Econ Theory (2012) 49:739–777
DOI 10.1007/s00199-011-0612-z

SYMPOSIUM

The dynamics of distributive politics

Marco Battaglini · Thomas R. Palfrey

Received: 22 September 2010 / Accepted: 4 March 2011 / Published online: 7 April 2011

Abstract We study dynamic committee bargaining over an infinite horizon with discounting. In each period, a committee proposal is generated by a random recognition rule, the committee chooses between the proposal and a status quo by majority rule, and the voting outcome in period t becomes the status quo in period $t + 1$. We study symmetric Markov equilibria of the resulting game and conduct an experiment to test hypotheses generated by the theory for pure distributional (divide-the-dollar) environments. In particular, we investigate the effects of concavity in the utility functions, the existence of a Condorcet winning alternative, and the discount factor (committee "impatience"). We report several new findings. Voting behavior is selfish and myopic. Status quo outcomes have great inertia. There are strong treatment effects that are in the direction predicted by the Markov equilibrium. We find significant evidence of concave utility functions.

We thank audiences at Northwestern University, NYU, MIT, the 2007 Constitutional Political Economy Conference in Torino, the 2007 Wallis Conference in Rochester, and the 2007 Workshop in Behavioral Public Economics in Innsbruck for comments. Anna Bassi, Kyle Mattes, Salvatore Nunnari, and Stephanie Wang provided valuable research assistance. Battaglini gratefully acknowledges financial support from a NSF CAREER Award (SES-0547748) and the hospitality of the Kellogg School's MEDS department for the academic year 2006–2007. Palfrey gratefully acknowledges financial support of grants from the National Science Foundation (SES-0617820 and SES-0962802) and the Gordon and Betty Moore Foundation.

M. Battaglini (✉)
Department of Economics, Princeton University, Princeton, NJ 08544, USA
e-mail: mbattagl@princeton.edu

T. R. Palfrey
Division of the Humanities and Social Sciences,
California Institute of Technology, Pasadena, CA 91125, USA
e-mail: trp@hss.caltech.edu

Keywords Dynamic bargaining · Voting · Experiments · Divide-the-dollar · Committees

JEL Classification D71 · D72 · C78 · C92

1 Introduction

The large redistributive programs that have characterized western democracies since the end of World War II—pension, health, and disability plans, for example—share a common characteristic. Once one of these programs is created by a legislature, it remains in force until explicitly revised. This feature makes the politics of redistribution an intrinsically dynamic game that cannot be studied as a simple static struggle for resources among different constituencies or even as a sequence of independent struggles. A policy chosen today will be the status quo tomorrow. In choosing the optimal policy, a policy maker should not only consider the direct effect of the policy today, but also the indirect effect that the policy has for future policy decisions. In the short run, the policy maker may prefer a policy that favors only his constituency; when the long run is considered, however, a more moderate policy may be preferred because a moderate status quo favorable to a larger constituency will be harder to overturn in the future.

Recent theoretical work has put particular emphasis on the dynamic structure of policy outcomes, producing a rich assortment of predictions.[1] This literature raises three natural questions: To what extent do game theoretic models accurately predict behavior in a dynamic policy game? Can the models be improved to better explain empirical evidence? If so, how?

In this paper, we take a step in answering these questions by studying equilibrium behavior in a simple dynamic model of committee bargaining with endogenous status quo and by presenting the first laboratory experiment on this class of games.[2] We consider an infinite horizon model in which a committee of three agents has to divide a dollar at every period. At the beginning of a period, a member of the committee is selected at random to propose a division to the committee. The committee then chooses by majority rule between the proposal and a given status quo. The selected policy is implemented, and it becomes the new status quo. With a positive probability, the game is repeated exactly as before, but with the new status quo; with the complementary probability, the game is terminated. The policy choice in period t, therefore, will affect the bargaining game at $t + 1$ and indirectly in the following period as well.

---

1 Among the most recent works, see Baron et al. (2011), Baron and Herron (2003), Battaglini and Coate (2006, 2007, 2008), Diermeier and Fong (2009), Duggan and Kalandrakis (2010), Kalandrakis (2004), Penn (2009).

2 Previous experimental work on legislative bargaining games is provided by McKelvey (1991) and, more recently, by Diermeier and Morton (2006), Diermeier and Gailmard (2006) and Frechette et al. (2003, 2005a,b,c) and Frechette et al. (2011). All these works, however, focus on static environments inspired by the seminal paper by Baron and Ferejohn (1989) in which a given amount of resources is allocated only once.

We study this model because similar models have been theoretically studied by a number of authors (Epple and Riordan 1987; Baron 1996; Baron and Herron 2003; Kalandrakis 2004; Duggan and Kalandrakis 2010), and it is therefore a natural starting point. Despite its simplicity, a complete understanding of behavior in this game has thus far proven elusive. When agents are risk neutral, Kalandrakis (2004) has shown by construction that this game has a symmetric Markov equilibrium in which committee members behave myopically, maximizing their current utility. In this equilibrium, therefore, proposers succeed in appropriating all or almost all the dollars in every period as if the game was a sequence of one-period games. Baron (1996) and Baron and Herron (2003) however have conjectured that with more general utilities, agents have stronger incentives for dynamic strategic behavior, suggesting that concavity in the utility function would lead to more equitable outcomes. However, equilibrium behavior in a legislative bargaining game with general utilities has not been characterized yet, so the importance of this phenomenon is not known in general.[3]

To investigate these issues and provide a theoretical benchmark for the experimental evidence, we proceed in two steps. First, we study a simple environment in which we can prove the existence of a unique equilibrium with certain desirable properties, and we can fully characterize it. In this environment, we can consider both situations in which the policy space admits a Condorcet winner and situations in which a Condorcet winner does not exist. A central principle of static models of committee decision making is that Condorcet winners will prevail. In the environment we study, even with a Condorcet winner, we predict the stability of dynamic regimes where non-Concorcet winners prevail indefinitely. Second, we extend the analysis to a "divide-the-dollar" game in which multiple equilibria may exist. We study this game by numerical methods, showing that Baron's conjecture is correct by computing an equilibrium in which as concavity increases, equilibrium outcomes become more equitable.

This theoretical analysis provides a rich set of predictions that we can test in the laboratory. We consider an experimental design that varies the environment across three dimensions. One dimension is whether the environment is a (nearly) continuous divide-the-dollar setting versus a more constrained set of allocations. The second dimension, applied to the finite environment, is the effect of the existence of a Condorcet winner. The third dimension, applied to the continuous environment, is the effect of long-run incentives, which we study by varying the discount factor of the committee members – comparing "patient" legislatures or committees with "impatient" ones.

Our experimental findings allow a clear evaluation of the ability of these complex theoretical models to predict empirical behavior. In environments where bargaining is over a limited set of states, the "standard" theoretical model assumed in the literature (in which utilities are linear and agents play according to Nash equilibrium) is consistent with many features of the data, but with some exceptions which we discuss. The model predicts, in particular, the difference in behavior that we observe between the case in which there is a Condorcet winner among the alternatives or not. When bargaining is over more complicated state spaces (as in the unit simplex), however, the standard model performs less well. The model predicts highly unequal outcomes

[3] Bowen and Zahran (2009) have constructed an example of an equilibrium with these properties for a non-degenerate interval of discount factors when then the number of agents is larger than four.

in which in each period one agent appropriates most of the resources: However, we rarely observe such outcomes, on the contrary, we observe a significant frequency of allocations in which resources are evenly distributed among all participants. We can however show that this type of behavior is not necessarily evidence of social preferences or non-strategic behavior. Indeed, there is little or no evidence in our data indicating a preference for fairness. Rather, our main experimental findings are consistent with selfish preferences and concave rather than linear utilities, and we fit such a model to the data. Players tend to make proposals that maximize their payoff at the expense of others when it is optimal to do so (as when they are favored by the status quo); and voting behavior is overwhelmingly myopic and selfish in all treatments.

The rest of the paper is organized as follows. The next section lays out the model. Section 3 characterizes the theoretical properties of the model. Section 4 describes the experimental design. Section 5 analyzes the results and findings of the experiment. We conclude in Sect. 6.

2 Model

We consider the problem faced by a set of N agents who repeatedly bargain over a set of outcomes X. In each period $t = 1, 2, 3\ldots$, a policy x_t is chosen by the agents. The bargaining protocol with which policy x_t is chosen is as follows. At the beginning of each period, an agent is chosen by nature as the proposer and proposes a policy, $y_t \in X$. The floor votes on this policy following a q rule, where $q \in [1, N]$. If the number voting in favor is greater than or equal to q, the proposal is accepted and $x_t = y_t$ is the implemented policy at t. If the proposal is voted by less than q agents, the proposal is rejected and a status quo policy $\overline{x}_t = x_{t-1}$ is implemented. The initial status quo $\overline{x}_1$ is exogenously specified. Each agent can be recognized as a policy proposer: The probability that agent i is recognized as proposer in period t is $\frac{1}{N}$, so the probabilities of being recognized are assumed here to be symmetric and history invariant.

Agents have a Von Neuman Morgenstern per period utility $U_i : X \to \mathbb{R}$, which is assumed to be continuous and (weakly) quasi concave. The policy implemented in period t, x_t, therefore induces an n-tuple of utilities $(U_i(x_t))_{i=1}^N$. The utility of an infinite sequence of policies, $x = \{x_1, \ldots, x_t, \ldots\}$, is given by $U_i^\delta(x) = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}U_i(x_t)$ where the non-negative discount factor, δ, is assumed to be strictly less than 1.

Many examples of this general framework can be constructed, such as the following two:

Example 1 (*Divide-the-dollar*) The agents have to divide a pie of size K. An allocation is vector $(x_t^i)_{i=1}^N$ where $x^i \geq 0 \forall i$ and $\sum_{i=1}^N x^i = K$. Each agent is interested only in the size of the pie that he receives.

Example 2 (*Public goods*) In this case, X is a collection of projects $\{x_1, \ldots, x_K\}$, where each project x_k gives utility payoffs to the agents, $(U_i(x_k))_{i=1}^N$.

An *outcome* in period t is defined by the current status quo, $\overline{x}_t$, nature's choice of a proposer $\iota_t \in \{1, \ldots, N\}$, the proposed policy $y_t \in X$, and the vector of votes

$\eta_t \in \{1, 0\}^N$. Let $\omega_t = \{\overline{x}_t, \iota_t, y_t, \eta_t\}$ be an outcome at time $t \geq 1$, and $x_t = \chi(\omega_t; q)$ is the policy implemented under the voting rule, q, if the outcome in period t is ω_t. A t-history h^t is defined as $h^t = \{\omega_1, \ldots .\omega_t\}$; the set of possible t-histories is H^t. A strategy for agent i is a set of functions $s_{it} = [\rho_{it}, \sigma_{it}]_{t=1}^{\infty}$ where $\rho_{it} : H^{t-1} \to \Delta X$ describes the proposal strategy of agent i in period t (ΔX is the set of randomizations over X); and $\sigma_{it} : H^{t-1} \times \{1, \ldots, N\} \times X \to [0, 1]$ associates to each history h^{t-1}, proposer, and proposal a probability to vote for the proposal.

In a sequential equilibrium, the strategies are measurable with respect to the entire history set H^t. A Markov strategy, on the contrary, is measurable only with respect to the status quo x_t and the payoff relevant events that occur in period t: $\rho_i : X \to \Delta X, \sigma_i : X \times X \to [0, 1]$. A *Markov equilibrium* is a subgame perfect Nash equilibrium in Markov strategies. The analysis in this paper focuses on Markov equilibria. To each Markov equilibrium and each agent i, we can associate a function $v_i(\theta)$, which represents the expected continuation value of agent i when the status quo (current policy) is θ before the proposer is randomly selected. Given this, we can define the function $u_i(\theta) = U_i(\theta) + \delta v_i(\theta)$ as the expected utility of agent i if policy θ is implemented in a representative period; and the function $u_i(\theta'; \theta)$ as the expected utility of agent i if he proposes θ' when the status quo is θ.

