Games and Economic Behavior 99 (2016) 275–294

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Gambler's fallacy and imperfect best response in legislative bargaining ☆

Salvatore Nunnari [a,*], Jan Zapal [b,c]

[a] Bocconi and IGIER, Via Roentgen 1, Milano, 20136, Italy
[b] CERGE-EI, a joint workplace of Charles University in Prague and the Economics Institute of the Czech Academy of Sciences, Politickych veznu 7, Prague, 11121, Czech Republic
[c] IAE-CSIC and Barcelona GSE, Campus UAB, Bellaterra, Barcelona, 08193, Spain

ARTICLE INFO

Article history:
Received 10 January 2014
Available online 28 June 2016

JEL classification:
D72
C92
C70

Keywords:
Legislative bargaining
Experiments
Quantal response
Gambler's fallacy

ABSTRACT

We investigate the implications of imperfect best response—in combination with different assumptions about correct (QRE) or incorrect beliefs (Quantal-Gambler's Fallacy or QGF)—in the alternating offer multilateral bargaining game. We prove that a QRE of this game exists and characterize the unique solution to the proposer's problem—that is, the proposal observed most frequently in a QRE. We structurally estimate this model on data from laboratory experiments, and show that it explains behavior better than the model with perfect best response: receivers vote probabilistically; proposers allocate resources mostly within a minimum winning coalition of legislators but do not fully exploit their bargaining power. Incorporating history-dependent beliefs about the future distribution of proposal power into the QRE model (QGF) leads to an even better match with the data, as this model implies slightly lower shares to the proposer, maintaining similar or higher frequencies of minimum winning coalitions and similar voting behavior.

1. Introduction

The Baron and Ferejohn (1989) model of alternating offer bargaining is the workhorse model of bargaining in legislatures and committees, and arguably one of the most influential models in political economy and formal political theory. In its simplest version, a committee of n legislators bargains over the allocation of a fixed budget. At the beginning of the game, one legislator is randomly assigned the ability to set the agenda; once a proposed agreement is on the floor, every legislator simultaneously votes on whether to accept or reject this proposal. If a simple majority of legislators supports the proposal, the bargaining process ends, and the resources are allocated according to the suggested agreement. If, instead, the committee rejects the proposal, the process repeats, with the random selection of a new agenda setter and the same bargaining protocol. In all stationary subgame perfect equilibria of this bargaining game, an agreement is reached immediately, with

☆ We are grateful to Guillaume Frechétte and Maria Montero for supplying the original data from the experiments in Frechette et al. (2003, 2005a, 2005b), and Drouvelis et al. (2010). We also thank Marina Agranov, Colin Camerer, Guillaume Frechétte, Steven Lehrer, Luis Miller, Massimo Morelli, Maria Montero, Thomas R. Palfrey, Matthew Rabin, Jean-Laurent Rosenthal and audiences at the California Institute of Technology, Barcelona Jocs, Universitat de Barcelona, ASSET 2013 Bilbao Annual Meeting, the Political Economy: Theory Meets Experiment Workshop at ETH Zurich and the Center for Experimental Social Science at New York University for comments and feedback. Zapal acknowledges financial support by the Czech Science Foundation, grant 14-27902P. The suggestions by two anonymous referees are also gratefully acknowledged.

* Corresponding author.
E-mail addresses: salvatore.nunnari@unibocconi.it (S. Nunnari), j.zapal@cerge-ei.cz (J. Zapal).

http://dx.doi.org/10.1016/j.geb.2016.06.008
0899-8256/

the agenda setter getting the lion's share of the resources and sharing the budget only with enough other legislators to guarantee passage on the floor (Eraslan, 2002).

A series of recent contributions tested the predictions of this basic model, and its many variations, using controlled laboratory experiments, which have a distinct advantage over field data in studying a complex institutional environment with many confounding factors (see, for instance, Frechette et al., 2003, 2005a, 2005b; Diermeier and Morton, 2005; Diermeier and Gailmard, 2006; Kagel et al., 2010; Drouvelis et al., 2010; Miller and Vanberg, 2013, 2015).

These studies produced several interesting regularities. Participants to legislative bargaining experiments learn over time to propose a positive allocation only to a *minimum winning coalition* of legislators, that is, only to the minimum number of agents required for passage of the proposal; they vote selfishly and do not always accept an advantageous offer: the probability they accept a proposal is increasing in the share offered to themselves but independent of the distribution of resources to other agents; and they do not fully exploit the first-mover advantage conferred by the power to set the agenda: while still proposing to themselves the largest share, they give to their coalition partners more than predicted by the theory.

In this paper, we propose a novel theoretical explanation that combines imperfect best response and mistaken beliefs on probabilistic events, and show that it can account for all the stylized facts from the laboratory.

The first contribution of this paper is to present two extensions of the benchmark model that have an independent theoretical interest. The first modification assumes that agents are subject to the *gambler's fallacy*. The gambler's fallacy—sometimes referred to as the belief in the law of small numbers (Tversky and Kahneman, 1971)—is the (erroneous) belief that small samples generated by a random variable should resemble large samples generated by the same random variable. In its most extreme form, agents who are prey of this cognitive fallacy and observe a series of fair coin flips, believe that the ratio of heads and tails in any sample of draws should be equal to one half—a fact that is true only in the limit, for a large sample.

There is a broad array of evidence, from laboratory experiments as well as field data, pointing to the relevance of this cognitive bias in the way people perceive stochastic processes (Bar-Hillel and Wagenaar, 1991; Rapoport and Budescu, 1992, 1997; Budescu and Rapoport, 1994; Clotfelter and Cook, 1993; Terrell, 1994). While the gambler's fallacy has successfully been incorporated in individual decision-making studies (Grether, 1980, 1992; Camerer, 1987; Rabin, 2002; Rabin and Vayanos, 2010), applications to game theoretical frameworks, where agents suffering from the same bias interact strategically, have not been explored.[1]

We introduce the gambler's fallacy in alternating offer bargaining by assuming that legislators (erroneously) believe that the stochastic process determining the identity of the agenda setter is history dependent. In particular, if a legislator sets the agenda in the current period and her proposal is rejected, the common belief in the committee is that the same legislator has a lower chance of making a second proposal. This, in turn, means that the other legislators believe they are more likely to be recognized in the next bargaining round, which increases their continuation value.[2] These beliefs do not affect the number of legislators who receive a positive allocation, but affect the allocation of resources among members of a minimum winning coalition: a higher continuation value of non-proposers implies that the proposer has to allocate a larger share of the pie to any coalition partner in order for her proposal to pass.

This suggests that incorporating history-dependent beliefs can explain a prominent feature of the experimental data. At the same time, this model does not capture other common patterns in the behavior of experimental committees. In particular, a model of alternating-offer bargaining with gambler's fallacy cannot account for proposals giving a positive allocation to more or fewer members than those required for passage[3] and cannot explain probabilistic voting strategies: legislators are still predicted to accept for sure any proposal above the continuation value of the game and to reject anything else.

The second modification we propose relaxes the assumption that agents perfectly best-respond to the strategies of other agents. The evaluation of theoretical models using data from field or laboratory experiments requires the incorporation of random errors representing omitted elements or noisy behavior. Probabilistic choice models have long been used in the analysis of individual decisions (Luce, 1959; Harless and Camerer, 1994; Hey and Orme, 1994; Camerer and Ho, 1994). Quantal Response Equilibrium (QRE, McKelvey and Palfrey, 1995, 1998) is the analogous way to model games with noisy players: it assumes that agents make mistakes in choosing their actions and that they make more expensive mistakes with smaller probability, but it also imposes the requirement that beliefs match equilibrium choice probabilities. A large body of

[1] Although this is not highlighted by the authors, Walker and Wooders (2001) find some evidence consistent with the gambler's fallacy in the play of zero-sum games by professional tennis players.

[2] Another well documented cognitive bias in perceiving stochastic processes is the *hot hand fallacy*, the (erroneous) belief that random sequences will exhibit excessive persistence rather than reversals (Rabin and Vayanos, 2010). In the alternating offer bargaining game, this means believing that the current proposer has a higher chance of making a second proposal. Ex-ante, it is not clear what bias is more relevant for predicting behavior in alternating-offer bargaining, as both biases are well documented in other setups, even for the same task and the same individual (see, for instance, Ayton and Fischer, 2004 and Sun and Wang, 2010). The structural estimation of a more general model—with agents subject to a bias in either direction—would point towards the gambler's fallacy. This is because the hot hand fallacy changes the equilibrium allocations in the opposite direction than what we observe in the laboratory: relative to the benchmark model with correct beliefs, the proposer's share is larger and agents in the winning coalition receive less.

[3] The incidence of these proposals decreases over time but remains positive, see Section 5 for details.

recent work has shown that this equilibrium concept can rationalize behavior in a variety of experimental settings where Nash equilibrium cannot.[4]

We show that, in a multilateral bargaining game, imperfect best response makes receivers sensitive to the amount proposed in a continuous way, reduces the share to the proposer and increases the incidence of non-minimum winning coalitions. Since receivers do not use a cutoff strategy but vote probabilistically, it is no longer optimal for the proposer to make a coalition partner indifferent between accepting and rejecting her proposal. When agents' best responses are, even minimally, imperfect, a proposer who does so faces a large probability of rejection. At the same time, the current proposer has a strong incentive to maximize the chance her proposal passes, because of the first-mover advantage she is sure to enjoy today but that she enjoys only stochastically tomorrow, if the game continues. As a result, the proposer is better off giving up some resources in order to increase the probability of passage. On the other hand, the optimal allocation gives a positive share only to a minimum winning coalition of players or fewer so any resource allocated to members in excess of a bare majority is attributable to mistakes.

A second contribution of the paper is to structurally estimate the model with imperfect best response, as well as the model which combines imperfect best response with erroneous beliefs, on the experimental data. We show that the predictions obtained from QRE fit remarkably well the proposing and voting decisions by laboratory subjects, significantly improving the explanatory power of the benchmark model, even when accounting for the additional degree of freedom. However, imperfect best response alone cannot account for all patterns in the data: reducing the coefficient of rationality, or the sensitivity to expected utilities, reduces the share to the proposer, but it also increases the resources allocated to legislators in excess of a minimum winning coalition. This means that there is a trade-off in how well the QRE can explain different patterns in the data: quantal response cannot account for both minimum winning coalitions and a more egalitarian split between coalition partners. Adding the gambler's fallacy—and keeping the responsiveness to payoffs in the QRE fixed—we predict an even lower share to the proposer but a similar or higher incidence of minimum winning coalitions. As a consequence, a model with both quantal response and gambler's fallacy (QGF) generates the best fit with the proposing and voting patters in the data.

2. Related literature

This paper contributes to the theoretical and experimental literature on alternating offer bargaining. The first wave of experimental works on structured bargaining in committees (Binmore et al., 1985; Ochs and Roth, 1989) studied the bilateral protocols introduced by Rubinstein (1982) and Stahl (1972). More related to this paper, a series of more recent contributions tested the predictions of the multilateral bargaining model in Baron and Ferejohn (1989).[5]

The mixed performance of the predictions from standard non-cooperative game theory has led researchers to propose alternative frameworks for the analysis of this data. Regarding the bilateral bargaining experiments, as summarized in Camerer (2003), "the basic finding from these studies is that offers and counteroffers are usually somewhere between equal split of the money being bargained over and the offer predicted by subgame-perfect game theory. Many of these results are approximately equilibrium outcomes when social preferences are properly specified" (page 469).

As discussed in the Introduction, subjects in multilateral bargaining experiments similarly fail to exploit their proposal power. However, behavior in multilateral bargaining experiments cannot be easily explained by other-regarding preferences. First, voting behavior in these games suggests subjects do not care about the share to the poorest member (Güth and van Damme, 1998; Bolton and Ockenfels, 1998) and proposers share resources only with the members of a minimum winning coalition. Second, theories with other regarding preferences actually predict a larger share to the proposer (Montero, 2007).[6] Other avenues have proven more successful. Frechette (2009) shows that belief-based learning models can help explain the experimental data from Frechette et al. (2003). In this learning model, subjects hold beliefs on the probability of each proposal being accepted or rejected and they update them on the basis of their experience in the experiment.[7] Diermeier (2011) suggests that moral considerations and framing play an important role in experimental behavior and argues that other factors—unmodeled in the benchmark framework, but relevant to sustained human interaction—could induce a bargainer to offer potential partners more than the absolute minimum amounts that they would accept. Agranov and Tergiman (2014) find that, when allowing for unrestricted cheap-talk communication, proposers extract rents in line with the theoretical predictions of the benchmark model. They conclude that deviations observed in the absence of communication could be due to uncertainty surrounding the amount a coalition member is willing to accept.

[4] These settings include, but are not limited to, two person zero sum games (McKelvey and Palfrey, 1995), signaling games (McKelvey and Palfrey, 1998), the centipede game (McKelvey and Palfrey, 1998), all-pay auctions (Anderson et al., 1998), first-price auctions (Goeree et al., 2002), bilateral bargaining (Goeree and Holt, 2000; Yi, 2005), coordination games (Anderson et al., 2001), and the "traveler's dilemma" (Goeree and Holt, 2001). We discuss the relationship with studies of imperfect best response in bilateral bargaining games in Section 2.

[5] See cites on page 276.

[6] The same is true of theories with risk aversion (Harrington, 1990) or present bias (note available from the authors).

[7] Similarly to our QRE model, Frechette (2009) assumes stochastic choice, with a parameter measuring the responsiveness of choice probabilities to expected utility. The main difference is that, in a QRE, beliefs on a proposal's probability of being accepted, as well as the continuation value of each proposal, are determined in equilibrium, assuming other agents imperfectly respond with the same responsiveness parameter. Moreover, our Gambler's Fallacy model impacts on continuation value and, indirectly, on beliefs over acceptance. Beliefs over proposal power process do not matter in Frechette (2009) because continuation values are assumed to be the same for all proposal types.

