

Gambler's Fallacy and Imperfect Best Response in Legislative Bargaining*

Salvatore Nunnari
Columbia University
sunnari@columbia.edu

Jan Zápál
IAE-CSIC, Barcelona GSE and CERGE-EI Prague
jan.zapal@iae.csic.es

January 9, 2014

Abstract

We analyze data from the legislative bargaining experiments reported in [Frechette, Kagel, and Lehrer \(2003\)](#), [Frechette, Kagel, and Morelli \(2005*a,b*\)](#), and [Drouvelis, Montero, and Sefton \(2010\)](#) in order to investigate the extent to which the observed deviations from the theoretical predictions can be explained by imperfect best response, in combination with different assumptions about correct (QRE) or incorrect beliefs (Quantal-Gambler's Fallacy or QGF) on future proposal power. The basic pattern in the data consists of proposers who allocate resources only within a minimal winning coalition of legislators but do not fully exploit their bargaining power and allocate to their coalition partner(s) a disproportionate share of the pie. The QRE model fits this pattern reasonably well. Incorporating history-dependent beliefs about the future distribution of proposal power into the QRE model 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 minimal winning coalitions and similar voting behavior.

JEL Classification: D72, C92, C70

Keywords: Legislative Bargaining; Experiments; Quantal Response; Gambler's Fallacy.

* We are grateful to Guillaume Frechette and Maria Montero for supplying the original data from the experiments in [Frechette, Kagel, and Lehrer \(2003\)](#), [Frechette, Kagel, and Morelli \(2005*a,b*\)](#), and [Drouvelis, Montero, and Sefton \(2010\)](#). We also thank Marina Agranov, Colin Camerer, Guillaume Frechette, 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 and ASSET 2013 Bilbao Annual Meeting for comments and feedback.

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 bargain 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 subgame perfect equilibria of this bargaining game, an agreement is reached immediately, with 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, Kagel, and Lehrer, 2003](#); [Frechette, Kagel, and Morelli, 2005a,b](#); [Diermeier and Morton, 2005](#); [Diermeier and Gailmard, 2006](#); [Kagel, Sung, and Winter, 2010](#); [Drouvelis, Montero, and Sefton, 2010](#); [Miller and Vanberg, 2013](#)).

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, accepting or rejecting a proposal only on the grounds of the share offered to themselves, but independently of the distribution of resources to other agents; and 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.

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. [Frechette \(2009\)](#) shows that belief-based learning models can help explain the experimental data from

Frechette, Kagel, and Lehrer (2003). 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 (2013) 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.¹

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. At this point, it is useful to summarize the basic features of the unique symmetric stationary subgame perfect equilibrium (*equilibrium* thereafter) in Baron and Ferejohn (1989). Consider a three-members committee, for simplicity. In the equilibrium, the proposer offers a positive share to one randomly selected legislator, making her just indifferent between accepting and rejecting; he offers nothing to the remaining legislator, and keeps the rest of the budget to herself. We call these allocations *minimum winning*. Legislators vote in favor of the proposed allocation if their share is larger than their *continuation value* and rejects the proposal otherwise. The continuation value is the expected utility a legislator gets from inducing the continuation of the bargaining process to a further round of negotiations.

The main contribution of this paper is to present two extensions of the benchmark model that have an independent theoretical interest and generate theoretical predictions in line with the experimental evidence. 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 (where the resemblance is measured by the mean or another summary statistic). In its most extreme form, agents who are prey of this cognitive fallacy and observe a series of

¹ The pattern in the data cannot be explained by bargaining theories with risk aversion (Harrington, 1990), present bias (note available from the authors), or other-regarding preferences (Montero, 2007). All these modifications of the basic framework increase proposal power, pushing the theoretical predictions further away from observed behavior.

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, 1997, 1992; 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.²

We introduce the gambler's fallacy in the alternating offer bargaining setup 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 are more likely to be recognized in the next bargaining round (or, at least, that they believe this is the case) which increases their continuation value. These beliefs do not affect the number of legislators who receive a positive allocation, but affects 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.

The second modification we propose relaxes the assumption that agents perfectly best-respond to the strategies of other agents and assumes that they make mistakes in choosing their actions. We focus on the logistic version of the Quantal Response Equilibrium (QRE) as defined by McKelvey and Palfrey (1998). The QRE assumes that agents make mistakes in choosing their actions and that they make more expensive mistakes with smaller probability.