A Markov equilibrium $s = \{(\rho_i, \sigma_i)\}_{i=1}^n$ is *symmetric* if $\{(\rho_i, \sigma_i)\}_{i=1}^n$ and if it has the following symmetry property. For any pair of agents i, j and for any pair of alternatives, $x, y \in X$, define x^{ij} (or y^{ij}) by switching the i^{th} and j^{th} components of x (or y), e.g., $x^{12} = (x_2, x_1, x_3)$. Then, we call s symmetric if $\rho_i(x|y) = \rho_j(x^{ij}|y^{ij})$ and $\sigma_i(x, y) = \sigma_j(x^{ij}, y^{ij})$ for any i, j and any $x, y \in X$. A Markov equilibrium is in *stage undominated* strategies if in each stage no agent uses a strategy that is weakly dominated given his equilibrium value function $v_i(\theta)$. From now on, we will focus on symmetric Markov equilibria in stage undominated strategies, and we refer to them simply as *equilibria*.

3 Theoretical predictions

In this section, we describe the equilibrium properties of the game described in Sect. 2 under additional assumptions on the policy space. We focus exclusively on the case of $N = 3$ and $q = 2$. In Sect. 3.1, we study a case with a finite set of alternatives in which the equilibrium is unique. In Sect. 3.2, we study equilibrium behavior in a standard divide-the-dollar game.

3.1 *Simplified* divide-the-dollar: coarse grid over allocations

In this section, we consider a relatively simple environment in which there is a unique equilibrium prediction. We focus on two possible cases. In both cases, the 3-way equal split that we call the *universal allocation*, is feasible. The two cases then differ in the other three allocations. In the first case, there is no *Condorcet winner* (NCW, which stands for “No Condorcet winner”): It only includes three *majoritarian allocations* where the pie is divided equally between two committee members, and the third committee member receives 0. In the second case (CW, which stands for "Condorcet winner"), the universal allocation is a Condorcet winner: that is, it is myopically preferred by a majority of voters in any pairwise comparison with the other three allocations.

3.1.1 No Condorcet winner (NCW)

Consider a bargaining game with a set of players $N = \{1, 2, 3\}$ and four states $X = \{x_0, x_1, x_2, x_3\}$ that induce payoffs described by the following matrix S:

$$\begin{bmatrix} & 1 & 2 & 3 \\ x_0 & 20 & 20 & 20 \\ x_1 & 0 & 30 & 30 \\ x_2 & 30 & 0 & 30 \\ x_3 & 30 & 30 & 0 \end{bmatrix} \tag{1}$$

where the rows describe the states and the column the players: The matrix specifies the per period utility of an agent for each state.[4] In this game, there are only two possibilities. Either the outcome is egalitarian if state x_0 is chosen; or the outcome is strictly majoritarian: A minimal wining coalition of players shares the dollar and leaves the remaining player with nothing. There is no Condorcet winner.

When the agents are identical, it is natural to consider equilibria in which agents behave and treat the other agents in the same way. We have already assumed symmetry of the strategies, but we can also consider a strong assumption that a strategy by i does not discriminate between other players j and k in terms of *outcomes*. We say an equilibrium is *neutral* if for any θ, x, y: $u_i(x; \theta) = u_i(y; \theta)$ implies $\rho_i(x|\theta) = \rho_i(y|\theta)$, and $u_i(x) = u_i(\theta)$ implies $\sigma_i(x|\theta) = \frac{1}{2}$ (where $\rho_i(x|\theta)$ is the probability that x is proposed by i in state θ, and $\sigma_i(x|\theta)$ is the probability that voter i votes for x when the status quo is θ). Intuitively, an agent cares only about his expected payoff, not about the particular state that achieves the payoffs. For the coarse grid divide-the-dollar games, this refinement of the symmetric Markov equilibrium will deliver uniqueness.[5]

Consider a strategy profile, in which voters vote "myopically" for the alternative that offers the highest per period payoff, mixing with equal probability when indifferent; and a proposer i proposes some $x_j \;\; j \neq i, 0$. We call a strategy profile with these characteristics a *myopic strategy*. The following result establishes that the symmetric forward-looking equilibrium strategies are myopic when payoffs are described by (1):

Proposition 1 *When payoffs are as* (1)*, with linear utilities, there is a unique neutral equilibrium. Each agent i proposes following history-independent strategy* $\rho_i\left(x_j\,|x\right) = \rho_i\left(x_k\,|x\right) = \frac{1}{2}\forall j, k \notin \{i, 0\}$*, and* $\forall x$ *and votes for the alternative that offers the highest immediate payoff, mixing with equal probability when indifferent. This remains an equilibrium if the agents have the same strictly increasing utility function.*

[4] In (1), the sum of payoffs is 60 because this is the size of the pie that we use in the experiments.

[5] Later, when we consider the continuous (or fine grid) divide-the-dollar game, we do not impose the condition of neutrality.

Proof See Appendix. □

This result provides clear-cut predictions that can be tested in the laboratory. The proposal behavior and the voting behavior do not depend on the initial status quo. Specifically, in every round, equilibrium proposal strategies are mixed, with equal probability on the two allocations that give the proposer 30. Voting behavior is myopic, with indifference leading to uniform mixing. This Markov equilibrium generates a unique transition probability function. From a status quo x_i, the state remains at x_i with probability $\frac{2}{3}$ and moves to a state x_j with probability $\frac{1}{3}$, and never moves to x_0. From status quo x_0, the state moves to each other state x_i with probability $\frac{1}{3}$. Proposals of (20, 20, 20), therefore, never occur in equilibrium, so votes involving (20, 20, 20) can only occur in the very first round, and only if (20, 20, 20) is the initial status quo.

It is useful to note that concavity ($\gamma > 0$) does not destroy this equilibrium in the present case. That is, regardless of the concavity of the utility function, the equilibrium identified in Proposition 1 persists.

3.1.2 Condorcet winner (CW)

With finite states, there can be a Condorcet winner. The bargaining game described by the following set of four allocations is one example:

$$\begin{bmatrix} & 1 & 2 & 3 \\ x_0 & 20 & 20 & 20 \\ x_1 & 30 & 15 & 15 \\ x_2 & 15 & 30 & 15 \\ x_3 & 15 & 15 & 30 \end{bmatrix} \tag{2}$$

As in (1), we have a symmetric policy (x_0); now, however, in policies x_i $i = 1, 2, 3$, one agent receives a payoff double to the payoff of the other players. The Condorcet winner is x_0.

Suppose that voters have linear utilities, and consider the value function implied by myopic strategies. The continuation value in state x_0 is the simplest to find, since—by virtue of it being the Condorcet winner—the outcome transitions out of this state with zero probability under myopic strategies. Therefore,

$$v_i(x_0) = \frac{20}{(1-\delta)} \quad \forall i. \tag{3}$$

Given this, the remaining value functions can also be easily found by backward induction. When the state is x_i, the value function $v_i(x_i)$ of agent i is

$$\begin{aligned} v_i(x_i) &= \frac{1}{3}(30 + \delta v_i(x_i)) + \frac{2}{3}\left[\frac{1}{2}(30 + \delta v_i(x_i)) + \frac{1}{2}\left(15 + \delta v_i(x_j)\right)\right] \\ &= \frac{2}{3}(30 + \delta v_i(x_i)) + \frac{1}{3}\left(15 + \delta v_i(x_j)\right) \end{aligned} \tag{4}$$

where $v_i(x_j)$ is the value function of the same agent i when the state is $j \notin \{0, i\}$. With probability 1/3, i is the proposer and he can guarantee that x_i is chosen; with probability 2/3 a different agent (say j) is proposer and proposes x_j: The proposal is accepted with probability 1/2, with probability 1/2 x_i is implemented again. The continuation value at x_j is for i is can be computed in a similar way:

$$v_i(x_j) = \frac{1}{6}(30 + \delta v_i(x_i)) + \frac{5}{6}\left(15 + \delta v_i(x_j)\right) \tag{5}$$

Solving equations (4)–(5) we have:

$$v_i(x_i) = \frac{50 - 30\delta}{(1-\delta)(2-\delta)}, \quad v_i(x_j) = \frac{35 - 15\delta}{(1-\delta)(2-\delta)}. \tag{6}$$

From these formulas, it is easy to verify that the strategies described elsewhere induce a value function that is monotonically increasing in the agent's one-shot payoff. The following proposition shows that not only these strategies are an equilibrium, but also they are the unique symmetric Markov equilibrium**:**

Proposition 2 *When the set of allocations is that in* (2) *and utilities are linear, there is a unique neutral equilibrium. In this equilibrium, players play myopic strategies.*

Proof See Appendix. □

The equilibrium strategies characterized in Proposition 2 imply a unique transition matrix. From a status quo x_i, the state remains at x_i with probability $\frac{2}{3}$ and moves to a state x_j with probability $\frac{1}{3}$, and never moves to x_0. From status quo x_0, the state remains at x_0 with probability 1.

As in the NCW case, voting behavior predicted by Proposition 3 is myopic; however, this leads here to an important difference. In the NCW case, (20, 20, 20) is defeated in any pairwise vote against any other allocation; but in the CW case, it will defeat any other allocation. Therefore, in the latter case, if the *initial status quo* is (20, 20, 20), it will remain the status quo forever. Proposals of (20, 20, 20), while off the equilibrium path in *both* the NCW and the CW case (except in the CW when it is the initial status quo), have a much different effect, since (20, 20, 20) is an absorbing state in the CW game, but it is always defeated the NCW game. As we show below this implies drastically different dynamics in Markov quantal response equilibrium (QRE) between the two tables, where because of the stochastic nature or propositions under QRE, CW case represented by 2 will alternate between epochs of "universal" regimes and "dictatorial" regimes, and the NCW environment represented by 1 will yield stable "majoritarian" regimes, where there is a random rotation of two players coalitions splitting the pie. The QRE dynamics are discussed in more detail in the results section.

3.2 The divide-the-dollar game

As we mentioned in Sect. 1, in a standard "divide-the-dollar" game, a feasible allocation is a vector $(x_t^i)_{i=1}^N$ where $x_t^i \geq 0$, and $\sum_{i=1}^N x_t^i$ is equal to a constant, the "size of the pie"; and each agent is interested only in the size of the pie that he receives. Despite its simplicity, there is no known characterization of the divide-the-dollar bargaining game described above for general utility functions.[6] To obtain predictions that we can test using our experimental data, we turned to numerical methods. Here, we describe properties of the numerically computed equilibrium under the parameter specifications used in the experiments: three agents, a pie of size 60 and a discount factor equal to either 0.83 or 0.75.

3.2.1 Numerical computation of the Markov equilibrium

We compute a Markov equilibrium for the family of utility functions with constant relative risk aversion:

$$U_i(x_i; \gamma) = \frac{1}{1-\gamma} (x_i)^{1-\gamma} \quad \forall i = 1, 2, 3$$

where x_i is the share received by agent i. The coefficient of relative risk aversion γ measures the curvature of the utility function: The higher is γ, the more concave is utility. For simplicity, in this section, we discuss two polar cases: the linear case, $\gamma = 0$, and a strictly concave case, $\gamma = 0.95$.[7]

Equilibria were computed as the limit of *Markov Logit Quantal Response Equilibria* (MLE) by gradually reducing noise in the agents reaction functions. This smooths out the best response correspondence, which is helpful in the numerical computation. In the logit version of quantal response equilibrium, as defined for extensive form games, each player at each information set uses a behavioral strategy where the log probability of choosing each available action is proportional to its continuation payoff, where the proportionality factor, λ, can interpreted as a responsiveness (or rationality) parameter. The continuation payoffs are computed using the MLE strategies of all future plays in the game, as, for example, in the definition of continuation payoffs in a sequential equilibrium (Kreps and Wilson 1982).[8] Markov perfect equilibria can be found as limits of MLE because for very high values of λ, players choose best responses with probability approaching 1, so limit points of the MLE correspondence, as $\lambda \to \infty$ are Markov equilibria. Moreover, Theorem 4.1 in McKelvey and Palfrey (1998) can be extended to the Markov equilibrium setting to show that for generic finite games in which a Markov perfect equilibrium exists, there is one Markov perfect equilibrium that is selected as the limit of the connected path in the equilibrium graph that has a solution for every value of $\lambda \geq 0$.

