The power of uncertainty in committees with unequal voting rights*

Georg D. Granic▽ Alexander Guggenberger◇ Alexander K. Wagner△

April 14, 2025

Abstract

We study strategic voting in a committee with unequal voting power and information asymmetries. The committee chair has tie-breaking authority, while the regular members have incomplete information about the chair's preferences. Characterizing all equilibria of the voting game reveals that, in contrast to the well-known 'chair's paradox' under complete information, in which tie-breaking authority hurts the chair, tie-breaking authority can also lead to the chair's preferred outcome under incomplete information. Which effect prevails depends crucially on the beliefs of the regular members and the coordination incentives they induce. Laboratory experiments broadly confirm our theoretical predictions, with behavior and voting outcomes closely matching equilibrium predictions. Empirical results also demonstrate that the strategic uncertainty of regular members, due to incomplete information, systematically shifts outcomes in favor of the chair. In support of the cognitive underpinnings of strategic voting, we find that participants with higher strategic sophistication are more likely to adopt equilibrium strategies.

JEL-classification: C72, C92, D71, D72, D82

Keywords: Strategic voting, committee, asymmetric information, tie-breaking power

---

* Granic and Wagner gratefully acknowledge financial support from the Diligentia Stiftung.

▽ Department of Applied Economics, Erasmus School of Economics, Erasmus University Rotterdam, PO Box 1739, 3000 DR Rotterdam, The Netherlands. Email: granic@ese.eur.nl

◇ Department of Economics, University of Salzburg, Austria. Email:alexander.guggenberger@plus.ac.at

△ Department of Economics, University of Salzburg, Austria, and VCEE, University of Vienna, Austria. Email: alexanderk.wagner@plus.ac.at

1 Introduction

Committee voting is common in many organizations, ranging from corporate boards to political bodies. Many such committees are presided by a chair with tie-breaking authority, including several constitutional courts in Europe (e.g., France, Italy, Spain) and the International Court of Justice of the United Nations in The Hague, among others.[1] Theoretical and experimental research has shown that power imbalances due to tie-breaking authority can lead to counterintuitive outcomes, weakening rather than strengthening the chair's position. The 'chair's paradox' is a case in point (Farquharson, 1969; Granic and Wagner, 2021; Alós-Ferrer, 2022). Although the chair is nominally the most powerful, in equilibrium the regular members coordinate against the chair, and the chair's least preferred option is implemented. The chair's paradox illustrates a more general theme in the strategic voting literature that power imbalances can bring about a "paradox of power": greater formal authority does not necessarily translate into more influence on the committee but may even backfire (see, e.g. Acemoglu et al. (2008) and Ke et al. (2022) for the case of coalition formation and Fréchette et al. (2005) and Maaser et al. (2019) for the case of legislative bargaining). This paper investigates how decisions in committees systematically shift as a function of preference uncertainty, revealing a new mechanism by which formal authority can become effective.

Whereas the existing literature emphasizes the effects of power imbalances under full information, real-world committees are also shaped by information asymmetries. Because of their position or selection, chairs often have access to privileged information than, or about, regular members, which can lead to more favorable outcomes for the chair (Blinder, 2004, 2007; Berry and Fowler, 2015, 2018). For example, in many corporate boards, the CEO serves as chair, holds more voting power than other board members, and has access to better information about the company's strategy or about other board members.[2] These examples of

[1] Another notable example is the U.S. Senate, where the Vice President, as its ex-officio chair, casts the tie-breaking vote in case of a 50:50 split. With the Senate almost evenly divided in Joe Biden's presidential term 2021-2025, then Vice President Kamala Harris held significant influence as chair with a total of 33 tie-breaking votes cast. Her predecessor, Mike Pence, cast 13 tie-breaking votes, as Vice President. Until January 2025, a total of 302 tie-breaking votes have been cast by the Vice President since 1789, see: https://www.senate.gov/legislative/TieVotes.htm

[2] Elon Musk is a point in case. As CEO, largest shareholder, and chair of the board of Tesla, he has often been accused of holding excessive influence over Tesla's board. Other examples include chair positions granted

asymmetric information capture what we refer to as preferences uncertainty, that is, situations in which regular members lack precise knowledge about the chair's preferences over possible outcomes. The central questions of our paper, which we investigate theoretically and experimentally, are how the interplay between tie-breaking authority and preference uncertainty shapes strategic voting in committees, and whether chairs can escape the 'chair's paradox' when other committee members are uncertain about the chair's preferences. Understanding how preference uncertainty interacts with power imbalances is critical to understanding how formal authority translates into real influence in committees.

Our vantage point is the model of strategic committee voting due to Farquharson (1969) which considers a committee with three members, a chair and two regular members. Each member votes for one of three alternatives and the winner is determined by plurality voting, with ties being broken in favor of the alternative voted for by the chair.[3] To answer our research questions, we generalize this committee voting model along two key dimensions: (1) by allowing arbitrarily strict preference structures, and (2) by introducing uncertainty about the chair's preferences, which is modeled as incomplete information about the chair's preference type in the game. That is, the chair knows the preferences of regular members, but regular members are uncertain about the true preference of the chair.

We derive the Bayesian Nash equilibria assuming iterative elimination of weakly dominated strategies (IEWDS) to characterize how preference uncertainty affects equilibrium behavior in the committee for different type distributions of the chair. We prove the existence of two classes of equilibria in the incomplete information game: one corresponding to the 'chair's paradox' and another, novel equilibrium where the chair successfully pushes through the preferred outcome. When regular members believe that their preferences are more aligned with each other than with the chair, they have incentives to tacitly coordinate against the chair, reinforcing the 'paradox of power'. This type of equilibrium nests the 'chair's paradox'

ex officio: the FED chair leads the FOMC, the president of the ECB chairs the ECB governing council, and the president of FIFA leads the FIFA council. It can also be argued that permanent members of the UN Security Council wield more power due to their veto rights while also having privileged access to intelligence and diplomatic channels, giving them an informational edge.

[3] In our theoretical setting, it is innocuous whether the chair's tie-breaking vote comes from a separate vote or not, provided that players in the game eliminate weakly dominated strategies, an assumption that is commonly assumed in the strategic voting literature (e.g. Moulin, 1979) and largely fulfilled in our experiment.

originally derived under complete information for a specific preference structure, as a special case. However, if a regular member expects the chair's preferences to be sufficiently aligned, they prefer to tacitly coordinate with the chair instead. Thus, under preference uncertainty, tie-breaking authority can be beneficial if there is sufficient overlap in preferences between the chair and regular members. Our theoretical results reveal how tie-breaking authority can either hurt or benefit the chair, depending on how regular members' beliefs shape their coordination incentives.

Because measuring the impact of preference uncertainty of committee members' behavior is challenging using field data, we collect empirical evidence to test our theoretical predictions using a preregistered laboratory experiment (with $N = 360$ participants). Implementing the theoretical model in the lab, we use the distribution of chair types in the committee, within a matching group, as the key treatment variable to test voting behavior in the different classes of equilibria in the incomplete information game. Different beliefs about the distribution of chair types are induced across treatments by reporting the fraction of participants assigned to different chair types. Our experimental results strongly confirm our theoretical predictions. Our first treatment comparison shows that behavior in the experiment changes with the distribution of chair types as predicted theoretically, going from outcomes that are favorable to disadvantageous for the chair as similarity in preferences between chair and regular members decreases. In each treatment, the corresponding equilibrium strategy profile constitutes the members' modal behavior in the experiment.

Our second treatment comparison provides causal evidence for the behavioral effects of strategic uncertainty about the chair's preference type on the voting behavior of regular committee members (the pure effect of preference uncertainty). Going from certainty to uncertainty about the chair's preferences, while keeping equilibrium predictions fixed, induces small but consistent deviations from equilibrium behavior of regular members, consistent with increased cognitive noise. Empirically, the higher complexity in voting that comes with a non-degenerate distribution of chair types in the game significantly changes committee outcomes in favor of the chair. In line with recent literature on complexity (e.g. Frydman and Nunnari, 2024), our results show, that increased strategic uncertainty, due to the introduction of incomplete information, leads to less successful equilibrium coordination against the chair. Additional analysis shows that participants' ability to perform IEWDS, a primary assumption in our theory, which we elicit in a separate task, significantly improves the propensity to play equilibrium strategies. This finding provides not only evidence for the internal validity of our theoretical modeling but is also a rare experimental validation of this commonly used equilibrium refinement in the voting literature. Taken together, the first-order mechanism of preference uncertainty we identify plausibly applies to many real-world committees. Belief-driven coordination considerations can explain real-world power consolidation in voting bodies with opaque preferences. Our results also imply that transparency about preferences might mitigate strategic biases in committee decisions where uncertainty undermines counterbalancing or coordination strategies.

We proceed as follows. Section 2 summarizes the related literature. In section 3, we develop the equilibrium model of committee decision making under asymmetric voting rights and asymmetric information from which we derive the main behavioral hypotheses, as well as a number of additional hypotheses, which are used to test the validity and the behavioral impact of the model's critical assumptions. Section 4 describes the experimental design. Section 5 reports the main behavioral results and treatment comparisons in detail, before section 6 concludes with a discussion of our main findings.

2 Related literature

By examining how information asymmetries interact with voting power imbalances, our study contributes to several strands of literature. First, while the theoretical literature on tie-breaking votes in committees typically examines specific preference structures under complete information (Farquharson, 1969; Brams et al., 1986; Alós-Ferrer, 2022), our model generalizes these results by considering asymmetric information and broader payoff structures. We are the first to show that equilibrium outcomes depend on the distribution of beliefs among regular members in these settings. Our characterization nests the classical chair's paradox, in which the chair's worst option is implemented in a committee under complete information, as a special case and additionally includes equilibria that implement the chair's preferred outcome, overturning the paradox of power result. Granic and Wagner (2021) investigate how the perception of the chair's tie-breaking power influences voting behavior in favor of the chair. Our paper, in contrast, focuses on the conditions under which chairs can leverage tie-breaking power in the presence of asymmetric information.

Our contribution is also closely related to the literature on weighted voting in committees. Differences in voting weights between members have been investigated, for instance, in the sequential-move Baron-Ferejohn model of legislative bargaining (Ansolabehere et al., 2005; Snyder et al., 2005; Ali et al., 2018). Fréchette et al. (2005) and Maaser et al. (2019) study experimentally the effects of purely nominal differences in voting weights on coalition bargaining. Maaser and Stratmann (2024) study a threshold public goods game with asymmetric voting power where committee members vote for a potentially immoral but for committee members beneficial policy. We show in section 3 that our model is strategically equivalent to a weighted voting model where the chair's weighted vote is greater than the weight of any single regular member but less than their combined weight.

The effects on strategic voting of tie-breaking power, or larger voting weights, has also been studied in the context of coalition formation (Acemoglu et al., 2008; Ke et al., 2022).[4] Ke et al. (2022) provide experimental evidence that, if the surplus in a coalition is negotiated after it is formed, the nominal strength (bargaining power) of a member can turn in a strategic disadvantage, and hence lead to a 'paradox of power' outcome. Our study contributes to a more realistic understanding of the paradox of power observed in strategic voting setting beyond the chair's paradox and games of complete information.

Our results also relate to the literature on information aggregation in committees. Blinder and Morgan (2005, 2008) investigate experimentally how leadership of the chair in monetary policy committees influences outcomes in information aggregation contexts. Bouton et al. (2018) examine the efficiency of different voting rules in the Condorcet jury model. In this stream of literature, uncertainty arises about the true state of the world and committee members are motivated by common interests to match the voting outcome with the correct

[4] The coalition formation model of Acemoglu et al. (2008) demonstrates that the majority coalition can be challenged by others, and that stronger members are more likely to be excluded by others as a preemptive measure to prevent them from gaining dominance at a later point.

state of the world. Since preferences among all committee members are fully aligned, the tie-breaking vote does not induce any strategic considerations. Hughes et al. (2023) investigate information aggregation in diverse committees, when preferences and information structures differ. In contrast to the literature on information aggregation, we do not consider efficiency gains or losses in voting outcomes but focus on purely private value settings with diverse preferences. The private value setting is better suited to studying whether power, such as tie-breaking authority, leads to more beneficial outcomes for the chair, because strategic voting behavior cannot be is independent of efficiency concerns. In this respect, our paper shares similar features to the literature on agenda-setting power in private interest settings and the power to exercise control over outcomes (e.g. Bernheim et al., 2006; Apesteguia et al., 2014). In settings where private and common interests coexist, our findings are likely to hold as long as private interests outweigh common interests.

Finally, our behavioral analysis also adds to the nascent literature on cognitive noise in strategic settings.[5] Our results are closely related to Frydman and Nunnari (2024), who study the impact of cognitive noise on equilibrium behavior in a two player coordination game by varying the complexity of the game. A crucial difference in their setting is that they manipulate the variability of a common payoff parameter in the game, which leads to imprecise payoff perceptions, whereas our uncertainty stems from incomplete information about the chair's preference type in the committee. However, our finding that the additional strategic uncertainty faced by regular members leads to more cognitive noise in individual voting decisions and thus to less successful equilibrium coordination against the chair is consistent with the findings in Frydman and Nunnari (2024).[6]

[5] In single-decision making settings, Enke and Graeber (2023) show that task complexity leads to higher self-reported cognitive uncertainty, measured in an independent task, and more noisy decisions. Enke et al. (2024) provide evidence for 'attenuated reactiveness' of participants to variations in the underlying parameters in a variety of single-decision making problems, suggesting large cognitive (information processing) limitations of human decision makers.

[6] In their setting, complexity is associated with a larger standard deviation of the prior distribution of the game's common payoff parameter. They show that coordination behavior in the high volatility setting leads to more cognitive noise than in the low volatility one.

3 Bayesian chair voting game

We now introduce the theoretical committee voting model, extending earlier work on the chair voting game to a Bayesian setting. Motivated by empirical evidence that shows that chairs are often more resourceful or better connected than regular members due to their position in the committee (Berry and Fowler, 2018, 2015), we assume that the chair has an informational advantage over regular members: the chair knows the preferences of regular members, but regular members are uncertain about the chair's preferences.

Following the setup of the classical chair's paradox (Farquharson, 1969; Alós-Ferrer, 2022), our Bayesian chair voting game is a three-member committee voting game. It involves three alternatives labeled A, B, C, and the outcome is implemented through plurality voting. All three members have regular votes and one member, the designated *chair* of the committee, holds tie-breaking authority: ties are broken in favor of the alternative the chair votes for. Member 2 (*m2*) and member 3 (*m3*) are referred to as "regular members" as they only hold regular votes. Let $I = \{chair, m2, m3\}$ denote the set of players, and let $A_i = \{A, B, C\}$ denote the (common) set of voting actions available to each member. We denote by T_{chair} the set of chair types, and impose a common prior and consistent beliefs over T_{chair}. Regular members have only one type (singleton type set), so we omit types in our notation for $m2$ and $m3$. Regular members and all types of the chair have strict preferences over the three alternatives, with different chair types having different preferences. Payoffs are defined over outcomes and reflect the members' preferences. For member i, the payoff from the outcome k is denoted $\pi_i(k)$. This setting defines a generic Bayesian chair voting game.

Table 1 shows the outcome space of the game from the viewpoint of a specific chair type (i.e., interim perspective). The chair type's vote determines the sub-matrix, and $m2$'s and $m3$'s votes determine the row and column of the sub-matrix, respectively. Voting outcomes of the game can be characterized along two cases. If two or more members agree and vote for the same alternative, it wins the election. For example, if $m2$ and $m3$ both vote for A, A is the outcome irrespective of the chair type's vote. If all members disagree and vote differently, the chair's vote determines the outcome. For example, if the chair type, $m2$, and $m3$ vote for A, B, and C, respectively, the tie-breaking authority of the chair implements outcome A.

Table 1: Generic outcome space for the Bayesian chair voting game.