Applied to the legislative bargaining setting, imperfect best response increases the incidence of non-minimum winning coalitions and reduces the share the proposer allocates to himself in any coalition. The intuition behind this latter result is that it is no longer optimal for the proposer to bring her coalition partner to indifference between accepting and rejecting her proposal. By doing so, the proposer faces a large probability that her coalition partner

² 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.

rejects the proposed agreement (when agents' best response are, even minimally, imperfect). 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 willing to give up some resources in order to increase the probability of her proposal passing in the current round.³

A second contribution of the paper is to structurally estimate these two models and their combination on the experimental data. We show that the predictions obtained when combining QRE with gambler's fallacy generates the best fit of the data: imperfect best response generates both lower shares to the proposer and a lower incidence of minimum winning coalitions with respect to the benchmark model. 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 experimental data.

Finally, while our main goal is to explain behavior in the laboratory implementation of a specific legislative bargaining protocol, we believe that the model with gambler's fallacy we introduce here is of independent interest and has important implications for real world legislatures. First, as we mentioned above, the gambler's fallacy is a cognitive bias that has been documented also outside the laboratory and also among professionals of the task being studied. This means that, in environments where the history-independent allocation of proposal rights is a good assumption, members of real world legislatures and committees are just as likely as laboratory subjects to be prey of this cognitive fallacy, especially when proposal rights are very valuable.

³ Quantal response has been applied to other bargaining frameworks (for example, to the ultimatum game by [Yi 2005](#)). The differences with the perfect best response case and the underlying intuition are similar to ours. Beyond applying quantal response to a different bargaining environment (multilateral bargaining with alternating offer), the novelty of our approach lies in suggesting this can fit the observed behavior in the lab, in combining QRE with gambler's fallacy, and in structurally estimating these models on laboratory data.

Second, while the Baron-Ferejohn bargaining model has many realistic features and appealing predictions, the key protocol feature that generates the cognitive bias we discuss—the random and independent selection of a proposer—is a stylized assumption not necessarily observed in real world legislatures or committees. In some political environments, the process determining proposal power is *de jure* or *de facto* history-dependent rather than being an independent and random draw in each period. The institutional protocol or the social norm adopted by a committee might prescribe that, following the failure to reach an agreement, a different committee member has the faculty to put forward a new proposal and attempt to form a decisive coalition. For example, in parliamentary systems and premier-presidential systems—where the government needs support by a majority in parliament (Shugart and Carey, 1992)—the head of state usually appoints a party leader, or *formateur*, to lead the formation of a coalition government and, if the first attempt is unsuccessful, he moves on to a different party leader. The model we introduce below can be interpreted as a model of a different bargaining protocol where the ability to set the agenda is history-dependent by institutional design, and agents have rational beliefs on the future distribution of proposal power.

The remainder of the paper proceeds as follows: Section 2 outlines the equilibrium analysis and theoretical predictions for the benchmark model and our proposed modifications. Section 3 describes the data from the laboratory experiments in Frechette, Kagel, and Lehrer (2003), Frechette, Kagel, and Morelli (2005a,b), and Drouvelis, Montero, and Sefton (2010). Section 4 presents the results of the structural estimation, and Section 5 concludes.

2 Theory

2.1 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. The bargaining proceeds over a potentially infinite number of rounds. In each round, one player $i \in N$ is selected to be a proposer, with probability $\rho_i = \frac{1}{n}$, and proposes a division of the pie $x \in X$ among N . The set of permissible allocations is the $n - 1$ dimensional simplex, $X = \{x \in \mathbb{R}_+^n \mid \sum_{i=1}^n x_i \leq 1\}$, 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*.⁴ If the number of the *yes* votes is equal to or larger than $\frac{n+1}{2}$, x passes, the game ends and the agents receive their share as determined by x . The utility player i derives from the 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 symmetric agents use the same strategies and proposers treat their potential coalition partners symmetrically (see Appendix A1 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 $\{2/3, 1/3, 0\}$.*

2.2 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

⁴ 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 provides her with higher utility relative to the continuation of the game. That is, i votes as if he is pivotal.

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 labelled ‘ 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 beliefs about the future recognition probabilities that are history dependent.⁵ For modeling convenience, we assume (as in [Rabin, 2002](#)) that the urn is refilled every two rounds.⁶

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 parametrized 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 converge to the benchmark model as ϕ goes to infinity.⁷

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).

Proposition 2.

In the alternating offer bargaining game where agents are prey to the gambler’s fallacy, there

⁵ 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.

⁶ 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.

⁷ 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.

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

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

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 A1.