We solved the game using discrete approximation of a unit simplex where allocations are in increments of 5.[9] This reduces the set of states to 91. Formally, the policy

6 The formal construction of an equilibrium in the divide-the-dollar game is available only for the case of linear utilities (see Kalandrakis 2004). We will discuss this equilibrium elsewhere.

7 In later sections, we will use experimental data to obtain a maximum likelihood estimate of γ.

8 A precise definition of this equilibrium concept is presented in Appendix 1.

9 For any discrete approximation, existence of a symmetric Markov equilibrium follows from standard fixed-point arguments.

space is as follows:

$$X := \left\{ x = (x_1, x_2, x_3) \text{ s.t. } \forall i \ \exists t \in \mathbb{N}, x_i = 5t, \ \sum_{i=1}^{3} x_i = 60 \right\}.$$

Given the smooth properties of this MLE path, there is a simple and (relatively) fast path-following algorithm, which will find this solution. It is simple because we know the solution at $\lambda = 0$: All (behavioral) strategies are chosen with equal probability, and this implies the unique value function. Hence, we begin with the solution at $\lambda = 0$ and can use that solution as the starting value to find the MLE for an incrementally larger value, say $\lambda = \epsilon$. Because we are guaranteed that for small enough ϵ, the starting value obtained from $\lambda = 0$ is very close to the solution at $\lambda = \epsilon$, so the fixed-point algorithm will find a solution at ϵ very quickly. Then, we use the solution at $\lambda = \epsilon$ to compute the solution at $\lambda = 2\epsilon$ and so forth, thereby tracing out the MLE path that converges to a Markov equilibrium of the game. There are some computational issues when λ becomes very large, and the algorithm takes several hours but conceptually it is quite simple, and convergence is not difficult to achieve.

3.2.2 Steady state equilibrium dynamics

A proposal strategy associates with each status quo a vector of probabilities of proposing each state. The voting strategies associate a probability of voting yes to each possible status quo–proposal pair. Because the equilibrium strategy space is so large, to describe the properties of equilibrium behavior, it is convenient to use the stationary distribution over outcomes induced by equilibrium strategies. The equilibrium strategies generate a Markov process with a stationary transition matrix. This transition matrix associates each state $x' \in X$ with a probability distribution $\varphi(x \,|x')$ over states $x \in X$ in the following periods. For a given initial distribution $P^0(x)$ over the status quo, we can therefore define the equilibrium distribution of states at t recursively as:

$$P^t(x) = \sum_{x' \in X} \varphi(x \,|x')\, P^{t-1}(x').$$

The probability function $P^t(x)$ converges to a stationary distribution $P^*(x)$ as $t \to \infty$. This distribution represents the frequency of the states that we would expect to observe in the long run, so it provides one of the fundamental properties of the Markov equilibrium.

For descriptive purposes, we cluster the states in coarser regions. Figure 1 describes a partition of the states in 7 regions. The D regions correspond to *dictatorial* allocations where one player receives the lion's share of the pie. The M regions correspond to *majoritarian* allocations where a coalition of two players receives most of the pie, with nearly equal shares, while the third player receives only a small amount or nothing. The U region consists of *universal* allocations, where the pie is equally, or nearly equally, shared. Conditional on being, say, in D1, we can use the stationary distribution of the computed Markov equilibrium to derive the probability of transition to state

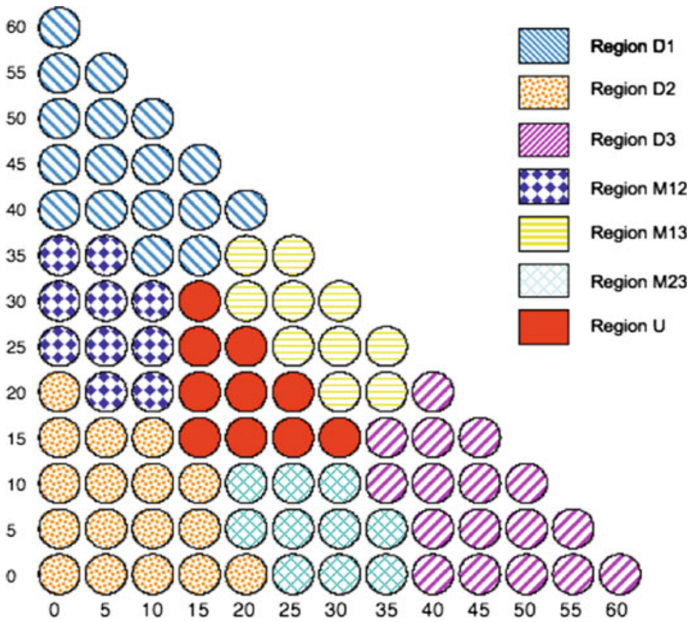

Fig. 1 Allocation regions. The *vertical coordinate* represents Agent 1's payoff, the *horizontal* represents Agent 3's payoff. The payoff of Agent 2 is the residual

$M12$ (the M region corresponding to the coalition of players 1 and 2). Doing this for all pairs of regions gives a representation of the steady state equilibrium dynamics of the infinitely repeated game in a simple 7×7 matrix. This allows one to describe the dynamics in a concise way.

Linear utilities We start with the discussion of the equilibrium with linear utilities and $\delta = 0.83$. In choosing how to allocate the pie, a proposer is faced with a trade-off between short-run and long-run effects of the allocation. In the *short run*, a proposer is facing a simple problem: If the proposer were completely myopic ($\delta = 0$), he would attempt to form a minimal winning coalition and maximize his immediate payoff. In the *long run*, however, the game is more complicated because a state that maximizes his payoff today may reduce his payoff in the future.

To see which effect dominates when agents are risk neutral, consider the equilibrium transition matrix, presented in Table 1, using the condensed states described in Fig. 1.[10] Given the symmetry of the equilibrium, we have only 3 regions to consider: if we are in D1, in M12, or in U: The remaining cases will be the same. The dynamics implied by Table 1 are therefore even more simply represented in Fig. 1, which describes the transition probabilities from these three states.

Table 1 makes clear that the short-run effect dominates. For example, suppose the initial state is in D1, where agent 1 receives most of the pie. In this case, the state

[10] In Table 1, transition probabilities may not sum to one due to rounding errors. SQ_t is the status quo in period t.

Table 1 Theoretical transition matrix of the 7 regions $\delta = .83, \gamma = 0$

SQ_t	SQ_{t+1}						
	D1	D2	D3	M12	M13	M23	U
D1	0.34	0.33	0.33	0	0	0	0
D2	0.33	0.34	0.33	0	0	0	0
D3	0.33	0.33	0.34	0	0	0	0
U	0.01	0.01	0.01	0.31	0.31	0.31	0.03
M12	0.33	0.33	0.25	0.08	0	0	0
M13	0.33	0.25	0.33	0	0.08	0	0
M23	0.25	0.33	0.33	0	0	0.08	0

will stay at D1 with 34% probability and move to D_j, $j = 2, 3$ with 33% probability; that is, with 100% probability, the state will remain in the extreme regions. This occurs because in D1 each agent will propose almost all the payoff for himself, with a minimal share going to a single coalition partner.

It is interesting to note the dynamics evolving from a status quo in U. In this case, the state does not jump directly to a region Di, $i = 1, 2, 3$ with high probability (in total, only 3% of the time). Much more likely the state will transition to a state Mij, $i, j = 1, 2, 3$. This is because that it is very difficult for i to convince any other agent to vote for a Di proposal. This can only happen if the state in U is bordering a region Dj, $j \neq i$, by offering to $k \notin \{i, j\}$ (a currently disadvantaged agent) a more advantageous payoff in Di. From a state Mij, however, the system moves to a D state with very high probability, more than 90% of the time. From U, the system moves with high probability to Mij.

In the long run, therefore we would expect the state to rotate around regions D1, D2 and D3. This myopic behavior can be clearly seen in the stationary distribution of outcomes represented in the top half of Fig. 3. In the long run, most of the mass of the distribution of states is on the extremes: that is on states in which a single agent receives a payoff between 50 and 60. When agents are risk neutral, therefore, they behave as if they were myopic, simply choosing allocations that maximize their current payoff.

This finding is consistent with the analysis in Kalandrakis (2004) who characterized an equilibrium of the bargaining game when the state space is the unit simplex (and so the unit of account is infinitesimal). There is, however, a slight difference. Kalandrakis (2004) shows that in the long run only the most extreme states are chosen (i.e., only states in which one agent receives 60). In the equilibrium presented earlier, this does not occur: indeed with a strictly positive probability at least one of the other agents receives a positive payment. This difference is due to the fact that in the model studied here the proposer must divide the pie in discrete units of 1/12 of the total size. With a continuum, the equilibrium must have voters voting in favor of the proposal when they are indifferent. (Otherwise, the proposer would have an incentive to sweeten the offer by an infinitesimal amount.) With a discrete pie, this no longer must be the case.[11]

[11] A similar property arises in looking at subgame perfect equilibrium of the ultimatum game. With a perfectly divisible pie, the only subgame perfect equilibrium is for the proposer to offer zero and the responder

Table 2 Theoretical transition matrix of the 7 regions $\delta = .83, \gamma = 0.95$

SQ_t	SQ_{t+1}						
	D1	D2	D3	M12	M13	M23	U
D1	0	0	0	0.27	0.27	0.27	0.19
D2	0	0	0	0.27	0.27	0.27	0.19
D3	0	0	0	0.27	0.27	0.27	0.19
U	0	0	0	0.01	0.01	0.01	0.98
M12	0	0	0.02	0.77	0.01	0.01	0.19
M13	0	0.02	0	0.01	0.77	0.01	0.19
M23	0.02	0	0	0.01	0.01	0.77	0.19

In the Markov equilibrium selected as the limit in our computations, the proposer does not want to make offers that leave the other players just indifferent, because in the quantal response equilibrium, they would vote in favor of the proposal only 1/2 of the time even for large values of λ, while they would accept all better offers (including the cheapest one) with probability 1. The proposer, therefore, has an incentive to offer something to his coalition partner. Of course, as the grid becomes arbitrarily fine, these equilibria become essentially identical.

It is interesting to compare these results with Proposition 1, which characterizes the equilibrium in the CW case. As in Proposition 1, the equilibrium proposals have minimal winning coalitions. As in the case with a coarse policy space, equilibrium behavior of agents mimics myopic behavior, behaving exactly as agents with a zero discount factor, even though all agents are in fact strategic and forward looking. The result, however, here is more extreme because the proposer is less constrained by the coarseness of the state space and can fully extract almost all the resources.

Concave utilities With strictly concave utility functions, agents are averse to sequences of outcomes in which the status quo—and hence their own share of the pie—changes at every period. Hence the incentives for more symmetric distributions are greater because such distributions generate less variance across time. Among the least efficient outcomes would be the one where a single voter, the proposer, appropriates the entire pie in each period. Though in this case, an agent is receiving 20 *on average* (60 one-third of the time and 0 two-thirds of the time), this gives a lower discounted utility than receiving exactly 20 in every period. Proposers can avoid such "rotating dictator" outcomes by choosing a division that is closer to the centroid of the simplex. By allocating a higher share to an agent, the proposer exposes himself *less* to expropriation in the future, because it makes it harder for a future proposer to extract a larger share of the surplus by forming a coalition with an excluded agent. A proposal close to the centroid is harder to overturn and reduces the volatility of the proposer's future payoffs.

To see that the myopic behavior is no longer optimal in equilibrium consider Table 2, which describes the 7 region transition matrix in the case with $\gamma = 0.95$, and Fig. 2

Footnote 11 continued
to accept any offer. However, with a discrete grid, there is also a subgame perfect equilibrium where the proposer makes the smallest positive offer, and the responder accepts only positive offers.