(i) chair type plays A

		$m3$: A	B	C
$m2$	A	A	A	A
	B	A	B	A
	C	A	A	C

(ii) chair type plays B

		$m3$: A	B	C
$m2$	A	A	B	B
	B	B	B	B
	C	B	B	C

(iii) chair type plays C

		$m3$: A	B	C
$m2$	A	A	C	C
	B	C	B	C
	C	C	C	C

It is important to mention that our Bayesian chair voting game is strategically equivalent to a broad class of Bayesian chair weighted voting games, using plurality voting without a tie-breaking rule.[7] For example, assume that regular members have the voting weights w_{m1} and w_{m2}, respectively, and that for the chair's voting weight w_{chair} the following holds: $w_{m1}, w_{m2} < w_{chair} < w_{m1} + w_{m2}$. The chair's weighted vote is greater than the weight of any single regular member but less than their combined weight $w_{m1} + w_{m2}$. It is straightforward to see that the outcome space of this weighted voting game is the same as the one of our Bayesian chair voting game presented in table 1. All theoretical results derived from our Bayesian chair voting game, thus, extend a broader class of voting games.

3.1 Model implemented in the experiment

Our main theoretical results are as follows. Bayesian chair voting games have two classes of Bayesian Nash Equilibria (BNE) in weakly undominated strategies: one class corresponds to the original chair's paradox equilibrium (the regular members tacitly gang up against the chair) and a novel class in which the committee always implements the most preferred alternative of the chair type. We, hereby, show that chair's paradox-type of equilibria emerge under a broader set of conditions than previously considered. Secondly, we establish the existence of another, novel class of equilibria that yields the opposite outcome of the chair's paradox. The

[7] We abstain from characterizing the entire class of weighted voting games for which the equivalence is true but instead present a simple subclass for which the equivalence holds.

Table 2: Preference profiles and types of chair in the specific model.

Player	Preference type	Preference ordering
Chair	$t = chair$ with p_A	$A \succ_{t=chair} B \succ_{t=chair} C$
Chair	$t = m3$ with p_B	$B \succ_{t=m3} C \succ_{t=m3} A$
Member 2	singleton type	$C \succ_{m2} A \succ_{m2} B$
Member 3	singleton type	$B \succ_{m3} C \succ_{m3} A$

main insight from the Bayesian framework is that regular members' beliefs about the chair's preferences determine if tie-breaking authority is detrimental (chair's paradox) or beneficial (novel equilibrium class) to the chair.

For expositional clarity, we begin with the specific model that we implemented in the experiment. This specific model takes some simplifying assumptions motivated by practical considerations of the laboratory experiment. Importantly, this model serves as a minimal example that contains the core mechanisms underlying our general theoretical results. Section 3.2 shows that these assumptions are innocuous and that our theoretical results generalize to all preference structures and all possible type spaces of the chair.

As in the classical complete-information chair voting game, we fix $m2$'s preferences to $C \succ_{m2} A \succ_{m2} B$, and $m3$'s preferences to $B \succ_{m3} C \succ_{m3} A$. The chair can be of two types. The first type, called $t = chair$, has preferences $A \succ_{t=chair} B \succ_{t=chair} C$. This type has alternative A as the most preferred alternative, and we denote the belief that the chair is of this type by $p_A \in [0, 1]$. In the classical complete-information chair voting game, this would be the only type of the chair, leading to the chair's paradox results known in the literature: in equilibrium $m2$ and $m3$ vote for C, implementing the chair's worst outcome (Farquharson, 1969; Alós-Ferrer, 2022).

Now suppose that regular members are uncertain about the chair's preference type. The chair can be of a second type, denoted as $t = m3$, with identical preferences as $m3$: $B \succ_{t=m3} C \succ_{t=m3} A$. B is the most preferred alternative of this type and the regular members' belief that the chair type is $t = m3$ is given by p_B. Naturally, we set $p_A + p_B = 1$. Table 2 summarizes the specific model.

The payoffs of the specific model are as follows: implementing a member's most preferred alternative, say k, yields a payoff of $\pi(k) = x$ for that member, implementing the second most preferred alternative yields a payoff of y, and implementing the least preferred alternative yields z, where $x > y > z$. For example, if A is implemented, $m2$ receives a payoff of $\pi_{m2}(A) = y$, $m3$ a payoff of $\pi_{m3}(A) = z$, and chair type $t = chair$ a payoff of $\pi_{t=chair}(A) = x$. Furthermore, we assume that the payoffs are equidistant between adjacently ranked alternatives ($x - y = y - z$).

Depending on the values of p_A (p_B), our main treatment variation in the experiment, the game's BNE surviving IEWDS (iterated elimination of weakly dominated strategies) can be characterized as follows:

EQ1: $s_{chair} = (A, B)$, $s_{m2} = C$, $s_{m3} = C$ if $p_B \leq p_A$.[8]

EQ2: $s_{chair} = (A, B)$, $s_{m2} = C$, $s_{m3} = B$ if $p_B \geq p_A$.

The equilibrium **EQ1** captures the spirit of the original chair's paradox, in which regular members gang up on the chair and vote for C. The chair's vote does not influence the outcome and C is implemented. In our Bayesian setting, C is the worst outcome for type $t = chair$, and the second best alternative of type $t = m_2$. There is a positive probability that in this equilibrium class chairs receive the worst or second best outcome, and they never receive their most-preferred outcome. Whether or not chair's-paradox-type equilibria emerge depends on the condition $p_B \leq p_A$, i.e. on the regular members' belief about the chair's type. Note that the original complete information chair voting game is nested in our Bayesian version. In the original game, the chair is only of type $t = chair$, so $p_A = 1$. In this case, (A, C, C) is the only equilibrium surviving IEWDS.

Conversely, if $p_B \geq p_A$, **EQ2** emerges. **EQ2** represents the opposite outcome of the chair's paradox. The regular members vote for different alternatives, and the chair's vote implements the type's most preferred alternative. In this equilibrium, the chair successfully implements

[8] In strategy notation, $s_{chair} = (A, B) = [s_{chair}(chair) = A, s_{chair}(m_3) = B]$.

its preferred outcome with certainty.

Comparing **EQ1** with **EQ2** reveals the strategic considerations of the different members in the game. The chair's respective type and $m2$ always vote for the most preferred alternatives in equilibrium. In contrast, member $m3$ faces the following trade-off. Voting for $s_{m3} = C$ implements C, $m3$'s second most preferred outcome, which yields a payoff of $\pi_{m3}(C) = y$ given the other players' equilibrium strategies. Voting $s_{m3} = B$, on the other hand, induces a lottery over the two alternatives with expected payoff $\pi_{m3}(p_A A p_B B) = p_A z + p_B x$. If the share of chair types $t = m3$ is high (p_B is high), the lottery is skewed towards $m3$'s most preferred alternative B and the expected payoff from the lottery exceeds the payoff of y which $m3$ receives from voting for C. If the share of chair types $t = chair$ is high (p_A is high), the lottery is skewed towards $m3$'s least preferred alternative A and the resulting expected payoff is below the payoff of y which $m3$ receives from voting for C.

3.2 General Bayesian chair voting game

In this section, we relax some of the assumptions made in section 3.1. Specifically, we now consider all possible strict preferences and all possible type spaces of the chair. We demonstrate that the equilibrium results presented in section 3.1 hold more generally. For our Bayesian Nash Equilibrium (BNE) analysis, we focus on pure strategy equilibria and equilibria in undominated strategies, i.e., strategies that are not weakly dominated for any player. This is common in the literature on voting games (Moulin, 1979; Kohlberg and Mertens, 1986; Dhillon and Lockwood, 2004; Granić, 2017). Furthermore, we eliminate all weakly dominated strategies of all players at each round of elimination to avoid order problems in equilibrium selection (see Fudenberg and Tirole, 1991, p. 461).

We begin by eliminating weakly dominated strategies. Our main argument links weak dominance to pivotality considerations, focusing on events in which members can change the implemented outcome by changing their vote (if they cannot change the outcome, they are indifferent between all strategies). The following two key observations underpin our analysis.

Observation 1. *If two members vote for the same alternative, that alternative is implemented with a majority, irrespective of the third member's vote. The remaining member is indifferent* *between all the strategies in this case.*

Observation 2. *If $m2$ and $m3$ vote for different alternatives, the chair's vote determines the outcome, either by breaking a tie or by creating a 2-to-1 majority for the alternative the chair votes for.*

These observations establish an important result for any chair type $t \in T_{chair}$. By observation 1, if $m2$ and $m3$ vote for the same alternative, chair type t is indifferent between all actions. If $m2$ and $m3$ vote for different alternatives, the chair's vote determines the outcome by observation 2. Chair type t can never do worse, but only do better by voting for the most preferred alternative. This establishes our first result.

Result 1. *Any chair type $t \in T_{chair}$ has exactly one undominated action in any Bayesian chair voting game: voting for their most preferred alternative.*

We now turn to regular members. We show that voting for the least preferred alternative is weakly dominated by voting for the most preferred alternative. To establish the latter, we note that, under consistent beliefs and using the Harsanyi transformation, the payoff of a regular member is just a convex combination of the type space distribution over the chair types' and regular members' actions. Therefore, if voting for the least preferred alternative is weakly dominated for a fixed type for the chair, it is also weakly dominated for the whole game. Consider w.l.o.g. $m2$ with preferences $A \succ_{m2} B \succ_{m2} C$. Compare voting for the most preferred alternative A with voting for the least preferred alternative C. Pivotality under plurality voting could induce outcome changes in two ways.

1. By switching the vote from C to A, alternative A now receives more votes and wins. According to the preferences of the regular member this is an improvement in the outcome.

2. Switching from C to A reduces the votes for C, potentially making B the winner or creating a tie. This can improve or worsen the outcome, depending on the alternatives previously winning.

We prove in appendix A that pivotality in the second way is precluded by the tie-breaking vote of the chair. Thus, a regular member can only do better, but never worse when voting for the most preferred alternative in comparison to voting for the least preferred alternative. This yields our second result.

Result 2. *For regular members $m2$ and $m3$, voting for the least preferred alternative is weakly dominated by voting for the most preferred alternative.*

These results hold for any preference structure, provided preferences are strict. Applying results 1 and 2 simplifies the search for BNE in undominated strategies: chairs types vote for their most preferred alternative and regular members do not vote for the least preferred alternatives. This reduces the original game to a 2×2 normal-form game between $m2$ and $m3$, treating chair types as automata who vote for their type's most preferred alternative. This simplification allows us to condense the beliefs, as regular members $m2$ and $m3$ only have to a form a belief about chair types' having certain most preferred alternatives. Denote by p_k the belief of regular members that the chair types' most preferred alternative is $k \in \{A, B, C\}$.

From the perspective of regular members, if they vote for the same alternative, this alternative is implemented by observation 1. If they vote for different alternatives, the two regular members face a lottery over outcomes as the chair type's vote determines the outcome in this case, see observation 2. In the latter case, regular members expect outcome A to be implemented with probability p_A, and outcome B and C with probability p_B and p_C, respectively. We denote this lottery outcome as $p_A A p_B B p_C C$.

Now consider regular member i. W.l.o.g. assume that i has preferences $A \succ_i B \succ_i C$, and that the payoffs associated with outcomes are $\pi_i(A) = x_i > \pi_i(B) = y_i > \pi_i(C) = z_i$. Table 3 summarizes the implemented outcomes in the reduced game after eliminating i's weakly dominated strategy voting for C, as per result 2, considering all possible votes of the other regular member j. At this stage, no assumptions are made about j's preferences, so the specific weakly dominated alternatives for j remain undefined.

Table 3 also allows us to characterize the best responses (BR) of i against j's possible strategies:

BR1 If the other member j votes $s_j = A$, member i's unique best response is to vote for $s_i = A$, assuming non-degenerative beliefs. Voting for $s_i = A$ induces payoff $\pi_i(A) = x_i$, whereas voting for $s_i = B$ induces $\pi_i(p_A A p_B B p_C C) = p_A x_i + p_B y_i + p_C z_i$.

Table 3: Outcomes given regular members i and j strategies treating the chair as an automaton, for preferences $A \succ_i B \succ_i C$.

	$s_j = A$	$s_j = B$	$s_j = C$
$s_i = A$	A	$p_A A p_B B p_C C$	$p_A A p_B B p_C C$
$s_i = B$	$p_A A p_B B p_C C$	B	$p_A A p_B B p_C C$

Notes: As an example, consider the upper middle entry: if $s_i = A$ and$s_j = B$, the outcome will be A or B by plurality rule if the chair votes for A or B, respectively, and the outcome will by C by tie-breaking rule if the chair votes for C. This leads to the outcome representation $p_A A p_B B p_C C$.

BR2 If member j votes $s_j = B$, which is i's second-best alternative, the best response of i depends on the expected payoffs of the lottery outcome. Voting $s_i = A$ induces the expected payoff $p_A x_i + p_B y_i + p_C z_i$. Voting $s_i = B$ implements B with payoff y_i. An increase in p_A (the probability that A wins) or in x_i (the payoff associated with outcome A) raises the expected payoff of voting $s_i = A$. Thus, i weighs the certainty of receiving the second-best outcome y_i against a lottery over all three alternatives. The more the lottery skews toward i's most preferred alternative A, the more likely it is that the lottery's expected payoff exceeds y_i.

BR3 If member j votes $s_j = C$, member i's worst alternative, i is indifferent between the two available strategies as they yield the same expected payoff.

The BNE in undominated strategies are given by $m2$ and $m3$ mutually best replying against each other. As the game is symmetric with regard to regular members, we can directly establish the two classes of equilibria that emerge in the game. Note that each class can contain several equilibria. We provide one example equilibrium for each class to demonstrate that it is not empty. Recall that the chairs' strategy is playing the most preferred alternative for each type in all equilibria, and we omit specifying the chair's equilibrium strategy. Again, BNE are formulated assuming w.l.o.g. that $A \succ_i B \succ_i C$.

1. Equilibria with regular members voting for the same alternative.

(a) BNE exist in which both regular members vote for A, the most preferred alternative of member i. For example, this is the case if A is the second most preferred alternative of member j and $\pi_j(A) \geq p_A\pi_j(A)+p_B\pi_j(B)+p_C\pi_j(C)$ (by **BR1** and **BR2**).

(b) BNE exist in which regular members vote for B, the second most preferred alternative of member i. For example, this is the case if B is the second most preferred alternative of member j and $\pi_r(B) \geq p_A\pi_r(A)+p_B\pi_r(B)+p_C\pi_i(C)$ holds for both members $r = m2, m3$ (by **BR2**).

2. BNE exist in which regular members vote for different alternatives. For example, this is the case if member j votes for the least preferred alternative of member i and vice versa (by **BR3**). This is possible if the least preferred alternative of one member is not the least preferred alternative of the other member (e.g., $m2$ has preferences $A \succ_{m2} B \succ_{m2} C$ and votes for $s_{m2} = A$, and $m3$ has preferences $C \succ_{m3} B \succ_{m3} A$ and votes for $s_{m3} = C$.)

Table 3 together with **BR1** to **BR3** provide a recipe for constructing the different equilibria of the game. Which equilibria (co)-exist in a given Bayesian chair voting game depends on the values of the payoff and belief parameters. The two classes of equilibria correspond to the two equilibria of the specific model discussed in section 3.1. The first class of equilibria captures the spirit of the complete-information chair's paradox akin to observation 1 where both regular members vote for the same alternative and the chair's vote does not influence the outcome. The second class of equilibria captures our novel insight and is related to observation 2: the chair type always receives the most preferred outcome.

To close this section, we reanalyze the specific model presented in section 3.1 to derive its equilibria. Applying results 1 and 2, we can reduce the game to a 2×2 game between $m2$ and $m3$ as shown in table 4 in terms of the implemented outcomes. Using all feasible combinations of **BR1**, **BR2**, and **BR3**, we obtain three pure action BNE in undominated strategies in the specific model.

