UCI	3-Partial-Incl	44.11	0.98
UCI	3-Partial-Excl	43.11	1.33
UCI	3-None	41.11	0.93

Table 7: t-statistics and p-values by Treatment and Location

Location	Treatment	t-stat	p-value $H_0 : \mu = 0.5$ $H_1 : \mu > 0.5$	p-value $H_0 : \mu = 0.5$ $H_1 : \mu \neq 0.5$
Caltech	1-Perfect	23.300	0.000	0.000
Caltech	2-Perfect	28.500	0.000	0.000
Caltech	3-Perfect	0.627	0.266	0.532
UCI	1-Perfect	42.100	0.000	0.000
UCI	2-Perfect	-5.650	1.000	0.000
UCI	3-None	-101.000	1.000	0.000
UCI	3-Partial-Excl	-34.700	1.000	0.000
UCI	3-Partial-Incl	-54.000	1.000	0.000
UCI	3-Perfect	-87.700	1.000	0.000

Table 8: Table for Figure 3

Location	Treatment	Grand Coalition	Dictatorial	MWC Weak	MWC Strong	MWC NA
Caltech	1-Perfect	11.36	2.27	86.36	–	–
Caltech	2-Perfect	12.50	1.14	76.14	10.23	–
Caltech	3-Perfect-Excl	23.08	–	57.69	19.23	–
Caltech	3-Perfect-Incl	13.46	1.92	–	–	84.62
UCI	1-Perfect	30.92	0.66	67.76	0.66	–
UCI	2-Perfect	25.66	–	25.66	48.68	–
UCI	3-Perfect-Excl	47.71	–	28.44	23.85	–
UCI	3-Perfect-Incl	37.25	1.96	–	–	60.78
UCI	3-None	64.29	–	–	–	35.71
UCI	3-Partial-Excl	41.82	–	–	–	58.18
UCI	3-Partial-Incl	40.00	0.95	18.10	40.95	–

Table 9: Table for Figure 4

Location	Treatment	Non Egalitarian	MWC Egalitarian	Grand Coalition Egalitarian
Caltech	1-Perfect	85.23	7.95	6.82
Caltech	2-Perfect	75.00	14.77	10.23
Caltech	3-Perfect-Excl	39.74	44.87	15.38
Caltech	3-Perfect-Incl	67.31	23.08	9.62
UCI	1-Perfect	86.18	7.89	5.92
UCI	2-Perfect	48.68	46.71	4.61
UCI	3-Perfect-Excl	17.43	46.79	35.78
UCI	3-Perfect-Incl	29.41	47.06	23.53
UCI	3-Partial-Excl	18.18	49.09	32.73
UCI	3-Partial-Incl	25.71	51.43	22.86
UCI	3-None	37.86	19.29	42.86

Table 10: Table for Figure 5

Location	Treatment	% Egalitarian Offers	Mean of proposer share
Caltech	1-Perfect	9.21	77.11
Caltech	2-Perfect	7.46	69.43
Caltech	3-Perfect-Excl	62.22	52.69
UCI	1-Perfect	11.65	69.78
UCI	2-Perfect	33.33	55.82
UCI	3-Perfect-Excl	90.32	51.61
UCI	3-Partial-Incl	84.21	50.94

C Additional Tables and Graphs

C.1 Distribution of Offers and Acceptances

Figure 6 omitted some treatments; we show results from those treatments below.