To the best of our knowledge, this is the first paper introducing imperfect best response in a multilateral bargaining environment. Yi (2005) and Goeree and Holt (2000) study QRE in simple one-period or two-period bilateral bargaining models. The differences with the multilateral bargaining game we study are technical, as well as substantial. From the technical point of view, both games have a finite horizon; this, together with a finite action space, means that existence of a QRE is guaranteed by results in McKelvey and Palfrey (1998). None of the results in McKelvey and Palfrey (1998) implies the existence of a QRE in the infinite horizon game we study. Moreover, while we provide a formal characterization of the QRE for any level of rationality, Yi (2005) only presents results in the limit, as irrationality vanishes. From the substantial point of view, both games involve only two players who decide by unanimity. Our multilateral bargaining game with a non-unanimous decision rule is a more complex strategic environment and generates richer predictions that are relevant for real world legislatures and committees and cannot be studied with a simpler bilateral game: the size of the proposer's coalition, how resources are allocated to different members of the committee, the extent of delay, and the impact of committee size.

3. Theory

3.1. Alternating-offer bargaining: benchmark model

Our benchmark model is the classic Baron and Ferejohn (1989) model of noncooperative legislative bargaining with alternating offers. In the original model, a group of $N = \{1, 2, \dots, n\}$ agents (where $n \geq 3$ and odd) bargains over a division of a pie of size 1. Bargaining proceeds over a potentially infinite number of rounds. In each round, one player $i \in N$ is randomly selected to be the agenda setter and proposes a division of the pie $x \in X$. The set of permissible allocations is the $n-1$ dimensional simplex, $X = \{x \in \mathbb{R}^n_+ \,|\, \sum_{i=1}^n x_i \leq 1\}$, where x_i denotes the share allocated to player i.

Once a proposal $x \in X$ is on the table, all agents vote either *yes* or *no*.[8] If the number of *yes* votes is equal to or larger than $\frac{n+1}{2}$, x passes, the game ends and agents receive their share as determined by x. The utility player i derives from agreement x, reached in round $s \geq 1$, is equal to $u_i(x) = \delta^{s-1} x_i$ where $\delta \in [0, 1]$ is a common discount factor. If x does not pass the voting stage, no agent receives a flow of utility, and the game moves to a further round of bargaining.

We focus on symmetric stationary subgame perfect Nash equilibria (SSPE) in which agents use the same strategies and proposers treat their potential coalition partners symmetrically (see Appendix A for a formal definition).

Proposition 1 *(Baron and Ferejohn, 1989). In the benchmark model, there exists a unique SSPE characterized by:*

1. *The agent recognized to be the proposer offers $1 - \frac{n-1}{2}\hat{x}$ to herself, $\hat{x}$ to $\frac{n-1}{2}$ randomly selected other agents and 0 to the remaining $\frac{n-1}{2}$ agents, where $\hat{x} = \frac{\delta}{n}$.*
2. *Agents vote in favor of a proposal $x \in X$ if and only if it satisfies $x_i \geq \hat{x}$.*
3. *The difference between the proposer's share and the share to agents with non-zero allocations (the proposer's advantage) equals $1 - \frac{n+1}{2}\hat{x} > 0$ and decreases with $\hat{x}$.*

Example 1. With $n = 3$ and $\delta = 1$, the unique equilibrium proposal is of the form $\{2/3, 1/3, 0\}$.

3.2. Gambler's fallacy

The belief in the law of small numbers, often called gambler's fallacy, is the erroneous belief that the distribution of a small random sample should closely resemble the distribution in the underlying population. In its most extreme form, after observing two consecutive heads in a series of four fair coin flips, an agent who is victim of the fallacy believes the next two coin flips will surely result in tails. In this example, the gambler's fallacy is driven by the belief that the ratio of heads and tails in any sample of flips should be equal to one half.

We model the gambler's fallacy as in Rabin (2002), adapting his framework to the alternating offer bargaining game. In particular, we assume that agents believe the identity of the proposer in every round of the game is determined by the draw from an urn. This urn contains ϕn balls where an integer $\phi \geq 1$ measures the extent of the gambler's fallacy. Initially, the urn contains an equal number of balls for each player, that is, it contains ϕ balls labeled 'i' for each $i \in N$. This means that, at the beginning of the game, agents share the common belief that they will be selected as the proposer for the first round with probability $1/n$. However, the draws from this urn are made without replacement, generating believes about the

[8] Voting can be either simultaneous as in Banks and Duggan (2000), who use the 'stage-undominated' voting strategies of Baron and Kalai (1993), or sequential as in Montero (2007). The resulting voting behavior is the same: $i \in N$ votes for x if and only if it gives him weakly more than the expected utility from continuing to another round of the game. That is, i votes as if he is pivotal.

future recognition probabilities that are history dependent.[9] For modeling convenience, we assume (as in Rabin, 2002) that the urn is refilled every two rounds.[10]

The game starts in round 1 with the urn in its initial state. This implies a probability of recognition equal to $\frac{\phi}{\phi n}$ for all the agents in round 1 and all subsequent odd rounds. In round 2 and all even rounds, the probability of recognition of the agents who have not been recognized in the previous round is equal to $\frac{\phi}{\phi n-1}$ and the probability of recognition of the former proposer is equal to $\frac{\phi-1}{\phi n-1}$. In other words, in even rounds, the former proposer believes he is less likely to propose than everyone else, while all the other agents believe they have a higher probability to set the agenda in this stage than the former proposer. The discrepancy between these even-round beliefs and the time-invariant beliefs of the benchmark model (an even chance of proposing for all agents in all rounds) is parameterized by ϕ. The highest degree of gambler's fallacy corresponds to $\phi = 1$, in which case the current proposer is certain not to set the agenda if we proceed to a further round of bargaining. On the other hand, the game with fallacious agents converges to the benchmark model as ϕ goes to infinity.[11]

Adding gambler's fallacy to the benchmark model makes even and odd rounds strategically different. However, all even rounds and all odd rounds are strategically equivalent. As a consequence, we extend the definition of strategies' stationarity to require that agents use the same strategy in all even rounds and the same strategy in all odd rounds (but strategies can potentially differ between even and odd rounds).[12]

Proposition 2. *In the alternating offer bargaining game where agents are prey to the gambler's fallacy, there exists a unique SSPE characterized by the following strategies:*

1. *The proposer offers $1 - \frac{n-1}{2}\hat{x}$ to herself, $\hat{x}$ to $\frac{n-1}{2}$ randomly selected other agents and 0 to the remaining $\frac{n-1}{2}$ agents, where*

$$\hat{x} = \begin{cases} \frac{\delta}{n}\frac{2n\phi-\delta}{2n\phi-2} & \text{in odd rounds} \\ \frac{\delta}{n} & \text{in even rounds} \end{cases} \tag{1}$$

2. *Agents vote for the proposal $x \in X$ if and only if it satisfies $x_i \geq \hat{x}$.*
3. *The difference between the proposer's share and the share to agents with non-zero allocations (the proposer's advantage) equals $1 - \frac{n+1}{2}\hat{x} > 0$ and decreases with $\hat{x}$. Relative to the benchmark model, the proposer's advantage is smaller in odd rounds and identical in even rounds.*
4. *Increasing the incidence of gambler's fallacy, by decreasing ϕ, strictly increases the amount $\hat{x}$ offered by the proposer to coalition partners in odd rounds.*

Proof. See Appendix A. □

Example 2. With $n = 3$ and $\delta = 1$, the equilibrium proposal will approach $\{8/12, 4/12, 0\}$ as $\phi \to \infty$ with the proposer receiving strictly less for any finite ϕ. For $\phi = 1$, the equilibrium allocation is $\{7/12, 5/12, 0\}$.

3.3. Quantal response equilibrium

Quantal response posits that agents are imperfect best responders: they choose better actions with higher probabilities than worse actions but do not choose best responses with probability one. Quantal Response Equilibrium (QRE) is an internally consistent equilibrium model: agents assume their opponents are imperfect best responders, and their response functions are based on the equilibrium probability distribution of the opponents' actions. We focus here on the logistic agent QRE as defined by McKelvey and Palfrey (1998), retaining the assumption of symmetry and stationarity.

More explicitly, as McKelvey and Palfrey (1998), we assume that the move of each player $i \in N$ at different information sets is done by different 'agents' of i. In the context of legislative bargaining, this implies that the proposing player i cannot control i in the subsequent voting stage. Moreover, in analogy with the concept of stage-undominated voting strategies, we assume that each player i votes as if pivotal.[13]

---

[9] Back to our example where a coin is flipped four draws, the urn initially contains two 'heads' and two 'tails'. After we draw two 'heads', only 'tails' will be drawn for sure in the remaining rounds.

[10] This assumption simplifies the analysis but is not necessary. What is important is that the act of recognition implies lower probability of recognition in the future.

[11] The limit case, when $\phi = \infty$, is equivalent to a situation where the draws from the urn are made with replacement and, thus, the probability each player is recognized to propose is time invariant.

[12] Departing from Rabin (2002), we can also use a stationary version of the Gambler's Fallacy model. Assume players believe that Nature excludes the former proposer with probability ϕ (in which case each of the other players is selected with probability $1/(n-1)$) and draws a proposer with the regular protocol with probability $1-\phi$ (in which case each player is selected with probability $1/n$). As we show in the Online Appendix, the equilibrium of this alternative model is very similar to the equilibrium of the model we discuss: the proposer allocates a positive and symmetric share only to the members of a minimum winning coalition and keeps the remainder for herself. The share to coalition partners is $\hat{x} = \frac{\delta}{n}\frac{2n-2+2\phi}{2n-2+\delta\phi}$ which equals $\frac{\delta}{n}$ when $\phi = 0$ (as in the benchmark model) and is clearly increasing in ϕ.

[13] This is a common assumption in the legislative bargaining literature, see Baron and Kalai (1993).

The concept of QRE in McKelvey and Palfrey (1998) is defined for discrete action spaces. For this reason, we change the space of permissible allocations in the bargaining game, X, to its discrete analog, $X' = \{x \in \mathbb{R}^n_+ | \sum_i^n x_i = 1 \wedge \text{mod}(x_i, d) = 0\ \forall i \in N\}$ for some $d \in (0,1)$ satisfying $\frac{1}{d} \in \mathbb{N}_{>0}$ in order for $X' \neq \emptyset$. In this specification, X' consists of allocations where each player receives an integer multiple of some d.[14] In the logit version of QRE, a legislator uses a behavioral strategy where the log probability of choosing each available action is proportional to its expected payoff, where the proportionality factor, λ, can be interpreted as a responsiveness (or rationality) parameter.

Using stationarity, we can think of each round of legislative bargaining as a one-shot game with disagreement payoffs given by the discounted continuation values, δv_i, for each $i \in N$. Symmetry further implies $v_i = v_j \equiv v$ for all $i, j \in N$. Given v, the probability i votes in favor of proposal $x \in X'$ is given by:

$$p^{\lambda}_{v,i}(x) = \frac{\exp \lambda\, x_i}{\exp \lambda\, x_i + \exp \lambda\, \delta\, v} \tag{2}$$

where $\lambda \geq 0$ measures the precision in i's best-response. With $\lambda = 0$, we have $p^{\lambda}_{v,i}(x) = \frac{1}{2}$, that is, a legislator randomizes in the voting stage, regardless of the proposal. On the other hand, with $\lambda \to \infty$, $p^{\lambda}_{v,i}(x) \to 1$ for $x_i > \delta v$ and $p^{\lambda}_{v,i}(x) \to 0$ for $x_i < \delta v$. That is, as in the benchmark model, legislators vote in favor of the proposal if it gives them strictly more than their discounted continuation value and reject the proposal when it gives them strictly less than their discounted continuation value. Note that any pair of voting strategies $(p^{\lambda}_{v,i}, p^{\lambda}_{v,j})_{i,j \in N}$ is symmetric.

Given the voting behavior described by $p^{\lambda}_{v,i}$, the probability $x \in X'$ is accepted, $p^{\lambda}_{v}(x)$, can be easily calculated as the probability that at least $\frac{n+1}{2}$ agents vote for x. Given this, proposer i proposes $x \in X'$ with probability

$$r^{\lambda}_{v,i}(x) = \frac{\exp \lambda\,(x_i\, p^{\lambda}_{v}(x) + \delta\, v\,(1 - p^{\lambda}_{v}(x)))}{\sum_{z \in X'} \exp \lambda\,(z_i\, p^{\lambda}_{v}(z) + \delta\, v\,(1 - p^{\lambda}_{v}(z)))} \tag{3}$$

which is symmetric across players and treats potential coalition partners equally.

Denote by $\sigma^{\lambda}(v) = (r^{\lambda}_{v,1}, \ldots, r^{\lambda}_{v,n}, p^{\lambda}_{v,1}, \ldots, p^{\lambda}_{v,n})$ the profile of proposal and voting strategies. Given v, λ and $\sigma^{\lambda}(v)$, we can calculate the expected utility of playing according to the profile of strategies $\sigma^{\lambda}(v)$, with disagreement utility v, which we denote with v'. With a slight abuse of notation, we denote this mapping by $v' = \sigma^{\lambda}(v)$. A v constitutes a QRE if and only if it is a fixed point of σ^{λ}.