EQ1: By **BR1** and **BR2**, the pure-strategy profile $s_{chair} = (A, B), s_{m2} = C, s_{m3} = C$ is a

Table 4: Outcomes in the specific model after elimination of weakly dominated strategies.

	$s_3 = B$	$s_3 = C$
$s_2 = A$	$p_A A p_B B$	$p_A A p_B B$
$s_2 = C$	$p_A A p_B B$	C

BNE in undominated strategies if $p_B \leq p_A$.

EQ2: By **BR1** and **BR2** the pure-strategy profile $s_{chair} = (A, B), s_{m2} = C, s_{m3} = B$ is a BNE in undominated strategies if $p_B \geq p_A$.

EQ3: By **BR3**, the pure-strategy profile $s_{chair} = (A, B), s_{m2} = A, s_{m3} = B$ is a BNE in undominated strategies, independent of any parameter values for payoffs and beliefs.

We further refine our equilibrium selection and apply iterative elimination of weakly dominated strategies (IEWDS). It is straightforward to see that voting C iteratively weakly dominates voting A for $m2$ in the reduced game by inspecting table 4. Suppose $m3$ votes B. In this case, $m2$ is indifferent between the two actions as they induce the same lottery over outcomes. If $m3$ votes C, $m2$ has a unique best reply: to vote C. The latter implements the most preferred outcome with a probability of 1. Applying IEWDS thus eliminates A for $m2$ and with it **EQ3**.

What are the surviving BNE in this game? We can see from table 4 that if $m2$ votes C, $m3$ weakly prefers playing B over C iff:

$$\begin{aligned} \pi_{m3}(p_A A p_B B) &> \pi_{m3}(C) \\ \Leftrightarrow p_A z + p_B x &\geq y \\ \Leftrightarrow p_B &\geq p_A (y - z)/(x - y) \\ \Leftrightarrow p_B &\geq p_A \end{aligned} \tag{1}$$

The latter inequality captures the trade-off $m3$ is facing in the game. If the share of chair types $t = m3$ is high (p_B is high), the lottery is skewed towards $m3$'s most preferred alternative B and the expected payoff from it exceeds the payoff of y which $m3$ receives from voting C. We thus obtain the two equilibria of the specific model presented in section 3.1.

4 Experimental design

We implement the Bayesian chair voting game presented in section 3.1 in a controlled laboratory experiment to test the model's main equilibrium predictions. In addition to the voting game, participants answer several incentivized tasks allowing us to track the effects of individual-level variation in strategic sophistication on voting behavior. The experimental design, hypotheses, and statistical analysis were preregistered at AsPredicted under project #120563, available at https://aspredicted.org/QSV_5NM.

Voting game. Following the specific model in section 3, the three committee members chair, $m2$, and $m3$ submit votes for one of the three alternatives A, B, or C (simultaneously and independently). Plurality voting decides which alternative is implemented. In case of a tie, the alternative voted for by the chair is implemented. The preferences of $m2$ and $m3$ are common knowledge. Incomplete-information in the experiment was induced by informing regular members about the probability of being matched with a chair type $t = chair$ (probability p_A) and a chair type $t = m3$ (probability $p_B = 1 - p_A$).

We induce preferences over alternatives using monetary incentives. Participants receive €15 if their most preferred outcome is implemented (corresponding to x in our model notation). They receive €10 if the second most preferred outcome (y), and €5 if the least preferred outcome (z) is implemented. Figure 1 shows the decision screen of how the payoffs are displayed to regular members. The payoffs induce the same ordinal preferences and fulfill the assumptions presented in section 3.1. The screens for chairs differ slightly. As chairs know their type, the irrelevant part of the payoff table is grayed out for chair participants.

Matching groups and rounds. Participants play 16 rounds of the game with random re-matching within a matching group of 12 participants. Before the first round, four participants from a matching group are each assigned randomly to the role of the chair, member 2, and member 3, respectively. Participants assigned to the chair role are randomly stratified into chair types $t = chair$ and $t = m3$, with the composition of the types determined by the treatment detailed below. In each round of the voting game, we randomly match one member 2, one member 3, and one chair of the matching group. A participant's role and preference type remains fixed for all rounds. This procedure was common knowledge among

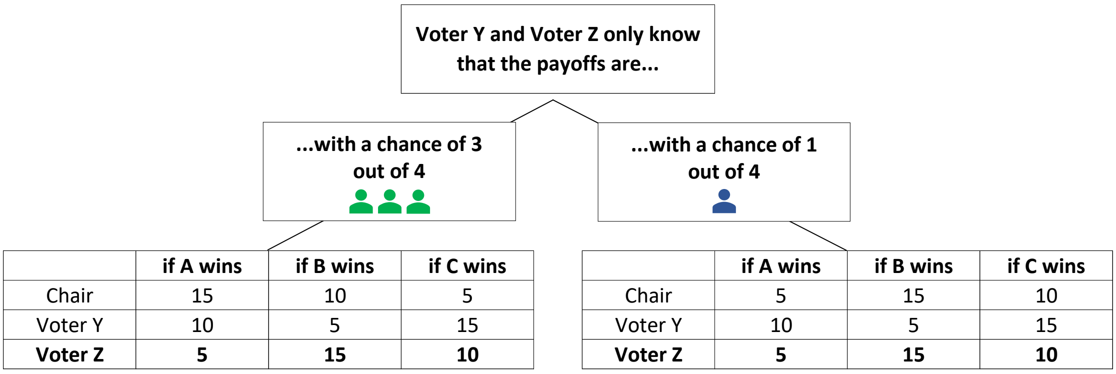

	if A wins	if B wins	if C wins
Chair	15	10	5
Voter Y	10	5	15
Voter Z	**5**	**15**	**10**

	if A wins	if B wins	if C wins
Chair	5	15	10
Voter Y	10	5	15
Voter Z	**5**	**15**	**10**

Figure 1: Bayesian voting game in the experiment (treatment $p_A = 3/4$).

Notes: We referred to the regular members as Voter Y ($m2$) and Z ($m3$) in the experiment.

all participants. At the end of a round, participants are provided with an overview of their own voting history, the number of votes each alternative received, the winning alternative, and the realized preference type of the chair (by highlighting the realized payoff table on the feedback screen similar to the one represented in figure 1). Feedback and repetition of the one-shot game allows participants to learn about the strategic environment.[9]

Treatments. We implement three between-subject treatments that vary the distribution of chair types in a matching group. Treatments, thus, differ only in the composition of chair types $t = chair$ and $t = m3$. Because participants maintain their assigned role (and type) over the course of all voting games, the belief p_A corresponds to the prior probability with which regular members are matched to chair type $t = chair$ in the matching group. This probability is explained intuitively using a frequency approach. We consider the following three treatments:

- Treatment $p_A = 1/4$: one of the four chairs in a matching group is of type $t = chair$, the remaining three chairs are of type $t = m3$. Theory predicts **EQ2** with $s = ((A, B), C, B)$.

- Treatment $p_A = 3/4$: three of the four chairs in a matching group are of type $t = chair$,

[9] Our design ensures that participants are matched with the same player on average four times over the course of the 16 games. This should minimize concerns about repeated-game effects in the voting data while allowing us to consider learning effects of participants in the experiment.

the remaining one chair is of type $t = m3$. Theory predicts **EQ1** with $s = ((A, B), C, C)$.

- Treatment $p_A = 1$: all four chairs in a matching group are of type $t = chair$. Theory predicts **EQ1** with $s = ((A, \cdot), C, C)$.

The treatment $p_A = 1$ corresponds to the original, complete information chair's paradox game. By comparing this treatment with the $p_A = 3/4$ treatment, we can investigate the effect of introducing uncertainty about the chair's preferences on committee members' behavior. Since the equilibrium is the same in both treatments, differences in behavior between treatments can be attributed to the complexity that the incomplete information about the chair type adds to the game.

Additional tasks and questionnaire. After the 16 rounds of the voting game, we further elicit individual-level characteristics in various incentivized tasks and a questionnaire, in the following sequence. *Strategic reasoning:* Participants in a matching group play a 2/3-beauty contest game (Nagel, 1995) and we use the guesses as a proxy for the participants' capability of strategic reasoning. *Iterative elimination of weakly dominated strategies (IEWDS):* In a sequence of 2 player normal-form games participants are asked to delete weakly dominated strategies iteratively. A game is counted as solved correctly if a participant correctly deletes iteratively weakly dominated strategies for both players in the simple matrix game. We use the number of correctly solved games (from zero to four) as a measure of the participant's ability to iteratively eliminate weakly dominated strategies. *Expected payoff maximization:* Participants are asked to choose payoff-maximizing actions in a simple decision problem under uncertainty. In five questions, the probability with which the computer chooses an action is varied and potential payoffs of different choices remain fixed; in five other questions, the participants' payoffs are varied but the probability with which the computer chooses them are fixed. The number of correct answers to these 10 questions informs us about whether voting behavior is partly driven by (in)ability to maximize expected payoffs but removes the *strategic uncertainty* players are facing in the Bayesian game about the strategy of other members. *Demographic questions:* Finally, participants answer a number of socioeconomic background questions, including age, gender (non-binary), field of study, and nationality which are used as demographic controls in the regression analysis. From the latter two, we construct a binary variable indicating whether the participant's field of study involves STEM skills, and the country of origin.

Further procedures and earnings. The experiment was programmed in oTree (Chen et al., 2016) and run at the experimental economics laboratory of the Vienna Center for Experimental Economics of the University of Vienna. A total of 360 individuals participated in 16 sessions in the experiment. For each treatment, we collected data from 10 independent matching groups, each with 12 participants. The randomization unit is a matching group. According to our power analysis, $N = 10$ independent observations on the matching-group level are sufficient (for 80% power and at a 5% significance level) to perform the planned statistical analyses. Before the start of the voting game, participants read the onscreen instructions and had to correctly answer a number of comprehension questions, see appendix E.1 for details. The total earnings from the voting game and additional tasks were disclosed to each participant privately and anonymously after completion of all tasks. In the voting game, one round of the game was randomly selected in a matching group to determine participants' earnings. The random payment mechanism ensures incentive compatibility (Azrieli et al., 2018) and minimizes possible confounds, such as carry-over effects from the repeated one-shot interactions. Participants received additional earnings from the incentivized tasks as well as €0.25 per correct comprehension questions. A session lasted about 55 minutes, including private and anonymous payments. Average earnings amounted to €17.03 (median €17.00, ranging between €7.50 and €32.75) per participant.

Participant characteristics. Before turning to the results, table 5 provides descriptive statistics for the various background characteristics of our $N = 360$ participants and the additional tasks they completed. Typical of a laboratory experiment, the participants are mainly students with an average age between 24 and 25. About 60% of our participants are female, the vast majority of them come from a European country, and about 30% of them study STEM fields. The average *guess* in the beauty game was between 42 and 46, on average they solved around 5 out of 10 expected payoff calculation problems correctly, and applied IEWDS correctly in 2 out of 4 normal-form matrices. Regarding the treatment balancing checks, none of these variables differ significantly across treatments at the 5 percent level.

Table 5: Background characteristics by treatment.

	$p_A = 1/4$		$p_A = 3/4$		$p_A = 1$		
Variable	mean	sd	mean	sd	mean	sd	p-value
Age	24.09	5.22	24.74	5.05	25.25	5.79	0.09
Sex female	0.58	0.50	0.58	0.50	0.62	0.49	0.83
Sex male	0.39	0.49	0.38	0.49	0.34	0.48	0.72
Prefer not to disclose/inter	0.03	0.16	0.04	0.20	0.04	0.20	0.73
Eastern europe	0.32	0.47	0.31	0.46	0.33	0.47	0.91
Western europe	0.56	0.50	0.57	0.50	0.52	0.50	0.71
STEM	0.32	0.47	0.25	0.43	0.23	0.42	0.30
Guess in beauty contest	42.68	21.66	43.66	22.18	46.42	23.27	0.38
Expected payoff optimizer	4.92	0.96	4.84	0.87	4.71	1.18	0.43
IEWDS	2.16	1.14	2.32	1.14	2.21	1.11	0.50

Notes: Reported $p-$values are based on Kruskal-Wallis tests for continuous variables and Chi-Squared for binary variables.

5 Experimental results

We now present the main results of our experiment, relating observed behavior to our equilibrium predictions. In section 5.1, we compare voting behavior and outcomes between the two incomplete information treatments which differ in the distribution of chair types ($p_A = 1/4$ and $p_A = 3/4$). In section 5.2, we investigate the effect of strategic uncertainty, caused by the complexity of incomplete information about the chair's type, on voting behavior of regular members. We further quantify non-equilibrium play of our participants to explain committee outcomes and provides evidence for our critical modeling assumptions in section 5.3. Unless otherwise noted, we restrict our main analysis to the last 8 (of the 16) rounds due to expected learning effects, as specified in the preregistration.[10]

[10] Appendix B provides evidence for learning effects in voting behavior over all rounds.

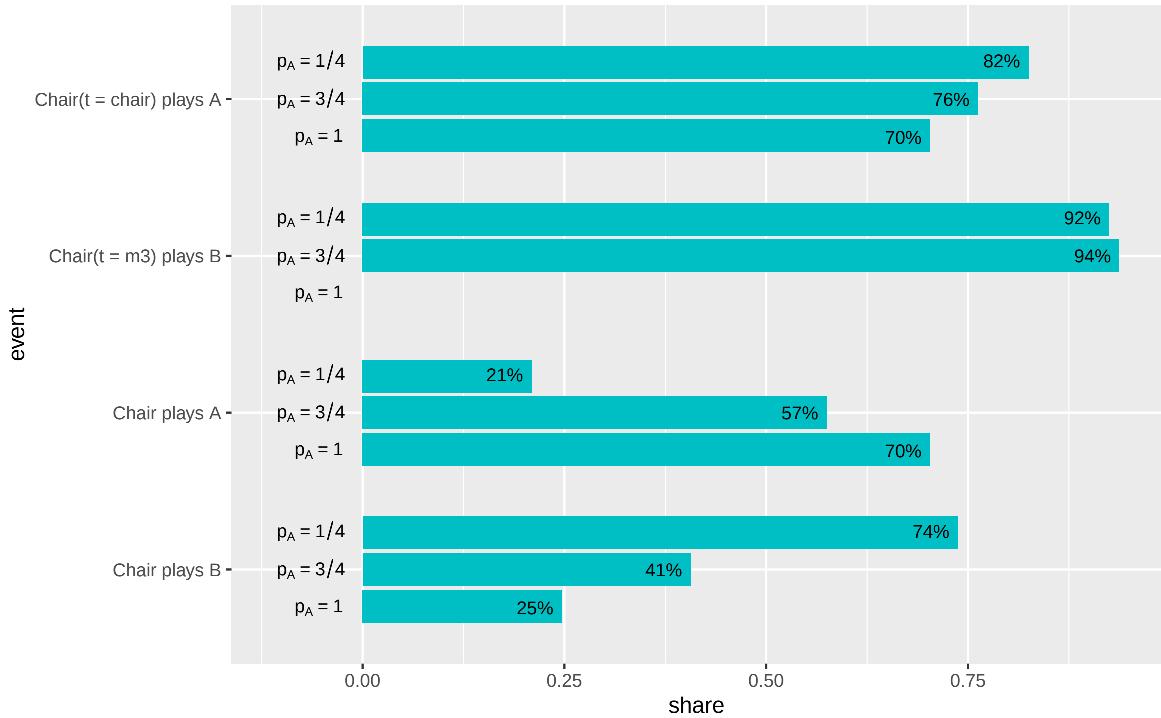

Figure 2: Voting behavior of chair types by treatment.