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

2.3 Quantal Response Equilibrium

Quantal response posits that agents make errors in choosing which pure strategy to play, despite knowing which strategy is their best-response to the strategies of the other agents. We focus here on the logistic agent Quantal Response Equilibrium (QRE) as defined by McKelvey and Palfrey (1998), retaining the assumption of symmetry and stationarity. We refer to this simply as QRE in the remainder.

More explicitly, as McKelvey and Palfrey (1998), we assume that the play for 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, we assume that agents use stage undominated voting strategies, that is, player i votes as if he is pivotal.⁸

⁸ 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, $X' = \{x \in \mathbb{R}_+^n \mid \sum_i^n x_i = 1 \wedge \text{mod}(x_i, d) = 0 \forall i \in N\}$ 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 .⁹ In the logit version of quantal response behavior, 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 the legislative bargaining as of 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:

$$p_{v,i}^\lambda(x) = \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v} \quad (2)$$

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

Given the voting behavior described by $p_{v,i}^\lambda$, the probability $x \in X'$ is accepted, $p_v^\lambda(x)$, 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

$$r_{v,i}^\lambda(x) = \frac{\exp \lambda (x_i p_v^\lambda(x) + \delta v (1 - p_v^\lambda(x)))}{\sum_{z \in X'} \exp \lambda (z_i p_v^\lambda(z) + \delta v (1 - p_v^\lambda(z)))} \quad (3)$$

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

Denote by $\sigma^\lambda(v) = (r_{v,1}^\lambda, \dots, r_{v,n}^\lambda, p_{v,1}^\lambda, \dots, p_{v,n}^\lambda)$ 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

⁹ 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.

a slight abuse of notation, we denote this mapping by $v' = \sigma^\lambda(v)$. The fixed point of σ^λ constitutes a QRE.

Proposition 3.

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

Proof. See Appendix A1.

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 imply legislators are more demanding when evaluating a proposed allocation of resources. In spite of this, we illustrate below that imperfect best response can easily generate a smaller average proposer’s share than the one from the benchmark model. However, little can be said in general about the distribution of proposed allocations, both in terms of their nature (resembling minimum winning coalitions or not) and in terms of share of the pie proposer allocates to herself.^{10,11}

Example 3. *Figure 1 shows the results of the numerical calculation of the QRE for $n = 3$ and $\delta = 1$ (see Appendix A2 for the details of the procedure). Figure 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.¹² Figure 1b shows the probability an agent votes in favor of the proposal, when it assigns him x_i .*


¹⁰ One can prove that there exists a unique QRE of the alternating offer bargaining game 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 (A3) in the proof of proposition 3 as $\sigma^\lambda(v) = \delta v + (\frac{1}{n} - \delta v) \sum_{x \in X'} r_{v,1}^\lambda(x) p_v^\lambda(x)$, 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).

¹¹ We have systematically (numerically) searched for counter example to uniqueness with no success. For all combinations of $\delta \in \{\frac{1}{6}, \frac{2}{6}, \frac{3}{6}, \frac{4}{6}, \frac{5}{6}, \frac{6}{6}\}$ 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 Supplementary Material available at http://www.columbia.edu/~sn2562/qregf_supplementary.pdf for the results of this exercise.

¹² 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 $\frac{n-1}{2}$.


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

(a) Proposing Behavior



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 vertical solid line is the benchmark SSPE prediction.

(b) Voting Behavior



Note: The lines represent the probability of voting *yes* given the offered share 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 $\lambda = 36$.

2.4 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 $\frac{1}{2}$ (since the legislator receives the same utility from accepting and rejecting, and the QRE strategies assign the same probability to payoff-equivalent actions). This makes it very risky for the proposing player 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.^{13,14}

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. On the other hand, the QRE, at least for moderate values of λ , predicts a non-negligible share of proposals that assign significant resources to a larger coalition of agents.

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.¹⁵ 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 generates a larger continuation value for the agents who are not proposing, which makes

¹³ Assuming that, in the benchmark model, an indifferent legislator rejects the proposal does *not* produce a similar reduction in proposal power. Even if an indifferent voter rejects with probability one, the proposer can induce her acceptance with probability one by offering $\delta v + \epsilon$. What generates a significant reduction in the proposer's share in the QRE is the fact that the acceptance probability changes smoothly on a non-trivial neighborhood of v .