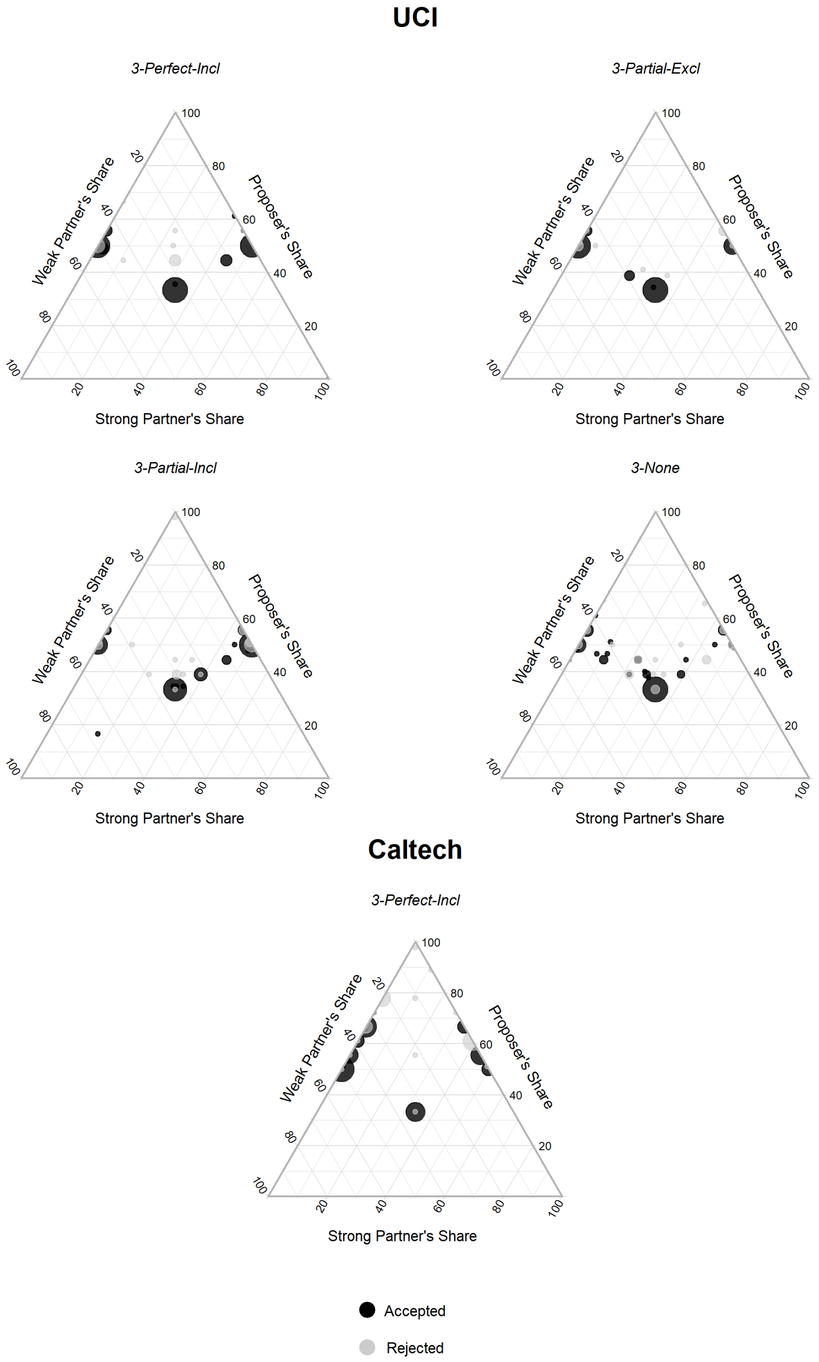

Figure 8: Distributions of Offers and Acceptances, Additional Treatments

C.2 Mapping from First Proposals to Expected Payoffs

Figure 7 omitted some treatments; we show results from those treatments below.

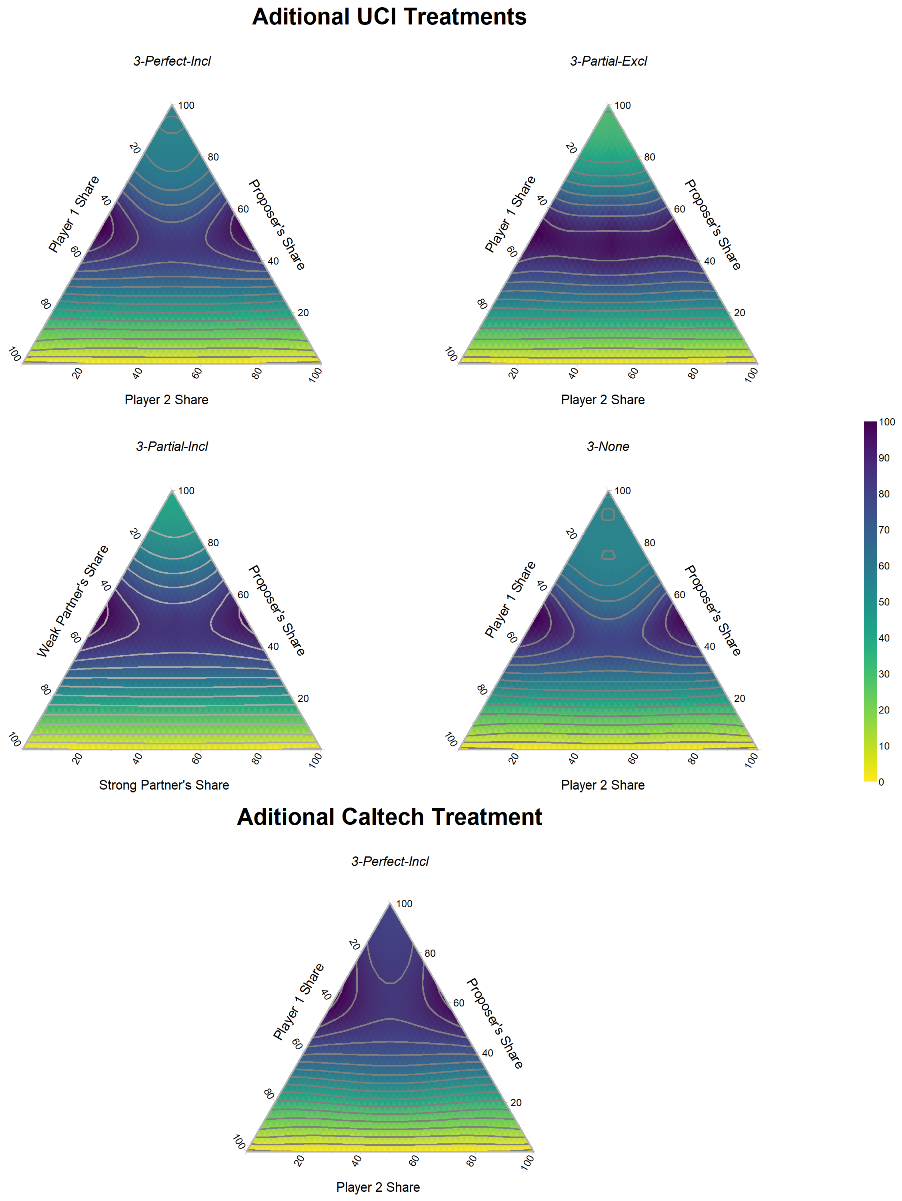

Figure 9: Mappings from First Proposals to Expected Payoffs, Additional Treatments