5.1 Voting under incomplete information

The two incomplete-information treatments we compare vary only in the distribution of the chair's type, with equilibrium prediction **EQ2** in treatment $p_A = 1/4$, and **EQ1** in $p_A = 3/4$. In both equilibria, chair types and regular member $m2$ vote for their most preferred alternative. The hypothesis regarding individual-level behavior follows directly from these equilibrium predictions.

Hypothesis 1 (Individual-level behavior)**.** *We expect that, (i) chairs, conditional on type, vote for A and B as often, (ii) members $m2$ vote for C as often, and (iii) members $m3$ vote for C more (B less) often in the $p_A = 3/4$ than in the $p_A = 1/4$ treatment.*

Our results show that the modal behavior in the two treatments closely matches equilibrium predictions. The top two bar charts in figure 2 present the share of votes cast for the theoretically predicted alternatives across treatments for each chair type. Consistent with theory, chair types mainly vote for the most preferred alternative, with no pronounced dif-

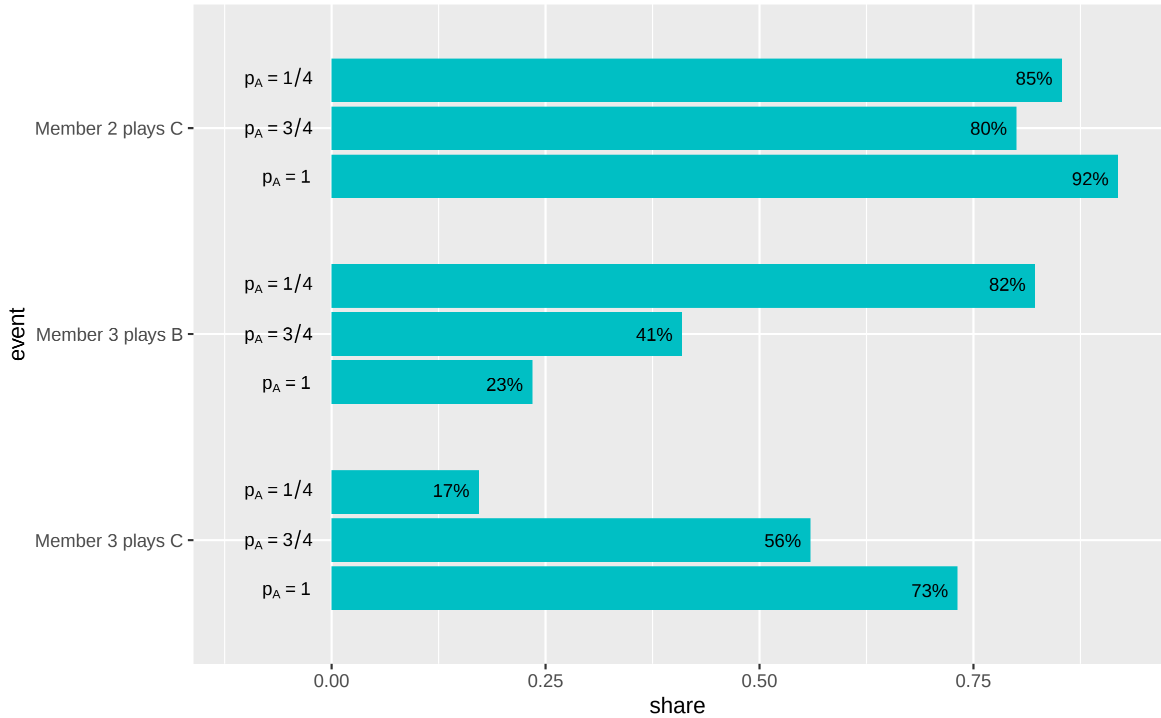

Figure 3: Voting behavior of regular members by treatment.

ferences in voting behavior between treatments. Chairs of type $t = chair$ vote for their most preferred alternative A in 82% ($p_A = 1/4$) and 76% ($p_A = 3/4$) of all elections, chair type $t = m3$ in 92% and 94% of the elections, respectively. The bottom two bar charts in figure 2 show chair behavior pooled across types. Note that the distribution of chair types differs between treatments: $t = chair$ is more frequent in $p_A = 3/4$, whereas $t = m3$ is more frequent in $p_A = 1/4$. Using the pooled chair types data, chair behavior varies between treatments in line with equilibrium predictions.

A similar pattern emerges for regular members. As shown in figure 3, the share of votes cast for theoretically predicted alternatives represents the modal behavior for regular members as well. Member $m2$ votes C in 85% and 80% of all elections in $p_A = 1/4$ and $p_A = 3/4$, respectively. As predicted, $m3$ displays high sensitivity to treatment-induced belief changes, reducing the vote for B from 82% of all elections in the $p_A = 1/4$ treatment to 41% in $p_A = 3/4$. This decrease is associated with a shift to voting C, which increases from 17% of all elections in $p_A = 1/4$ to 56% in $p_A = 3/4$. It appears that members $m2$'s behavior in the experiment matches equilibrium predictions more closely than $m3$'s behavior. We will come back to this observation when discussing non-equilibrium play in detail in section 5.3.

Overall, our equilibrium predictions are strongly supported by observed individual level behavior in the experiment. All the treatment comparisons reported above are validated by a series of Wilcoxon-Mann-Whitney (WMW) tests, using independent observations at the matching group level, with $n = 10$ observations per treatment, which are corrected for multiple hypothesis testing using Holm-Bonferroni.[11] The results are summarized in table 6 for chairs and in table 7 for regular members. We also run probit regressions in appendix D, table D1 and D2, which corroborate all our main results.

Next, we turn to the analysis of treatment differences in strategy profiles and outcomes, with the hypothesis again following directly from the theory.

Hypothesis 2 (Equilibrium strategy profiles and outcomes)**.** *We expect that (i) the equilibrium strategy profile* $((A, B), C, C)$ *is played more and* $((A, B), C, B)$ *less often, (ii) alternative* C *wins more (A and B less) often, and (iii) chairs receive on average lower payoffs in the* $p_A = 3/4$ *than in* $p_A = 1/4$ *treatment.*

Recall that chair participants in our experiment are assigned a fixed chair type, either $t = chair$ or $t = m3$. This implies that we do not observe a chair's strategy, which would be the behavior of a chair participant in both type roles, only the action chosen by a given type. To account for this fact, we define an election to be *consistent* with a strategy profile if the three voting actions observed in an election are induced by the strategy profile. For example, the observed action profile (A, C, B) is consistent with the predicted equilibrium strategy profile $((A, B), C, B)$ if the chair in the committee was of type $t = chair$, and inconsistent with the strategy profile if the chair was of type $t = m3$. For ease of exposition, we will say that a strategy profile is played if the observed election is consistent with an equilibrium strategy profile.

Figure 4 presents the share of elections consistent with equilibrium strategy profiles across

[11] We jointly test 9 null-hypotheses for hypothesis 1, 6 tests for hypothesis 2, 9 tests for hypothesis 3, and 6 tests for hypothesis 4. While we acknowledge that an insignificant test cannot confirm the truth of the null hypothesis, it is notable that insignificant tests arise precisely in cases where our model predicts no treatment differences. This observation lends further credence to the robustness of our equilibrium predictions.

Table 6: Voting behavior of chair types by treatment.

comparison	variable	H_0	share 1	share 2	p-value	consistent
3/4 vs 1/4	Chair plays A	<	0.58	0.21	0.00	Yes
3/4 vs 1/4	Chair plays B	>	0.41	0.74	0.00	Yes
3/4 vs 1/4	Chair(t=chair) plays A	=	0.76	0.82	0.25	Yes
3/4 vs 1/4	Chair(t=chair) plays B	=	0.23	0.17	0.25	Yes
3/4 vs 1/4	Chair(t=m3) plays A	=	0.01	0.00	0.94	Yes
3/4 vs 1/4	Chair(t=m3) plays B	=	0.94	0.92	0.66	Yes
3/4 vs 1	Chair plays A	>	0.58	0.70	0.06	No
3/4 vs 1	Chair plays B	<	0.41	0.25	0.00	Yes
3/4 vs 1	Chair(t=chair) plays A	=	0.76	0.70	0.60	Yes
3/4 vs 1	Chair(t=chair) plays B	=	0.23	0.25	0.60	Yes
3/4 vs 1	Chair(t=m3) plays A	=	0.01	-	-	-
3/4 vs 1	Chair(t=m3) plays B	=	0.94	-	-	-

Notes: H_0 refers to the null hypothesis of the WMW test. Column "consistent" indicates whether the test result is in line with theory. P-values are corrected for multiple testing within each hypothesis using the Holm-Bonferroni correction. Note that chair type $t = m3$ behavior cannot be compared between the $p = 3/4$ and the $p = 1$ treatments because this type does not exist in the later treatment.

treatments. As predicted, we find that equilibrium strategy profile $((A, B), C, C)$ is played more frequently and $((A, B), C, B)$ less frequently in treatment $p_A = 3/4$ than in $p_A = 1/4$. Specifically, equilibrium profile $((A, B), C, B)$ is played in 60% of elections in the $p_A = 1/4$ treatment, compared to only 27% in the $p_A = 3/4$ treatment. Recall that in this equilibrium the chair type's most preferred alternative is implemented. In contrast, the equilibrium profile $((A, B), C, C)$ is played in 35% of elections in the $p_A = 3/4$ treatment and in 15% of elections in treatment $p_A = 1/4$. In this chair's paradox equilibrium $((A, B), C, C)$, the chair's vote does not influence the outcome. As suggested by our analysis of individual behavior, the higher degree of non-equilibrium play in treatment $p_A = 3/4$ is largely driven by regular $m3$ deviating from the equilibrium strategy.

Table 7: Voting behavior of regular member by treatment.

comparison	variable	H_0	share 1	share 2	p-value	consistent
3/4 vs 1/4	M2 plays C	=	0.80	0.85	0.66	Yes
3/4 vs 1/4	M3 plays B	>	0.41	0.82	0.00	Yes
3/4 vs 1/4	M3 plays C	<	0.56	0.17	0.00	Yes
3/4 vs 1	M2 plays C	=	0.80	0.92	0.06	Yes
3/4 vs 1	M3 plays B	=	0.41	0.23	0.60	Yes
3/4 vs 1	M3 plays C	=	0.56	0.73	0.60	Yes

Notes: H_0 refers to the null hypothesis of the WMW test. Column "consistent" indicates whether the test result is in line with theory. P-values are corrected for multiple testing within each hypothesis using the Holm-Bonferroni correction.

Notably, the equilibrium strategy profiles represent the modal observation in both treatments. However, compared to individual-level behavior, we observe a lower degree of congruence between theory and experimental outcomes. This is because outcomes are a noisy mapping of individual votes, and a single deviation by any committee member results in a non-equilibrium strategy profile. We also observe the theoretically predicted differences between treatments in implemented outcomes, see figure 5. As expected, alternative C wins more often (45% vs 20%), and B wins less often (23% vs 63%) in the $p_A = 3/4$ than in the $p_A = 1/4$ treatment. Interestingly, alternative A does not win less often in $p_A = 3/4$ but, in fact, wins more frequently. This observation is driven by differences in non-equilibrium behavior of committee members. Our theory predicts that alternative A wins 25% of elections in $p_A = 1/4$, all due to ties being broken by the chair type $t = chair$ (equilibrium action profile (A, B, C)). No ties are predicted in $p_A = 3/4$ and alternative A is predicted to win 0% of election. However, the experiment reveals a smaller fraction of elections being decided by a tie break in $p_A = 1/4$ (15%) compared to $p_A = 3/4$ (27.5%). Tie breaks in favor of A occur almost exclusively when the chair is of type $t = chair$, so in favor $t = chair$'s most preferred alternative. In this sense, chairs benefit benefit substantially from non-equilibrium behavior

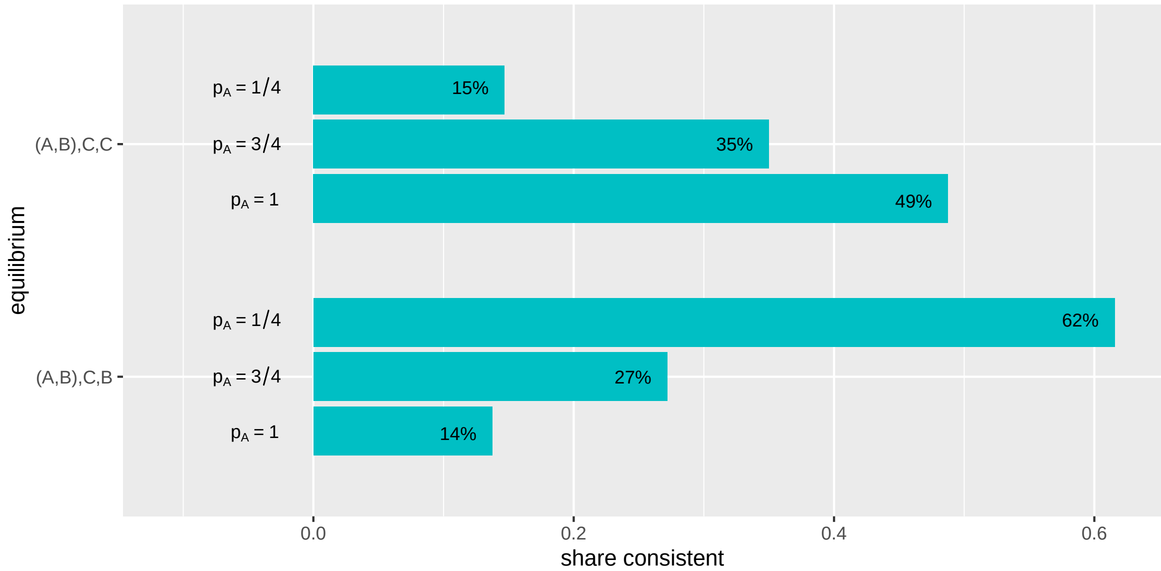

Figure 4: Share of elections consistent with equilibrium predictions.

in treatment $p_A = 3/4$. All above observations are corroborated by a series of WMW tests summarized in table 8.

As a direct consequence of the observed strategy profiles in the experiment, we find that chairs fare better in $p_A = 1/4$ than in $p_A = 3/4$, with average payoffs of €13.52 versus €10.53, respectively (WMW test, p-value < 0.01). This payoff observation is particularly important because it provides a comprehensive measure of a participant's overall 'performance' in the experiment. In particular, the winning frequency of the most preferred alternative, as analyzed above, does not account for which alternative (second or least preferred) wins in other elections. Thus, we conclude that chairs benefit significantly from the change in other members' beliefs about the type distribution, both in terms of the frequency with which their most-preferred alternative wins and their overall payoff. As predicted, chairs fare better in comparison to regular members in $p_A = 1/4$ ($m2$ earns €7.81, $m3$ earns €12.33). For $p_A = 3/4$, chair earnings are between the ones of regular members ($m2$ earns €11.08, $m3$ earns €9.55). The latter result is mainly due to unsuccessful coordination of regular members against the chairs, the reason for which we will discuss in section 5.3.