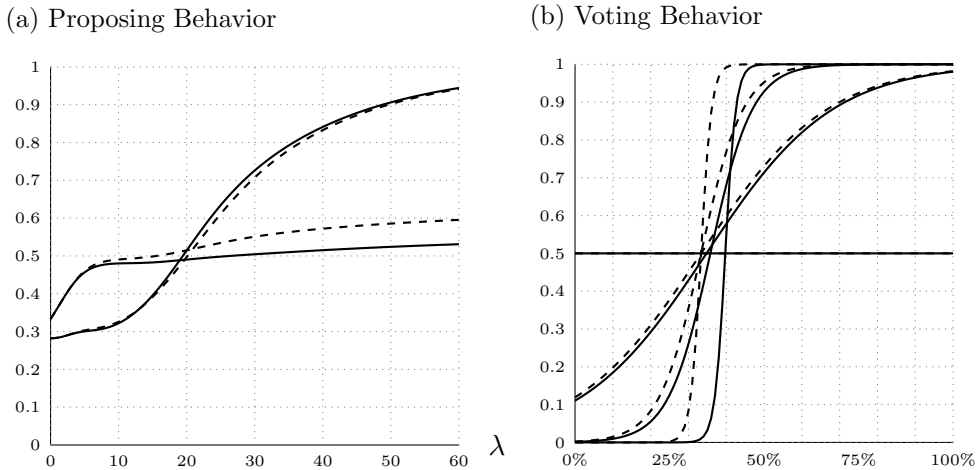
¹⁴ Additionally, imperfect best response changes the proposer's behavior by introducing mistakes in her choice of proposed allocations. Not only is she (optimally) more generous, but she also trembles and proposes, with positive probability, more or less generous allocations both in terms of the shares offered and in terms of the number of coalition partners.

¹⁵ 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.

non-minimum winning coalitions more costly to maintain relative to the minimum winning ones.

Figure 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.¹⁶ 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 for the same proposal.

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



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 $\lambda = 72$.

3 Data

We analyze data from six legislative bargaining experiments testing the [Baron and Ferejohn \(1989\)](#) model in the laboratory. Experiment 1 comes from [Frechette, Kagel, and Lehrer \(2003\)](#) (their ‘closed rule’ treatment), experiments 2 through 4 from [Frechette, Kagel, and Morelli \(2005b\)](#) (their ‘EWES’, ‘UWES’ and ‘EWES with $\delta = \frac{1}{2}$ ’ treatments), experiment 5

¹⁶ 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 Supplementary Material, the continuation values in the QRE and in the even rounds of the QGF are, if not equal, then very similar.

from Frechette, Kagel, and Morelli (2005a) (their ‘Baron and Ferejohn equal weight’ treatment) and experiment 6 from Drouvelis, Montero, and Sefton (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.¹⁷ 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.¹⁸

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.¹⁹ Focusing on all rounds could confound the data with repeated play effects and focusing only on approved allocations would cut the data size by a factor of three or five, depending on the experimental committee size.^{20,21} 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

¹⁷ Appendix A3 reports detailed information regarding each experiment.

¹⁸ 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.

¹⁹ 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.

²⁰ 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, Kagel, and Lehrer, 2003), all rounds proposals (Frechette, Kagel, and Morelli, 2005b) or all rounds minimum winning coalition proposals (Frechette, Kagel, and Morelli, 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 Supplementary Material and differ little from the results presented below.

²¹ 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).

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.

Figures 3a and 3b show the main features of our data.²² 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.²³

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).²⁴

4 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.²⁵

In Table 1, we show the average share to the proposer (X_{PR}), the quartiles of its distribution, and the incidence of minimum winning coalitions in the 6 experimental datasets, and we compare them with the predictions from the benchmark model (X_{PR}^*)²⁶ and with

²² Appendix A3 includes more detailed information in a non-graphical form.