Proposition 3. *There exists a QRE of the alternating offer bargaining game. In any QRE, the associated continuation value v^* satisfies $v^* \leq \frac{1}{n}$. If $\delta = 1$, there exists a unique QRE with $v^* = \frac{1}{n}$, as in the unique SSPE of the benchmark model. If $\delta < 1$, $v^* < \frac{1}{n}$ in any QRE.*

Proof. See Appendix A. □

Any QRE is characterized by $v^* \leq \frac{1}{n}$, that is, by a continuation value weakly smaller than in the benchmark model. This means that, relaxing perfect response, does not lead to legislators who are more demanding when evaluating a proposed allocation of resources. In spite of this, the example below illustrates that imperfect best response can easily generate a smaller average proposer's share than the one from the benchmark model.[15,16]

Example 3. Fig. 1 shows the results of the numerical calculation of the QRE for $n = 3$ and $\delta = 1$ (see Appendix B for the details of the procedure). Fig. 1a shows the mean share the proposer allocates to herself in this equilibrium along with its standard deviation and the percentage of proposed allocations that give a positive share to a minimum winning coalition.[17] Fig. 1b shows the probability an agent votes in favor of the proposal, when it assigns him x_i.

---

[14] We do not allow for non-exhaustive allocations as this would greatly complicate the computation of the QRE and would change little in terms of the results.

[15] In addition to the result in the second part of Proposition 3, one can prove that there exists a unique QRE of the alternating offer bargaining game for any δ for sufficiently small values of λ using the same technique as in the uniqueness proof of Lemma 1 in McKelvey and Palfrey (1998). In fact, this result is almost immediate if we rewrite (A.3) in the proof of Proposition 3 as $\sigma^{\lambda}(v) = \delta v + (\frac{1}{n} - \delta v) \sum_{x \in X'} r^{\lambda}_{v,1}(x) p^{\lambda}_{v}(x)$, with the sum equal to $\frac{1}{2}$ for $\lambda = 0$, after noting that $\sigma^{\lambda}(v)$ is continuous in λ for fixed v. However, in general, $\sigma^{\lambda}(v)$ fails to be monotonic, convex or contraction (in v), any of which would suffice for uniqueness (with some additional observations).

[16] We have systematically (numerically) searched for counter example to uniqueness with no success. For all combinations of $\delta \in \{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}\}$ and $\lambda \in \{0, 2, 6, 10, 18, 36, 72, 144\}$ we have calculated and plotted $\sigma^{\lambda}(v)$ for $v \in [0,1]$ (on a grid of 101 equally spaced values) and searched, unsuccessfully, for multiple solutions to $v = \sigma^{\lambda}(v)$. For this reason, we refer to the equilibrium from Proposition 3 as *the QRE*. See the Online Appendix for the results of this exercise.

[17] In any QRE with $\lambda < \infty$, proposal and voting strategies place a positive probability on every available action. For this reason, we show summary statistics of the probability distribution over actions associated with a QRE rather than point estimates. Both here and in the remainder of the paper, we classify an allocation x as 'minimum winning' if at most a minimum winning coalition of agents receive a non-negligible share of the pie, that is, if the number of agents receiving less than 5% of the pie is weakly larger than $\frac{n-1}{2}$.

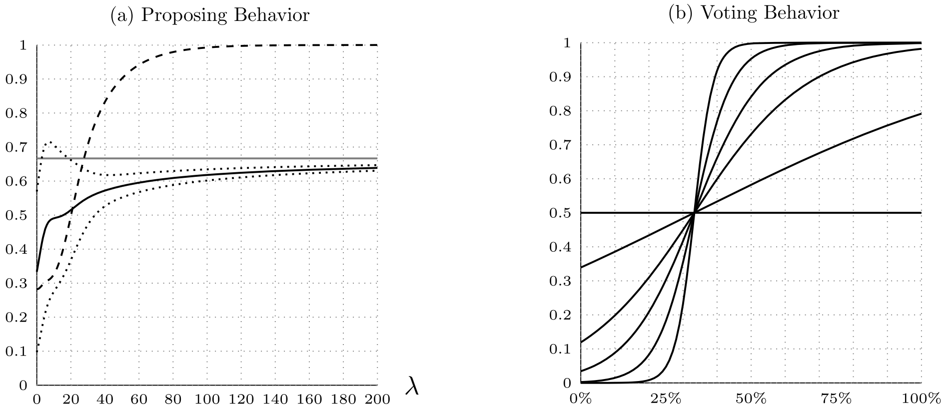

Note: The solid line represents the mean share to the proposer; the dotted lines represent $\pm$ one standard deviation; the dashed line represents the frequency of approximate minimum winning coalitions (that is, proposals where only a simple majority of legislators receives more than 5% of the budget); the horizontal solid line is the benchmark SSPE prediction.

Note: The lines represent the probability of voting *yes* (on the y-axis) given the offered share (on the x-axis) for different values of λ equal to $\{0, 2, 6, 10, 18, 36\}$. This probability is constant for $\lambda = 0$ and most responsive to the offered share for $\lambda = 36$.

Fig. 1. QRE in Baron and Ferejohn (1989), $n = 3$ and $\delta = 1$.

A full characterization of the equilibrium strategies and outcomes is difficult to obtain. In the remainder of this section, we offer a partial characterization of the proposals observed in a QRE, focusing on the allocation that is proposed most frequently.

Proposition 4. *Let $eu^{\lambda}_{v^*,i}(x) = x_i p^{\lambda}_{v^*}(x) + \delta v^*(1 - p^{\lambda}_{v^*}(x))$ be the expected utility of player i from proposing allocation x in a QRE with associated v^*. Let x^* be an allocation that maximizes $eu^{\lambda}_{v^*,i}$ subject to $x \in X'$. Then $x^*_i \geq x^*_j$ for $\forall j \in N \setminus \{i\}$ and $x^*_i > \delta v^*$.*

Proof. See Appendix A. □

In a QRE, proposers choose any $x \in X'$ with positive probability, placing more weight on proposals which give a higher expected utility. Proposition 4 shows that the allocation that maximizes player i's expected utility—that is, the allocation player i proposes with the highest probability—gives her a weakly higher share than to any other player; this allocation also gives her strictly more than her discounted continuation value. Proposition 5 shows that, in the preferred allocation, proposers treat coalition partners equally.

Proposition 5. *In any solution x^* to the proposer's problem, $\max_{x \in X} eu^{\lambda}_{v^*,i}(x)$, the proposer allocates x^*_i to herself, x^*_l to n^* legislators, with $x^*_i > x^*_l > 0$, and 0 to the remaining $n - n^* - 1$ legislators.*

Proof. See Appendix A. □

Proposition 5 is shown by proving that if player i proposes a strictly positive share to two other players in the allocation that maximizes $eu^{\lambda}_{v^*,i}$, then these two shares are equal.[18] This means that the allocation chosen most frequently by a proposer in the QRE can be fully characterized by two numbers: the share she allocates to herself, x^*_i, and the number of players among which she distributes evenly the remainder, n^*. Proposition 6 establishes that, with three players, the allocation proposed most frequently in the QRE never gives a positive share to a supermajority of committee members.

Proposition 6. *If $n = 3$, then any solution x^* to $\max_{x \in X} eu^{\lambda}_{v^*,i}(x)$ satisfies $x^*_j = 0$ for at least one $j \in N \setminus \{i\}$.*

Proof. See Appendix A. □

[18] Proposition 5 maximizes $eu^{\lambda}_{v^*,i}$ over X rather than X'. Working with a continuous proposal space allows us to use calculus in the proof of this proposition.

In the Online Appendix, we present numerical computations of the optimal and most frequent proposals in the QRE for games with $n \in \{3, 5\}$ and $\delta \in \{\frac{1}{2}, 1\}$. For low values of λ, the proposer gets the whole pie (with n^* equal to zero) or a share larger than the benchmark SSPE prediction (with $n^* < \frac{n-1}{2}$). The intuition is that, when voting behavior is sufficiently noisy, the proposed allocation has small or no impact on the probability of acceptance. As we increase λ, x_i^* first decreases below the benchmark SSPE prediction and then approaches it from below (with $n^* = \frac{n-1}{2}$). This implies that the share to the non-proposing members of a minimum winning coalition, x_l^*, approaches the benchmark SSPE prediction from above.

3.4. Combining quantal response and gambler's fallacy

Combining QRE with proposal recognition probabilities subject to the gambler's fallacy is now straightforward. Just as for the benchmark model, gambler's fallacy makes the QRE model strategically equivalent every two rounds. Once we extend the QRE mapping σ^λ to cover two rounds of play and we map the continuation value at the end of an even round, v, into $v' = \sigma^{\lambda,\phi}(v)$, the proof of equilibrium existence proceeds along similar lines to the proof of Proposition 3.

In this section, we want to highlight the different ways through which imperfect best response (in the QRE sense) and the gambler's fallacy change the equilibrium predictions relative to the benchmark model. Based on Proposition 2 and Example 3, we can see that both decrease a proposer's share. However, they do so through a different underlying mechanism.

The gambler's fallacy increases the continuation value of all non-proposing agents in the odd rounds and, thus, it makes them more expensive coalition partners, decreasing the proposer's advantage. Imperfect best-response, on the other hand, (weakly) decreases the continuation value of all agents and, thus, makes them (weakly) cheaper coalition partners with respect to the benchmark model with perfect best response. Note, however, that, given some continuation value $v \leq \frac{1}{n}$, allocating δv to a given player makes her vote for the proposed allocation with probability $\frac{1}{2}$. This is because the legislator receives the same utility from accepting and rejecting, and QRE strategies assign the same probability to payoff-equivalent actions. In turn, this makes it very risky for the proposer to offer the allocation $x = \{1 - \delta v, \delta v, 0\}$ since the expected utility following a rejection of x is $\delta v \ll 1 - \delta v$. A proposer maximizing her expected utility then needs to be more generous to her coalition partners, at the expense of her own share.[19,20] The fact that imperfect best-response implies more generous proposers does not, however, imply that it eliminates proposal power completely. In fact, Propositions 4 and 5 show that the proposing player allocates to herself more than to any other player.[21]

Regarding the predicted share to the proposer, gambler's fallacy and imperfect best response have a similar effect with respect to the benchmark model: they both contribute to reduce the advantage of the agenda setter and the resources she can assign to himself. On the other hand, imperfect best response and gambler's fallacy have different effects on the predicted frequency of minimum winning coalitions. For any value of ϕ, the gambler's fallacy predicts that any equilibrium proposal will allocate zero share of the pie to $\frac{n-1}{2}$ agents. That is, all coalitions are predicted to be minimum winning, as in the benchmark model. To the contrary, the QRE, at least for moderate values of λ, predicts a non-negligible share of proposals that assign significant resources to a coalition of agents which is larger or smaller than a minimum winning. For $n = 3$ and $n = 5$, the numerical computations in the Online Appendix show that, in the optimal proposal—that is, the proposal most frequently observed in a QRE—, the proposer gives a positive allocation only to the members of a minimum winning coalition for most values of λ. The only exception is for low values of λ, when voting behavior is very noisy: in this case, the proposer is better off giving a positive share to a coalition smaller than minimum winning.[22]

We are now ready to discuss what happens when we combine the two models. Increasing the degree of imperfect best response (that is, decreasing λ), while maintaining the degree of gambler's fallacy constant (that is, holding ϕ fixed), will reduce the predicted share to the proposer. The same will happen when we increase the extent of gambler's fallacy (that is, decrease ϕ), while maintaining the degree of imperfect best response constant. On the other hand, decreasing λ makes minimum winning coalitions less frequent while decreasing ϕ makes them more frequent.[23] The intuition behind the former effect is that, with lower λ, the probability of a mistake is larger. The intuition behind the latter effect is that a lower ϕ

---

[19] This is not the case if behavior is irresponsive to the share offered, which happens for very noisy voters, that is, very low λ.

[20] An alternative model which differs from the benchmark model only by assuming that an indifferent legislator rejects a proposal would *not* produce a similar reduction in proposal power. In this case, the proposer can ensure acceptance with an offer of $\delta v + \epsilon$. QRE generates a significant reduction in the proposer's share because the acceptance probability changes smoothly on a non-trivial neighborhood of δv.

[21] Additionally, imperfect best response changes the proposer's behavior by introducing mistakes in her choice of proposed allocations. Not only is she (optimally) more generous, but she also trembles and proposes, with positive probability, more or less generous allocations both in terms of the shares offered and in terms of the number of coalition partners. To show that mistakes do not eliminate the proposal power, we calculated the QRE expected payoff of a proposing and of a non-proposing player for $n \in \{3, 5\}$ and $\delta \in \{\frac{1}{2}, 1\}$. The results, presented graphically in the Online Appendix, show a proposer's advantage even after the introduction of mistakes.

[22] Note that, in a QRE, all feasible allocations are proposed with positive probability and the optimal proposal only informs us on the proposal offered with the highest chance. The probability the proposer offers the optimal proposal is increasing in λ. This means that, for those values of λ which make it advantageous to propose a minority coalition, the proposer is also very prone to 'mistakes' and likely to propose a positive amount to a supermajority of agents.

[23] This second statement is true for any finite value of λ. Reducing ϕ in the best-response case has no effect on the incidence of minimum winning coalition.