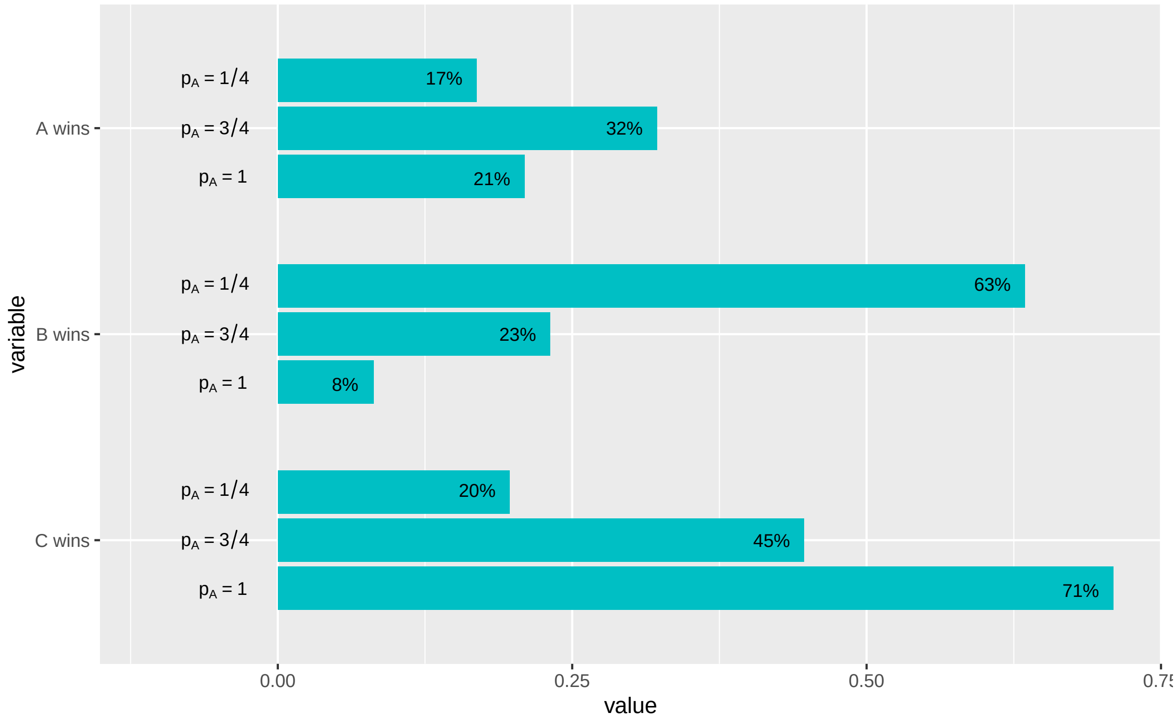

Figure 5: Share of winning alternatives by treatment.

5.2 The effects of preference uncertainty on voting

We now address whether uncertainty among regular members about the preferences of the chair confers an advantage on the chair of a committee. To this end, we compare behavior in treatment $p_A = 1$, where regular members know with certainty that the chair is of type $t =$ chair, against treatment $p_A = 3/4$, where the chair is of type $t =$ chair with probability $3/4$ and of type $t = m3$ with probability $1/4$. In both treatments, the equilibrium predictions prescribe chairs of type $t =$ chair to vote for A, and both regular members to vote for C. Any treatment differences in observed behavior of regular members can therefore be attributed to the added strategic uncertainty, introduced by the complexity of the non-degenerate type distribution of the chair in the $p_A = 3/4$ treatment. Strategic uncertainty here reflects uncertainty about the behavior of other committee members in the game. Although it can have many different sources in general, such as uncertainty about the degree of rationality of the other members, our treatment comparison quantifies specifically the effect of an increase

Table 8: Strategy profiles and outcomes by treatment.

comparison	variable	H_0	share 1	share 2	p-value	consistent
3/4 vs 1/4	((A, B), C, C)	$<$	0.35	0.15	0.00	Yes
3/4 vs 1/4	((A, B), C, B)	$>$	0.27	0.62	0.00	Yes
3/4 vs 1/4	A wins	$>$	0.32	0.17	0.99	No
3/4 vs 1/4	B wins	$>$	0.23	0.63	0.00	Yes
3/4 vs 1/4	C wins	$<$	0.45	0.20	0.00	Yes
3/4 vs 1	((A, B), C, C)	$=$	0.35	0.49	0.15	Yes
3/4 vs 1	((A, B), C, B)	$=$	0.27	0.14	0.15	Yes
3/4 vs 1	A wins	$=$	0.32	0.21	0.15	Yes
3/4 vs 1	B wins	$=$	0.23	0.08	0.00	No
3/4 vs 1	C wins	$=$	0.45	0.71	0.08	Yes

Notes: H_0 refers to the null hypothesis of the WMW test. Column "consistent" indicates whether the test result is in line with theory. P-values are corrected for multiple testing within each hypothesis using the Holm-Bonferroni correction.

in strategic uncertainty, or complexity, caused by the introduction of incomplete information about the chair's preferences.[12] We set out to test the following hypothesis regarding rational individual behavior of the committee members.

Hypothesis 3 (Individual-level behavior)**.** *We expect that (i) chairs of* $t = chair$ *play A as often, (ii) chairs unconditional on type choose A more (B less) often, (iii) member 2s and 3s choose C as often in the* $p_A = 1$ *than in* $p_A = 3/4$ *treatment.*

Figure 2 reveals a small difference in voting behavior for chairs of type t = chair when comparing behavior across treatments $p_A = 3/4$ and $p_A = 1$, conditional on the type. Specifically, the share of votes for A drops from 76% in $p_A = 3/4$ to 70% in $p_A = 1$, but the difference is not significant (WMW test, p-value = 0.5980). It is not surprising that the introduction

[12] In a related a approach, Frydman and Nunnari (2024) investigate the effect of cognitive noise and strategic uncertainty by varying a common payoff parameter in a two player coordination game.

of incomplete information about the chair's preference does not affect voting behavior of the chair type $t = chair$'s significantly because chairs know their type. Analyzing chair behavior with chair types pooled, the bottom two bar charts in figure 2 show that chairs play A less frequently and B more frequently in the $p_A = 1$ treatment compared to the $p_A = 3/4$ treatment. This pattern arises mechanically from the different distribution of chair types in the treatments. However, only the difference in the frequency of playing B is significant, as indicated by the WMW tests in table 6.

Turning to regular members, figure 3 shows the voting shares for alternative C which is the theoretical prediction for regular members in both treatments. Certainty about the chair's type in treatment $p_A = 1$ appears to increase regular members' votes for the equilibrium action C. Specifically, $m2$ members increase the vote share of C from 80% to 92%, and $m3$ from 56% to 73%, when moving from treatment $p_A = 3/4$ to $p_A = 1$. Although these difference in voting shares are not significant as shown in table 7, the impact of the introduction of preference uncertainty is clearly more pronounced for regular members than for the chairs.

In the following, we will show how the above changes in individual voting by regular members induced by preference uncertainty are large enough to significantly change the realized outcomes and strategy profiles. The corresponding hypothesis 4 summarizes the equilibrium strategies and outcomes predicted by the theory.

Hypothesis 4 (Equilibrium strategy profiles and outcomes)**.** *We expect that (i) the equilibrium strategy* $((A, B), C, C)$ *is played equally often, (ii) alternative* C *wins equally often, and (iii) type* $t = chair$ *chairs receive the same average payoff in* $p_A = 1$ *as in* $p_A = 3/4$*.*

Observed strategy profiles and outcomes indicate a sizable increase in equilibrium play when removing the preference uncertainty. The elections consistent with the equilibrium strategy profile $((A, B), C, C)$ rises from 35% in $p_A = 3/4$ to 49% in $p_A = 1$ (p-value = 0.15) and alternatives C's winning share increases from 45% to 71% (p-value = 0.08). At the same time, both A's and B's winning share decreases from 32% to 21% (p-value = 0.15) and from 23% to 8% (p-value < 0.01) between $p_A = 3/4$ and $p_A = 1$, respectively. A notable consequence of these differences in outcomes is that chairs receive a significantly higher payoff in the $p_A = 3/4$ treatment compared to the $p_A = 1$ treatment (p-value $= 0.011$), with average payoffs of €10.53 and €7.50 respectively. In summary, chairs clearly benefit from the introduction of outcome uncertainty, as shown by the higher payoffs in $p_A = 3/4$ than in $p_A = 1$. This result, which contradicts our theory-based hypothesis 4(iii) derived under full rationality, is primarily due to more non-equilibrium behavior by regular members in the uncertainty treatment $p_A = 3/4$, which results in more favorable outcomes for the chair. Our finding that strategic, respectively cognitive, uncertainty increases the share of decision errors, leading to less optimal decisions, is consistent with recent evidence by Enke and Graeber (2023) in single decision-making tasks and Frydman and Nunnari (2024) in a strategic setting. Next, we take a closer look at the drivers of non-equilibrium committee member behavior and their implications.

5.3 Deviations from equilibrium and individual characteristics

Although voting behavior in the experiment is broadly consistent with our theoretical model, some members deviate more from equilibrium actions than others. In this section, we show how chair types systematically benefit from the combination of non-equilibrium behavior of regular members and their power to break ties. In addition, we provide evidence on the ability of members to iteratively eliminate weakly dominated strategies, a key assumption in section 3. We also show how individual-level characteristics, elicited in separate tasks, influence the propensity to vote as predicted by equilibrium, providing evidence for the internal validity of our modeling approach.

Weak domination and non-equilibrium behavior. Voting behavior of different committee members' is largely consistent with the assumptions of eliminating weakly dominated strategies (WDS) and its iterated elimination (IEWDS) used to derive the Bayesian Nash Equilibria (BNE) in section 3. Table 9 summarizes voting behavior according to consistency with (IE)WDS by committee member and treatment. It can be seen that members $m2$ and $m3$ violate WDS at extremely low rates, at most 2-3%, which we attribute to decision errors.

For chairs the picture looks different and WDS violations are substantial for some chair types. Chair types $t = m3$ with at most 7% exhibit a comparable low level of WDS violations as regular members, but chairs of type $t =$ chair violate WDS at much higher rates with 24%

Table 9: Share of equilibrium, IEWDS, and weakly dominant actions.

treatment	player	A	B	C	survives first iteration
$p_A = 1/4$	Member 2	0.12	0.02	0.85	0.98
$p_A = 1/4$	Member 3	0.01	0.82	0.17	0.99
$p_A = 1/4$	Chair(t = chair)	0.82	0.17	0.00	0.83
$p_A = 1/4$	Chair(t = m3)	0.00	0.92	0.07	0.93
$p_A = 3/4$	Member 2	0.17	0.03	0.80	0.97
$p_A = 3/4$	Member 3	0.03	0.41	0.56	0.97
$p_A = 3/4$	Chair(t = chair)	0.76	0.23	0.01	0.76
$p_A = 3/4$	Chair(t = m3)	0.01	0.94	0.05	0.94
$p_A = 1$	Member 2	0.06	0.02	0.92	0.98
$p_A = 1$	Member 3	0.03	0.23	0.73	0.97
$p_A = 1$	Chair(t = chair)	0.70	0.25	0.05	0.70

Notes: Green entries denote the share of equilibrium actions, red entries denote the share of weakly dominated actions, and the last column shows the share of actions surviving the first iteration of IEWDS.

in treatment $p_A = 3/4$ and 30% in treatment $p_A = 1$. These rates of WDS violations can be explained by the fact that, the chair's voting choice is irrelevant when both regular members coordinate on C as predicted by the equilibrium strategy profile in treatment $p_A = 3/4$ and $p_A = 1$. Indeed, the bulk of WDS violations are stemming from chairs of type t = chair playing B. By voting B, chairs of type $t = chair$ attempt to coordinate with $m3$ on outcome B, which would yield a higher payoff for this chair type and for $m3$ that they get from C winning. However, whenever $m3$ votes for B, the best response for a chair of type t = chair is to play A and secure the win via tie-breaking power for their most-preferred alternative. This reasoning is reflected in our data, and figure 6, with the corresponding action profile (B, C, B) and (A, C, B) occurring in 13% and 20% of the elections in treatment $p_A = 3/4$.

The share of equilibrium actions played varies between 56% and 92% across members and treatments (highlighted in green), see Table 9. This large variation is best understood by considering differences in strategic incentives and the complexity of reasoning required for each member to vote optimally. For member $m3$, the decision-making process is particularly challenging as $m3$'s equilibrium strategy depends on the distribution of chair types. This complexity is best illustrated in the $p_A = 3/4$ treatment, where $m3$ chooses the equilibrium action C in only 56% of elections, significantly lower than the 73% and 82% in the $p_A = 1$ and $p_A = 1/4$ treatments. The reason for this pronounced difference is that strategic sophistication is most demanding for $m3$ in the incomplete information treatment $p_A = 3/4$, and thus some $m3$ members try, mostly unsuccessfully, to coordinate with the chair on a better outcome using off-equilibrium actions.[13]

Why chairs benefit from non-equilibrium behavior. We have shown that deviations from equilibrium behavior occur mostly for member $m3$, which is directly linked to the uncertainty $m3$ faces about the chair's preference type. The resulting non-equilibrium behavior of $m3$ benefits the chair by increasing the likelihood of a tie, thereby allowing the chair's vote to determine the committee outcome.

Comparing behavior in the $p_A = 1$ and $p_A = 3/4$ treatments allows to identify the effects of the strategic uncertainty caused by incomplete information about the chair's type, as both predict the same equilibrium strategy. As mentioned before, the most frequent deviation from equilibrium occurs when member $m3$ votes B instead of the equilibrium vote for C in the $p_A = 3/4$ treatment in an attempt to coordinate with the chair. This non-equilibrium behavior is well reflected in figure 6, which shows that (A, C, B) is the second most frequently observed action profile in the $p_A = 3/4$ treatment. In (A, C, B), a tie occurs and the outcome is A due to the tie-breaking authority of the chair. This non-equilibrium behavior of $m3$ is a key driver for the significantly higher share of tie-breaks in the $p_A = 3/4$ treatment (28%) compared to the $p_A = 1$ treatment (17%) (two-sided WMW, p-value = 0.036). In the

[13] As mentioned already above, some $m3$ members may find it appealing to deviate from equilibrium by switching from C to B in an attempt to coordinate with the chair on B, which would be a more profitable outcome for both players, regardless of whether the chair is of type t = chair or $t = m3$. However, this reasoning fails to anticipate that a rational chair of type t = chair would then have an incentive to play A instead of B, creating a tie (A, C, B) and ultimately implementing A via tie-breaking power. Figures C6 and C7 in the appendix confirms that the 56% of C votes cast by $m3$ in $p_A = 3/4$ elections stems from a high fraction of $m3$ participants 'mixing' between B and C.

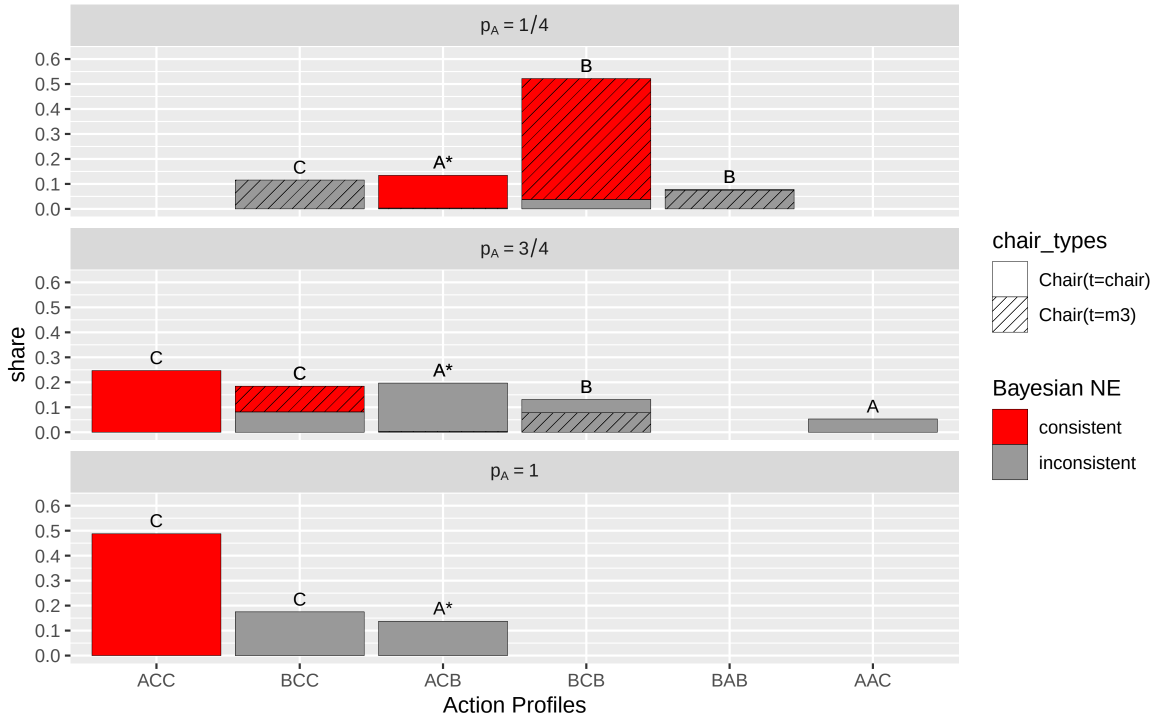

Figure 6: Most frequent action profiles by treatment.