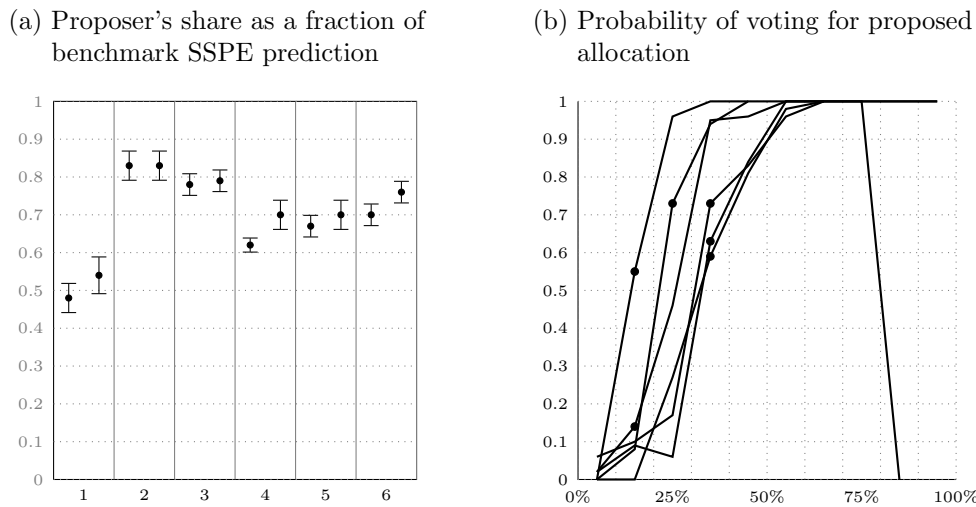
²³ The difference between experienced and unexperienced subjects in the mean proposer’s share in round 1 proposals is significant at the 1% level, using a standard t -test, in experiments 1, 4 and 6. This does not give a clear indication of whether to analyze the (un)experienced data separately or not. For space considerations, we present below results from pooled data. The Supplementary Material includes separate results for inexperienced and experienced subjects data and shows that the results do not depend on this choice.

²⁴ Frechette, Kagel, and Lehrer (2003); Frechette, Kagel, and Morelli (2005a,b) and Drouvelis, Montero, and Sefton (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.

²⁵ See Appendix A4 for the details of the procedure.

²⁶ In the equilibrium of the benchmark model, with perfect best response and no gambler’s fallacy, all

Figure 3: Experimental Results from Legislative Bargaining Experiments



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 A3 for details.

the predictions from the QRE and QGF models (using the estimated parameters).²⁷

proposals give the proposer exactly X_{PR}^* and distribute benefits only to a minimum winning coalitions. We omit this from Table 1 in the interest of space.

²⁷ While here we focus on the most important summary statistics, the Supplementary Material 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.

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_{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
	QRE					
$\hat{\lambda}$	20.2	22.1	23.4	10.6	33.5	18.4
$Avg(\hat{X}_{PR})$.441	.527	.532	.627	.400	.512
$Q1(\hat{X}_{PR})$.350	.490	.500	.540	.350	.460
$Q2(\hat{X}_{PR})$.450	.550	.550	.640	.400	.540
$Q3(\hat{X}_{PR})$.500	.600	.600	.730	.450	.600
% MWC	.547	.605	.635	.424	.773	.516
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					
$\hat{\lambda}$	21.7	21.7	22.5	13.3	33.8	18.8
$\hat{\phi}$	1	3	1	1	1	1
$Avg(\hat{X}_{PR})$.421	.518	.497	.594	.387	.491
$Q1(\hat{X}_{PR})$.350	.480	.460	.530	.350	.440
$Q2(\hat{X}_{PR})$.450	.540	.510	.610	.400	.510
$Q3(\hat{X}_{PR})$.500	.590	.560	.680	.450	.570
% MWC	.553	.601	.637	.449	.794	.541
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, Kagel, and Lehrer \(2003\)](#); Experiments 2 through 4 from [Frechette, Kagel, and Morelli \(2005b\)](#); Experiment 5 from [Frechette, Kagel, and Morelli \(2005a\)](#); Experiment 6 from [Drouvelis, Montero, and Sefton \(2010\)](#). X_{PR} , X_{PR}^* , and \hat{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 minimal 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.

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. 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, $\hat{\phi} < \infty$).

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 1, 3 and 5, only the incidence of minimum winning coalitions 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.

5 Conclusions

The interesting empirical results reported by [Frechette, Kagel, and Lehrer \(2003\)](#), [Frechette, Kagel, and Morelli \(2005a,b\)](#), and [Drouvelis, Montero, and Sefton \(2010\)](#) 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.