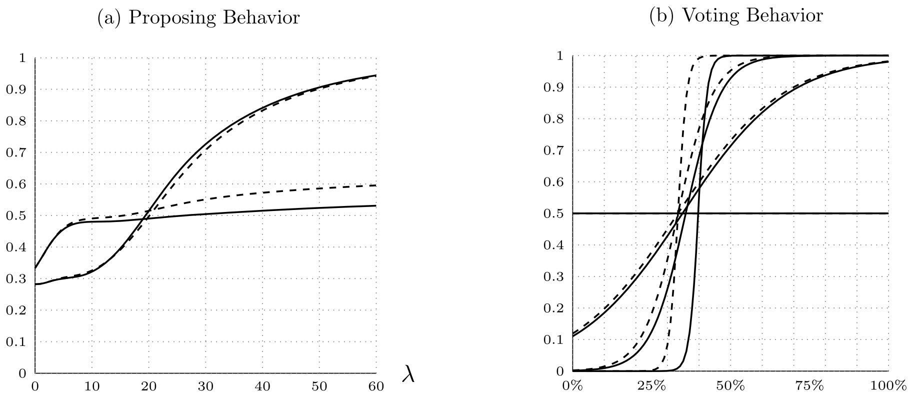

Note: Odd round predictions from QRE (dashed lines) and QGF with $\phi = 1$ (solid lines). Left panel: mean share to proposer (lines below 0.6) and frequency of approximate minimum winning coalitions (lines converging to 1). Right panel: probability of voting *yes* given offered share; values of $\lambda \in \{0, 6, 18, 72\}$; probability constant for $\lambda = 0$, most responsive for $\lambda = 72$.

Fig. 2. QRE vs. QGF, $n = 3$ and $\delta = 1$.

generates a larger continuation value for the agents who are not proposing, which makes non-minimum winning coalitions more costly to maintain relative to the minimum winning ones.

Fig. 2 shows the key differences between the QRE model without gambler's fallacy and the QRE model with gambler's fallacy (QGF). For the QGF, the figure uses the predictions for the odd rounds.[24] For $n = 3$ and $\delta = 1$, the left panel shows that the effect of introducing the gambler's fallacy in the imperfect best response model is a decrease in the mean proposer's share and an increase in the frequency of proposals that are minimum winning. For the same parameters, the right panel shows the impact on the voting behavior: in the QGF equilibrium, we have a higher continuation value for non-proposing agents in odd rounds or, equivalently, a lower probability of accepting the same proposal.

4. Data

We analyze data from six legislative bargaining experiments testing the Baron and Ferejohn (1989) model in the laboratory. Experiment 1 comes from Frechette et al. (2003) (their 'closed rule' treatment), experiments 2 through 4 from Frechette et al. (2005b) (their 'EWES', 'UWES' and 'EWES with $\delta = \frac{1}{2}$' treatments), experiment 5 from Frechette et al. (2005a) (their 'Baron and Ferejohn equal weight' treatment) and experiment 6 from Drouvelis et al. (2010) (their 'symmetric' treatment).

The experiments use either three-members ($n = 3$) or five-members ($n = 5$) committees, with or without discounting, and our data include all possible size-discounting combinations.[25] All the experiments implement a symmetric version of the Baron and Ferejohn (1989) model with equal recognition probabilities, and majoritarian voting, with the bargaining proceeding to a further round until an agreement is reached.[26]

In these experiments, subjects play 10 (experiment 6), 15 (experiment 1), or 20 (experiments 2–5) legislative bargaining games, that we call periods. For a given period and round, all the experiments asks all the subjects to propose a division of the pie. For each committee, one of the proposals is then selected and voted on.

From these data, we select round 1 proposals, selected or not, and votes for or against the selected round 1 proposals.[27] Focusing on all rounds could confound the data with repeated play effects and focusing only on approved allocations

[24] From Proposition 2 we know that, with gambler's fallacy, equilibrium proposals in even rounds are the same as in the benchmark model. This is because, in even rounds, the equilibrium continuation values are unchanged. We do not have a similar result comparing QRE and QGF. However, as illustrated in the Online Appendix, the continuation values in the QRE and in the even rounds of the QGF are, if not equal, then very similar.

[25] Appendix C reports detailed information regarding each experiment.

[26] There are two exceptions, none of which changes our equilibrium predictions. In experiment 3, the three members of each committee have, respectively, 45, 45 and 9 votes, and a proposal needs at least 50 votes for passage. This model is isomorphic to the model with majoritarian voting and each subject controlling one vote. Similarly, in experiment 6, the three members of each committee control, respectively, 3, 2 and 2 votes, and 4 votes are required for approval. Finally, experiment 6 limits the number of bargaining rounds to 20. This limit is never reached in the data and the predictions we use for this experiment assume the bargaining process can last for an infinite number of rounds. To verify that this does not play a significant role, we calculated the QRE for the finite model and for $\lambda = 0$, which is the value of λ that makes the effect of truncation at 20 rounds most relevant. The equilibrium continuation value in round 10 differs, relative to the infinite model, only at the fourth decimal place.

[27] In the context of data selection, round 1 proposals and all rounds proposals refers to all proposals, whether selected or not. Approved proposals on the other hand are subset, by definition, of selected proposals.

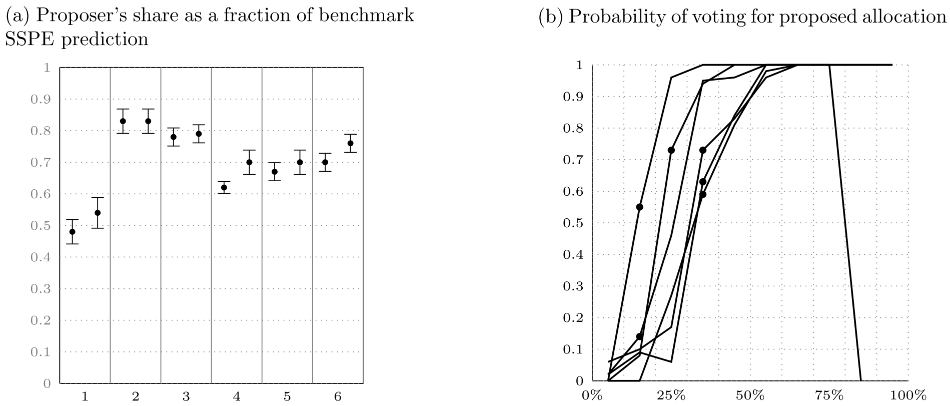

Note: Data from Frechette, Kagel and Lehrer (2003); Frechette, Kagel and Morelli (2005a,b); Drouvelis, Montero and Sefton (2010). Left panel: mean and 99% confidence interval for round 1 proposals. Inexperienced (experienced) subjects on left (right). Right panel: probability of voting *yes* given offered share. Based on response to selected round 1 proposals. • shows probability of voting *yes* at offered share predicted by the benchmark SSPE. See Appendix C for details.

Fig. 3. Experimental results from legislative bargaining experiments.

would cut the data size by a factor of three or five, depending on the experimental committee size.[28,29] We rescale all the allocations in the data to be shares of the pie rather than absolute experimental units (that is, we normalize all pies to size 1). One observation of proposing behavior is then a triplet (for three-members committees) or a quintuplet (for five-members committees) of shares in a given proposal; one observation of voting behavior is the share offered to a subject and her corresponding vote. Since subjects propose and vote exactly once in every round of every period, the number of proposing and voting observations is the same.

Figs. 3a and 3b show the main features of our data.[30] The left panel shows the mean proposer's share in the six experiments as a fraction of the benchmark SSPE prediction, along with a 99% confidence interval. For each experiment, we show separate results for experienced and unexperienced subjects in order to illustrate that, even with experience, proposers receive lower shares than what the benchmark model predicts. Because of these small differences between experienced and unexperienced subjects, we pool the data together in the remainder of the analysis.[31]

The right panel shows the probability of approving a proposal as a function of the share offered to oneself. We divided the voting data into ten 10% wide bins and, for each bin, calculated the probability of a favorable vote. Not surprisingly, the probability of approval is increasing in the share offered, and, for most experiments, larger than $\frac{1}{2}$, when the offered share is the one predicted by the SSPE in the benchmark model (the exception here being experiment 4).[32]

5. Estimation results

For each experiment in our data, we structurally estimate, using maximum likelihood techniques, the best-fitting λ for the QRE model and the best-fitting $\{\lambda, \phi\}$ pair for the QGF model.[33]

---

[28] The papers from which we draw our data use different data selection methods. All papers analyze approved allocations, adding the analysis of either round 1 proposals (Frechette et al., 2003), all rounds proposals (Frechette et al., 2005b) or all rounds minimum winning coalition proposals (Frechette et al., 2005a). To check that our estimation results do not depend on a particular data selection method, we repeated all the MLE structural estimations using either proposals from all rounds or only approved proposals. The full results of these estimations are in the Online Appendix and differ little from the results presented below.

[29] The difference in the mean proposer's share between round 1 proposals and approved proposals is, using a standard t-test, significant at the 1% level in experiments 2 and 5. The difference in the mean proposer's share between round 1 proposals and all rounds proposals is, using again a standard t-test, never significant at conventional levels (the smallest p-value being 0.45).

[30] Appendix C includes more detailed information in a non-graphical form.

[31] The difference between experienced and unexperienced subjects in the mean proposer's share in round 1 proposals is significant at the 1% level, using a standard t-test, in experiments 1, 4 and 6. This does not give a clear indication of whether to analyze the (un)experienced data separately or not. For space considerations, we present below results from pooled data. The Online Appendix includes separate results for inexperienced and experienced subjects data and shows that the results do not depend on this choice.

[32] Frechette et al. (2003, 2005a, 2005b) and Drouvelis et al. (2010) observe several other common patterns in their data. Experimental subjects require time to learn to play, mainly in the proposer role. The experiments start with a low fraction of minimum winning proposals but this fraction increases as proposers learn to play. The share offered to a subject in the selected proposal determines her voting behavior, but the shares allocated to the other subjects (and their distribution) do not matter.

[33] See Appendix D for the details of the procedure. Table 1 presents summary statistics for convenience. The maximum likelihood procedure does not fit the summary statistics but rather the complete proposing (proposed vectors of shares offered to all players) and voting (shares offered and votes cast) data.

Table 1
MLE Estimation Results for QRE and QGF.

Experiment	1	2	3	4	5	6
N	5	3	3	3	5	3
δ	4/5	1	1	1/2	1	1
Observations	275	330	411	420	450	480
$X^\star_{PR}$	.680	.666	.666	.833	.600	.666
Data						
$Avg(X_{PR})$	.338	.553	.522	.537	.409	.486
$Q1(X_{PR})$	.250	.500	.500	.500	.350	.420
$Q2(X_{PR})$	.350	.530	.500	.530	.400	.500
$Q3(X_{PR})$	.400	.600	.570	.600	.500	.570
% MWC	.422	.727	.793	.695	.853	.573
$\mathbb{P}$[delay]	.036	.318	.219	.093	.422	.269
QRE						
$\widehat{\lambda}$	20.2	22.1	23.4	10.6	33.5	18.4
$Avg(\widehat{X}_{PR})$	.441	.527	.532	.627	.400	.512
$Q1(\widehat{X}_{PR})$	.350	.490	.500	.540	.350	.460
$Q2(\widehat{X}_{PR})$	.450	.550	.550	.640	.400	.540
$Q3(\widehat{X}_{PR})$	.500	.600	.600	.730	.450	.600
% MWC	.547	.605	.635	.424	.773	.516
$\mathbb{P}$[delay]	.324	.345	.313	.179	.353	.443
Ln(L) – Overall	−2434.1	−2380.4	−2903.3	−3432.5	−3012.7	−3849.9
Ln(L) – Proposing	−2369.3	−2300.1	−2800.7	−3346.2	−2920.9	−3662.4
Ln(L) – Voting	−64.8	−80.2	−102.7	−86.3	−91.8	−187.6
QGF						
$\widehat{\lambda}$	21.7	21.7	22.5	13.3	33.8	18.8
$\widehat{\phi}$	1	3	1	1	1	1
$Avg(\widehat{X}_{PR})$	.421	.518	.497	.594	.387	.491
$Q1(\widehat{X}_{PR})$	.350	.480	.460	.530	.350	.440
$Q2(\widehat{X}_{PR})$	.450	.540	.510	.610	.400	.510
$Q3(\widehat{X}_{PR})$	.500	.590	.560	.680	.450	.570
% MWC	.553	.601	.637	.449	.794	.541
$\mathbb{P}$[delay]	.301	.339	.273	.152	.315	.382
Ln(L) – Overall	−2374.5	−2377.7	−2833.1	−3309.2	−2945.4	-3766.0
Ln(L) – Proposing	−2308.8	−2296.4	−2722.4	−3237.2	−2857.5	-3560.9
Ln(L) – Voting	−65.7	−81.3	−110.6	−72.1	−87.9	-205.2
Likelihood ratio test (p-values)						
Overall	0.0000	0.0213	0.0000	0.0000	0.0000	0.0000
Proposing	0.0000	0.0066	0.0000	0.0000	0.0000	0.0000
Voting	1.0000	1.0000	1.0000	0.0000	0.0053	1.0000

Note: Experiment 1 from Frechette et al., 2003; Experiments 2 through 4 from Frechette et al., 2005b; Experiment 5 from Frechette et al., 2005a; Experiment 6 from Drouvelis et al., 2010. X_{PR}, $X^\star_{PR}$, and $\widehat{X}_{PR}$ refer, respectively to the proposer's allocation observed in the data, the proposer's allocation predicted by the benchmark model, and the proposer's allocation predicted by the MLE estimates. % MWC refers to the incidence of minimum winning coalitions, defined as proposals where at least 1 member (for $n = 3$), or at least 2 members (for $n = 5$), receive less than 5% of the pie. All data and estimates refer to round 1 behavior.

In Table 1, we show the average share to the proposer (X_{PR}), the quartiles of its distribution, the incidence of minimum winning coalitions, and the frequency of delays in the 6 experimental datasets, and we compare them with the predictions from the benchmark model ($X^\star_{PR}$)[34] and with the predictions from the QRE and QGF models (using the estimated parameters).[35]

The second panel in Table 1 shows the summary statistics from the experimental data and allows us to gauge the distance with the predictions of the benchmark model: the median share to the proposer is between 51% (in experiment 1) and 80% (in experiment 2) of the equilibrium predictions; and the incidence of minimum winning coalitions—which is predicted to be 100% by the benchmark model—is between 42% (in experiment 1) and 85% (in experiment 5).