Notes: Share of action profiles predicted by the BNE conditional on chair type in each treatment. Conditional on the chair type implies that, for instance, action profile (B, C, C) is only counted as consistent with BNE if the chair in the election is $t = m3$, otherwise it is counted as inconsistent. Only action profiles occurring with a share of more than 5% of all elections in a treatment are displayed inf figure for readability.

$p_A = 1/4$ treatment, tie-breaks occur in only 15% of cases which is a significantly lower rate than in $p_A = 3/4$ (two-sided WMW, p-value $= 0.021$). Unlike in the other treatments, where tie-breaks result from non-equilibrium behavior, tie-breaks in the $p_A = 1/4$ treatment are part of equilibrium play.[14] Overall, our results indicate that the complexity introduced by preference uncertainty in the $p_A = 3/4$ treatment increases the frequency of tie-breaks, which in turn lead to favorable outcome for the chair compared to the $p_A = 1$ treatment. Thus, we conclude that chairs systematically benefit from the increase in non-equilibrium behavior induced by incomplete information of regular members.

[14] Note that the tie-breaking shares in the statistical tests differ slightly from those in Figure 6, as the figure omits action profiles observed in fewer than 5% of cases for readability, some of which include tie-breaking scenarios.

Table 10: Impact of sophistication measures on equilibrium voting behavior.

DV: equilibrium behavior	(1)	(2)	(3)	(4)
Guess in beauty contest	0.0005	0.0003	0.001	0.0004
	(0.001)	(0.001)	(0.001)	(0.001)
Expected payoff maximizer	−0.010	−0.014	−0.001	0.001
	(0.022)	(0.022)	(0.018)	(0.019)
IEWDS	0.039**	0.043**	0.049***	0.050***
	(0.016)	(0.017)	(0.015)	(0.018)
Decision time	−0.524***	−0.475***	−0.493***	−0.477***
	(0.177)	(0.178)	(0.133)	(0.152)
Round number	0.007**	0.007**	0.007**	0.007**
	(0.003)	(0.003)	(0.003)	(0.003)
Demographic controls	No	Yes	Yes	Yes
Role type x treatment FE	No	No	Yes	Yes
Matching group FE	No	No	No	Yes
Observations	1,920	1,920	1,920	1,920
R^2	0.052	0.081	0.130	0.155

Notes: Linear regressions with standard errors clustered on the matching-group level in parentheses. The dependent variable is an indicator whether the equilibrium voting action was played. We use data from rounds 9-16 and treatments $p_A = 1/4$ and $p_A = 3/4$. Significance levels: ** $p < 0.05$; *** $p < 0.01$

Individual-level determinants of voting behavior. Finally, we provide evidence how individual-level characteristics related to participants' strategic sophistication determine the propensity to vote according to equilibrium predictions. In particular, recall that we imposed a number of rationality and sophistication assumptions, such as the use of IEWDS (Moulin, 1979) to derive the equilibria in section 3. Investigating the relationship between these measures and optimal behavior provides empirical evidence on the validity of some of our theoretically imposed key assumptions regarding sophisticated voting behavior.

Hypothesis 5 (Sophistication and equilibrium behavior)**.** *Participants to vote for the equilibrium alternative more often if they are able to (i) perform IEWDS, (ii) maximize expected payoffs, and (iii) exhibit a high degree of strategic reasoning.*

To test this hypothesis, we regress the propensity to vote for the equilibrium alternative on the different individual-level measures using committee members' voting decisions. Our individual-level measures were elicited independently after the voting game in various incentivized decisions tasks (see section 4 for details). Table 10 summarizes the results of linear model on the propensity to play the equilibrium action; model (2) adds demographic controls, model (3) player-role dummies and treatment fixed effects and their interactions, and model (4) matching-group fixed effects to the baseline model.

In line with our sophisticated voting assumption in section 3, we find that a participant's ability to perform IEWDS has a significant positive effect on the probability of voting optimally in all model specifications as hypothesized ($p-$values of 0.018, 0.020, 0.017 and 0.020 in models (1) to (4) in table 10). That is, a participant's performance on the IEWDS task is strongly associated with a higher propensity to play the equilibrium action as hypothesized. Other variables related to strategic sophistication, such as a participant's *guess* in the beauty contest and the ability to maximize expected payoffs, see *payoff maximizer*, have no significant effect on the propensity to play optimally in any of the models. In addition to IEWDS, decision time and round number are also significant. Probit models in table D4 in the appendix confirm the results of the above linear models.

We also fit the full model from column 4 in table D5 in the appendix separately for each member and type. Table D5 shows that the IEWDS coefficient is significant only for member 2 and member 3. This is intuitive, since we already argued above that finding the optimal strategy is much more complex for regular players than for chairs. In particular, member 3's optimal voting and naively choosing the most preferred outcome varies across all treatments.[15] Overall, the regression results on equilibrium voting behavior on these individual characteristics confirm the rationality assumptions underlying our voting model.

6 Conclusion

In this paper, we investigated the interplay between voting power imbalances and informational asymmetries in a small committee. Our theoretical model generalizes the classical

[15] Note that the insignificant coefficient for the chair member regressions could also be due to the smaller number of observations for the $t = chair$ and $t = m3$ types, but the point estimate for IEWDS is also smaller in size in these regressions.

chair's paradox by introducing incomplete information about the chair's preferences, considering the effect of different chair types on equilibrium outcomes. We found that information asymmetries significantly influence how tie-breaking power can be leveraged by the chair, giving rise to novel equilibria that do not exist in complete information settings. Our results contribute to the broader literature on committee voting by extending previous models to account for incomplete information and general payoff structures. Our experimental results largely corroborated the theoretical equilibrium predictions. We also found that the response of regular members to changes in the distribution of chair preferences was somewhat attenuated. This attenuated response is consistent with an increase in non-equilibrium behavior of regular members, which in turn systematically favors the committee chair, also because it increases the number of ties produced in the committee. Future research could explore the welfare implications of these power imbalances, particularly in mixed-interest environments where both private and common interests are at play. In addition, studying the effects of more complex voting systems and institutional designs on power and information asymmetries could further shed light on the conditions under which power imbalances are either mitigated or exacerbated. Modeling the incentives to strategically conceal one's preferences as a chair may be another avenue for future research, as our results suggest that this may be a profitable rationale for those in positions of power.

References

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin, "Coalition formation in non-democracies," *The Review of Economic Studies*, 2008, *75* (4), 987–1009.

Ali, S Nageeb, B Douglas Bernheim, and Xiaochen Fan, "Predictability and Power in Legislative Bargaining," *The Review of Economic Studies*, 02 2018, *86* (2), 500–525.

Alós-Ferrer, Carlos, "The Trembling Chairman Paradox," *Games and Economic Behavior*, 2022, *131*, 51–56.

Ansolabehere, Stephen, James M. Snyder, Aaron B. Strauss, and Michael M. Ting, "Voting Weights and Formateur Advantages in the Formation of Coalition Governments," *American Journal of Political Science*, 2005, *49* (3), 550–563.

Apesteguia, Jose, Miguel A. Ballester, and Yusufcan Masatlioglu, "A Foundation for Strategic Agenda Voting," *Games and Economic Behavior*, 2014, *87*, 91–99.

Azrieli, Yaron, Christopher P Chambers, and Paul J Healy, "Incentives in experiments: A theoretical analysis," *Journal of Political Economy*, 2018, *126* (4), 1472–1503.

Bernheim, B. Douglas, Antonio Rangel, and Luis Rayo, "The Power of the Last Word in Legislative Policy Making," *Econometrica*, 2006, *74* (5), 1161–1190.

Berry, Christopher R. and Anthony Fowler, "Cardinals or Clerics? Congressional Committees and the Distribution of Pork," *American Journal of Political Science*, 2015, *60* (3), 692–708.

_ and _ , "Congressional Committees, Legislative Influence, and the Hegemony of Chairs," *Journal of Publich Economics*, 2018, *158*, 1–11.

Blinder, Alan S., *The quiet revolution: central banking goes modern*, Yale University Press, 2004.

_ , "Monetary policy by committee: Why and how?," *European Journal of Political Economy*, 2007, *23* (1), 106–123.

_ and John Morgan, "Are Two Heads Better than One? Monetary Policy by Committee," *Journal of Money, Credit, and Banking*, 2005, *37* (5), 798–811.

_ and _ , "Do Monetary Policy Committees Need Leaders? A Report on an Experiment," *American Economic Review, Papers & Proceedings*, 2008, *98* (2), 224–229.

Bouton, Laurent, Aniol Llorente-Saguer, and Frédéric Malherbe, "Get rid of unanimity rule: The superiority of majority rules with veto power," *Journal of Political Economy*, 2018, *126* (1), 107–149.

Brams, Steven J., Dan S. Felsenthal, and Zeev Maoz, "New Chairman Paradoxes," in Andreas Diekmann and Peter Mitter, eds., *Paradoxical Effects of Social Behavior: Essays in Honor of Anatol Rapoport*, Physica-Verlag, 1986.

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

Dhillon, Amrita and Ben Lockwood, "When are plurality rule voting games dominance-solvable?," *Games and Economic Behavior*, 2004, *46* (1), 55–75.

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

_ , _ , Ryan Oprea, and Jeffrey Yang, "Behavioral attenuation," 2024. working paper.

Farquharson, Robin, *Theory of Voting*, Yale University Press, 1969.

Fréchette, 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.

Frydman, Cary and Salvatore Nunnari, "Coordination with Cognitive Noise," 2024. working paper.

Fudenberg, Drew and Jean Tirole, *Game theory*, MIT press, 1991.

Granić, Dura-Georg, "The Problem of the Divided Majority: Preference Aggregation and Uncertainty," *Journal of Economic Behavior & Organization*, 2017, *133*, 21–38.

Granic, Georg D. and Alexander K. Wagner, "Where power resides in committees," *The Leadership Quarterly*, 2021, *32* (4), 101285.

Hughes, Niall, Zia-Ul-Hassan Khan, and Friederike Mengel, "Diversity in Committees," 2023. working paper.

Ke, Changxia, Florian Morath, Anthony Newell, and Lionel Page, "Too big to prevail: The paradox of power in coalition formation," *Games and Economic Behavior*, 2022, *135*, 394–410.

Kohlberg, Elon and Jean-Francois Mertens, "On the strategic stability of equilibria," *Econometrica*, 1986, *54* (5), 1003–1037.

Maaser, Nicola and Thomas Stratmann, "Costly voting in weighted committees: The case of moral costs," *European Economic Review*, 2024, *162*, 104652.

_ **, Fabian Paetzel, and Stefan Traub**, "Power illusion in coalitional bargaining: An experimental analysis," *Games and Economic Behavior*, 2019, *117*, 433–450.

Moulin, Hervé, "Dominance Solvable Voting Schemes," *Econometrica*, 1979, *47* (6), 1137–1151.

Nagel, Rosemarie, "Unravelling in Guessing Games: An Experimental Study," *American Economic Review*, 1995, *85* (5), 1313–1326.

Snyder, James M., Michael M. Ting, and Stephen Ansolabehere, "Legislative Bargaining under Weighted Voting," *American Economic Review*, 2005, *95* (4), 981–1004.

Online Appendix
for
"How asymmetric information shapes decisions
in committees with unequal voting rights"

A Proofs

We want to establish that voting for the least preferred alternative is weakly dominated by voting for the most preferred alternative for regular members fixing a chair type. Since we assume consistent beliefs and can use Harsanyi transformation, the payoff of a regular members is just a convex combination of the type space distribution over the chair types' and regular members' actions. In this sense, if voting for the least preferred alternative is weakly dominated fixing a specific type for the chair, it is also weakly dominated for the whole game. This follows immediately from our assertions.

Result 3. *Voting for the least preferred alternative is weakly dominated by voting for the most preferred alternative for member 2.*

Proof. To simplify notation call $m2$'s most preferred alternative $BEST$, the second most preferred alternative $MIDDLE$, and the least preferred alternative $LEAST$. We have to consider three cases:

case (i) $s_{chair}(t) = BEST$, i.e. the chair type votes for $m2$'s most preferred alternative. If $s_{m2} = BEST$, this alternative is implemented as $BEST$ receives a majority of 2 votes. If $s_{m2} = LEAST$ two sub-cases emerge. Either $m2$ and $m3$ disagree, i.e. $s_{m3} \neq LEAST$. In this case $BEST$ wins by Observation 2. Or $m2$ and $m3$ agree, i.e. $s_{m2} = s_{m3} = LEAST$. In this case $LEAST$ is implemented by *Observation 1*. $s_{m2} = BEST$ thus induces either the same or a strictly higher payoff than $s_{m2} = LEAST$.

case (ii) $s_{chair}(t) = Second$. If $s_{m2} = BEST$ either $s_{m3} = BEST$ and $BEST$ is implemented by *Observation 1*, or $s_{m3} \neq BEST$ in which case $MIDDLE$ is implemented by *Observation 2*. $s_{m2} = BEST$ at least implements $MIDDLE$ and sometimes $BEST$. If $s_{m2} = LEAST$ either $s_{m3} = Bottom$ and $LEAST$ is implemented by *Observation 1*, or $s_{m3} \neq LEAST$ and $MIDDLE$ is implemented by *Observation 2*. $s_{m2} = LEAST$ at best implements $MIDDLE$ and sometimes $LEAST$. $s_{m2} = BEST$ thus induces either the same or a strictly higher payoff than $s_{m2} = LEAST$.

case (iii) $s_{chair}(t) = LEAST$. If $s_{m2} = LEAST$, this alternative is implemented as $LEAST$ receives at least a majority of 2 votes. If $s_{m2} = BEST$ two sub-cases emerge. Either $s_{m3} = BEST$ and $BEST$ is implemented by *Observation 1*. Or $s_{m3} \neq BEST$ and $LEAST$ is implemented by *Observation 2*. $s_{m2} = BEST$ thus induces either the same or a strictly higher payoff than $s_{m2} = LEAST$.

□

Since the game is symmetric with respect to the regular members $m2$ and $m3$ (i.e. the voting rule is anonymous with regard to these two players), we immediately obtain:

Result 4. *Voting for the least preferred alternative is weakly dominated by voting for the most preferred alternative for member 3.*

B Additional analysis: learning behavior

In our main analysis we only considered rounds 9-16 in the voting game.The regressions in table 10 show that the round number coefficient is significant and positive, indicating an increasing propensity of playing optimal over time in the second half of rounds. Considering voting behavior in all rounds (including rounds 1-8), we provide further evidence on participants' learning to vote as described by equilibrium. The figures in appendix C plot the convergence to equilibrium actions over time for each player role and show that learning under incomplete information can be more difficult. In particular, consider the behavior of member 3s in treatment $p_A = 3/4$, for which we have found a large share of votes deviating from equilibrium actions in the main text. Figure C5 plots member 3s' actions over all 16 rounds of the committee game. The pattern of behavior indicates a strong convergence to the equilibrium actions in all treatments. In the incomplete-information treatment $p_A = 3/4$, member 3s vote for the non-equilibrium action B and the equilibrium action C almost equally often in the first half of the experiment and then learn to play the equilibrium action C more frequently. Comparing the equilibrium convergence between treatment $p_A = 3/4$ and the complete information treatment $p_A = 1$, where member 3 faces no uncertainty about the chair's type in the committee, results are in line with our interpretation that treatment $p_A = 3/4$ is cognitively more challenging for member 3s because they learn to play equilibrium actions in $p_A = 3/4$ at a lower rate as compared to $p_A = 1$.