References

- Agranov, Marina, and Chloe Tergiman. 2013. "The Effect of Negotiations on Bargaining in Legislatures: An Experimental Study." *mimeo* .
- Banks, Jeffrey S., and John Duggan. 2000. "A Bargaining Model of Collective Choice." *American Political Science Review* 94(1): 73–88.
- Bar-Hillel, Maya, and Willem A. Wagenaar. 1991. "The Perception of Randomness." *Advances in Applied Mathematics* 12(4): 428–454.
- Baron, David P., and Ehud Kalai. 1993. "The Simplest Equilibrium of a Majority-Rule Division Game." *Journal of Economic Theory* 61(2): 290–301.
- Baron, David P., and John A. Ferejohn. 1989. "Bargaining in Legislatures." *American Political Science Review* 83(4): 1181–1206.
- Budescu, David V., and Amnon Rapoport. 1994. "Subjective Randomization in One- and Two-Person Games." *Journal of Behavioral Decision Making* 7(4): 261–278.
- Camerer, Colin F. 1987. "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence." *American Economic Review* 77(5): 981–997.
- Clotfelter, Charles T., and Philip J. Cook. 1993. "Notes: The 'Gambler's Fallacy' in Lottery Play." *Management Science* 39(12): 1521–1525.
- Diermeier, Daniel. 2011. "Coalition Experiments." In *Handbook of Experimental Political Science*, ed. James N. Druckman, Donald P. Green, James H. Kuklinski, and Arthur Lupia. New York, NY: Cambridge University Press pp. 399–412.
- Diermeier, Daniel, and Rebecca Morton. 2005. "Experiments in Majoritarian Bargaining." In *Social Choice and Strategic Decisions*, ed. David Austen-Smith, and John Duggan. Berlin: Springer pp. 201–226.
- Diermeier, Daniel, and Sean Gailmard. 2006. "Self-Interest, Inequality, and Entitlement in Majoritarian Decision-Making." *Quarterly Journal of Political Science* 1(4): 327–350.
- Drouvelis, Michalis, Maria Montero, and Martin Sefton. 2010. "Gaining Power Through Enlargement: Strategic Foundations and Experimental Evidence." *Games and Economic Behavior* 69(2): 274–292.

- Eraslan, Hülya. 2002. “Uniqueness of Stationary Equilibrium Payoffs in the Baron-Ferejohn Model.” *Journal of Economic Theory* 103(1): 11–30.
- Frechette, Guillaume R. 2009. “Learning in a Multilateral Bargaining Experiment.” *Journal of Econometrics* 153(2): 183–195.
- Frechette, Guillaume R., John H. Kagel, and Massimo Morelli. 2005a. “Behavioral Identification in Coalitional Bargaining: An Experimental Analysis of Demand Bargaining and Alternating Offers.” *Econometrica* 73(6): 1893–1937.
- Frechette, Guillaume R., John H. Kagel, and Massimo Morelli. 2005b. “Nominal Bargaining Power, Selection Protocol, and Discounting in Legislative Bargaining.” *Journal of Public Economics* 89(8): 1497–1517.
- Frechette, Guillaume R., John H. Kagel, and Steven F. Lehrer. 2003. “Bargaining in Legislatures: An Experimental Investigation of Open versus Closed Amendment Rules.” *American Political Science Review* 97(2): 221–232.
- Grether, David M. 1980. “Bayes Rule as a Descriptive Model: The Representativeness Heuristic.” *Quarterly Journal of Economics* 95(3): 537–557.
- Grether, David M. 1992. “Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence.” *Journal of Economic Behavior & Organization* 17(1): 31–57.
- Harrington, Joseph E. 1990. “The Power of the Proposal Maker in a Model of Endogenous Agenda Formation.” *Public Choice* 64(1): 1–20.
- Kagel, John, Hankyoung Sung, and Eyal Winter. 2010. “Veto Power in Committees: An Experimental Study.” *Experimental Economics* 13(2): 167–188.
- McKelvey, Richard, and Thomas R. Palfrey. 1998. “Quantal Response Equilibria for Extensive Form Games.” *Experimental Economics* 1(1): 9–41.
- Miller, Luis, and Christoph Vanberg. 2013. “Decision Costs in Legislative Bargaining: An Experimental Analysis.” *Public Choice* 155(3-4): 373–394.
- Montero, Maria. 2007. “Inequity Aversion May Increase Inequity.” *Economic Journal* 117(519): C192–C204.

- Rabin, Matthew. 2002. “Inference by Believers in the Law of Small Numbers.” *Quarterly Journal of Economics* 117(3): 775–816.
- Rabin, Matthew, and Dimitri Vayanos. 2010. “The Gambler’s and Hot-Hand Fallacies: Theory and Applications.” *Review of Economic Studies* 77(2): 730–778.
- Rapoport, Amnon, and David V. Budescu. 1992. “Generation of Random Series in Two-Person Strictly Competitive Games.” *Journal of Experimental Psychology: General* 121(3): 352–363.
- Rapoport, Amnon, and David V. Budescu. 1997. “Randomization in Individual Choice Behavior.” *Psychological Review* 104(3): 603–617.
- Shugart, Matthew Soberg, and John M. Carey. 1992. *Presidents and Assemblies: Constitutional Design and Electoral Dynamics*. New York, NY: Cambridge University Press.
- Terrell, Dek. 1994. “A Test of the Gambler’s Gallacy: Evidence from Pari-Mutuel Games.” *Journal of Risk and Uncertainty* 8(3): 309–317.
- Tversky, Amos, and Daniel Kahneman. 1971. “Belief in the Law of Small Numbers.” *Psychological Bulletin* 76(2): 105–110.
- Walker, Mark, and John Wooders. 2001. “Minimax Play at Wimbledon.” *American Economic Review* 91(5): 1521–1538.
- Yi, Kang-Oh. 2005. “Quantal-response Equilibrium Models of the Ultimatum Bargaining Game.” *Games and Economic Behavior* 51(2): 324–348.