The third panel in Table 1 shows, for each experiment, the estimated λ from the QRE model, the corresponding predictions, and the corresponding log-likelihoods (separately, for the whole dataset, for the proposing behavior, and for the voting behavior). The estimated QRE model fits the experimental data better than the equilibrium from the benchmark model.[36]

[34] In the equilibrium of the benchmark model, with perfect best response and no gambler's fallacy, all proposals give the proposer exactly $X^\star_{PR}$ and distribute benefits only to a minimum winning coalitions. We omit this from Table 1 in the interest of space.

[35] While here we focus on the most important summary statistics, the Online Appendix shows the whole distribution of the share to the proposer, as well as the probability of approving a proposed allocation for the experimental data and for the QRE and QGF estimates.

[36] The predictions from the estimated QRE are significantly superior to the predictions from the QRE restricted to $\lambda = 500$ (a close approximation of the benchmark SSPE). See the Online Appendix for the results of the likelihood-ratio tests comparing these two models.

The estimated QRE, however, cannot explain at the same time the two main stylized facts from the experiments: a high incidence of minimum winning coalitions, and a more egalitarian split between coalition partners. The typical problem of the QRE estimates is that they over-predict the share to the proposer while under-predicting the frequency of minimum winning coalitions. This is the case in experiments 3, 4, and 6. This is because, as we mentioned above, changing the degree of rationality in the QRE entails a trade-off in the predicted behavior: increasing λ increases both the predicted proposer's share and the predicted frequency of minimum winning coalition (hence, worsening the fit in terms of proposer's share); on the other hand, decreasing λ decreases both the predicted proposer's share and the predicted frequency of minimum winning coalitions (hence, worsening the fit in terms of minimum winning coalitions).

The fourth panel in Table 1 shows, for each experiment, the estimated λ and ϕ from the QGF model, the corresponding predictions, and the corresponding log-likelihoods. First, we note that the estimated degree of gambler's fallacy is in line with the presence of a significant cognitive bias: in 5 out of 6 experiments, the best fitting ϕ is 1, the highest degree of gambler's fallacy; the best fitting ϕ is consistent with the presence of a bias also in the remaining experiment (that is, $\widehat{\phi} < \infty$).

Adding gambler's fallacy to imperfect best response helps to reconcile the theoretical predictions with the observed data on both counts. The QGF estimates generally predict a lower proposer's share (with respect to the QRE) and a higher frequency of minimum winning coalitions. For the best fitting parameters, this is the case for experiments 4 and 6. In experiments 3 and 5, only the incidence of minimum winning coalitions gets closer to the data and in experiment 1, only the proposer's share gets closer to the data. This is because the maximum likelihood estimates are chosen to fit best the complete proposing and voting data (that is, proposal vectors that specify allocations to every committee member, and voting probabilities as a function of proposed allocations), not just the share to the proposer and whether an allocation is minimum winning or not (the summary statistics that we present here). However, as shown in the fifth panel of Table 1, the predictions from the QGF model are significantly superior to the predictions of the QRE model for all experiments. The better fit is achieved mostly through a better prediction of proposal behavior.

Contrary to the benchmark model, the QRE and QGF models predict occurrence of delays with strictly positive probability. Table 1 shows that the frequency of delays, defined as the frequency of first-round rejections, in the experimental data ranges from 4% (in experiment 1) to 42% (in experiment 5). The QRE model predicts fairly well the probability of delays, although it generally overestimates it (this is true for all data but experiment 5, and in particular for experiment 1). Adding gambler's fallacy to the QRE model lowers the predicted probability of delays and moves the predictions closer to the experimental data for all experiments (except for experiment 5).[37]

6. Conclusions

The interesting empirical results reported by Frechette et al. (2003, 2005a, 2005b), and Drouvelis et al. (2010) on alternating offer bargaining games show some consistent deviation from the predictions of standard non-cooperative game theory: proposers do not fully exploit the advantage conferred by their agenda setting power, but still get the largest share in the coalition supporting the agreement, and do not allocate resources to committee members in excess of the minimal coalition needed for approval; voters vote selfishly and care about the allocation to themselves, but not about the allocation to others.

In this paper, we show that a model with imperfect best response (logit QRE) fits these data rather well. This model fits the incomplete exploitation of proposal power and the occasional occurrence of proposals that distribute benefits to a coalition larger than minimum winning. However, imperfect best response alone cannot account for all patterns in the data: reducing the coefficient of rationality in the QRE, reduces the share to the proposer, but it also increases the resources allocated to legislators in excess of a minimum winning coalition.

The fit is further improved significantly by assuming erroneous beliefs on the stochastic allocation of future proposal power, in the form of the gambler's fallacy. The addition of this well documented cognitive bias decreases the share to the proposer to match the observed proposing behavior more closely, without decreasing the incidence of minimum winning coalitions or changing the fit of voting behavior. We conclude that the suboptimal behavior in these bargaining environments is plausibly attributable to mistaken beliefs on probabilistic events combined with imperfect best response.

[37] As discussed above, we estimated the models using alternative datasets, that is, using proposals from all rounds, only approved proposals, only unexperienced subjects or only experienced subjects. The results of these robustness checks are reported in the Online Appendix and are very similar to the estimates from Table 1. In particular, with respect to the estimates we present here, the estimates of λ are slightly lower for proposals from all rounds, lower for unexperienced subjects, and higher for experienced subjects. All the estimates of λ are higher for the datasets including only experienced subjects relative to the datasets including only unexperienced subjects. The estimated ϕ for the QGF models is generally equal to 1, except in experiment 2 where it varies between 1 and 6, depending on the dataset used.

Appendix A. Proofs

A.1. SSPE definition

A proposal strategy of $i \in N$ is $r_i : X \to \Delta(X)$, where $\Delta(X)$ is the set of distributions defined on X; and a voting strategy of $i \in N$ is $p_i : X \to \Delta(\{\text{accept}, \text{reject}\})$. Stationarity means $i \in N$ uses p_i and r_i in every round of the game and we define symmetry to mean $p_i = p_j$ and $r_i = r_j$ for $\forall i, j \in N$, where the need to permute entries in arguments of the strategies is implicitly understood. Symmetric treatment of potential coalition partners means that if $i \in N$ is recognized to be the proposer and both $x \in X$ and $y \in X$ maximize her expected utility given the strategies of the other agents and x can be obtained from y by re-labeling of entries, then $r_i(x) = r_i(y)$. SSPE is a profile of voting and proposal strategies that are stationary, symmetric, treat potential coalition partners symmetrically and constitute subgame perfect Nash equilibrium.

A.2. Proof of Proposition 2

Proof. Denote by v^o odd round s continuation value of the agents who have not proposed in s. There is no need to index v^o by i due to symmetry. For the same odd round s, v^{op} denotes continuation value of the agent who proposed in s. Finally denote by v^e continuation value of the agents in even rounds. There is no need to index v^e differently for the proposing player; the recognition probabilities, and hence the continuation values, are equal for all the agents every two rounds. Similarly, x^e and x^o are shares offered in even and odd rounds respectively.

Given the recognition probabilities

$$\begin{aligned} v^e &= \frac{\phi}{\phi n}\left(1 - \frac{n-1}{2}x^o\right) + \frac{\phi n - \phi}{\phi n}\left(\frac{1}{2}x^o\right) \\ v^o &= \frac{\phi}{\phi n - 1}\left(1 - \frac{n-1}{2}x^e\right) + \frac{\phi n - 1 - \phi}{\phi n - 1}\left(\frac{1}{2}x^e\right) \\ v^{op} &= \frac{\phi - 1}{\phi n - 1}\left(1 - \frac{n-1}{2}x^e\right) + \frac{\phi n - 1 - (\phi - 1)}{\phi n - 1}\left(\frac{1}{2}x^e\right) \end{aligned} \tag{A.1}$$

which simplifies to $v^e = \frac{1}{n}$, $v^o = \frac{\phi - x^e/2}{\phi n - 1}$ and $v^{op} = \frac{\phi - \left(1 - \frac{n-1}{2}x^e\right)}{\phi n - 1}$.

Proposing player i, if she finds inducing acceptance of her proposal optimal, will propose positive shares to $\frac{n-1}{2}$ other agents with the shares being just sufficient to guarantee their agreement. This implies $x^e = \delta v^e$ and $x^o = \delta v^o$. This implies that there exists unique $x^e = \frac{\delta}{n}$ which in turn implies existence of unique $x^o = \frac{\delta}{n}\frac{2n\phi - \delta}{2n\phi - 2}$. What remains to be shown is that proposer's payoff from inducing approval of her proposal, $1 - \frac{n-1}{2}x^o$, is larger than payoff from inducing rejection, δv^{op}. This rewrites using $x^o = \delta v^o$ as $1 \geq \delta[v^{op} + \frac{n-1}{2}v^o]$ and certainly holds as $v^{op} + (n-1)v^o = 1$.

Given the continuation values voting behavior described in the proposition is clearly optimal. Finally, comparative statics on ϕ is immediate due to $\frac{2n\phi - \delta}{2n\phi - 2}$ strictly decreasing in ϕ and limit equal to 1 as $\phi \to \infty$. □

A.3. Proof of Proposition 3

Proof. First, we prove that there exists a QRE of the alternating offer bargaining game and that, in any QRE, the associated continuation value v^* satisfies $v^* \leq \frac{1}{n}$.

Take σ^λ as defined in the text and restrict its domain to $[0, 1]$. It is easy to see $\sigma^\lambda : [0, 1] \to [0, 1]$. All we need to show is that σ^λ has fixed point $v^* = \sigma^\lambda(v^*)$, due to $\sigma^\lambda(v)$ being symmetric for any $v \in [0, 1]$ and $\{\sigma^\lambda(v^*), \sigma^\lambda(v^*), \ldots\}$ constituting stationary strategy profile.

Fixing $v \in [0, 1]$, by Theorem 3 in McKelvey and Palfrey (1998) there exists unique $v' = \sigma^\lambda(v)$ so that we can think of σ^λ as of a function. With the proposal and voting strategies continuous in v, σ^λ is continuous in v as well. By Brouwer fixed point theorem, there exists $v^* = \sigma^\lambda(v^*)$.

To show that $v^* \leq \frac{1}{n}$ for any $v^* = \sigma^\lambda(v^*)$ and $v^* < \frac{1}{n}$ when $\delta < 1$, first notice

$$\sigma^\lambda(v) = \sum_{j \in N} \frac{1}{n} \sum_{x \in X'} r^\lambda_{v,j}(x) \left[x_i\, p^\lambda_v(x) + \delta\, v\, (1 - p^\lambda_v(x))\right] \tag{A.2}$$

for any $i \in N$.

In order to proceed we need extra notation. Take arbitrary allocation $x \in X'$, $x = \{x_1, \ldots, x_n\}$, and corresponding $p^\lambda_v(x)$ and $r^\lambda_{v,1}(x)$. Define m-time circular shift operator $\sigma^m(i)$ as $\sigma^m(i) = i - m$ modulo n for $m \in \{0, \ldots, n-1\}$. Applied to x, let $x_{\sigma^m} = \{x_{\sigma^m(1)}, \ldots, x_{\sigma^m(n)}\}$ ordering the entries by their new index. For example, $x_{\sigma^0} = x$ and using the original indexes $x_{\sigma^1} = \{x_n, x_1, \ldots, x_{n-1}\}$. Due to symmetry $p^\lambda_v(x_{\sigma^m}) = p^\lambda_v(x_{\sigma^0})$ for any $m \in \{0, \ldots, n-1\}$. Denoting by $r^\lambda_{v,\sigma^m(1)}(x_{\sigma^m})$ probability of player $\sigma^m(1)$ proposing x_{σ^m} we further have $r^\lambda_{v,\sigma^m(1)}(x_{\sigma^m}) = r^\lambda_{v,1}(x_{\sigma^0})$ for all $m \in \{0, \dots, n-1\}$. Using this along with $\sum_{m=0}^{n-1} x_{\sigma^m} = \{1, \dots, 1\}$ we can rewrite (A.2) as

$$\sigma^\lambda(v) = \sum_{x \in X'} \frac{1}{n} r^\lambda_{v,1}(x) \left[p^\lambda_v(x) + n\,\delta\, v\,(1 - p^\lambda_v(x)) \right]. \tag{A.3}$$

For $n\delta v = 1$ (A.3) implies $\sigma^\lambda(v) = \frac{1}{n} = \delta v$ so that for and only for $\delta = 1$, $v = \frac{1}{n\delta}$ is fixed point of σ^λ. For $n\delta v > 1$ (A.3) implies $\sigma^\lambda(v) < \delta v$ so that σ^λ cannot have fixed point when $v > \frac{1}{\delta n}$. For $n\delta v < 1$ (A.3) implies $\sigma^\lambda(v) < \frac{1}{n}$ so that any fixed point of σ^λ has to occur for $v < \frac{1}{n}$.