C Additional figures

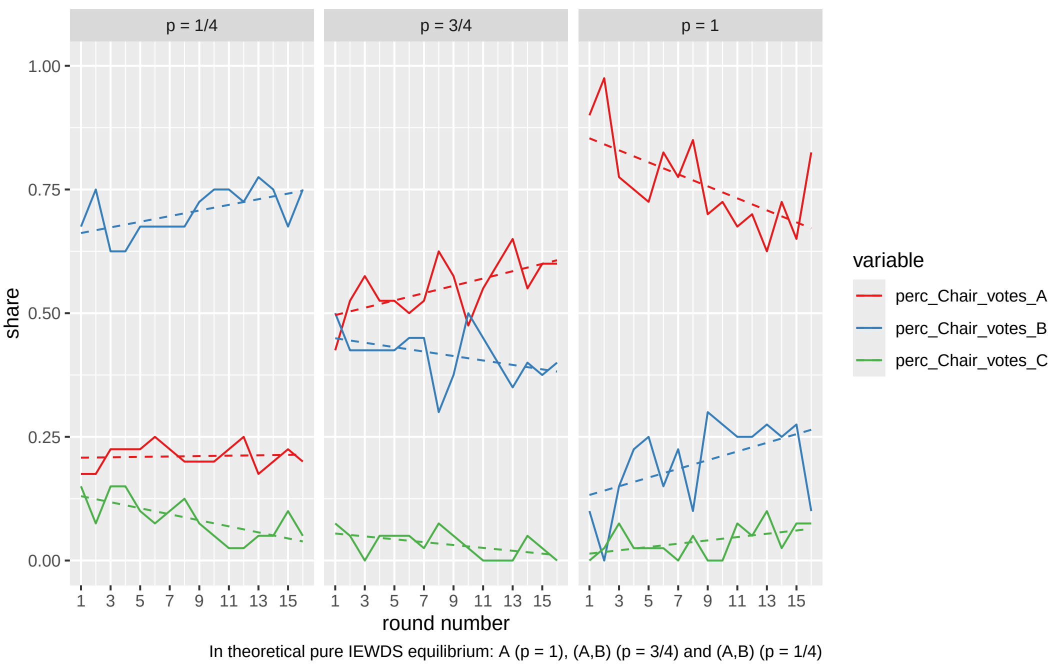

Figure C1: Equilibrium convergence of chair (types pooled) voting over all rounds.

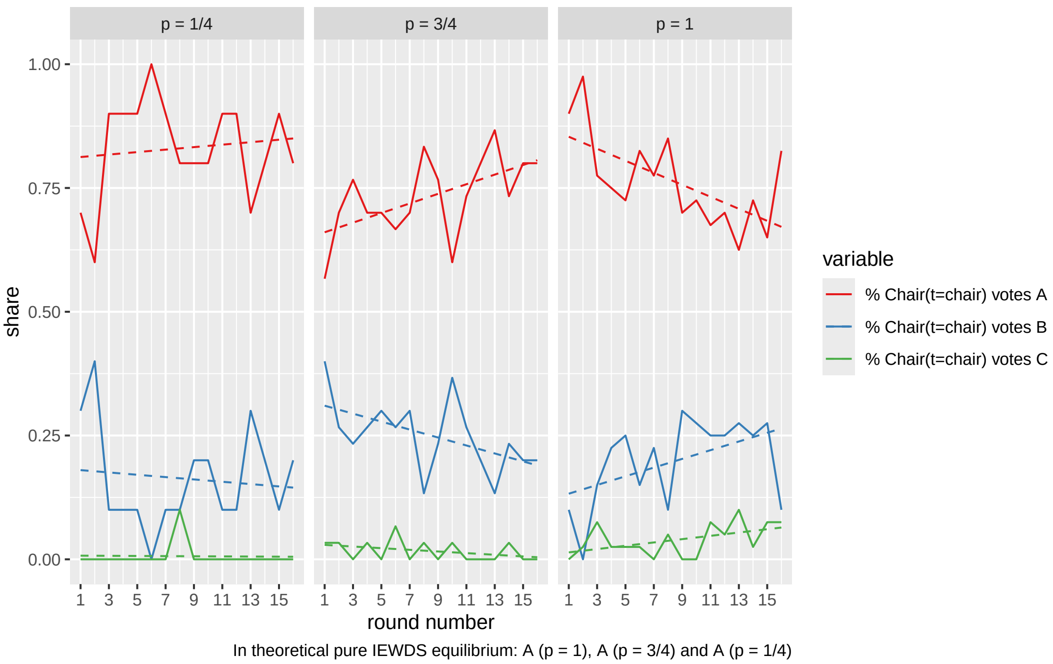

Figure C2: Equilibrium convergence of chair (t=chair) voting over all rounds.

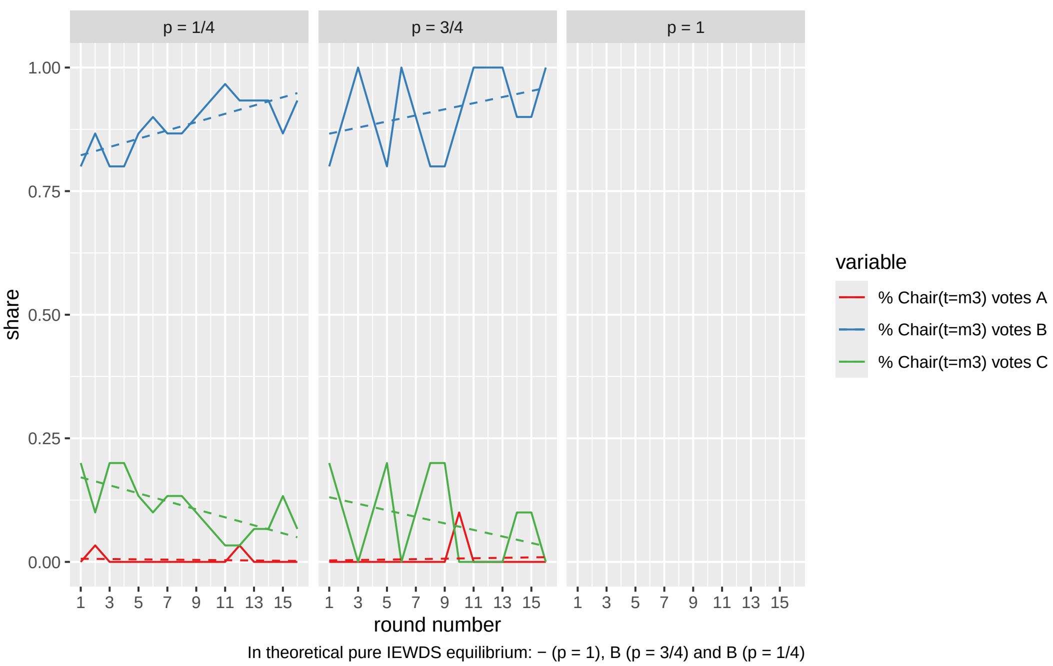

Figure C3: Equilibrium convergence of chair (t=m3) voting over all rounds.

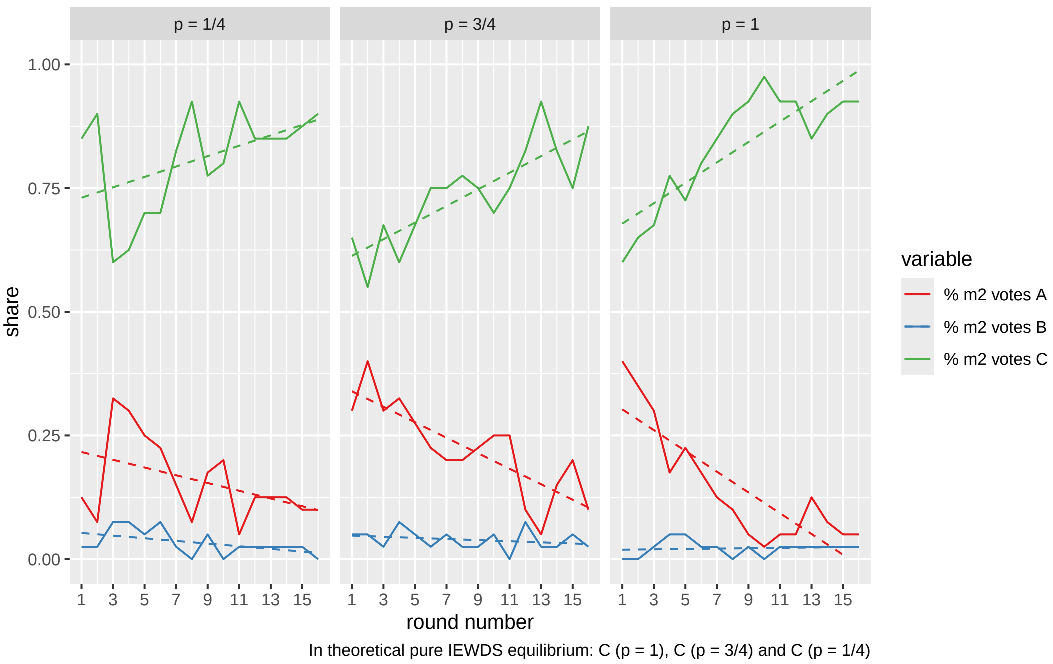

Figure C4: Equilibrium convergence of member 2 voting over all rounds.

Figure C5: Equilibrium convergence of member 3 voting over all rounds.

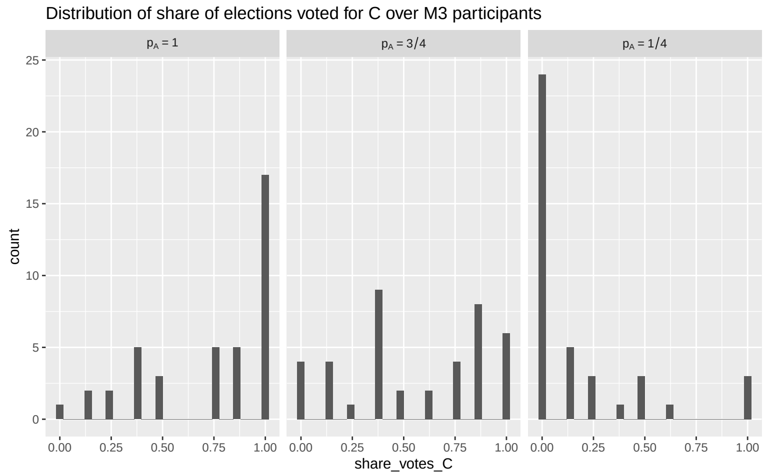

Figure C6: Distribution of share of elections voted for C by member 3.

Notes: For each participant in role of member 3, the share of voting for C is calculated. Only the second half of rounds is used.

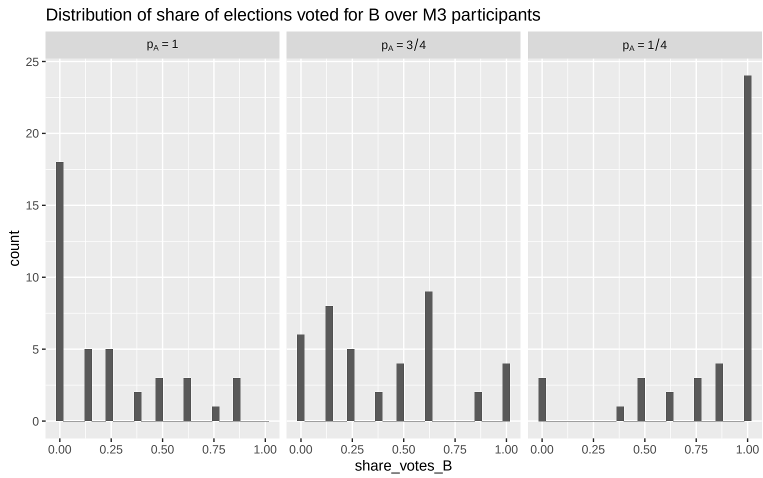

Figure C7: Distribution of share of elections voted for B by member 3.

Notes: For each participant in role of member 3, the share of voting for B is calculated. Only the second half of rounds is used.

D Additional tables

Table D1: Voting differences in uncertainty treatments ($p = 1/4$ vs $p = 3/4$).

	Chair		M2	M3	
DV:	votes A	votes B	votes C	votes C	votes B
	(1)	(2)	(3)	(4)	(5)
Treatment p = 3/4	2.166***	−1.907***	−1.348**	0.987**	−1.336***
	(0.515)	(0.502)	(0.581)	(0.386)	(0.387)
Demographic controls	Yes	Yes	Yes	Yes	Yes
Sophistication Tasks	Yes	Yes	Yes	Yes	Yes
Matching group FE	Yes	Yes	Yes	Yes	Yes
Observations	640	640	640	640	640

Notes: Probit regressions with standard errors clustered at the matching-group level in parentheses. The dependent variable is an indicator variable of the respective equilibrium action being played. We use data from rounds 9-16 and treatments $p_A = 1/4$ and $p_A = 3/4$; the reference category of the treatment dummy is treatment $p_A = 1/4$. Significance levels: ** $p < 0.05$; *** $p < 0.01$

Table D2: Voting differences in uncertainty treatments ($p = 1/4$ vs $p = 3/4$) by chair type.

	Chair (t = chair)		Chair (t = m3)	
DV:	votes A	votes B	votes A	votes B
	(1)	(2)	(3)	(4)
Treatment p = 3/4	0.374***	−0.372***	0.010	−0.046
	(0.119)	(0.116)	(0.025)	(0.159)
Demographic controls	Yes	Yes	Yes	Yes
Sophistication Tasks	Yes	Yes	Yes	Yes
Matching group FE	Yes	Yes	Yes	Yes
Observations	320	320	320	320
R^2	0.296	0.271	0.107	0.324

Notes: Linear regressions with standard errors clustered at the matching-group level in parentheses. The dependent variable is an indicator variable of the action being played. We use data from rounds 9-16 and treatments $p_A = 1/4$ and $p_A = 3/4$; the reference category of the treatment dummy is treatment $p_A = 3/4$. Significance levels: ** $p < 0.05$; *** $p < 0.01$

Table D3: Voting differences in certainty and uncertainty treatments ($p = 1$ vs $p = 3/4$).

	Chair		M2	M3	
DV:	votes A	votes B	votes C	votes C	votes B
	(1)	(2)	(3)	(4)	(5)
Treatment p = 3/4	0.580***	−0.547***	−0.042**	−0.547***	0.465***
	(0.051)	(0.050)	(0.020)	(0.111)	(0.111)
Demographic controls	Yes	Yes	Yes	Yes	Yes
Sophistication Tasks	Yes	Yes	Yes	Yes	Yes
Matching group FE	Yes	Yes	Yes	Yes	Yes
Observations	640	640	640	640	640
R^2	0.266	0.220	0.173	0.273	0.245

Notes: Linear regressions with standard errors clustered at the matching-group level in parentheses. The dependent variable is an indicator variable of the equilibrium action. We use data from rounds 9-16 and treatments $p_A = 1$ and $p_A = 3/4$; the reference category of the treatment dummy is treatment $p_A = 1$.. Significance levels: ** $p < 0.05$; *** $p < 0.01$

Table D4: Impact of sophistication measures on equilibrium behavior.