C.3 Alternate Figures for Section 5.1.1

Figures 1 and 2 used the data from experienced matches in which the first offer was accepted. Below, we produce the corresponding figures of final payoffs across *all* experienced matches, *including matches where the first proposal was rejected*.

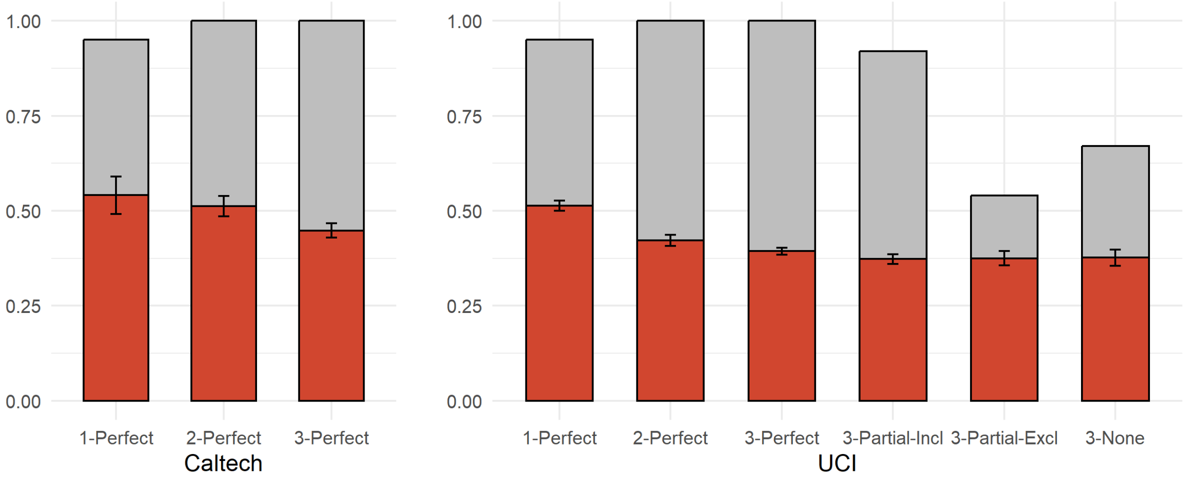

Figure 10: Realized Final Share of First Proposer Across All Experienced Matches (Including Matches Where the First Proposal was Rejected).

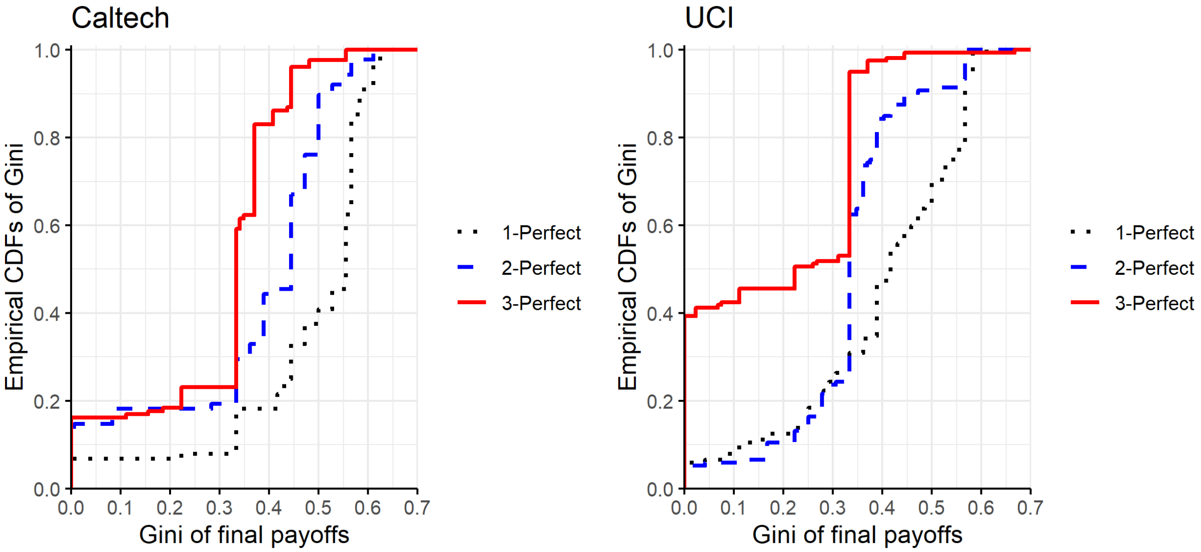

Figure 11: CDFs of Gini Coefficients of Final Payoffs Across All Experienced Matches (Including Matches Where the First Proposal was Rejected)

C.4 Tables for Section 5.2.3

First, we display the mean rejection payoffs for each treatment, alongside the rejection rates of the first offer. These are used in the computation of our two measures of proposer optimization.