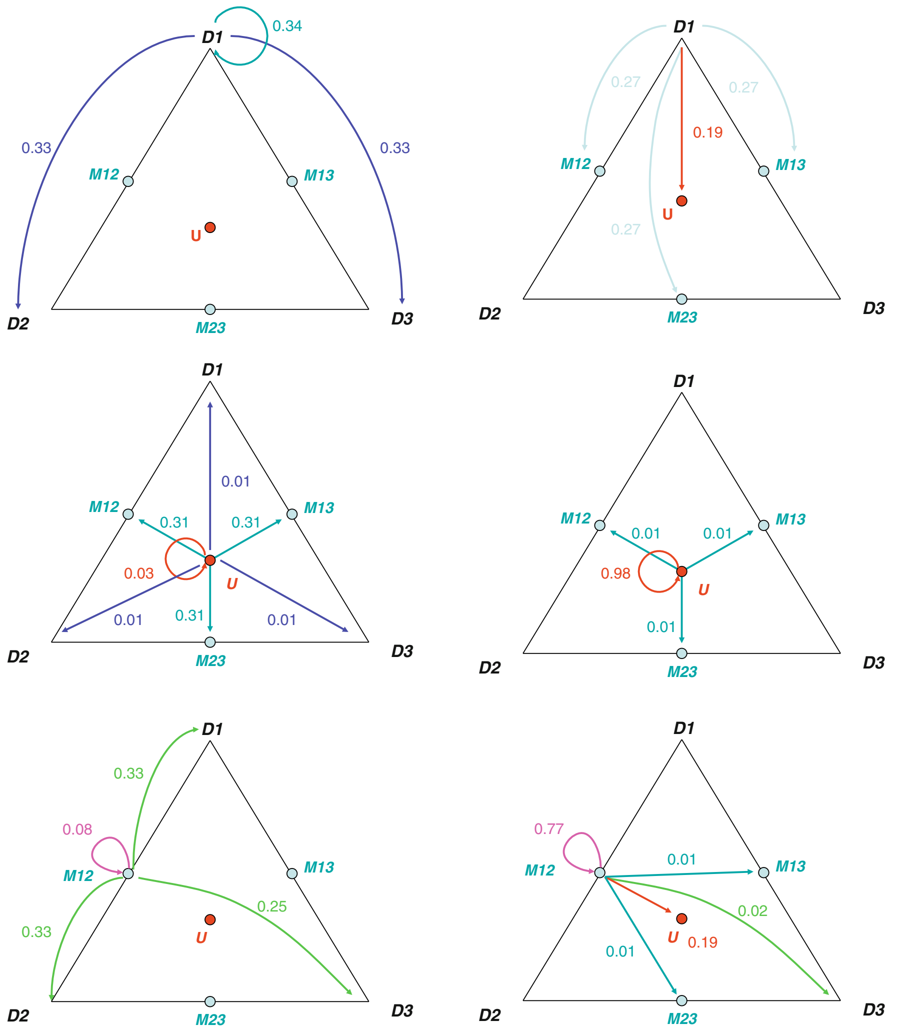

Fig. 2 $\delta = 0.83$. The *left graphs* (**a**) are for the linear utility equilibrium and the *right graphs* (**b**) are for the concave utility (=.95) equilibrium. The *top graphs* illustrate transitions from D regions; the *middle graphs* from U; the *bottom graphs* from M

which represents the 3 representative cases. The differences between Fig. 2b and Fig. 2a ($\gamma = 0$) are striking. Starting from a D state, when $\gamma = 0.95$ we *never* stay in a D state, but usually move to a state M12, M13, or M23 (over 80% of the time) and occasionally to region U. When the state is in a majoritarian region, we usually remain in the same region but again move to U with significant probability, 20% of the time. Once region U is reached, it is essentially an absorbing state, staying there 98% of the time, rarely moving to the M region. Also of some independent interest is that the only M_{ij} states that are visited with positive probability divide the pie equally between i and j, with the third voter receiving zero. Such states also have high persistence, not

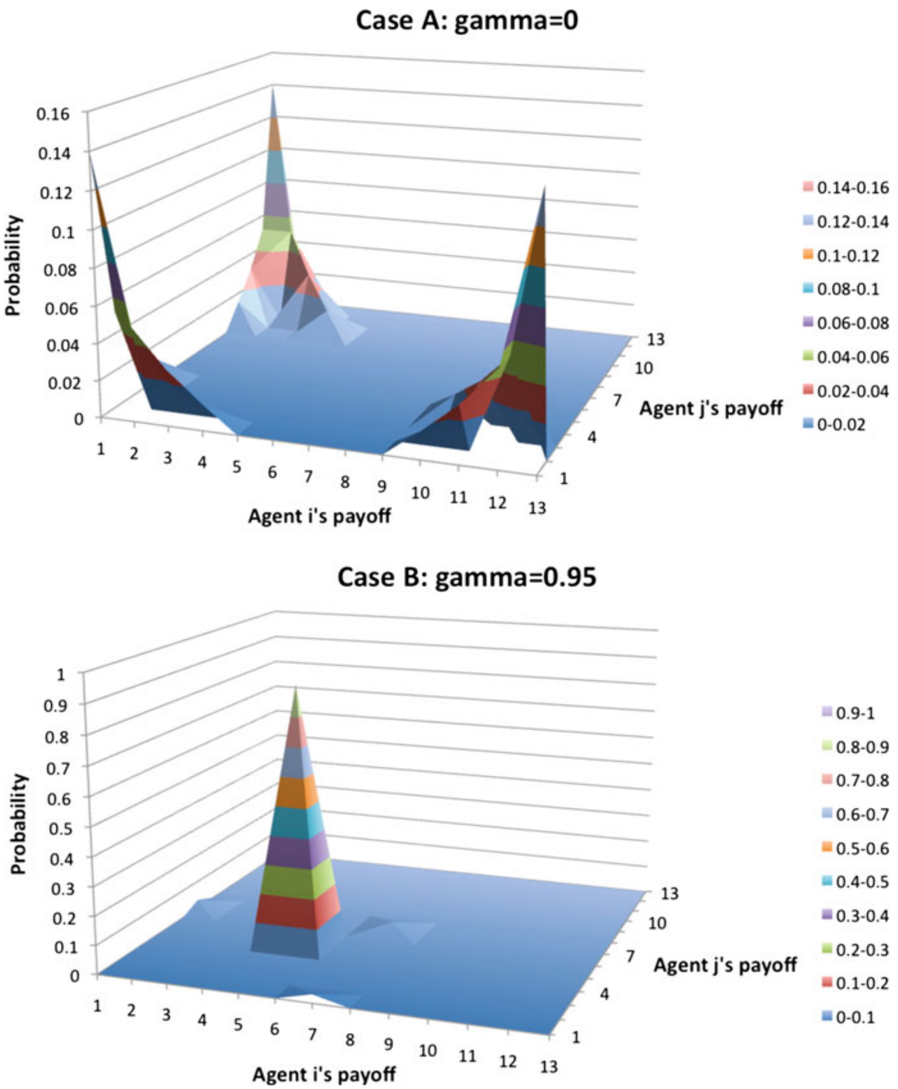

Fig. 3 Stationary distribution of allocations ($\delta = .83$). The *top graph* is for $\gamma = 0$ and the *bottom graph* is for $\gamma = 0.95$

only in the sense of usually staying in an M region (i.e., rarely moving to a D_i region or U), but also in the stronger sense that there is essentially no transition probability from M_{ij} to M_{ik} or M_{kj}.

The long-run incentives to move toward the center are clear in the stationary distribution represented in Fig. 3. While with linear utilities the probability that agent i receives a payoff between 20 and 40 in the stationary distribution is .00, it is over .96 with $\gamma = 0.95$.

The tendency of outcomes to cluster around the centroid confirms the phenomenon identified in Baron (1996) for a unidimensional case.

Table 3 Experimental design

Session	δ	Set of feasible allocations		# Subjects
		Matches 1–10	Matches 11–20	
1	.75	NCW	CW	9
2	.75	CW	NCW	12
3	.83	Continuous		12
4	.75	Continuous		12
5	.75	Continuous		12

4 Laboratory experiment

We use controlled laboratory experiments to study behavior in these dynamic committee bargaining environments with endogenous status quo allocations. In our experimental design, we vary the discount factor and the set of feasible allocations. We conduct two sessions with discrete allocations, as in the theoretical section and three sessions with allocations with a very fine grid, as a finite approximation to the continuous-state divide-the-dollar game. We refer to the fine grid sessions as "continuous." In all sessions, the sum of the three agents' allocations equals 60 (Table 3).

4.1 Procedures

Discount factors were induced by a probabilistic endpoint. After each t, a fair die was rolled by the experimenter at the front of the room, and the game continued to period $t+1$ if and only if the die roll belonged to a preannounced subset of the possible faces. For example, to implement $\delta = .75$, we rolled a twelve-sided die, and the game ended if and only if a 10, 11, or 12 was rolled.[12] If the die came up with a number less than 10, the game proceeded to round $t + 1$, with the status quo being determined by the majority rule winner in round t. In all except one session, we used a discount factor of $\delta = .75$. The remaining session used $\delta = .83$ and a six-sided die.

The experiments were all conducted at the Princeton Laboratory for Experimental Social Science and used registered students from Princeton University. Each subject participated in exactly one session.

Each of the discrete allocation sessions was divided into two subsessions, each of which lasted for 10 matches. Each match corresponded to one play of the infinitely repeated game, using the die-termination rule described earlier.[13] The set of feasible allocations was different in the two subsessions. The two sets of allocations are exactly the ones previously described as the NCW and CW cases in (1) and (2).

In the continuous allocation sessions, proposals could be any non-negative integer division of the 60-unit pie. Because this was a more difficult task, subjects took a

[12] In some sessions, we used an 8-sided die to implement $\delta = .75$. In the $\delta = .83$ session, we used a 6-sided die.

[13] As a result, there was a lot of variance in the length of the matches, which ranged from 1 round to 23 rounds.

longer time deciding on their proposals. Consequently, we only ran 10 matches in each of the three continuous sessions.

Instructions were read aloud, and subjects were required to correctly answer all questions on a short comprehension quiz before the experiment was conducted. Subjects were also provided a summary sheet about the rules of the experiment which they could consult. The experiments were conducted via computers.[14]

At the beginning of each match, subjects were randomly divided into committees of 3 members each. In each committee, members were assigned to be either Committee Member 1, Committee Member 2, or Committee Member 3 and this member assignment remained the same for all rounds of a match. An initial status quo was randomly chosen by the computer, using a uniform distribution of the set of feasible allocations. Initial status quo assignments were independent across matches and across committees.

After being informed of the initial status quo, each committee member was prompted by the computer to enter a "provisional proposal". After all members had entered a provisional proposal, one was selected at random to become the "active proposal". The active proposal was then voted on against the status quo, which was referred to as the "standing alternative". Whichever received more votes was the policy that was implemented in that round, and each member received earnings accordingly. After all committees had finished the round, a die was rolled to determine whether to continue. If the match continued, then the winning proposal in the previous round became the standing alternative for the new round. This continued until a die roll terminated the match.

This was repeated with the group membership shuffled randomly after each match. Each subject was paid the sum of his or her earnings over all rounds of all matches in cash at the end of the experiment. Average earnings were approximately $30 (including a $10 show up fee), with each session lasting about 90 min.

5 Experimental results

We analyze the results separately for the discrete allocation sessions and the continuous allocation sessions.

5.1 Coarse grid bargaining committees

5.1.1 State transition probabilities

State transition probabilities provide a clear summary of the dynamics of *outcomes* since they provide a synthetic description of aggregate behavioral data on both *proposal making* and *voting*.[15] The transitions and outcomes for the two coarse grid games

[14] Sample instructions and the computer program used for the experiment are available from the authors. The computer program was an extension to the open source multistage game software. See http://multistage.ssel.caltech.edu.

[15] Note that this does not obviously follow, unless voting and proposal-making strategies are stationary.