	DV: equilibrium behavior			
	(1)	(2)	(3)	(4)
Guess in beauty contest	0.001	0.001	0.002	0.001
	(0.003)	(0.003)	(0.003)	(0.004)
Expected payoff maximizer	−0.033	−0.049	−0.011	−0.003
	(0.079)	(0.081)	(0.067)	(0.074)
IEWDS	0.143**	0.164**	0.195***	0.201***
	(0.067)	(0.077)	(0.067)	(0.078)
Decision time	−1.742***	−1.641***	−1.778***	−1.788***
	(0.515)	(0.542)	(0.500)	(0.533)
Round	0.024**	0.026**	0.028**	0.029**
	(0.011)	(0.011)	(0.011)	(0.012)
Demographic controls	No	Yes	Yes	Yes
Role type x treatment FE	No	No	Yes	Yes
Matching group FE	No	No	No	Yes
Observations	1,920	1,920	1,920	1,920

Notes: Probit regressions with standard errors clustered at the matching-group level in parentheses. The dependent variable is an indicator variable of the equilibrium action. We use data from rounds 9-16 and treatments $p_A = 1/4$ and $p_A = 3/4$. Significance levels: ** $p < 0.05$; *** $p < 0.01$

.

Table D5: Impact of sophistication measures on equilibrium behavior by role.

DV: equilibrium behavior	(1)	(2)	(3)	(4)
	M2	M3	Chair (t=chair)	Chair (t=m3)
Guess in beauty contest	−0.002	0.0005	0.009*	0.002
	(0.001)	(0.002)	(0.005)	(0.002)
Expexted payoff opt.	−0.009	−0.010	0.096	−0.031
	(0.031)	(0.033)	(0.079)	(0.039)
IEWDS	0.073**	0.086***	0.037	0.032
	(0.029)	(0.029)	(0.040)	(0.057)
Decision time	−0.687***	−0.243	0.228	0.008
	(0.244)	(0.337)	(0.593)	(0.203)
Round number	0.012**	0.003	0.012	0.001
	(0.005)	(0.008)	(0.008)	(0.005)
Demographic controls	Yes	Yes	Yes	Yes
Treat FE	Yes	Yes	Yes	Yes
Matching group FE	Yes	Yes	Yes	Yes
Observations	640	640	320	320
R^2	0.202	0.320	0.296	0.324

Notes: Linear regressions with standard errors clustered on the matching-group level in parentheses. The dependent variable is an indicator whether the equilibrium voting action was played. We use data from rounds 9-16 and treatments $p_A = 1/4$ and $p_A = 3/4$. Significance levels: ** $p < 0.05$; *** $p < 0.01$

E Details experimental design

E.1 Comprehension questions

The following quiz questions are for treatment $p = 3/4$. For quiz questions referring to payoff matrices, the voting screen as in figure 1 was shown to the participants.

1. Assume that Member Y votes for B, Member Z votes for A and the Chair votes for C. Which alternative wins the election?
   - Alternative C is the winner of the election because the chair has tie-breaking power and has voted for C.
   - A coin toss decides which alternative is the winner.
   - No alternative wins the election because there is a tie.
2. Before the first election, each participant is assigned to a Member role (Chair, Member Y or Member Z) and then keeps the role in every election of the experiment.
   - True
   - False
3. You will be randomly matched with two other members...
   - ...only once in beginning of the experiment.
   - ...in every election.

4. Assume your payoffs are 15 if C wins, 10 if B wins, 5 if A wins. You will receive 15 Euro...
    - ...whenever C wins the election.
    - ...only if C wins and you voted for C.
5. Why do Member Y and Member Z see both the LEFT and the RIGHT payoff table in an election?
    - Because the Chair is indecisive of what to choose.
    - Because Member Y and Member Z only know the probability that the Chair (drawn from the 4 Chairs in the experiment) has the payoffs in the LEFT table or the payoffs in the RIGHT table.
6. Is it correct that if A wins, then the Chair receives 15 for sure?
    - Yes
    - Only if the Chair has the payoffs shown in the LEFT table
    - Only if the Chair has the payoffs shown in the RIGHT table
7. Is it correct that the Chair is more likely to have the payoffs in the LEFT table than in the RIGHT table?
    - Yes
    - No
8. Is it correct that the Chair knows for sure what her/his payoff is if alternative C wins?
    - Yes
    - No
9. Assume you are Member Y. Do you know for sure what the Chair's payoff is if alternative C wins?
    - Yes
    - No
10. Assume you are Member Z. Is it correct that your payoff is 10 if alternative C wins?.
    - Yes
    - No
    - Only if the Chair has the payoffs in the RIGHT table.

E.2 Onscreen instructions

The following instructions are for treatments $p = 3/4$. Instructions are adapted accordingly for treatments $p = 1$ and $p = 1/4$.

General Information

Welcome! The experiment in which you are about to participate is part of a research project on decision-making.

Please remain silent during the experiment, do not speak to other participants, and switch off all communication devices NOW. If you have any question or need assistance of any kind, please raise your hand, and an experimenter will come to you and answer your question in private.

You will be asked to make various decisions and you can earn money for your decisions. How much you can earn in a task will be announced before you make your decisions.

All participants and their decisions will remain anonymous to other participants during the experiment. You will neither learn the identity of the participants you will interact with, nor will others find out about your identity.

At the end of the experiment, **you will be privately and anonymously paid in cash** the amount you earned in the experiment.

The experiment consists of several independent parts. We now describe Part 1. The other parts will be explained later.

Next

Part 1: Instructions

Overview of Part I

- You will do this part of the experiment in a group of 12 participants.
- 16 times in a row, the computer will randomly select **groups of 3** participants from these 12 participants to do an election.
- In each election, there is **one Chair, one Voter Y, and one Voter Z.**
- **Each participant has to vote for one of three alternatives called A, B, and C.**
- **The alternative with the most votes wins the election and determines the payoff of the voters.**
- At the end of the experiment, one of the 16 elections will be randomly selected, and your earnings from this part will be your payoff from that election.

Next, we will explain the rules and details of Part 1. Click to proceed.

Next Page

Part 1: Instructions

Voter roles

There are 12 participants in your group. **Before the first election, the computer assigns randomly**

- **4** participants to the role of **a Chair**,
- **4** participants to the role of **a Voter Y**, and
- **4** participants to the role of **a Voter Z**.

Each participant keeps her/his assigned voter role for the entire experiment.

Back Next Page

Part 1: Instructions

Possible payoffs of each voter role

The payoff of a Voter Role depends on which of the alternatives A, B, or C wins the election. The payoffs for each voter role are as follows:

- Each of the 4 **Voter Y** ( ) participants receives 10€ if alternative A wins, 5€ if B wins, and 15€ if C wins the election.
- Each of the 4 **Voter Z** ( ) participants receives 5€ if alternative A wins, 15€ if B wins, and 10€ if C wins the election.

In our experiment in general, there can be **two types of Chairs**, green ( ) and blue ( ), **which differ in their payoffs**. For your session today, there will be **3 Chairs of type green and 1 Chairs of type blue**.

- **3 of the 4 Chairs** are of type green. They receive 15€ if alternative A wins, 10€ if B wins, and 5€ if C wins the election.
- **1 of the 4 Chairs** are of type blue. They receive 5€ if alternative A wins, 15€ if B wins, and 10€ if C wins the election.

The possible payoffs remain the same in all 16 elections for each voter.

From the 12 voters in the experiment, 3 are type-green Chairs, 1 are type-blue Chairs, 4 are Voters Y in grey and 4 are Voters Z in black, as shown below.

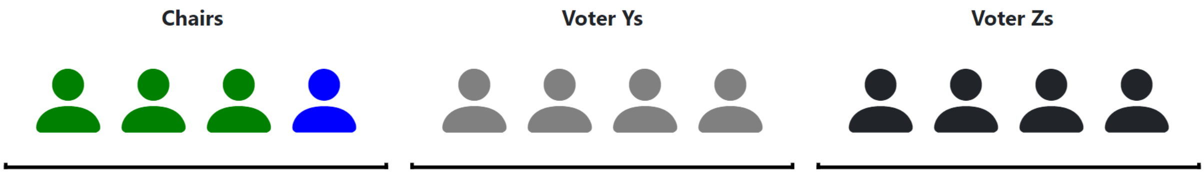

Back Next Page

Part 1: Instructions

Voting and election winners

In each election:

- There is always one Chair, one Voter Y, and one Voter Z.
- Each voter **has to vote for one of the three alternatives A, B, or C.**
- When voting, you do not know how the other two voters vote and the other two do not know your choice when they are voting.

Which alternative wins the election?

- **The alternative (A, B, or C) that receives the most votes wins the election.**
- **In case of a tie between alternatives, the alternative the Chair has voted for in the election wins.**
- At the end of an election, **each voter receives feedback** about the alternative she/he voted for, how many votes each alternative received, which alternative won the election, and her/his payoffs from the current election.

Examples: Assume you are Voter Y and your payoff is such that you receive 10 if A wins, 5 if B wins and 15 if C wins.

- Example 1: You and Voter Z vote for A and the Chair Voter votes for B. A wins the election so you get 10.
- Example 2: You vote for A, Voter Z votes for B and the Chair votes for C. There is a tie (every alternative gets one vote), so the Chair's vote decides. The Chair has voted for C, so C is the winner. You receive 15 (your payoff if C wins).

Back Next Page

Part 1: Instructions

Matching voters in an election

You will participate in a total of 16 independent elections. **Before the start of each election**,

- The computer **randomly selects 3 voters to participate in the election** from the 12 voters in the experiment.
- This is done by randomly drawing **1 of the 4 Chairs, 1 of the 4 Voter Ys, and 1 of the 4 Voter Zs in the experiment.** That is, the voters you will interact with will change from election to election!

The example below shows how 3 voters, one of each voter role, are randomly selected to interact in an election. Recall that different colors represent different payoffs.

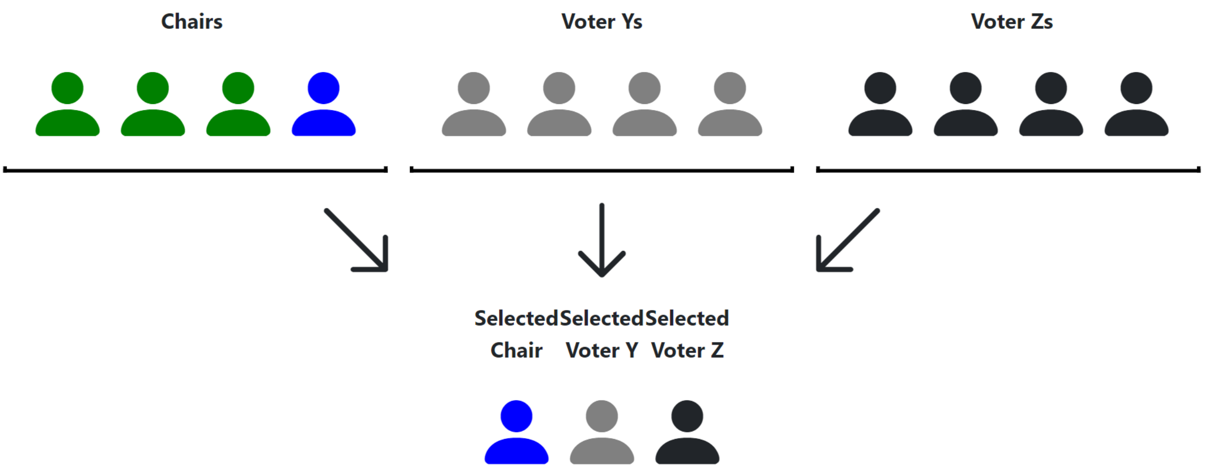

You see that the randomly selected Chair can be either of the green or of the blue type, depending on the random draw.

Back Next Page

Part 1: Instructions

Information of Voter Y and of Voter Z

In each election, 3 voters are randomly selected to participate.

Every voter knows the payoff of Voter Y and Voter Z when voting. **However, Voter Y and Voter Z do not know of which type the Chair is.** That is, Voter Y and Voter Z do not know for sure if the Chair in their election is of the green-type or the blue-type when casting their vote.

On the decision screen, a payoff table similar to the one below is shown to Voter Y and Voter Z. The payoff table summarizes the possible payoffs of Voter Y, Voter Z, and of each type of the Chair in the election.

Voter Y and Voter Z only know that the payoffs are...

...with a chance of 3 out of 4:

	if A wins	if B wins	if C wins
Chair	15	10	5
Voter Y	10	5	15
Voter Z	5	15	10

...with a chance of 1 out of 4:

	if A wins	if B wins	if C wins
Chair	5	15	10
Voter Y	10	5	15
Voter Z	5	15	10

Note that the payoffs for Voter Y and Voter Z are the same in the LEFT and in the RIGHT table. For example, Voter Z receives 5€ if A wins, and Voter Y receives 10€ if A wins.

Only the Chair's payoff is different in the LEFT and RIGHT table as each table shows the payoffs of the chair for each of her/his possible payoff types green and blue.

Voter Y and Voter Z do not know for sure whether the Chair in the current election has the payoff profile shown in the LEFT or RIGHT table. But, they know that

- the **Chair** will have the **payoffs from the LEFT table with probability 75%**. Because there are 4 Chairs in the experiment, 3 type-green ( ) and 1 type-blue ( ), the **chance is 3 out of 4** that the randomly selected Chair is of the green-type with the payoffs given in the LEFT table.
- the **Chair** will have the **payoffs in the RIGHT table with probability 25%**. That is, the **chance is 1 out of 4** that the randomly selected Chair is a blue-type with the payoffs given in the RIGHT table.

Back Next Page

Part 1: Instructions

Information of the Chair

- The Chair knows the payoffs of Voter Y and Voter Z and also knows her/his own payoff type.
- The payoff table summarizes the possible payoffs of Voter Y, Voter Z, and of the Chair in the election.
- In the example payoff table below, the Chair is a green-type. For this reason, the LEFT table is highlighted. The RIGHT table is not applicable and therefore greyed out.

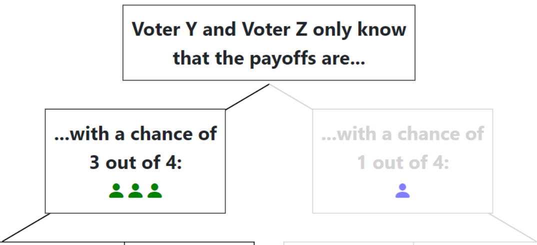

	if A wins	if B wins	if C wins
Chair	15	10	5
Voter Y	10	5	15
Voter Z	5	15	10

	if A wins	if B wins	if C wins
Chair	5	15	10
Voter Y	10	5	15
Voter Z	5	15	10

Back Next Page

Part 1: Instructions

Summary of the Instructions for Part 1

- **3 participants** are randomly selected to vote in an election. In each election, there is **one Chair, one Voter Y, and one Voter Z.**
- **Each participant has to vote for one of three alternatives called A, B, and C.**
- **The alternative with the most votes wins the election.** If there is tie, such that A, B, and C receive exactly one vote, the Chair has tie-breaking power. That is, the alternative the Chair has voted for in the election wins.
- **The payoff you (and other participants) receive depend on which alternative wins the election.** Participants in different Voter Roles may receive different payoffs for a given winning alternative.
- After the winning alternative is determined, the election ends. Then, 3 participants of the 12 participants are **randomly selected** to cast votes in the next election. This procedure is repeated until you have participated in 16 elections. Then Part 1 of the experiment is over.
- Note that the participants you interact with may change from election to election. Their identity is not revealed. You will only learn the Voter Role of the other participants in an election, but not which randomly selected participant is assuming a voter role in an election.
- After all 16 elections have been played, **one of the elections you participated in will be selected randomly**.
- **The Euro payoff you received in this randomly selected election will be paid out in cash** to you at the end of the experiment (together with additional earnings you may receive from other parts in the experiment).

Once you have read the instructions, you are asked to answer a short quiz. These questions ensure that everyone understands the instructions. Then, Part I of the experiment will start. You will receive instructions for Part 2 of the experiment after Part 1 has ended.

Back Start the quiz!