Second, we prove that, if $\delta = 1$, there exists a unique QRE with $v^* = \frac{1}{n}$. Rewriting (A.3) as

$$\sigma^\lambda(v) = \frac{1}{n} \sum_{x \in X'} r^\lambda_{v,1}(x) p^\lambda_v(x) + \delta v \sum_{x \in X'} r^\lambda_{v,1}(x)(1 - p^\lambda_v(x)) \tag{A.4}$$

shows that $\sigma^\lambda(v)$ is a weighted average of $\frac{1}{n}$ and δv, with a weight $\sum_{x \in X'} r^\lambda_{v,1}(x) p^\lambda_v(x) \in (0,1)$. When $\delta = 1$, this implies that

$$\sigma^\lambda(v) \in \begin{cases} \left(v, \frac{1}{n}\right) & \text{if } v \in \left[0, \frac{1}{n}\right) \\ \left(\frac{1}{n}, v\right) & \text{if } v \in \left(\frac{1}{n}, 1\right] \end{cases}$$

so that $v^* = \sigma^\lambda(v^*)$ cannot hold unless $v^* = \frac{1}{n}$. □

A.4. Proof of Proposition 4

Proof. Let x^* be an allocation that maximizes $eu^\lambda_{v^*,i}$ subject to $x \in X'$. We first show that $x^*_i \geq x^*_j$ for $\forall j \in N \setminus \{i\}$. Suppose $x^*_i < x^*_j$ for some $j \in N \setminus \{i\}$. Let x' be identical to x^* except for swapping allocations to i and j. That is, $x'_k = x^*_k$ for $k \in N \setminus \{i, j\}$, $x'_i = x^*_j$ and $x'_j = x^*_i$. Because $x^*_i < x^*_j$, $x'_i > x'_j$. Furthermore, $p^\lambda_{v^*}(x^*) = p^\lambda_{v^*}(x')$. Hence $eu^\lambda_{v^*,i}(x^*) < eu^\lambda_{v^*,i}(x')$, which contradicts x^* maximizing $eu^\lambda_{v^*,i}$.

We now show that $x^*_i > \delta v^*$. Suppose $\delta < 1$ and $x^*_i \leq \delta v^*$. Since $v^* \leq \frac{1}{n}$, $x^*_i < \frac{1}{n}$ and thus $nx^*_i < 1$. From the previous paragraph $x^*_i = \max_{j \in N} x^*_j$ and hence $nx^*_i \geq \sum_{j \in N} x^*_j = 1$, a contradiction to $nx^*_i < 1$.

Suppose $\delta = 1$ and $x_i \leq v^*$. If $x^*_i < \frac{1}{n}$, then we obtain a contradiction by the argument used in the previous paragraph. Hence $x^*_i \geq \frac{1}{n}$. Since $v^* \leq \frac{1}{n}$, it must be the case that $x^*_i = v^*$ and hence $eu^\lambda_{v^*,i}(x^*) = v^*$. Let $x' \in X'$ be such that $x'_i = 1$ and $x'_j = 0$ for $\forall j \in N \setminus \{i\}$. Then $eu^\lambda_{v^*,i}(x') = (1 - v^*) p^\lambda_{v^*}(x') + v^* > eu^\lambda_{v^*,i}(x^*)$, where the inequality follows from $1 - v^* > 0$ and $p^\lambda_{v^*}(x') > 0$. But $eu^\lambda_{v^*,i}(x') > eu^\lambda_{v^*,i}(x^*)$ contradicts x^* maximizing $eu^\lambda_{v^*,i}$. □

A.5. Proof of Proposition 5

Proof. To prove the proposition, we will show that if x^* solves $\max_{x \in X} eu^\lambda_{v^*,i}(x)$ with $x^*_j \in (0,1)$ and $x^*_k \in (0,1)$ for some $j, k \in N \setminus \{i\}$, then $x^*_i > x^*_j = x^*_k$. If $\lambda = 0$, x^* is unique, $x^*_i = 1$ and $x^*_j = 0$ for any $j \in N \setminus \{i\}$, so consider $\lambda > 0$ throughout.

The proposer's problem is

$$\max_x eu^\lambda_{v^*,i}(x) \text{ s.t. } x_j \in [0,1] \text{ for } \forall j \in N \text{ and } \sum_{j \in N} x_j \leq 1 \tag{OP}$$

where

$$eu^\lambda_{v^*,i}(x) = x_i p^\lambda_{v^*}(x) + \delta v^*(1 - p^\lambda_{v^*}(x)) = (x_i - \delta v^*) p^\lambda_{v^*}(x) + \delta v^*. \tag{A.5}$$

Because the objective function in (OP) is continuous and the constraint set is compact, the solution exists. Furthermore, the constraints are cut-out by affine functions and hence any solution x^* to (OP) has to satisfy the standard Kuhn–Tucker necessary conditions without need for constraint qualification. It is also clear that the $\sum_{i \in N} x_i \leq 1$ constraint has to bind at any solution. We also note that Proposition 4 continues to hold if X' is replaced by X. Hence any solution x^* to (OP) has to satisfy $x^*_i > \delta v^*$ and $x^*_i \geq x^*_j$ for $\forall j \in N \setminus \{i\}$.

Before writing the Lagrangian associated with (OP), we note that $p^\lambda_{v^*}(x)$ is the probability that at least $\frac{n+1}{2}$ out of n players vote for x. Denote by $p(x)_{q|n}$ the probability that at least q out of n players vote for x and by $p(x)^*_{q|n}$ the probability that exactly q out of n voters vote for x, where the dependence on v^* and λ is understood but omitted in order to economize on notation. Setting $q = \frac{n+1}{2}$, $p^\lambda_{v^*}(x) = p(x)_{q|n}$. Similarly we drop v^* and λ from $p^\lambda_{v^*,i}(x)$ for $\forall i \in N$ which becomes $p_i(x)$ for $\forall i \in N$. Since $p_i(x) = \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^*}$, simple algebra shows that $\frac{\partial p_i(x)}{\partial x_i} = \lambda p_i(x)(1 - p_i(x))$. Denote by x^{ij} the allocation x with x_i and x_j dropped. The following way of writing $p(x)_{q|n}$ will be useful below

$$\begin{aligned} p(x)_{q|n} &= p_i(x)p_j(x)p(x^{ij})_{q-2|n-2} + (1-p_i(x))(1-p_j(x))p(x^{ij})_{q|n-2} \\ &\quad + \left[p_i(x)(1-p_j(x)) + (1-p_i(x))p_j(x)\right] p(x^{ij})_{q-1|n-2} \\ &= p(x^{ij})^*_{q-2|n-2}p_i(x)p_j(x) - p(x^{ij})^*_{q-1|n-2}(1-p_i(x))(1-p_j(x)) + p(x^{ij})_{q-1|n-2}. \end{aligned} \tag{A.6}$$

Taking derivative with respect to x_i gives

$$\frac{\partial p(x)_{q|n}}{\partial x_i} = \lambda p_i(x)(1-p_i(x))\left[p_j(x)p(x^{ij})^*_{q-2|n-2} + (1-p_j(x))p(x^{ij})^*_{q-1|n-2}\right]. \tag{A.7}$$

With the notation established the Lagrangian associated with (OP) is

$$L(x,\mu,\mu_-,\mu_+) = (x_i - \delta v^*)p(x)_{q|n} + \delta v^* - \mu\left[\sum\nolimits_{i\in N} x_i - 1\right] - \mu_- \cdot (-x) - \mu_+ \cdot (x-1) \tag{A.8}$$

where μ_- and μ_+ are vectors of multipliers associated with the $x_j \geq 0$ and $x_j \leq 1$ constraints and $\cdot$ is the standard inner product.

Fix x^* that solves (OP) and suppose, towards a contradiction, that $x_j^* \in (0,1)$ and $x_k^* \in (0,1)$ for some $j,k \in N \setminus \{i\}$ and $x_j^* \neq x_k^*$. We also have $x_i^* \in (0,1)$ since $x_i^* > \delta v^* \geq 0$ and $x_i^* \leq 1 - x_j^* - x_k^*$. The first order conditions associated with x_i^*, x_j^* and x_k^* thus are

$$\begin{aligned} \frac{\partial L(x^*,\mu^*,\mu_-^*,\mu_+^*)}{\partial x_i} &= p(x^*)_{q|n} + (x_i^* - \delta v^*)\frac{\partial p(x^*)_{q|n}}{\partial x_i} - \mu^* = 0 \\ \frac{\partial L(x^*,\mu^*,\mu_-^*,\mu_+^*)}{\partial x_j} &= (x_i^* - \delta v^*)\frac{\partial p(x^*)_{q|n}}{\partial x_j} - \mu^* = 0 \\ \frac{\partial L(x^*,\mu^*,\mu_-^*,\mu_+^*)}{\partial x_k} &= (x_i^* - \delta v^*)\frac{\partial p(x^*)_{q|n}}{\partial x_k} - \mu^* = 0. \end{aligned} \tag{A.9}$$

Since $x_i^* > \delta v^*$, we have $\frac{\partial p(x^*)_{q|n}}{\partial x_j} - \frac{\partial p(x^*)_{q|n}}{\partial x_k} = 0$. Using (A.7), this implies

$$\begin{aligned} &\left[\frac{\partial p(x^*)_{q|n}}{\partial x_j} - \frac{\partial p(x^*)_{q|n}}{\partial x_k}\right]\left[\lambda(p_j(x^*) - p_k(x^*))\right]^{-1} \\ &= \left[p(x^{jk,*})^*_{q-1|n-2}(1-p_j(x^*))(1-p_k(x^*)) - p(x^{jk,*})^*_{q-2|n-2}p_j(x^*)p_k(x^*)\right] \\ &= \frac{1}{d(x_j^*,x_k^*)}\left[p(x^{jk,*})^*_{q-1|n-2}\exp 2\lambda\delta v^* - p(x^{jk,*})^*_{q-2|n-2}\exp\lambda(x_j^*+x_k^*)\right] \\ &= \frac{1}{d(x_i^*,x_j^*,x_k^*)}\left[\begin{array}{l} p(x^{ijk,*})^*_{q-1|n-3}\exp\lambda 3\delta v^* - p(x^{ijk,*})^*_{q-3|n-3}\exp\lambda(x_i^*+x_j^*+x_k^*) \\ + p(x^{ijk,*})^*_{q-2|n-3}\left[\exp\lambda(2\delta v^* + x_i^*) - \exp\lambda(\delta v^* + x_j^* + x_k^*)\right] \end{array}\right] \end{aligned} \tag{A.10}$$

where $d(x_j^*,x_k^*) = \prod_{z\in\{x_j^*,x_k^*\}}(\exp\lambda z + \exp\lambda\delta v^*)$, similarly for $d(x_i^*,x_j^*,x_k^*)$ and, by convention, $p(x^{ijk,*})^*_{0|0} = 1$ and $p(x^{ijk,*})^*_{q|0} = 0$ for $q \neq 0$. Use $T_1(x^*,j,k)$, $T_2(x^*,j,k)$ and $T_3(x^*,i,j,k)$ as the shorthand for the term in the square brackets on the second, third and fourth line of (A.10) respectively. Because $x_j^* \neq x_k^*$, $p_j(x^*) \neq p_k(x^*)$. Hence $T_1(x^*,j,k) = T_2(x^*,j,k) = T_3(x^*,i,j,k) = 0$. Moreover, using $T_1(x^*,j,k)$ in (A.6), we have $p(x^*)_{q|n} = p(x^{jk,*})_{q-1|n-2}$.

Denote $X^* = \{x \in X | x_l = x_l^* \ \forall l \in N \setminus \{i,j,k\}, x_i = x_i^*, x_j + x_k = x_j^* + x_k^*\}$ and note $x^* \in X^*$. Since $T_2(x,j,k)$ depends on x_j and x_k only through $x_j + x_k$, $T_2(x',j,k) = 0$ for any $x' \in X^*$. Hence $T_1(x',j,k) = 0$ and thus $p(x')_{q|n} = p(x^{jk,'})_{q-1|n-2} = p(x^{jk,*})_{q-1|n-2}$ for any $x' \in X^*$. Therefore, any $x' \in X^*$ solves (OP) since $x_i' = x_i^*$ and $p(x')_{q|n} = p(x^*)_{q|n}$. Since any $x' \in X^*$ solves (OP), it satisfies (except at boundaries of X^* when $x_j' \in \{0, x_j^* + x_k^*\}$)

$$(x_i' - \delta v^*)\left[\frac{\partial p(x')_{q|n}}{\partial x_i} - \frac{\partial p(x')_{q|n}}{\partial x_j}\right] + p(x')_{q|n} = 0 \tag{A.11}$$

implied by the first order conditions associated with x_i' and x_j'. Since $x_i' = x_i^*$ and $p(x')_{q|n} = p(x^*)_{q|n}$, $\frac{\partial p(x')_{q|n}}{\partial x_i} - \frac{\partial p(x')_{q|n}}{\partial x_j}$ is constant in $x' \in X^*$.