Table 4 Empirical transition matrices in the NCW and CW games

SQ_t	SQ_{t+1}			
	30–30–0	30–0–30	0–30–30	20–20–20
NCW committees				
30–30–0	0.77 (14)	0.16 (16)	0.05 (5)	0.02 (2)
30–0–30	0.13 (7)	0.79 (73)	0.05 (5)	0.02 (2)
0–30–30	0.11 (5)	0.11 (7)	0.75 (46)	0.02 (1)
20–20–20	0.19 (1)	0.16 (6)	0.14 (5)	0.51 (19)
Freq (SQ)	0.36 (27)	0.35 (102)	0.21 (61)	0.08 (24)
	30–15–15	15–30–15	15–15–30	20–20–20
CW committees				
30–15–15	0.44 (14)	0.09 (3)	0.09 (3)	0.38 (12)
15–30–15	0.17 (7)	0.54 (22)	0.02 (1)	0.27 (11)
15–15–30	0.17 (5)	0.03 (1)	0.53 (16)	0.27 (8)
20–20–20	0.01 (1)	0.03 (3)	0.03 (3)	0.94 (105)
Freq (SQ)	0.13 (27)	0.13 (29)	0.11 (23)	0.63 (136)

are summarized in Table 4. For each table, the last row is obtained by summing each of the columns corresponding to the $t + 1$ status quo. These frequencies give the overall outcome frequencies, excluding the initial round 0 status quos, which were decided randomly by the computer to start each match. Because the game is symmetric, it is useful to look at the transitions from the two key possible cases: when the status quo is the universal allocation or not. Figure 4 represents these transition probabilities.

Several features of the outcome data are noteworthy. First, there is a striking difference between the outcomes of the two coarse grid treatments: In the NCW treatment, the universal outcome was the committee decision only 24 out of 291 times (8%). In contrast, majoritarian outcomes prevailed nearly always 92% of the time. For the CW treatment, this is reversed. There were 136 out of 215 (63%) universal outcomes, while non-universal outcomes were chosen 37% of the time. The theoretical prediction of more universal outcomes in CW than NCW, therefore, is strongly supported by the data, and is significant at any conventional level.

Second, as also predicted by the theory, there is strong “persistence of regimes”. For the CW treatment, non-universal allocations should usually map into non-universal allocations and universal allocations map into universal allocations. This we find, with persistence rates of 70% in the case of the non-universal regime and 94% in the universal regime. In the NCW committees, the universal regime is not part of the equilibrium, so there should be less persistence, which is what we find: universal outcomes map into universal outcomes significantly less often (51%) than in the CW committees. And we find nearly 100% persistence of majoritarian allocations in the NCW committees, as predicted. Hence, in all cases, allocations are significantly more likely to persist than not, except in the case where they are not part of the Markov equilibrium (the case of universal allocation in NCW committees).

Figure 4, however, suggests a potentially interesting systematic departure from equilibrium behavior. In particular, the non-universal regime is somewhat less stable

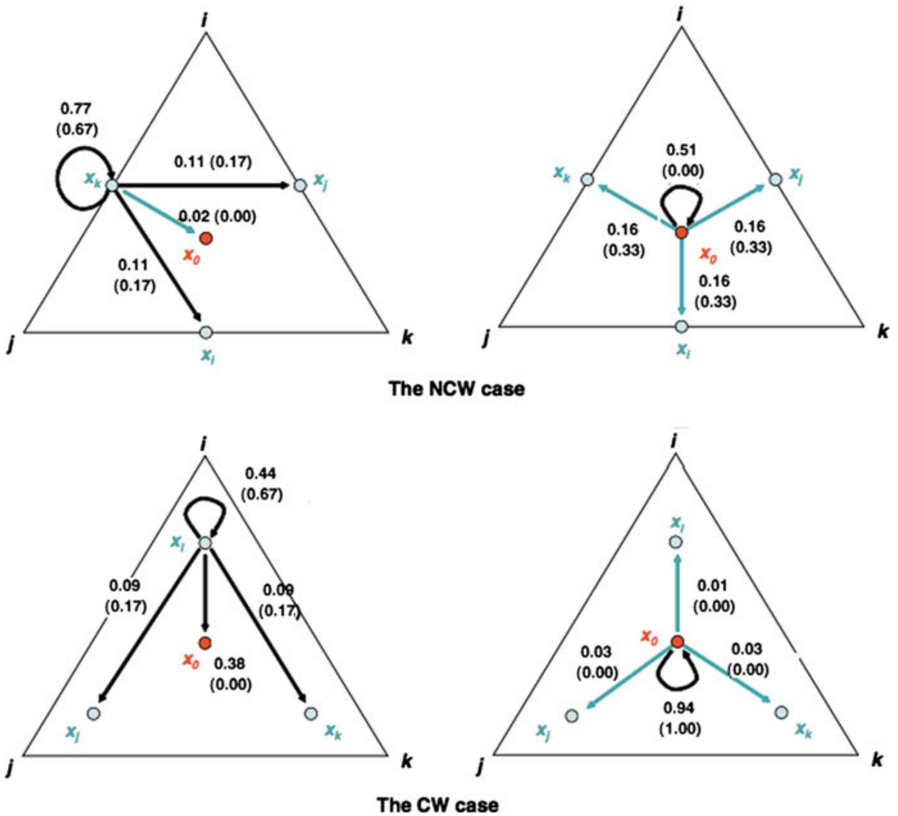

Fig. 4 Empirical transition probabilities for the NCW and CW game from the universal and majoritarian allocations. The *numbers* in *parenthesis* are the predictions of the equilibrium with linear utilities

than predicted in the CW committees, and universal allocations are somewhat more stable than predicted in the NCW committees. However, the latter is only 50% meaning that in NCW committees, universal allocations are as likely to be replaced by non-universal allocations as they are to persist.

To investigate the origin of these dynamic patterns, in the next two sections, we decompose the determinants of the transition probabilities by analyzing in detail proposal and voting behavior. As we will show, the departure from the equilibrium observed above appears to be mainly due to small deviations from equilibrium in *proposal* strategies rather than *voting* strategies.

5.1.2 Proposal making and outcomes

Table 5 displays the aggregate proposal frequencies as a function of the status quo and the position of the player, for the CW (left) and NCW (right) treatments, respectively. The table is "anonymized" for data-pooling purposes, in the sense that a proposal

Table 5 Empirical proposal strategies (theoretical strategies in parenthesis)

SQ	Proposal			
	30	0	20	# Obs
NCW				
30	0.94 (1.00)	0.01 (0.00)	0.04 (0.00)	508
0	0.89 (1.00)	0.01 (0.00)	0.10 (0.00)	254
20	0.70 (1.00)	0.01 (0.00)	0.29 (0.00)	111
	30	15	20	# Obs
CW				
30	0.94 (1.00)	0.00 (0.00)	0.07 (0.00)	103
15	0.57 (1.00)	0.06 (0.00)	0.37 (0.00)	206
20	0.47 (0.00)	0.05 (0.00)	0.48 (1.00)	336

of (20, 20, 20) by member 1 when the status quo is (30, 30, 0) is treated the same as a proposal of (20, 20, 20) by member 2 when the status quo is (0, 30, 30), and so forth. Furthermore, allocations that give an agent the equivalent share are combined. For example, in the CW treatment, observations of 15–15–30 and 15–30–15 are merged together for player 1 into the category "15". This anonymization leads to a very simple 3×3 matrix representation of the aggregate proposing data. For example, the entry 0.07 (7) in the CW table for row 30 and column 20 indicates that 7% of subjects who receive 30 in the status quo under the CW treatment propose 20–20–20 (7 observations).

These tables show several features. First, subjects almost never offer to receive the lowest payoff (either 15 or 0). The rarity of these events (2%) suggests that subjects understand the basic task, and such events are simply trembles that can be ignored.

Second, the proposal strategies in the NCW treatment closely track the theoretical predictions. If the status quo is majoritarian, then proposals are majoritarian more than 90% of the time (100% is the prediction). In the relatively rare instances where the universal allocation is the status quo, that probability is more than 70%.

Third, in contrast, the strategies in the CW treatment are not as close to the linear utility Markov equilibrium predictions. In one respect, it is consistent with the theory: Just as in NCW, a favored committee member receiving 30 almost always (93%) proposes to continue being the favored committee member. Both members receiving low payoffs in a non-universal status quo also usually (57%) propose to be the new favored member, but still propose the universal allocation more than one-third of the time (37%)[16]. However, in the universal status quo, subjects are equally likely to propose their favorable allocation as they are to propose to stay at the universal outcome. This is inconsistent with optimal behavior since in state x_0 a proposer is indifferent between proposals since x_0 defeats all proposals.

[16] The remaining 6% is accounted for by a few cases where a member proposed to be on the short end of an unequal allocation. This also happened in 5% of cases when the status quo was universal. In the NCW treatment, there were also a few cases (<5%) where a subject proposed an allocation where they would receive 0.

Table 6 Voting behavior Entries are Pr (vote for proposal). Theoretical strategies in parenthesis

SQ	Standing proposal			
	30	0	20	# Obs
NCW				
30	0.25 (0.50)	0.03 (0.00)	0.17 (0.00)	508
0	1.00 (1.00)	0.41 (0.50)	1.00 (1.00)	254
20	0.86 (1.00)	0.00 (0.00)	0.28 (0.50)	111
	30	15	20	# Obs
CW				
30	0.44 (0.50)	0.00 (0.00)	0.06 (0.00)	103
15	1.00 (1.00)	0.42 (0.50)	0.96 (1.00)	206
20	0.93 (1.00)	0.06 (0.00)	0.23 (0.50)	336

Thus, there are some sharp differences between the two treatments, the most important being the greater frequency of proposing 20 in the CW treatment, which is predicted by theory. However, in both treatments, subjects propose their myopically best allocation, but the probability of making such a proposal varies with the status quo, being most likely if it is already the current status quo, and least likely if the status quo is 20–20–20.[17]

5.1.3 Voting decisions

Voting decisions are overwhelmingly myopic and selfish for both treatments (meaning that agents vote for the alternative that offers the highest short-run payoff), as predicted. Overall, in the two coarse grid sessions, voters voted myopically 96% of the time (723/751). This is broken down in more detail, as follows. See Table 6.

In the CW treatment, when a member is faced with a choice a large share of a non-universal allocation (30) versus any other allocation (where their share would be either 20 or 15), they vote for the favorable allocation 96% of the time (180/187), and this does not depend on whether the favorable allocation was the status quo or the new proposal. When faced with a choice between the universal allocation and the smaller share of the non-universal allocation (15), a member voted for the universal allocation 95% of the time (195/206).[18]

In the NCW treatment, when faced with a choice between the alternative where they were out of the coalition and received 0, versus any other alternative, the member voted for the other alternative 99% of the time (272/276). When faced with a

[17] While not apparent from the anonymized tables, we found statistically significant evidence of non-anonymous proposal making in the NCW treatment. This arises when there is a majoritarian status quo. The two current coalition members are more than three times as likely (76% vs. 24%) to propose the same coalition, rather than proposing to switch partners. In contrast, no such asymmetry is observed in the proposal strategy of the "out" member (or in the case of majoritarian proposals, when the status quo is universal) with respect to which partner they propose to form a coalition with in a majoritarian outcome.

[18] This is evidence against a hypothesis of pro-social preferences playing a role in behavior here. On the whole, subjects in this experiment are not motivated by concerns for fairness.

choice between receiving 30 in a majoritarian allocation versus receiving 20 in the universal allocation, they voted myopically for the majoritarian allocation 85% of the time (68/80).

The similarity between the voting behavior in the two treatments is remarkable. The only real differences between the two tables is in the number of observations, not in the frequency of voting for one's myopically preferred alternative. Hence, this leads us to conclude unambiguously that *proposal behavior, not voting behavior, is what drives the differences in outcomes between the CW and NCW treatments.*

5.1.4 Quantal response equilibrium: Markov logit equilibrium (MLE)

As discussed earlier, because there are two regimes in the CW treatment, a quantal response equilibrium analysis will produce a somewhat different dynamics compared with the Markov Nash equilibrium. In the Markov Nash equilibrium, depending on the randomly assigned initial status quo, each committee will find itself stuck forever in exactly one regime, either the regime of the static Condorcet winner (where the universal outcome occurs every round) or the rotation regime, where outcomes randomly rotate around the three non-universal allocations. That is, they "lock in" on one regime from the very start and stay there forever, and this is entirely determined by the initial status quo. In a quantal response equilibrium, stochastic choice will result in long-run alternation between the two regimes, with the expected duration of a regime depending on the response parameter.