Table 11: Mean Rejection Payoffs

Location	Treatment	Rejection rate	Mean Rejection Payoff
Caltech	1-Perfect	22.73	5.00
Caltech	2-Perfect	27.27	27.24
Caltech	3-Perfect-Excl	29.49	26.09
Caltech	3-Perfect-Incl	38.46	43.00
UCI	1-Perfect	15.79	5.00
UCI	2-Perfect	31.58	26.28
UCI	3-Perfect-Excl	21.10	26.96
UCI	3-Perfect-Incl	21.57	25.76
UCI	3-None	20.71	24.52
UCI	3-Partial-Excl	18.18	12.22
UCI	3-Partial-Incl	23.81	15.47

The table below displays what the main text refers to as the second measure, capturing the ratio (in percentages) of the actual first proposer payoff to the predicted optimal payoff, both on the aggregate ("Agg" column) and per coalition type. Some entries exceed 100% because the actual payoff accrued by proposers for some coalition types exceeded the predicted optimal payoffs.

Table 12: Proposer optimization rates: Measure two

Location	Treatment	Optimal Payoff	MWC Weak	MWC Strong	MWC NA	Grand Coalition	Dictatorial	Agg
Caltech	1-Perfect	64	93	–	–	49	8	86
Caltech	2-Perfect	60	94	63	–	60	56	86
Caltech	3-Perfect-Excl	46	104	62	–	75	–	89
Caltech	3-Perfect-Incl	55	–	–	96	62	89	91
UCI	1-Perfect	65	86	8	–	66	140	80
UCI	2-Perfect	49	92	86	–	78	–	86
UCI	3-Perfect-Excl	44	97	97	–	78	–	88
UCI	3-Perfect-Incl	47	–	–	94	75	0	85
UCI	3-None	45	–	–	98	75	–	83
UCI	3-Partial-Excl	41	–	–	104	76	–	92
UCI	3-Partial-Incl	42	105	100	–	71	0	88

The following table presents the fraction of the optimal payoff an equal split offer to a grand coalition would produce on average for each location-treatment.

Table 13: Optimization rates for Equal-split Grand Coalition offers

Location	Treatment	Optimal Payoff	GC Egalitarian
Caltech	1-Perfect	64	50
Caltech	2-Perfect	60	55
Caltech	3-Perfect-Excl	46	70
Caltech	3-Perfect-Incl	55	66
UCI	1-Perfect	65	50
UCI	2-Perfect	49	65
UCI	3-Perfect-Excl	44	75
UCI	3-Perfect-Incl	47	70
UCI	3-None	45	72
UCI	3-Partial-Excl	41	80
UCI	3-Partial-Incl	42	75

Finally, the following table presents the fraction of the optimal payoff that the first proposer obtains if her offer is rejected.

Table 14: Optimization rates for mean rejected payoff

Location	Treatment	Optimal Payoff	MRP
Caltech	1-Perfect	64	8
Caltech	2-Perfect	60	46
Caltech	3-Perfect-Excl	46	56
Caltech	3-Perfect-Incl	55	78
UCI	1-Perfect	65	8
UCI	2-Perfect	49	53
UCI	3-Perfect-Excl	44	61
UCI	3-Perfect-Incl	47	55
UCI	3-None	45	54
UCI	3-Partial-Excl	41	30
UCI	3-Partial-Incl	42	37

D Results from all matches

The following figures replicate Figures 1-6 using data from all matches.

Figure 12: Realized Proposer Share of Accepted First-round Offers

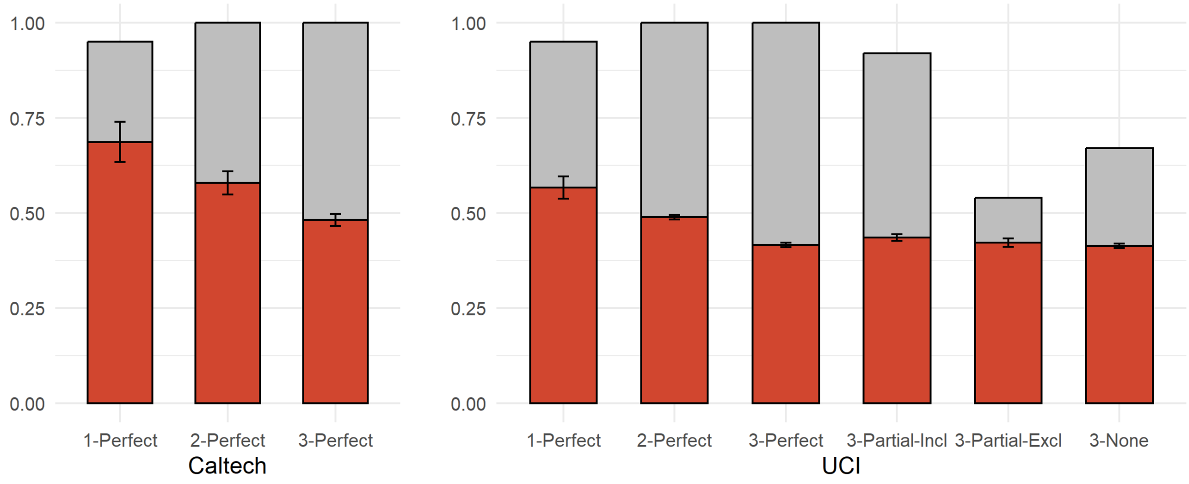

Figure 13: CDFs of Gini Coefficients of Final Payoffs of First-round Accepted Offers