A1 Proofs

SSPE definition

A proposal strategy of $i \in N$ is $r_i : X \rightarrow \Delta(X)$, where $\Delta(X)$ is the set of distributions defined on X ; and a voting strategy of $i \in N$ is $p_i : X \rightarrow \Delta(\{\text{accept}, \text{reject}\})$. 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-labelling 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.

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

$$\begin{aligned} v^e &= \frac{\phi}{\phi n} \left(1 - \frac{n-1}{2} x^o \right) + \frac{\phi n - \phi}{\phi n} \left(\frac{1}{2} x^o \right) \\ v^o &= \frac{\phi}{\phi n - 1} \left(1 - \frac{n-1}{2} x^e \right) + \frac{\phi n - 1 - \phi}{\phi n - 1} \left(\frac{1}{2} x^e \right) \\ v^{op} &= \frac{\phi - 1}{\phi n - 1} \left(1 - \frac{n-1}{2} x^e \right) + \frac{\phi n - 1 - (\phi - 1)}{\phi n - 1} \left(\frac{1}{2} x^e \right) \end{aligned} \quad (\text{A1})$$

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

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 behaviour 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 $\phi \rightarrow \infty$. \square

Proof of Proposition 3

Proof. Take σ^λ as defined in the text and restrict its domain to $[0, 1]$. It is easy to see $\sigma^\lambda : [0, 1] \rightarrow [0, 1]$. 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 $\{\sigma^\lambda(v^*), \sigma^\lambda(v^*), \dots\}$ 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 $v^* = \sigma^\lambda(v^*)$, first notice

$$\sigma^\lambda(v) = \sum_{j \in N} \frac{1}{n} \sum_{x \in X'} r_{v,j}^\lambda(x) \left[x_i p_v^\lambda(x) + \delta v (1 - p_v^\lambda(x)) \right] \quad (\text{A2})$$

for any $i \in N$.

In order to proceed we need extra notation. Take arbitrary allocation $x \in X'$, $x = \{x_1, \dots, x_n\}$, and corresponding $p_v^\lambda(x)$ and $r_{v,1}^\lambda(x)$. Define m -time circular shift operator $\sigma^m(i)$ as $\sigma^m(i) = i + m$ modulo n for $m \in \{0, \dots, n-1\}$. Applied to x , let $x_{\sigma^m} = \{x_{\sigma^m(1)}, \dots, x_{\sigma^m(n)}\}$ ordering the entries by their new index. For example, $x_{\sigma^0} = x$ and using the original indexes $x_{\sigma^1} = \{x_n, x_1, \dots, x_{n-1}\}$. Due to symmetry $p_v^\lambda(x_{\sigma^m}) = p_v^\lambda(x_{\sigma^0})$ for any $m \in \{0, \dots, n-1\}$. Denoting by $r_{v,\sigma^m(1)}^\lambda(x_{\sigma^m})$ probability of player $\sigma^m(1)$ proposing x_{σ^m} we further have $r_{v,\sigma^m(1)}^\lambda(x_{\sigma^m}) = r_{v,1}^\lambda(x_{\sigma^0})$ 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 (A2) as

$$\sigma^\lambda(v) = \sum_{x \in X'} \frac{1}{n} r_{v,1}^\lambda(x) \left[p_v^\lambda(x) + n \delta v (1 - p_v^\lambda(x)) \right]. \quad (\text{A3})$$

For $n\delta v = 1$ this implies $\sigma^\lambda(v) = \frac{1}{n} = \delta v$ so that for $\delta = 1$, $v = \frac{1}{n}$ is fixed point of σ^λ . For $n\delta v > 1$ (A3) implies $\sigma^\lambda(v) < \delta v$ so that σ^λ cannot have fixed point when $v > \frac{1}{\delta n} \geq \frac{1}{n}$. For $n\delta v < 1$ (A3) implies $\sigma^\lambda(v) < \frac{1}{n}$ so that any fixed point of σ^λ has to occur for $v \leq \frac{1}{n}$. \square

A2 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 $v^{e'}$. 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 $|v^{e'} - v^e| \leq 10^{-6}$.