Table 7 reports the fitted and actual choice probabilities for all the different behavior strategies of the players, under the assumption of anonymity (e.g., the probability voter 1 votes for (30, 15, 15) over (20, 20, 20) equals the probability voter 2 votes for (15, 30, 15) over (20, 20, 20), etc.). The fitted choice probabilities are the MLE choice probabilities at the maximum likelihood value of lambda, estimated separately for the two treatments.[19] The estimated choice probabilities track the data reasonably well in the following sense. Figure 5 presents a scatter plot of the predicted and actual choice probabilities, and also shows the regression line, which has slope close to 1 and intercept close to 0, and $R^2 = .85$.[20]

5.2 Continuous allocation sessions

We analyze continuous Session 1 separately from Sessions 2 and 3, because the discount rate was different. Several interesting comparisons emerge, in spite of the fact

[19] Recall that each treatment obtained data from two separate sessions, one in which it was the first subsession and another where it was the second subsession. Lambdas estimated separately for each session are not significantly different.

[20] We also estimated a model that included a concavity parameter for the utility function. That led to a small improvement in fit, but the predicted strategies were very close to the QRE estimated strategies with linear utility. This is not surprising. Recall from Proposition 1 that the equilibrium with in the NCW treatment does not depend on concavity of the utility function.

Table 7 Empirical vs. MLE proposing and voting probabilities

SQ	Prop	Pr (proposal)		Pr (yes)	
		Fitted	Data	Fitted	Data
NCW treatment					
Lambda: 0.19					
Log likelihood: −867.33					
0	0	0.03	0.02		
0	20	0.11	0.15	0.82	1.00
0	30	0.86	0.85	0.97	1.00
20	0	0.01	0.01	0.18	0.00
20	20	0.19	0.31		
20	30	0.81	0.68	0.81	0.84
30	0	0.02	0.02	0.03	0.03
30	20	0.35	0.08	0.19	0.17
30	30	0.63	0.90		
CW treatment					
Lambda: 0.24					
Log likelihood:					−905.05
30	30	0.77	0.93		
30	20	0.14	0.07	0.21	0.06
30	15	0.09	0.00	0.10	0.00
20	30	0.70	0.47	0.79	0.93
20	20	0.19	0.48		
20	15	0.11	0.05	0.32	0.06
15	30	0.74	0.57	0.89	1.00
15	20	0.18	0.37	0.68	0.96
15	15	0.09	0.06		

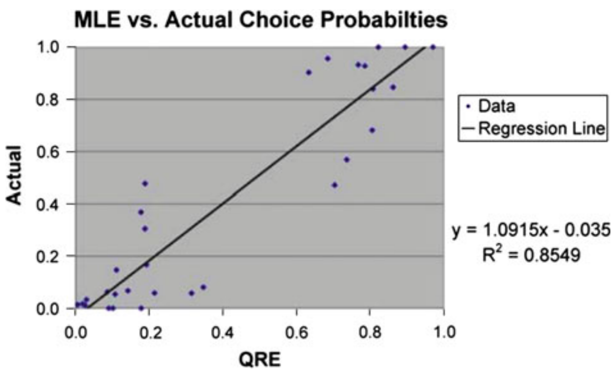

Fig. 5 MLE vs. actual choice probabilities

that the Markov equilibria are identical with linear preferences. As we will see, many of these differences can be explained extending the analysis to Markov QRE and considering concave utilities.

Table 8 State transition probabilities

SQ_t	SQ_{t+1}						
	D1	D2	D3	M12	M13	M23	U
$\delta = 0.83$							
D1	0.33 (3)	0.22 (2)	0.11 (1)	0.00 (0)	0.00 (0)	0.33 (3)	0.00 (0)
D2	0.14 (1)	0.14 (1)	0.14 (1)	0.00 (0)	0.43 (3)	0.14 (1)	0.00 (0)
D3	0.06 (1)	0.06 (1)	0.44 (8)	0.33 (8)	0.00 (0)	0.00 (0)	0.11 (2)
M12	0.02 (1)	0.06 (3)	0.00 (0)	0.55 (28)	0.14 (7)	0.12 (6)	0.12 (6)
M13	0.06 (3)	0.00 (0)	0.04 (2)	0.10 (5)	0.67 (34)	0.08 (4)	0.06 (3)
M23	0.00 (0)	0.04 (1)	0.11 (3)	0.11 (3)	0.18 (5)	0.39 (11)	0.18 (5)
U	0.00 (0)	0.01 (1)	0.00 (0)	0.07 (9)	0.02 (3)	0.02 (2)	0.88 (107)
Freq (SQ)	0.03	0.03	0.05	0.18	0.18	0.09	0.43
$\delta = 0.75$							
D1	0.33 (7)	0.10 (2)	0.14 (3)	0.05 (1)	0.00 (0)	0.33 (7)	0.05 (1)
D2	0.18 (3)	0.18 (3)	0.35 (6)	0.18 (3)	0.12 (2)	0.00 (0)	0.00 (0)
D3	0.13 (5)	0.16 (6)	0.42 (16)	0.16 (6)	0.00 (0)	0.00 (0)	0.13 (5)
M12	0.05 (4)	0.10 (8)	0.00 (0)	0.31 (25)	0.27 (22)	0.25 (20)	0.02 (2)
M13	0.06 (4)	0.00 (0)	0.03 (2)	0.22 (15)	0.46 (31)	0.18 (12)	0.04 (3)
M23	0.00 (0)	0.11 (7)	0.17 (11)	0.21 (13)	0.22 (14)	0.27 (17)	0.02 (1)
U	0.00 (0)	0.00 (0)	0.00 (0)	0.21 (11)	0.28 (15)	0.09 (5)	0.42 (22)
Freq (SQ)	0.07	0.08	0.11	0.22	0.25	0.18	0.10

Continuous data converted to 7-region grid

5.2.1 Empirical distribution and transition probabilities

Table 8 shows the transition frequencies, in percentages for each status quo, using the 7-region grid described in Sect. 3.1.2 (see Fig. 1).[21] Regions D1, D2, and D3 are the dictatorial regions; M12, M13, and M23 are the majoritarian regions, and U is the universal region. The top half is for patient committees and the bottom half for impatient committees. The last row in each half gives the relative frequencies of outcomes in each region. The left panel of Fig. 6 represents the empirical transition probabilities (as done in Fig. 2).

These tables yield the following results. From the last rows of the upper and lower half of Table 8, *this clearly refutes the Markov equilibrium prediction based on linear utilities that the committees outcomes will approximate a rotating dictatorship* (Table 2, Fig. 3a). We observe outcomes in the D regions only 10% of the time in

[21] The actual implementation in the laboratory was finite, with 1891 possible states (proposals to divide 60 into three non-negative integer allocations). For computational reasons, the Markov QRE and the Nash equilibrium benchmarks were computed assuming the grid described in Fig. 1 where the unit of measure is 5 (which implies 91 states). To compare the results, the states in the experimental state space were aggregated associating each of them to the closest state (in euclidean norm) in the coarser 91 state space. Nearly all (over 90%) proposed allocations observed in the experiment are divisible by 5, so the impact of this approximation is minimal.

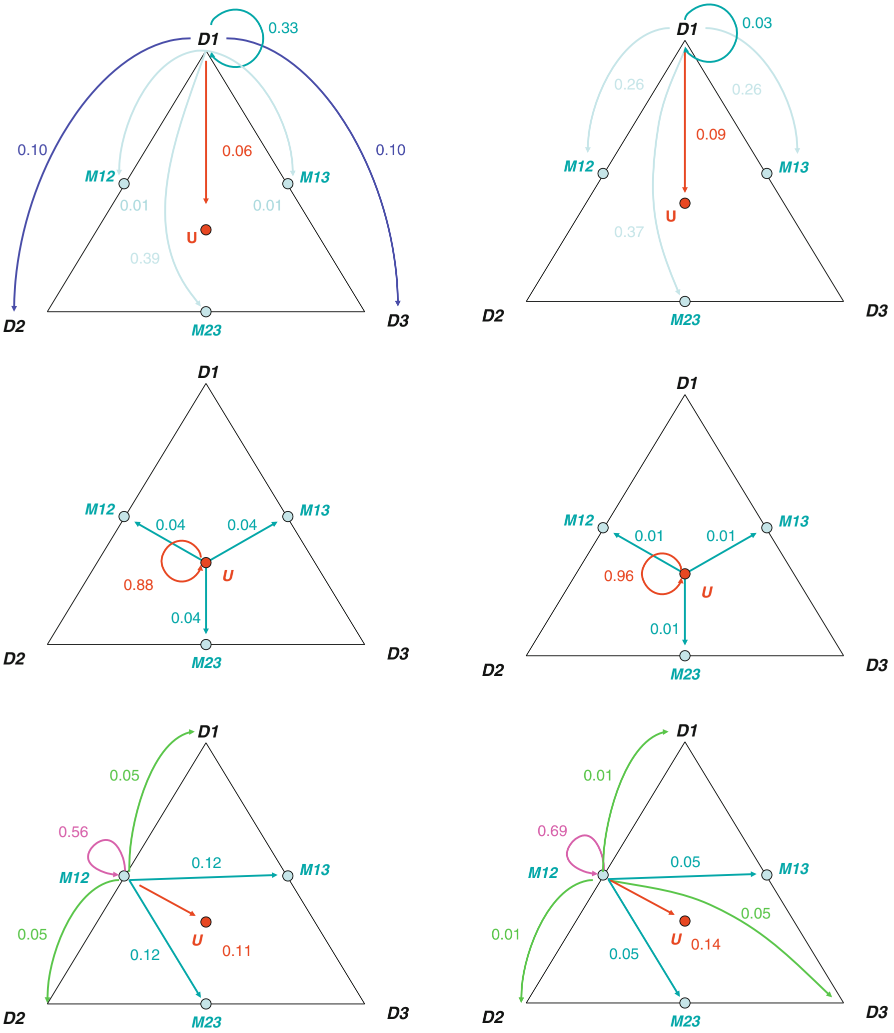

Fig. 6 Empirical vs. theoretical transition matrices, $\delta = 0.83$

the patient committees, and only 26% of the time in impatient committees. It is worth noting, however, that there is still a lot of persistence to dictatorial outcomes. Conditional on the status quo being in a D region, the outcome in the next period is almost twice as likely to be in one of these regions compared with non-D outcomes, with this persistence strongest for impatient committees. The reason so few outcomes are observed in these regions overall is that they are only occasionally reached from any other regions (less than 11% of the time).

In the *patient committees*, the remaining outcomes are divided equally between majoritarian outcomes and the universal outcome (approximately 45% of the time each). The transition matrix gives some indication of the dynamics of the patient committees. The universal outcomes function nearly, but not quite, as an absorbing state.

The empirical probability of moving away from U is only 12%. In the linear utility equilibrium, this probability is 100%, and it should be defeated by majoritarian proposals. Indeed, the latter property of equilibrium is observed to an extent in our data: i.e., when U is defeated, it is almost always by a majoritarian (2-person coalition) proposal. Majoritarian status quos are less stable than U, with the probability of transitioning to a new region being 44%. The theoretical (with linear utilities) probability of a majority status quo being defeated is 100%, and it should be defeated only by dictatorial proposals. In fact, we find the opposite: When an M status quo is defeated, it is usually by U, not by a dictatorial proposal. The D status quos are by far the least stable, in contrast to the linear utility model, which predicts them to be the *most* stable (being defeated only 33% of the time, always by another D outcome). In the experiment, they are defeated almost two-thirds of the time—essentially whenever one of the non-dictators is chosen as the proposer.