Writing $\frac{\partial p(x')_{q|n}}{\partial x_i} - \frac{\partial p(x')_{q|n}}{\partial x_j}$, for any $x' \in X^*$ we have

$$\begin{aligned}
&\left[\frac{\partial p(x')_{q|n}}{\partial x_i} - \frac{\partial p(x')_{q|n}}{\partial x_j}\right]\left[\lambda(p_i(x') - p_j(x'))\right]^{-1} \\
&= \frac{1}{d(x'_i, x'_j, x'_k)}\left[\begin{array}{l} p(x^{ijk,'})^*_{q-1|n-3}\exp\lambda 3\delta v^* - p(x^{ijk,'})^*_{q-3|n-3}\exp\lambda(x'_i + x'_j + x'_k) \\ + p(x^{ijk,'})^*_{q-2|n-3}\left[\exp\lambda(2\delta v^* + x'_k) - \exp\lambda(\delta v^* + x'_i + x'_j)\right] \end{array}\right] \\
&= \frac{p(x^{ijk,'})^*_{q-2|n-3}}{d(x'_i, x'_j, x'_k)}\left[\begin{array}{l} -\exp\lambda(2\delta v^* + x'_i) + \exp\lambda(\delta v^* + x'_j + x'_k) \\ \exp\lambda(2\delta v^* + x'_k) - \exp\lambda(\delta v^* + x'_i + x'_j) \end{array}\right] \\
&= p(x^{ijk,'})^*_{q-2|n-3}\left[\begin{array}{l} -p_i(x')(1 - p_j(x'))(1 - p_k(x')) + (1 - p_i(x'))p_j(x')p_k(x') \\ (1 - p_i(x'))(1 - p_j(x'))p_k(x') - p_i(x')p_j(x')(1 - p_k(x')) \end{array}\right]
\end{aligned} \tag{A.12}$$

where going from the second to the third line uses $T_3(x', i, j, k) = 0$ for any $x' \in X^*$. Routine algebra shows that

$$\frac{\partial p(x')_{q|n}}{\partial x_i} - \frac{\partial p(x')_{q|n}}{\partial x_j} = -\lambda p(x^{ijk,'})^*_{q-2|n-3}(p_i(x') - p_j(x'))(p_i(x') - p_k(x')). \tag{A.13}$$

Since $\lambda p(x^{ijk,'})^*_{q-2|n-3} > 0$ is independent of $x' \in X^*$, $(p_i(x') - p_j(x'))(p_i(x') - p_k(x'))$ has to be as well. Evaluating the expression at x'_i, x'_j and $x'_k = x^*_j + x^*_k - x'_j$, we have

$$(p_i(x'_i) - p_j(x'_j))(p_i(x'_i) - p_k(x^*_j + x^*_k - x'_j)) \tag{A.14}$$

and its derivative with respect to x'_j writes, after some algebra, as

$$-\lambda[p_j(x'_j) - p_k(x'_k)][p_i(x'_i)(1 - p_j(x'_j))(1 - p_k(x'_k)) + p_j(x'_j)p_k(x'_k)(1 - p_i(x'_i))] \tag{A.15}$$

which is non-zero unless $x'_j = x'_k$, a contradiction.

Finally, we show $x^*_i > x^*_l$ for any $l \in N \setminus \{i\}$. By Proposition 4 it suffices to show $x^*_i \neq x^*_l$. Suppose, towards a contradiction, that $x^*_i = x^*_l$ for some $l \neq i$. This implies that $x^*_i \in (0, 1)$ since $x^*_i > \delta v^* \geq 0$ and $x^*_i \leq 1 - x^*_l < 1 - \delta v^*$. Hence, the first order conditions associated with x^*_i and x^*_l are

$$\begin{aligned}
\frac{\partial L(x^*, \mu^*, \mu^*_-, \mu^*_+)}{\partial x_i} &= p(x^*)_{q|n} + (x^*_i - \delta v^*)\frac{\partial p(x^*)_{q|n}}{\partial x_i} - \mu^* = 0 \\
\frac{\partial L(x^*, \mu^*, \mu^*_-, \mu^*_+)}{\partial x_l} &= (x^*_i - \delta v^*)\frac{\partial p(x^*)_{q|n}}{\partial x_l} - \mu^* = 0
\end{aligned} \tag{A.16}$$

which, because $\frac{\partial p(x^*)_{q|n}}{\partial x_i} = \frac{\partial p(x^*)_{q|n}}{\partial x_l}$ when $x^*_i = x^*_l$ and because $p(x^*)_{q|n} > 0$, cannot be satisfied simultaneously. □

A.6. Proof of Proposition 6

Proof. Fix $n = 3$ and let player i be the proposer and players j and k the responders. From Proposition 4, by the remark we made in the proof of Proposition 5, any solution x^* to (OP) satisfies $x^*_i > \delta v^*$. If $\lambda = 0$, (OP) has unique solution x^* with $x^*_i = 1$ so that the proposition holds and hence we only need to consider $\lambda > 0$.

Suppose first that $x^*_i > \frac{1-\delta v^*}{2}$ and, towards a contradiction, that $x^*_j > 0$ and $x^*_k > 0$. Since x^* solves (OP), (x^*_j, x^*_k) has to be a solution to $\max_{x_j \in [0,1], x_k \in [0,1]}(x^*_i - \delta v^*)p((x^*_i, x_j, x_k))_{2|3}$ subject to $x_j + x_k = 1 - x^*_i$. Since $x^*_i - \delta v^* > 0$, we can suppress the $(x^*_i - \delta v^*)$ term. Using (A.6) and $x' = (x^*_i, x_j, x_k)$, we can write $p((x^*_i, x_j, x_k))_{2|3} = p(x')_{2|3}$ as

$$\begin{aligned}
p(x')_{2|3} &= (1 - p_i(x^*_i))p_j(x_j)p_k(x_k) - p_i(x^*_i)(1 - p_j(x_j))(1 - p_k(x_k)) + p_i(x^*_i) \\
&= \frac{\exp(\lambda(\delta v^* + x_j + x_k)) - \exp(\lambda(x^*_i + 2\delta v^*))}{\prod_{z \in \{x^*_i, x_j, x_k\}}\exp(\lambda\delta v^*) + \exp(\lambda z)} + p_i(x^*_i).
\end{aligned} \tag{A.17}$$

Since $x^*_i > \frac{1-\delta v^*}{2}$ and $x_j + x_k = 1 - x^*_i$, $\delta v^* + x_j + x_k < x^*_i + 2\delta v^*$ and hence the numerator in (A.17) is strictly negative. Maximization of $p(x')_{2|3}$ is thus equivalent to the maximization of the denominator in (A.17). This problem after some algebra reads (ignoring strictly positive constants) $\max_{x_j \in [0, 1-x^*_i]}\left[\exp(\lambda(x_j + \delta v^*)) + \exp(\lambda(1 - x^*_i - x_j + \delta v^*))\right]$ and has two solutions $x_j^* = 0$ and $x_j^* = 1 - x_i^*$ since the objective function of the problem is strictly decreasing at $x_j = 0$ and is strictly convex. But this contradicts $x_j^* > 0$ and $x_k^* = 1 - x_i^* - x_j^* > 0$.

Suppose now that $x_i^* \leq \frac{1-\delta v^*}{2}$. We will show that $x_i^* \leq \frac{1-\delta v^*}{2}$ leads to a contradiction. Denote by $\bar{x}$ any x with x_i restricted to $x_i \leq \frac{1-\delta v^*}{2}$. Now note that because x_i is restricted to $x_i \leq \frac{1-\delta v^*}{2}$, $eu^{\lambda}_{v^*,i}(\bar{x}) = (\bar{x}_i - \delta v^*) p^{\lambda}_{v^*}(\bar{x}) + \delta v^* \leq \left[\frac{1-\delta v^*}{2} - \delta v^*\right] \cdot 1 + \delta v^*$. Consider now $x' = (1 - \delta v^*, 0, \delta v^*)$. Since

$$p(x')_{2|3} = \frac{1}{2} \frac{\exp(\lambda) + \exp(\lambda 2\delta v^*) + \exp(\lambda(1+\delta v^*)) + \exp(\lambda)}{\exp(\lambda) + \exp(\lambda 2\delta v^*) + \exp(\lambda(1+\delta v^*)) + \exp(\lambda 3\delta v^*)} \geq \frac{1}{2} \tag{A.18}$$

we have $eu^{\lambda}_{v^*,i}(x') \geq [1 - \delta v^* - \delta v^*] \frac{1}{2} + \delta v^*$. Because

$$eu^{\lambda}_{v^*,i}(x') \geq \left[1 - \delta v^* - \delta v^*\right] \frac{1}{2} + \delta v^* > \left[\frac{1-\delta v^*}{2} - \delta v^*\right] \cdot 1 + \delta v^* \geq eu^{\lambda}_{v^*,i}(\bar{x}) \tag{A.19}$$

holds for $\delta v^* > 0$, what remains is the case when $\delta v^* = 0$. When $\delta v^* = 0$, $eu^{\lambda}_{v^*,i}(x') \geq \frac{1}{2}$. For $\bar{x}$, $eu^{\lambda}_{v^*,i}(\bar{x}) \leq \frac{1}{2} p(\bar{x})_{2|3} < \frac{1}{2}$ where the last inequality follows from $p(\bar{x})_{2|3} \in (0, 1)$. Thus $eu^{\lambda}_{v^*,i}(x') > eu^{\lambda}_{v^*,i}(\bar{x})$. In summary, $x_i^* \leq \frac{1-\delta v^*}{2}$ yields a contradiction, just as $x_i^* > \frac{1-\delta v^*}{2}$ along with $x_j^* > 0$ and $x_k^* > 0$ does. Hence either $x_j^* = 0$ or $x_k^* = 0$ (or both). □

Appendix B. Numerical computation of QRE

Here we explain in sufficient detail numerical computation of the QRE. We do so for the version with gambler's fallacy, which is the more sophisticated procedure. Adapting the computation to the QRE without gambler's fallacy is straightforward.

First we compute the space of all permissible allocations X'. We set value of d parameterizing fineness of X' to $d = 0.001$ for the computations underlying Example 3 and to $d = 0.01$ ($d = 0.05$) for the maximum likelihood estimations with $n = 3$ ($n = 5$). This produces space of 501501, 5151 and 10626 distinct allocations respectively.

The computation uses iteration on the continuation values. Given continuation value of all the agents in an even round s, v^e, it calculates proposing and voting probabilities in s, continuation value for the proposing and for the non-proposing agents in an odd round $s-1$ and new even round continuation value $v^{e\prime}$. Iterating on this procedure gives us equilibrium continuation value. We used $v^e = \frac{1}{n}$ as a starting value, experienced no problems achieving convergence and stopped the iterations when $|v^{e\prime} - v^e| \leq 10^{-6}$.

Fix λ and ϕ. Denote probabilities of recognition by $\rho^o = \frac{1}{n}$ in odd rounds, by $\rho^e = \frac{\phi}{\phi n - 1}$ in even rounds for the previously non-proposing agents and by $\rho^{ep} = \frac{\phi - 1}{\phi n - 1}$ in even rounds for the previously proposing player.

Take even round s with continuation value v^e. Then $i \in N$ will vote for $x \in X'$ with probability

$$p^{\lambda,\phi}_{v^e,i}(x) = \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^e} \tag{B.1}$$

which allows us to calculate probability of $x \in X'$ being accepted, $p^{\lambda,\phi}_{v^e}(x)$, in a natural way. Given $p^{\lambda,\phi}_{v^e}(x)$ player $i \in N$ will propose $x \in X'$ with probability

$$r^{\lambda,\phi}_{v^e,i}(x) = \frac{\exp \lambda \left(x_i p^{\lambda,\phi}_{v^e}(x) + \delta v^e (1 - p^{\lambda,\phi}_{v^e}(x))\right)}{\sum_{z \in X'} \exp \lambda \left(z_i p^{\lambda,\phi}_{v^e}(z) + \delta v^e (1 - p^{\lambda,\phi}_{v^e}(z))\right)} \tag{B.2}$$

which is easy to compute.

Given $p^{\lambda,\phi}_{v^e}(x)$ and $r^{\lambda,\phi}_{v^e,i}(x)$, odd round $s-1$ continuation values v^{op} and v^o for the agents proposing and not proposing in $s-1$ can be calculated as

$$v_i^{\pi} = \sum_{j \in N} (\mathbb{I}_j^{\pi} \rho^{ep} + (1 - \mathbb{I}_j^{\pi}) \rho^e) \sum_{x \in X'} r^{\lambda,\phi}_{v^e,j}(x) \left[x_i p^{\lambda,\phi}_{v^e}(x) + \delta v^e (1 - p^{\lambda,\phi}_{v^e}(x))\right] \tag{B.3}$$

where the $\mathbb{I}_j^{\pi}$ indicator function equals unity if and only if $j = \pi$. Index $\pi \in N$ denotes player proposing in $s-1$ so that $v^o = v_i^{\pi}$ by setting $i \neq \pi$ (it is immediate $v_i^{\pi} = v_j^{\pi}$ if $i \neq \pi$ and $j \neq \pi$) and $v^{op} = v_i^{\pi}$ by setting $i = \pi$.

Given the odd period continuation values v^o and v^{op} player $i \in N$ votes for $x \in X'$ proposed in $s-1$

$$\begin{aligned} p^{\lambda,\phi}_{v^{op},i}(x) &= \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^{op}} \\ p^{\lambda,\phi}_{v^{o},i}(x) &= \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^{o}} \end{aligned} \tag{B.4}$$

if she has and has not proposed x, respectively. Using $p^{\lambda,\phi}_{v^{op},i}(x)$ and $p^{\lambda,\phi}_{v^{o},i}(x)$ we can again calculate probability of x being approved, $p^{\lambda,\phi}_{v^{o}}(x)$, and probability of i proposing x as

$$r^{\lambda,\phi}_{v^{o},i}(x) = \frac{\exp\lambda\,(x_i\, p^{\lambda,\phi}_{v^{o}}(x) + \delta\, v^{op}\,(1 - p^{\lambda,\phi}_{v^{o}}(x)))}{\sum_{z\in X'} \exp\lambda\,(z_i\, p^{\lambda,\phi}_{v^{o}}(z) + \delta\, v^{op}\,(1 - p^{\lambda,\phi}_{v^{o}}(z)))}. \tag{B.5}$$

Having $p^{\lambda,\phi}_{v^{o}}(x)$ and $r^{\lambda,\phi}_{v^{o},i}(x)$ we can calculate $v^{e\prime}$ from

$$v^{e\prime} = \sum_{j\in N} \rho^o \sum_{x\in X'} r^{\lambda,\phi}_{v^{o},j}(x) \left[x_i\, p^{\lambda,\phi}_{v^{o}}(x) + \delta(\mathbb{I}^{\pi}_{i}\, v^{op} + (1 - \mathbb{I}^{\pi}_{i})\, v^{o})(1 - p^{\lambda,\phi}_{v^{o}}(x))\right] \tag{B.6}$$

(it is immediate $v^{e\prime}$ does not depend on choice of i and π) concluding one step of the continuation value iteration.