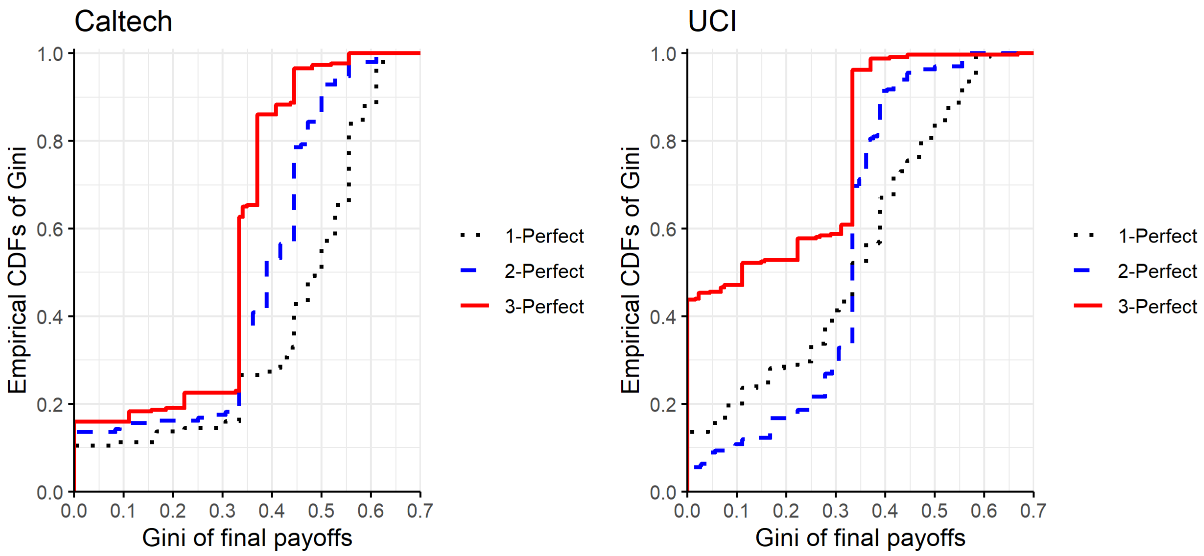

Figure 14: Coalition types in first-round proposals

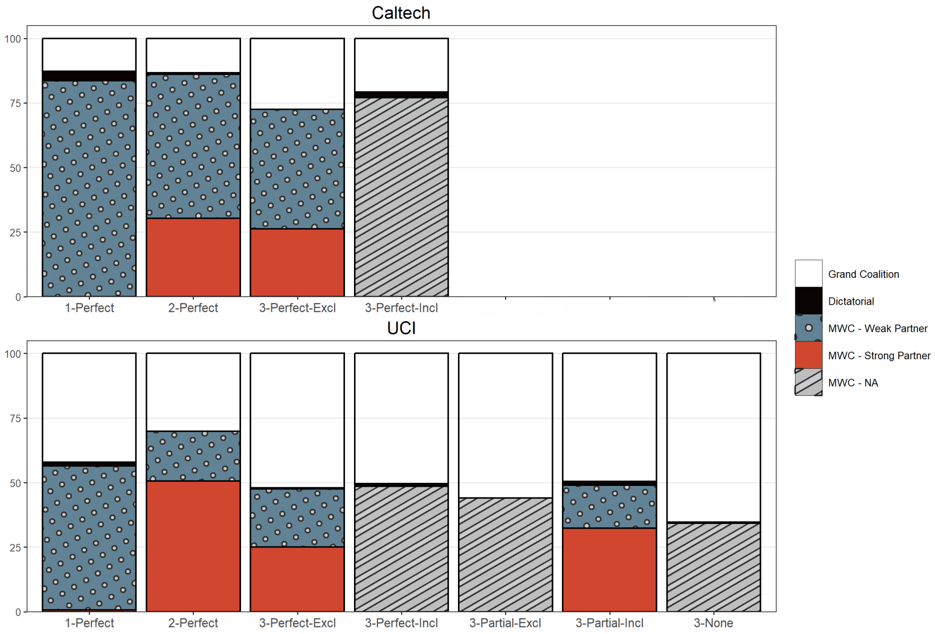

Figure 15: The Frequency of Egalitarian Offers

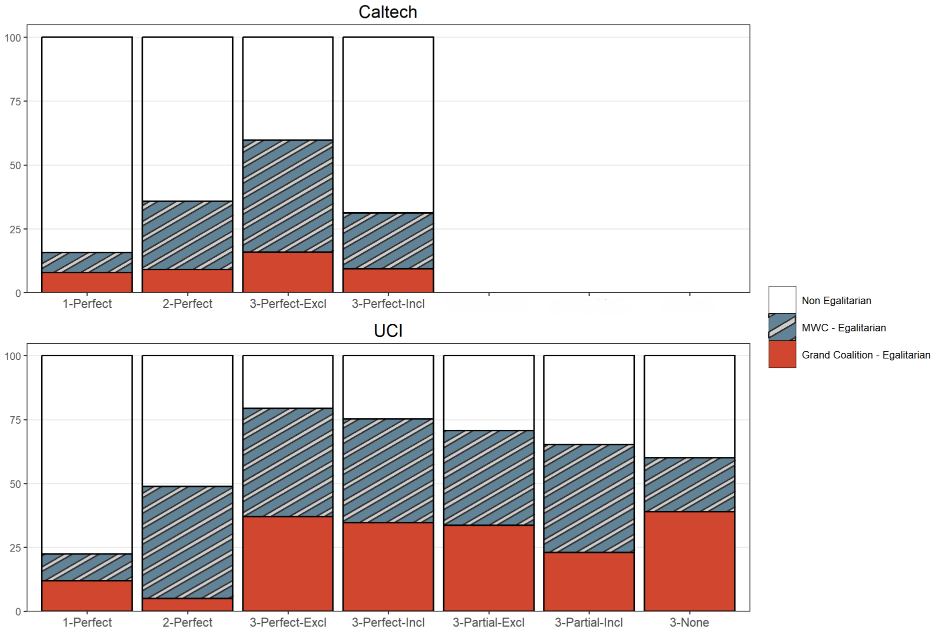

Figure 16: Frequency of Egalitarianism Among MWC Offers to Weak Partners

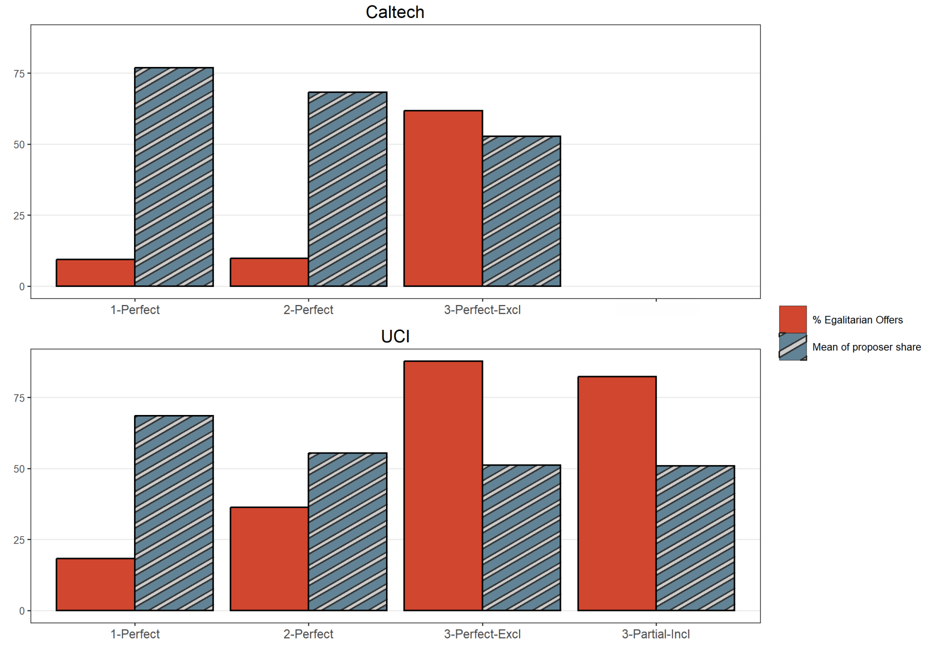

Figure 17: Distribution of Offers and Acceptances

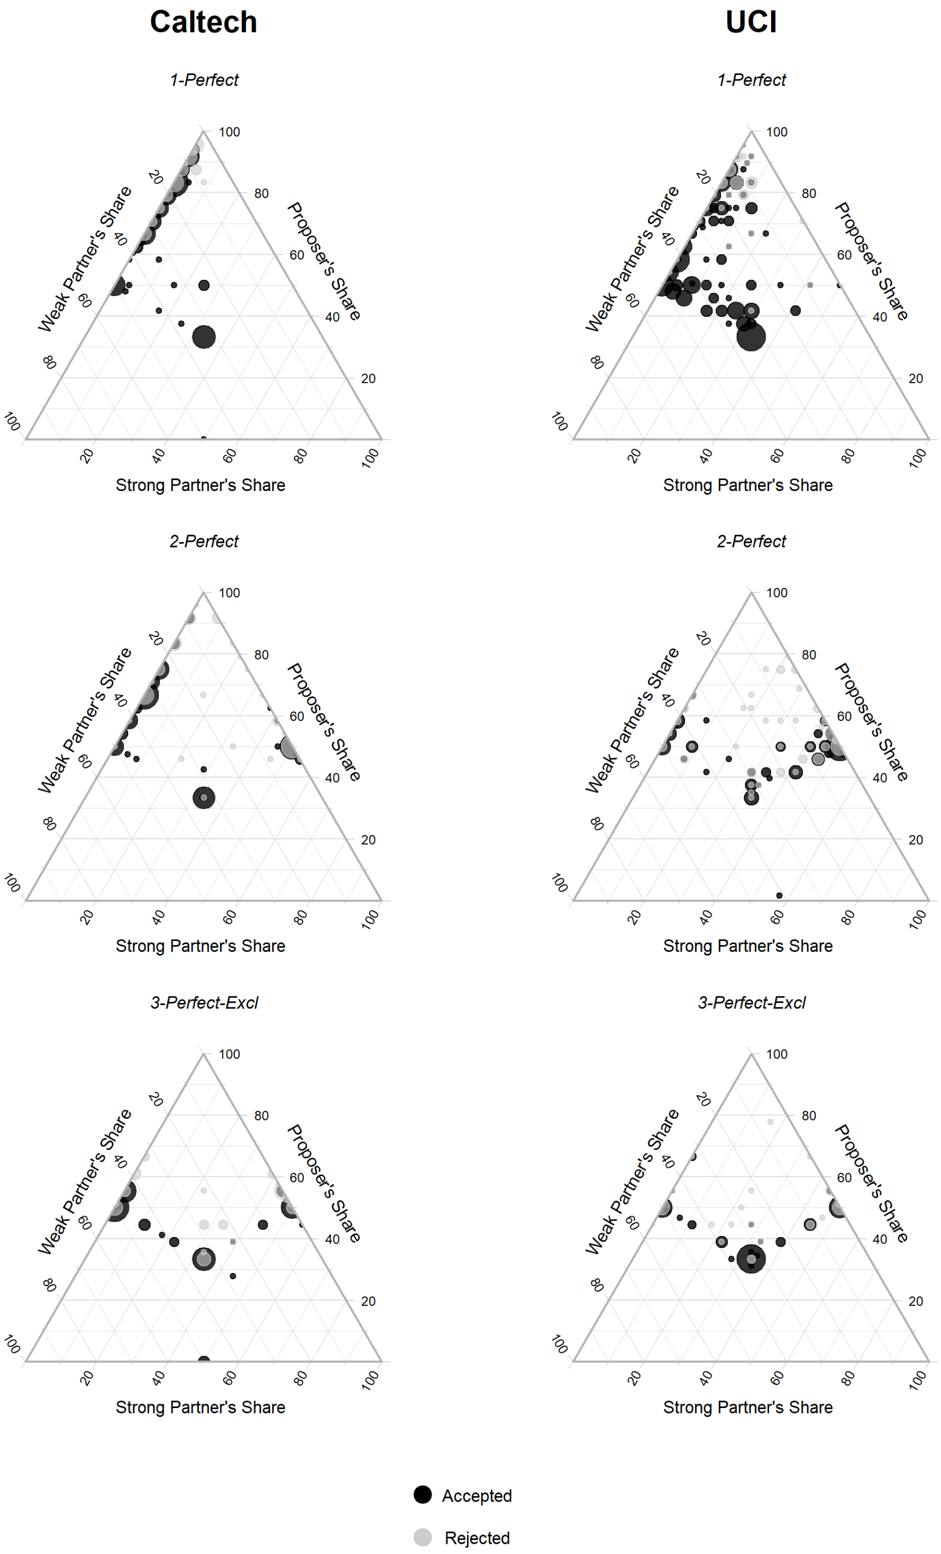

E Sample Instructions (3-Perfect treatment)

The following instructions were read aloud to subjects at the beginning of a 3-Perfect session:

Thank you for agreeing to participate in this group decision-making experiment. During the experiment we require your complete, undistracted attention, and ask that you follow instructions carefully. You may not open other applications on your computer, talk with other students, or engage in other distracting activities, such as using your phone, reading books, etc.

You will be paid for your participation in cash, at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. Your earnings during the experiment are denominated in points. At the end of the experiment, the points that you earn will be converted into US dollars using the rate: 1 point = 10 cents.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear. If you have any questions after the experiment has begun, raise your hand, and an experimenter will come and assist you.

You will make choices over a sequence of 15 matches. At the end of the session, the computer will randomly select one match for payment. Each match is equally likely to be selected for payment. You will be paid what you have earned in this selected match, plus the show-up fee of $15. Everyone will be paid in private and you are under no obligation to tell others how much you earned.

WHAT HAPPENS IN EVERY MATCH

At the beginning of every Match, all subjects are randomly divided into 3-member groups. In addition, each member is randomly assigned an ID number, either 1, 2, or 3. The group assignment and the ID assignments are random, so you will not know the identity of the subjects you are matched with and your group-members will not know your identity. The groups are completely independent of each other and payoffs and decisions in one group have no effect on payoffs and decisions in other groups.

Each group will decide how to allocate budget among the 3 members in your group. Proposals will be voted up or down (accepted or rejected) by majority rule. That is, whenever at least 2 out of 3 members approve a proposal, it passes.