The distribution of outcomes for the *impatient committees* is different. D outcomes are somewhat more common, but still occur only slightly more than one-quarter of the time. Outside the D regions, outcomes are significantly more likely to lie in M than in U (64% vs. 10%, p<0.01), compared with the patient committee outcomes where there is almost no difference (45.5% vs. 43.1%). The reason for a difference in transition probabilities between the treatments is that U region is essentially absorbing in the patient committees, while in the impatient committees, a status quo of U is usually defeated immediately by an M proposal. The D and M regions all have strong persistence for impatient committees, stronger than in the patient committees. Conditional on being in a D region, the empirical probability of staying in some D region is 67% (for the patient committees, it was 53%). Conditional on being in a M region, the empirical probability of staying in that region or transitioning to a different 2-person coalition is over 80% (for the patient committees, it was 79%, so the difference in this case is not significant). In contrast, the U region is significantly more ($p < 0.01$) unstable: The probability of transitioning away from U is nearly 60%, compared with 12% in patient committees.

We now decompose the data into proposal and voting behavior. To compare the evidence with the coarse grid treatments discussed earlier, it is useful to represent the 7 regions of Fig. 1 in terms of just 5 "anonymously equivalent" regions. For example, for player 1, D2 and D3 are combined and M12 and M13 are combined. We have relabeled the 5 regions: D, DX, M, MX, and U. If we take agent i as a reference, D is the region where a proposer i receives the most and corresponds to Di in the 7 region partition; DX is composed by Dj and Dk for $j, k \neq i$; M is Mij and Mik; MX is Mjk; and U is the same as in the seven region partition.

5.2.2 *Proposal making and voting behavior*

Table 9 shows the aggregate proposing behavior for continuous committees for each status quo. The top half of the table is for the "patient" committees ($\delta = .83$). The bottom half of the table is for the "impatient" committees ($\delta = .75$). The entries in the table are the relative frequencies of proposal in that category, given the status quo, pooling across all committees, rounds, and members.

Table 9 Proposal probabilities, conditional on status quo

Status quo	Provisional proposal					
	D	DX	M	MX	U	# Obs
$\delta = 0.83$						
D	0.88	0.00	0.09	0.00	0.03	34
DX	0.34	0.01	0.51	0.00	0.13	68
M	0.19	0.00	0.66	0.01	0.14	260
MX	0.00	0.08	0.71	0.00	0.21	130
U	0.01	0.01	0.33	0.00	0.66	366
$\delta = 0.75$						
D	0.78	0.00	0.18	0.01	0.03	76
DX	0.58	0.00	0.35	0.01	0.07	152
M	0.28	0.01	0.65	0.00	0.06	422
MX	0.00	0.05	0.90	0.00	0.05	211
U	0.03	0.01	0.78	0.01	0.18	159

There are several notable features. First, as in the coarse grid data, subjects almost never propose to have the smallest allocation. That is, we see essentially no observations of a subject proposing an allocation in DX or in MX. Second, if a member is receiving a dictatorial allocation (i.e., the status quo is in D), then they almost always propose to stay in D, as predicted by the theory, and also consistent with the evidence from the discrete treatment. Interestingly, U proposals are almost never made by a subject in the D, DX, or M states. There is a lot of persistence of proposals in both the M and the D regions in both kinds of committees; subjects usually propose to stay in the same region, about two-thirds of the time.

In the *patient committees*, when the status quo is in the U region, U proposals are proposed two-thirds of the time and M proposals are made one-third of the time. While the linear utility equilibrium predicts almost all proposals in the U regions should be outside U, this is no longer true for $\gamma > 0$. Depending on the concavity of the utility function the equilibrium proposal when the status quo is in U can be in either region U or M. In the DX region (as seen in the analysis of the cases with $\gamma = 0$ and $\gamma = 0.95$ of Sect. 3.2.2), the Nash proposals are either in D or M, which is true 93% of the time in our data. For MX status quos, we never see D proposals, which is consistent with the equilibrium behavior because D proposals will always be defeated. In the M region, the non-M proposals are equally split between D and U.

The proposing behavior is different in the *impatient committees* compared with the patient committees, especially in the U region. In impatient committees, U allocations are rarely proposed at *any* status quo. Intuitively, this makes sense because the U allocation only is valuable under two conditions, first utilities must be sufficiently concave, and second, the continuation probability must be high enough to make the value of an equal share tomorrow be higher than a more than equal share today. To see this, consider the case where δ is very small, so that value is determined almost entirely by your current share in the allocation. Then, you would most prefer D ($\approx$60), next most prefer M ($\approx$30), and U ($\approx$20) is least preferred. This is reflected in the data

Table 10 Prob (vote for proposal) (# of observations)

Status quo	Active proposal				
	D	DX	M	MX	U
$\delta = 0.83$					
D	0.91 (11)	0.00 (7)	0.00 (2)	0.08 (12)	0.50 (2)
DX	1.00 (7)	0.83 (29)	1.00 (26)	0.50 (2)	1.00 (4)
M	0.89 (18)	0.11 (18)	0.78 (135)	0.09 (47)	0.45 (42)
MX	− (0)	0.83 (18)	1.00 (47)	0.57 (44)	1.00 (21)
U	0.75 (4)	0.13 (8)	0.66 (70)	0.09 (35)	0.86 (249)
$\delta = 0.75$					
D	0.91 (22)	0.00 (27)	0.40 (5)	0.13 (16)	0.33 (6)
DX	1.00 (27)	0.70 (71)	0.97 (37)	0.60 (5)	1.00 (12)
M	1.00 (39)	0.05 (41)	0.80 (204)	0.09 (116)	0.32 (22)
MX	1.00 (1)	0.92 (39)	1.00 (116)	0.41 (44)	1.00 (11)
U	1.00 (1)	0.00 (2)	0.84 (90)	0.04 (45)	0.67 (21)

where over 90% of proposals are in M or D. Clearly, the M proposals are most likely in the MX region, because these are the only proposals that can make the proposer better off and also make one other coalition partner better off.

Voting decisions are somewhat less myopic in the continuous allocation committees than in the CW and NCW committees with only 4 states. While 96% of voting decisions were myopic in those coarse grid treatments, only 83% of voting decisions were myopic for $\delta = .83$ and 92% of the decisions are myopic for $\delta = .75$. Still, an overwhelming fraction of voting decisions is myopic. Part of the lower frequency of myopic decision making in the continuous allocation committees is probably attributable to the fact that often subjects were faced with a choice between two alternatives that were quite close in terms of their share of the allocation. In the coarse grid treatments, no such pairs exist; a voter is either completely indifferent or receives a significantly higher share in one of the alternatives. But there is another factor that is also important. In contrast to the discrete sessions, the value functions in the continuous treatments are *not monotone increasing* in a member's own allocation.

The top half of Table 10 shows the aggregate voting behavior for (status quo, active proposal) pairs with patient committees, for each status quo, using the 5-region grid. The entries in the table are the empirical fraction of yes votes (in favor of the active proposal), pooling across all committees, rounds, and members. The bottom table shows the relative frequencies of the active proposals for each status quo category.

There are several results that follow from these tables. First, members almost always vote for D outcomes over any alternative. Second, members tend to vote for majoritarian outcomes, against any alternative, except when the status quo is in D. Even in that case, members vote for the majoritarian outcome one-third of the time, but these kinds of elections are rare events. Third, members tend to vote for U over other alternatives, except majoritarian outcomes, where the two members of the majoritarian coalition usually vote against U. Indeed, the predominance of U votes in the U state accounts for the bulk of the non-myopic voting behavior in the session.

The bottom half of Table 10 shows the aggregate voting behavior for (status quo, active proposal) pairs in impatient committees for each status quo, again using the 5-region grid. The results are much the same as the voting behavior in patient committees. The only difference is that impatient committees tended to vote for M proposals over U proposals with slightly greater frequency, but the difference is rather small in magnitude. This finding is similar to what we observed in the discrete allocation committees, where the main differences in behavior across treatments was in proposal behavior, not voting behavior.

5.3 Comparison with theory: QRE and concave utilities

For the continuous allocation environments, the predictions of the Markov equilibrium with linear preferences ($\gamma = 0$) are clearly rejected by the data, for at least three reasons. First, in both treatments, the dictatorial outcomes are infrequently observed (26% in the .75 treatment and 12% in the .83 treatment). Overall, more than 80% of outcomes are either majoritarian or universal. Second, as it can be seen by comparing Fig. 2 with the left panel of Figs. 6 and 7, the transition probabilities are quite different for the two sessions: The .83 treatment has many more U observations and more persistence in the U state but less persistence in the D and M states. But the Markov equilibrium with linear preferences predicts no treatment effect. Third, we observe many transitions (and proposals) that are predicted never to happen according to the Markov equilibrium with linear preferences.

A possible explanation for this finding is that agents have altruistic preferences. This hypothesis, however, seems contradicted by voting behavior, which is predominantly myopic and selfish, and proposing behavior: As we have seen in the previous section, proposers generally take advantage of proposal power when the status quo allows them to do it.

Based on the theoretical analysis earlier in the paper, we can suggest an alternative explanation for the findings, which is based on more standard assumptions on utility: that individuals have strictly concave utility functions. To test this hypothesis, we estimate the value of γ in each treatment, and also estimate a constrained value of γ, assuming it to be the same in both treatments, and assuming it to be the same across subjects. In order to obtain such an estimate, we use the Logit equilibrium as a structural model of the errors and therefore, simultaneously estimate λ and γ. The QRE model is a natural one to use for the error structure to estimate γ, because evidence from past experiments indicates a significant stochastic component of choice, which is correlated with equilibrium expected payoffs (McKelvey and Palfrey 1995, 1998). Because the stochastic choice *affects* expected payoffs, it will generally have additional equilibrium effects. The estimation is done using standard maximum likelihood methods. Using the path-following algorithm used earlier to compute Nash equilibria, we trace out the logit solution to the game; that is, we trace out a unique connected family of MLEs, $\{\rho(\lambda, \gamma), \sigma(\lambda, \gamma)\}$ for increasing values of lambda and γ ranging from 0 to 1. This defines a likelihood function $L(\widehat{\rho}, \widehat{\sigma}; \gamma, \lambda)$, where $(\widehat{\rho}, \widehat{\sigma})$ are the observed proposal and voting choice frequencies in the data, using the 91 state grid defined earlier in the paper.