Appendix C. Summary statistics of experimental data

Table C.1
Experimental results from legislative bargaining experiments: Proposer's share.

Experiment	1	2	3	4	5	6
N	5	3	3	3	5	3
δ	4/5	1	1	1/2	1	1
pie	\$25	\$30	\$30	\$30	\$60	£3.60
SSPE	0.68	0.67	0.67	0.83	0.60	0.67
Proposer's share						
average	0.34	0.55	0.52	0.54	0.41	0.49
s.d.	(0.11)	(0.12)	(0.10)	(0.13)	(0.11)	(0.12)
obs.	275	330	411	420	450	480

Note: Experiment 1 from Frechette et al. (2003), experiments 2 through 4 from Frechette et al. (2005b), experiment 5 from Frechette et al. (2005a), experiment 6 from Drouvelis et al. (2010). Based on round 1 proposals. Number of rounds is 10 (experiments 6), 15 (experiment 1) and 20 (experiments 2–5). Difference between experiment 2 and 3 are different voting shares (with identical equilibrium prediction).

Table C.2
Experimental results from legislative bargaining experiments: Probability of voting for proposed allocation given offered share.

Experiment	1	2	3	4	5	6
Offered share						
[0.0, 0.1)	0.00	0.02	0.00	0.02	0.00	0.06
s.d.	(0.00)	(0.15)	(0.00)	(0.14)	(0.00)	(0.24)
obs.	57	92	109	108	164	111
[0.1, 0.2)	0.55*	0.09	0.00	0.14*	0.08	0.10
s.d.	(0.50)	(0.30)	(0.00)	(0.36)	(0.28)	(0.31)
obs.	47	11	9	14	37	20
[0.2, 0.3)	0.96	0.06	0.27	0.46	0.73*	0.17
s.d.	(0.21)	(0.24)	(0.47)	(0.51)	(0.45)	(0.38)
obs.	112	18	11	24	95	29
[0.3, 0.4)	1.00	0.63*	0.59*	0.95	0.94	0.73*
s.d.	(0.00)	(0.49)	(0.50)	(0.23)	(0.24)	(0.44)
obs.	33	43	58	74	96	83
[0.4, 0.5)	1.00	0.84	0.81	0.96	1.00	0.83
s.d.	(0.00)	(0.37)	(0.39)	(0.20)	(0.00)	(0.38)
obs.	26	57	64	71	28	80
[0.5, 0.6)		1.00	0.98	1.00	1.00	0.96
s.d.		(0.00)	(0.12)	(0.00)	(0.00)	(0.20)
obs.		73	129	86	23	139
[0.6, 0.7)		1.00	1.00	1.00	1.00	1.00
s.d.		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
obs.		27	29	37	6	16
[0.7, 0.8)		1.00	1.00	1.00		
s.d.		(0.00)	(0.00)	(0.00)		
obs.		2	1	1		

Table C.2 (*continued*)

Experiment	1	2	3	4	5	6
[0.8, 0.9]		1.00	0.00	1.00		1.00
s.d.		(0.00)	(0.00)	(0.00)		(0.00)
obs.		2	1	1		1
[0.9, 1.0]		1.00		1.00	1.00	1.00
s.d.		(0.00)		(0.00)	(0.00)	(0.00)
obs.		5		4	1	1

Note: * denotes interval with benchmark SSPE prediction. Based on response to selected round 1 proposals.

Appendix D. Estimation methodology

For each experiment in our data, we structurally estimate, using maximum likelihood techniques, the best-fitting λ for the QRE model and the best-fitting $\{\lambda, \phi\}$ pair for the QGF model. We outline here the estimation procedure for the QGF model.[38] For a given value of λ and ϕ, we first numerically calculate the equilibrium probability of $i \in N$ approving $x \in X'$, $p^{\lambda,\phi}_{v*,i}(x) = p^{\lambda,\phi}_{v*,i}(x_i)$, and the equilibrium probability of proposing $x \in X'$, $r^{\lambda,\phi}_{v*,i}(x)$. We set the parameter determining the coarseness of the space of allocations, X', to $d = 0.01$ ($d = 0.05$) for experiments with $n = 3$ ($n = 5$). The resulting X' contains 5151 (10626) distinct permissible allocations. We round the experimental data to multiples of d in order to match the values in X'.

Recall that each observation s of proposing behavior, $x^{p,s} \in X'$, consists of a triplet (quintuplet) of shares, $x^{p,s} = \{x^{p,s}_1, \ldots, x^{p,s}_n\}$ and each observation of voting behavior consists of a share $y^{v,s}$ offered to $i \in N$ and her vote $u^{v,s} \in \{0, 1\}$. Denote by S_e the number of observations in experiment $e \in \{1, \ldots, 6\}$. For given e, the corresponding dataset is then $(x^{p,s}, y^{v,s}, u^{v,s})_{s \in S_e}$ where we order entries in $x^{p,s}$ such that the first entry corresponds to the share the proposer allocates to herself.

For a given s, λ and ϕ, the likelihood of observing $(x^{p,s}, y^{v,s}, u^{v,s})$ in the QGF is:

$$L(s, \lambda, \phi) = r^{\lambda,\phi}_{v*,1}(x^{p,s}) + u^{v,s}\, p^{\lambda,\phi}_{v*,i}(y^{v,s}) + (1 - u^{v,s})\,(1 - p^{\lambda,\phi}_{v*,i}(y^{v,s})) \tag{D.1}$$

so that the log-likelihood of observing the whole dataset S_e is

$$l(\lambda, \phi) = \sum_{s \in S_e} \ln\left(L(s, \lambda, \phi)\right). \tag{D.2}$$

The maximum likelihood estimate of $\{\lambda, \phi\}$ is then $\{\hat{\lambda}, \hat{\phi}\} = \arg\max_{\lambda,\phi} l(\lambda, \phi)$. We identified $\{\hat{\lambda}, \hat{\phi}\}$ searching over a grid of $\lambda \in \{0, 0.1, \ldots\}$ and $\phi \in \{1, 2, \ldots\}$.

Appendix E. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.geb.2016.06.008.

References

Agranov, M., Tergiman, C., 2014. Communication in multilateral bargaining. J. Public Econ. 118, 75–85.
Anderson, S.P., Goeree, J.K., Holt, C.A., 1998. Rent seeking with bounded rationality: an analysis of the all-pay auction. J. Polit. Economy 106, 828–853.
Anderson, S.P., Goeree, J.K., Holt, C.A., 2001. Minimum-effort coordination games: stochastic potential and logit equilibrium. Games Econ. Behav. 34, 177–199.
Ayton, P., Fischer, I., 2004. The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? Mem. Cogn. 32, 1369–1378.
Banks, J.S., Duggan, J., 2000. A bargaining model of collective choice. Amer. Polit. Sci. Rev. 94, 73–88.
Bar-Hillel, M., Wagenaar, W.A., 1991. The perception of randomness. Adv. Appl. Math. 12, 428–454.
Baron, D.P., Ferejohn, J.A., 1989. Bargaining in legislatures. Amer. Polit. Sci. Rev. 83, 1181–1206.
Baron, D.P., Kalai, E., 1993. The simplest equilibrium of a majority-rule division game. J. Econ. Theory 61, 290–301.
Binmore, K., Shaked, A., Sutton, J., 1985. Testing noncooperative bargaining theory: a preliminary study. Amer. Econ. Rev. 75, 1178–1180.
Bolton, G.E., Ockenfels, A., 1998. Strategy and equity: an ERC-analysis of the Güth–van Damme game. J. Math. Psychol. 42, 215–226.
Budescu, D.V., Rapoport, A., 1994. Subjective randomization in one- and two-person games. J. Behav. Decis. Mak. 7, 261–278.
Camerer, C.F., 1987. Do biases in probability judgment matter in markets? Experimental evidence. Amer. Econ. Rev. 77, 981–997.
Camerer, C.F., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press, Princeton, NJ.
Camerer, C.F., Ho, T.H., 1994. Violations of the betweenness axiom and nonlinearity in probability. J. Risk Uncertainty 8, 167–196.
Clotfelter, C.T., Cook, P.J., 1993. Notes: the 'gambler's fallacy' in lottery play. Manage. Sci. 39, 1521–1525.
Diermeier, D., 2011. Coalition experiments. In: Druckman, J.N., Green, D.P., Kuklinski, J.H., Lupia, A. (Eds.), Handbook of Experimental Political Science. Cambridge University Press, New York, NY, pp. 399–412.
Diermeier, D., Gailmard, S., 2006. Self-interest, inequality, and entitlement in majoritarian decision-making. Q. J. Polit. Sci. 1, 327–350.

[38] The estimation procedure for the QRE model is similar and simpler. Essentially, dropping ϕ from the notation of this section describes the estimation methodology for the QRE model.

Diermeier, D., Morton, R., 2005. Experiments in majoritarian bargaining. In: Austen-Smith, D., Duggan, J. (Eds.), Social Choice and Strategic Decisions. Springer, Berlin, pp. 201–226.
Drouvelis, M., Montero, M., Sefton, M., 2010. Gaining power through enlargement: strategic foundations and experimental evidence. Games Econ. Behav. 69, 274–292.
Eraslan, H., 2002. Uniqueness of stationary equilibrium payoffs in the Baron–Ferejohn model. J. Econ. Theory 103, 11–30.
Frechette, G.R., 2009. Learning in a multilateral bargaining experiment. J. Econ. 153, 183–195.
Frechette, G.R., Kagel, J.H., Lehrer, S.F., 2003. Bargaining in legislatures: an experimental investigation of open versus closed amendment rules. Amer. Polit. Sci. Rev. 97, 221–232.
Frechette, G.R., Kagel, J.H., Morelli, M., 2005a. Behavioral identification in coalitional bargaining: an experimental analysis of demand bargaining and alternating offers. Econometrica 73, 1893–1937.
Frechette, G.R., Kagel, J.H., Morelli, M., 2005b. Nominal bargaining power, selection protocol, and discounting in legislative bargaining. J. Public Econ. 89, 1497–1517.
Goeree, J.K., Holt, C.A., 2000. Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. Europ. Econ. Rev. 44, 1079–1089.
Goeree, J.K., Holt, C.A., 2001. Ten little treasures of game theory and ten intuitive contradictions. Amer. Econ. Rev. 91, 1402–1422.
Goeree, J.K., Holt, C.A., Palfrey, T.R., 2002. Quantal response equilibrium and overbidding in private-value auctions. J. Econ. Theory 104, 247–272.
Grether, D.M., 1980. Bayes rule as a descriptive model: the representativeness heuristic. Quart. J. Econ. 95, 537–557.
Grether, D.M., 1992. Testing Bayes rule and the representativeness heuristic: some experimental evidence. J. Econ. Behav. Organ. 17, 31–57.
Güth, W., van Damme, E., 1998. Information, strategic behavior, and fairness in ultimatum bargaining: an experimental study. J. Math. Psychol. 42, 227–247.
Harless, D.W., Camerer, C.F., 1994. The predictive utility of generalized expected utility theories. Econometrica 62, 1251–1289.
Harrington, J.E., 1990. The power of the proposal maker in a model of endogenous agenda formation. Public Choice 64, 1–20.
Hey, J.D., Orme, C., 1994. Investigating generalizations of expected utility theory using experimental data. Econometrica 62, 1291–1326.
Kagel, J., Sung, H., Winter, E., 2010. Veto power in committees: an experimental study. Exper. Econ. 13, 167–188.
Luce, D.R., 1959. Individual Choice Behavior: A Theoretical Analysis. Wiley, New York, NY.
McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10, 6–38.
McKelvey, R.D., Palfrey, T.R., 1998. Quantal response equilibria for extensive form games. Exper. Econ. 1, 9–41.
Miller, L., Vanberg, C., 2013. Decision costs in legislative bargaining: an experimental analysis. Public Choice 155, 373–394.
Miller, L., Vanberg, C., 2015. Group size and decision rules in legislative bargaining. Europ. J. Polit. Economy 37, 288–302.
Montero, M., 2007. Inequity aversion may increase inequity. Econ. J. 117, C192–C204.
Ochs, J., Roth, A.E., 1989. An experimental study of sequential bargaining. Amer. Econ. Rev. 79, 355–384.
Rabin, M., 2002. Inference by believers in the law of small numbers. Quart. J. Econ. 117, 775–816.
Rabin, M., Vayanos, D., 2010. The gambler's and hot-hand fallacies: theory and applications. Rev. Econ. Stud. 77, 730–778.
Rapoport, A., Budescu, D.V., 1992. Generation of random series in two-person strictly competitive games. J. Exp. Psychol. Gen. 121, 352–363.
Rapoport, A., Budescu, D.V., 1997. Randomization in individual choice behavior. Psychol. Rev. 104, 603–617.
Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. Econometrica 50, 97–109.
Stahl, I., 1972. Bargaining theory. Ph.D. thesis. Stockholm School of Economics.
Sun, Y., Wang, H., 2010. Gambler's fallacy, hot hand belief, and the time of patterns. Judgm. Decis. Mak. 5, 124–132.
Terrell, D., 1994. A test of the gambler's Gallacy: evidence from pari-mutuel games. J. Risk Uncertainty 8, 309–317.
Tversky, A., Kahneman, D., 1971. Belief in the law of small numbers. Psychol. Bull. 76, 105–110.
Walker, M., Wooders, J., 2001. Minimax play at Wimbledon. Amer. Econ. Rev. 91, 1521–1538.
Yi, K.O., 2005. Quantal-response equilibrium models of the ultimatum bargaining game. Games Econ. Behav. 51, 324–348.