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

$$p_{v^e, i}^{\lambda, \phi}(x) = \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^e} \quad (\text{A4})$$

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

$$r_{v^e, i}^{\lambda, \phi}(x) = \frac{\exp \lambda (x_i p_{v^e}^{\lambda, \phi}(x) + \delta v^e (1 - p_{v^e}^{\lambda, \phi}(x)))}{\sum_{z \in X'} \exp \lambda (z_i p_{v^e}^{\lambda, \phi}(z) + \delta v^e (1 - p_{v^e}^{\lambda, \phi}(z)))} \quad (\text{A5})$$

which is easy to compute.

Given $p_{v^e}^{\lambda, \phi}(x)$ and $r_{v^e, i}^{\lambda, \phi}(x)$, 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

$$v_i^\pi = \sum_{j \in N} (\mathbb{I}_j^\pi \rho^{ep} + (1 - \mathbb{I}_j^\pi) \rho^e) \sum_{x \in X'} r_{v^e, j}^{\lambda, \phi}(x) \left[x_i p_{v^e}^{\lambda, \phi}(x) + \delta v^e (1 - p_{v^e}^{\lambda, \phi}(x)) \right] \quad (\text{A6})$$

where the \mathbb{I}_j^π 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 $s - 1$

$$\begin{aligned} p_{v^{op},i}^{\lambda,\phi}(x) &= \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^{op}} \\ p_{v^o,i}^{\lambda,\phi}(x) &= \frac{\exp \lambda x_i}{\exp \lambda x_i + \exp \lambda \delta v^o} \end{aligned} \tag{A7}$$

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

$$r_{v^o,i}^{\lambda,\phi}(x) = \frac{\exp \lambda (x_i p_{v^o}^{\lambda,\phi}(x) + \delta v^{op} (1 - p_{v^o}^{\lambda,\phi}(x)))}{\sum_{z \in X'} \exp \lambda (z_i p_{v^o}^{\lambda,\phi}(z) + \delta v^{op} (1 - p_{v^o}^{\lambda,\phi}(z)))}. \tag{A8}$$

Having $p_{v^o}^{\lambda,\phi}(x)$ and $r_{v^o,i}^{\lambda,\phi}(x)$ we can calculate $v^{e'}$ from

$$\begin{aligned} v^{e'} &= \\ &\sum_{j \in N} \rho^o \sum_{x \in X'} r_{v^o,j}^{\lambda,\phi}(x) \left[x_j p_{v^o}^{\lambda,\phi}(x) + \delta (\mathbb{I}_j^\pi v^{op} + (1 - \mathbb{I}_j^\pi) v^o) (1 - p_{v^o}^{\lambda,\phi}(x)) \right] \end{aligned} \tag{A9}$$

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

A3 Summary Statistics of Experimental Data

Table A1: 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, Kagel, and Lehrer \(2003\)](#), experiments 2 through 4 from [Frechette, Kagel, and Morelli \(2005b\)](#), experiment 5 from [Frechette, Kagel, and Morelli \(2005a\)](#), experiment 6 from [Drouvelis, Montero, and Sefton \(2010\)](#). 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 A2: 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		
[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.

A4 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.²⁸ For a given value of λ and ϕ , we first numerically calculate the equilibrium probability of $i \in N$ approving $x \in X'$, $p_{v^*,i}^{\lambda,\phi}(x) = p_{v^*,i}^{\lambda,\phi}(x_i)$, and the equilibrium probability of proposing $x \in X'$, $r_{v^*,i}^{\lambda,\phi}(x)$. 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, $x^{p,s} = \{x_1^{p,s}, \dots, x_n^{p,s}\}$ 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 $e \in \{1, \dots, 6\}$. 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:

$$L(s, \lambda, \phi) = r_{v^*,1}^{\lambda,\phi}(x^{p,s}) + u^{v,s} p_{v^*,i}^{\lambda,\phi}(y^{v,s}) + (1 - u^{v,s}) (1 - p_{v^*,i}^{\lambda,\phi}(y^{v,s})) \quad (\text{A10})$$

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

$$l(\lambda, \phi) = \sum_{s \in S_e} \ln(L(s, \lambda, \phi)). \quad (\text{A11})$$

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 $\lambda \in \{0, 0.1, \dots\}$ and $\phi \in \{1, 2, \dots\}$.

²⁸ 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.