Each Match consists of potentially several Rounds. Your ID number will stay the same in all Rounds of a Match and will not be re-assigned until the next Match. At the beginning of Round 1, your group has a budget of 240 points and needs to decide how to allocate these 240 points among the 3 members in your group. One of 3 members in your group will be randomly chosen to make a proposal of how to allocate 240 points among the 3 members. We will call this member the proposer. Each member has the same chance of being selected as the proposer. Allocations to each member must be between 0 and 240 points. All allocations must add up to 240 points.

After the selected proposer has made his/her proposal, this proposal will be posted on your computer screens with the proposed allocation to you and the other members clearly indicated. You will then have to vote either to accept or to reject this proposal. If the proposal passes (gets 2 or more yes votes), the proposed allocation is implemented and the Match is over. If the proposal is defeated (gets 1 or less votes), there will be another Round of the same Match. However, the available budget will be reduced by 5

In every Round, each group member is equally likely to be selected as the proposer. However, at the beginning of each Round will be told both the ID number of the proposer in the current Round and the ID number of the proposer in the next Round should your group reach the next Round. This process will repeat itself until a proposed allocation passes (gets 2 or more yes votes).

To summarize, the steps of the process will work as follows: 1. At the beginning of each Round, one member is randomly selected to be the proposer. The proposer makes a proposal of how to split available budget. In the very first Round, the budget is 240 points. 2. A vote is held (each member of the group votes to accept or reject the proposal). 3. If 2 or more members of the group accept it (vote yes), then the proposal passes and the Match is over. If the proposal is rejected, then the budget shrinks by 5

Recall, that in each Match, you will be randomly re-matched into groups of 3 members each. Each member of the group will be assigned an ID number (from 1 to 3), which is displayed on the top of the screen. Once the Match is over, you will be randomly re-matched to form new groups of 3 members and you will be assigned a NEW ID for the next Match. Please make sure you know your ID number when making your decisions. Since ID numbers will be randomly assigned prior to the start of each Match, everyone's ID number in a Match is only temporary, and will usually change from Match to Match.

COMMUNICATION

In each Round, after the proposer is selected but before he/she submits his/her proposal, members of a group will have the opportunity to communicate with each other using the chat box. The communication is structured as follows. On the top of the screen, each member of the group will be told his/her ID number. You will also know the ID number of the member who is currently selected to make a proposal (proposer).

Below you will see a box, in which you will see all messages sent to either all members of your group or to you personally. You will not see the chat messages that are sent privately to other members. In the box below that one, you can type your own message and send it either to the entire group or to particular members of your group. To select members that will receive your message, simply click on the buttons that correspond to the ID numbers of the members who you want to receive this message and hit SEND. You can send message to all members of your group by clicking the SELECT ALL button.

The chat option will be available until the proposer submits her proposal. At this moment the chat option will be disabled. SCREEN LAYOUT

On the left side of your screen, above the chat box, there is a table, which displays information about the current Match. In this table, you will be shown information about the past Rounds of the current Match, the current round of this Match as well as information about the next Round proposer, which will be relevant if the current Round proposal fails to obtain a majority of votes.

Take a look at this table. The first row of the table refers to Round 1. Specifically, the first row of the table displays the available budget in Round 1 (240 points) as well as the ID number of the proposer. The table also displays some information about Round 2 of the game (the second row of the table) in case the Round 1 proposal fails. In particular, you can see that the budget that would be available in Round 2 is 228 points and you can also see the ID number of the second round proposer should the group reach second round. Recall that in each Round, the proposer is chosen randomly among the three members of your group. However, in every Round you will be told which member was selected to be the proposer in the next Round should the current Round proposal fail. The ID number of the next round proposer is displayed in the last column of the table labeled Next Round Proposers.

After the Round 1 proposer submits his/her proposal, this proposal will be displayed in the table in the column “Proposal” with your own allocation highlighted in RED. Similarly, after all members of your group submit their votes, all the votes will be listed in this table in the column “Votes” with your vote highlighted in RED.

At the beginning of Round 2, a third row will appear in this table with information about the available budget and potential proposers in Round 3 should you reach this Round. Similarly, at the beginning of Round 3, a fourth row will appear in the table, etc...

In addition, there is a history box at the bottom of the screen. This history box lists the final allocation reached by your group in each Match as well as votes of all group members.

REVIEW

1. The experiment will consist of 15 Matches. There may be several Rounds in each Match. 2. Prior to each Match, you will be randomly divided into groups of 3 members each. Each member in a group will be assigned an ID number. 3. At the start of each Match, one member in your group will be randomly selected to make a proposal of how to allocate 240 points among the three of you. Before he/she submits his/her proposal, members of the group can use the chat box to communicate with each other. You may send public messages that will be delivered to all members of your group as well private messages that will be delivered to members that you specify explicitly. 4. Proposals to each member must be greater than or equal to 0 points. 5. If a simple majority accepts the proposal (2 or more members), the Match ends. 6. If a simple majority rejects the proposal then the group moves on to Round 2, in which the available budget is reduced by 57. Each of the three group members is equally likely to be chosen as the proposer in each Round of each Match. However, at the beginning of each Round, the group members will be told the ID number of the current Round proposer as well as the ID number of the next Round proposer should your group reach the next Round.

Are there any questions? We will now go slowly through one practice Match to familiarize you with the screen. After the practice Match is over, we will start the experiment, in which you will play 15 Matches.