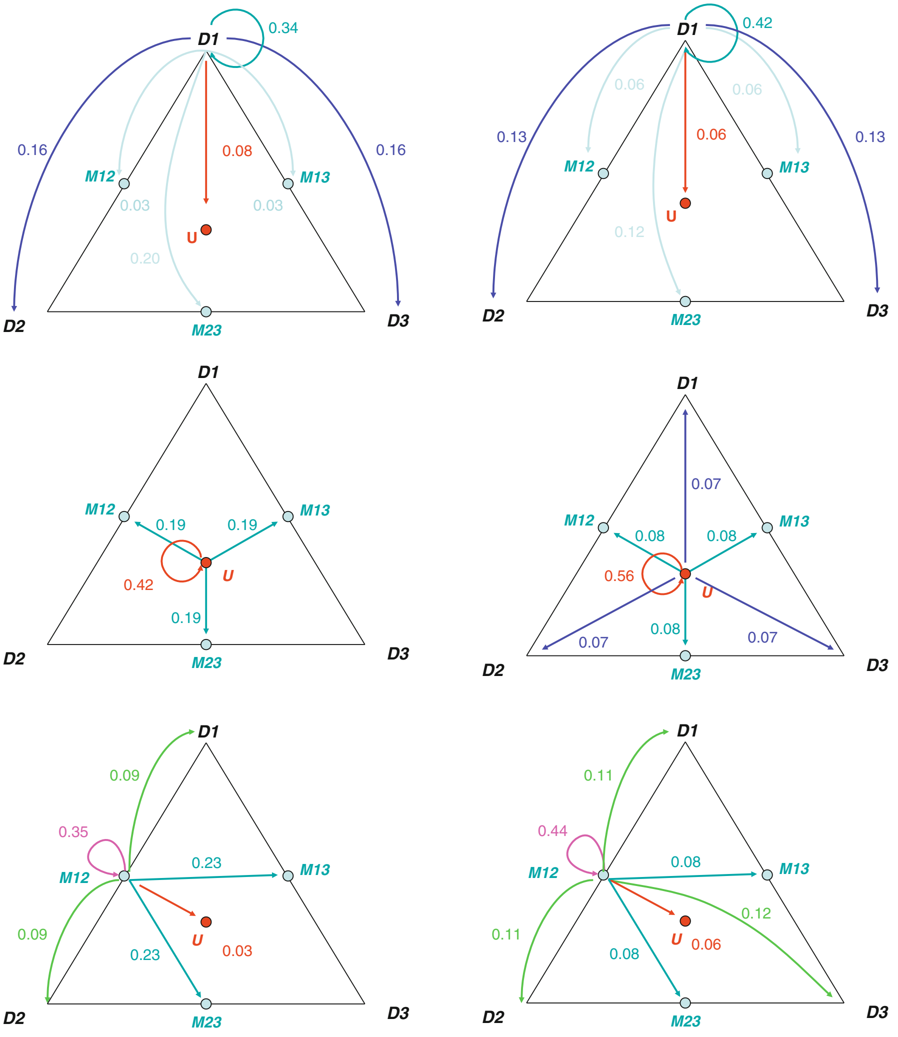

Fig. 7 Empirical vs. theoretical transition matrices, $\delta = 0.75$

Table 11 MLE parameters

δ	$\widehat{\gamma}$	$\widehat{\lambda}$	$-\ln L$
.83	0.70	8.01	3, 827
.75	0.40	1.20	3, 792
Pooled	0.50	1.68	7, 643

Table 11 below gives the results of the estimation for the continuous sessions. In both treatments, the concavity parameter is highly significant.[22] The estimates we

[22] A likelihood ratio test rejects the $\widehat{\gamma} = 0$ model in both treatments, and in the pooled data, at very high significance ($p < 0.001$).

obtain are also in the same range (between 0.40 and 0.70) as concavity estimates from many other sources of data, including auction experiments (Goeree et al. 2002), abstract game experiments (Goeree et al. 2003), lottery choice experiments (Holt and Laury 2002), and field data from auctions (Campo et al. 2011). The Chi-square test shows a significant difference between $(\widehat{\lambda}, \widehat{\gamma})$ for the two treatments. However, the differences between the constrained (pooled) estimates and the separate estimates in terms of the fit to voting and proposing behavior are small. The improvement in likelihood with the separate estimates is less than one-third of one percent, compared with the constrained (pooled) estimates. While the estimates of $\widehat{\lambda}$ appear quite different across the two treatments, this is exaggerated by a ridge in the likelihood functions, where slightly higher values of $\widehat{\gamma}$ lead to much higher estimates of $\widehat{\lambda}$, especially with the .83 data, but with virtually no change in the likelihood function. For example, in the .83 data, $-\ln L[\widehat{\gamma} = .65, \widehat{\lambda} = 3.01] = 3{,}832$ and $-\ln L[\widehat{\gamma} = .70, \widehat{\lambda} = 8.01] = 3{,}827$.

Figure 6 compares the empirical transition matrices for the two treatments, using the 7-state grid with the theoretical transition matrices implied at the estimated values of $\widehat{\gamma}$ and $\widehat{\lambda}$. The empirical transitions are obtained by using symmetry to condense the 7 regions in Table 8 into five regions (as in Table 10). Thus, for example in the upper left triangle diagram, we represent all transitions from D as transitions from D1. Transition frequencies from D to DX are equal to the sum of the two numbers in the arrows of that diagram from D1 to D2 and D1 to D3, and these two numbers are always equal by symmetry. Empirical transitions elsewhere in the left side triangles of that figure are represented similarly.

The theoretical transitions for this condensed grid are derived in the following way. First, using the 91-state grid, the estimated parameters directly imply estimated proposal and voting strategies in the $(\widehat{\gamma}, \widehat{\lambda})$ Markov Logit equilibrium. We then use those estimated proposal and voting strategies to obtain a 91×91 transition matrix, for the states shown in Fig. 1. In order to condense this into a 7×7 transition matrix, we need to weight the states within each of the 7 coarsened states according to the probability those states would theoretically occur. The apparently obvious way to do this by using the stationary distribution implied by the theoretical transition matrix turns out to be incorrect, because in our experiment we used a random stopping rule rather than playing the game an infinite number of times. Hence, the actual distribution is influenced by the initial status quo allocations, which were uniformly distributed on the simplex. Thus, we compute the expected distribution, given that we start round 1 with a uniform distribution over the allocations, and then compute the implied distribution over future allocations, given the stationary stopping rule (either .83 or .75). These weights are then used to coarsen the transition matrix to 7×7.[23]

[23] To avoid clutter, the figure only shows the transitions from D1 (top), U (middle), and M12 (bottom). By symmetry, the theoretical transitions (right half of figure) from D2 and D3 are identical to the one illustrated for D1; transitions from M13 and M23 are identical to the one illustrated for M12. The empirical transitions illustrated from D1 (top left) represent the average across all three D regions, similarly for the empirical transitions from M12 (bottom left).

There are some apparent similarities between the theoretical transition matrices and the empirical ones, but also some differences. First, for both the theoretical and empirical transitions, the status quo has a lot of persistence in all regions of both treatments, with the single exception being the D regions of the .83 treatment. Specifically, the probability of staying in a D region is 53%, while the theoretical probability is only 3%. This is the one case where the fitted QRE transitions track the empirical transitions poorly, but it is based on very few observations. In .75 treatment, for the theoretical transitions at the $(\widehat{\lambda}, \widehat{\gamma})$ estimates, the probability of staying in a D region, given the status quo is a D outcome is nearly 70%; it is 60% for the M regions; and 56% for the U region. The corresponding empirical probabilities are 66, 80, and 42%, respectively. For the .83 treatment, these theoretical persistence probabilities for the M and U regions are 79 and 96%, respectively, compared with the empirical findings of 78 and 88%, respectively. For the exceptional case of the D regions in the .83 treatment, the theoretical persistence probability is 3% and the observed finding is 55%. This is the one case where the fitted QRE transitions track the empirical transitions poorly, but it is based on very few observations.[24]

There are also similarities between the empirical and theoretical non-persistent transitions (D $\rightarrow$ M or U, M $\rightarrow$ D or U, and U $\rightarrow$ D or M). For the .83 committees, when the status quo transitions out of a D region, we find that it goes to an M allocation 88% of the time, identical to the theoretical transition probability from D to M; when the status quo leaves the U region, we observe it going to an M region over 90% of the time, where the theoretical conditional transition probability is 100% to an M region; from an M region we observe it occasionally going to both U and D, but somewhat more frequently to U, consistent with the theoretical transition probabilities. For the .75 committees, the non-persistent transitions also track the theoretical transitions reasonably well: From D, we observe transitions to M three times more often than transitions to U, where the theoretical ratio is two to one; from M, transitions to D account for 81% of the non-persistent transitions in our data, compared the theoretical transition probability of 79%; from U, transitions to M are more likely than transitions to D theoretically, and we observe this in the data. However, in the latter case, we actually observe zero transitions from U, which is not consistent with the fitted estimate of 81%.

From this estimation, we conclude two main results. First, we can reject the hypothesis that the utility of the agents is linear with very high significance. In the distributive problem under analysis, concavity of the utility function seems to be an essential ingredient for an explanation that accounts for the empirical behavior that we observe in the laboratory. Second, the MLE model with concave preferences can account for much of the the qualitative empirical evidence, on the basis of comparing the observed transitions with the fitted theoretical transitions implied by the estimated MLE model.

[24] There were only 36 observations of a status quo in D, out of a total of 288 observations in the $\delta = .83$ treatment.

6 Concluding remarks

We studied dynamic committee behavior in multiperiod environments that are linked dynamically by a sequence of endogenous status quo outcomes, where the status quo outcome in period t is determined by the committee decision in period $t + 1$. Such models of repeated bargaining dynamics are intended to study the inertia that is a natural consequence of procedures used by many standing committees or legislatures. Specifically, outcomes in earlier periods have long-run effects that are created simply by the powerful and special status quo position embodied in past decisions.

Several parametric versions of these environments were studied in the laboratory, and the environments were varied along three dimensions. First, we consider both continuous divide-the-dollar environments and highly constrained allocations. Second, for highly constrained settings, we consider Condorcet and non-Condorcet environments. Third, we look at the effect of committee "patience", or long-run incentives, by varying the effective discount factor.

Many specific findings were reported in the data analysis. These findings can be boiled down to six main results.

- Result 1: Outcomes have strong persistence, and this persistence roughly follows the theoretical prediction, but with significant stochastic variation, as captured in the QRE model.
- Result 2: Voting behavior in all the different environments is selfish and myopic.
- Result 3: To the extent that the outcomes deviate from equilibrium predictions, it is mainly due to differences between actual proposal behavior and observed proposal behavior.
- Result 4: Patient committees exhibit substantially different proposal behavior than impatient committees.
- Result 5: We observe more universal and majoritarian outcomes than predicted by the Markov equilibrium with linear preferences.

The fact that in all treatments, we see more universal outcomes than expected based on the "standard" Markov equilibrium theory with linear preferences might tempt some to interpret as evidence of fairness and pro-social preferences. However, evidence from voting and proposal behavior does not seem to support such an interpretation. Voters consistently vote selfishly for outcomes that are extremely lopsided and unfair—provided they gain personally by doing so. Proposers tend to take advantage of their proposal power whenever the status quo allows them to do it.

We are able to show theoretically that universal outcomes will arise more often over time if preferences are concave. Using quantal response equilibrium, we estimate the concavity of the utility function and find it to be a significant factor. The coefficient of constant relative risk aversion is estimated on the pooled continuous data to be .5, similar to estimates of utility function curvature in a variety of other economics experiments. Based on a comparison of the empirical state transitions and the theoretical transition probabilities generated by the model estimates, we find

- Result 6: The data reject the model of linear utility, and the concave utility model accounts for several prominent features of behavior and the trajectory of allocations in our dynamic distributive politics environment.

The specific framework in this paper focused on dynamic distributive politics, but in principle it could be generalized to more complex political environments. In practice, in addition to purely distributive allocations, there are public good decisions, some of which have important dynamic components. Legislative committees make important decisions about the production and accumulation of real resources, for example public infrastructure and other durable public goods, and may finance these expenditures by running public debt, all of which has long-run implications for future policy decisions.[25] The next step of our research program is to extend the theoretical model in this direction and to collect and study experimental data in such environments.

7 Appendix

7.1 Markov quantal response equilibrium

In this appendix, we define a *Markov quantal response equilibrium* for the case of J feasible alternatives (states) and N legislators. We will focus on a particular version in which the quantal response function is logit, which we call *Markov Logit equilibrium* (*MLE*). As in the Markov Nash equilibrium, the state of the system is given by the status quo, and hence the state space is equal to the set of feasible alternatives.[26]

7.1.1 Expected utilities

We define v_{ji} to be the expected continuation utility of agent i in state j (so before the proposer is chosen) and a $J \times N$ matrix collecting all these values. All the following equations take V as inputs.

We define $\sigma_k^{ji}(V)$ to be the probability of voting yes to a proposal x_k in state x_j by agent i. Let $P_y^i(V)$ and $F_y^i(V)$ be the probabilities that the proposal, respectively, passes or fail if i votes *yes* in an equilibrium in which the value function is V. When, as in the experiment, $N = 3$, we have $P_y^i(V) = 1 - \Pi_{l \neq i}(1 - \sigma_k^{jl}(V))$ and $F_y^i(V) = \Pi_{l \neq i}(1 - \sigma_k^{jl}(V))$. The expected utility from voting yes is $(x_{ki} + \delta v_{ki}) P_y^i(V) + \left(x_{ji} + \delta v_{ji}\right) F_y^i(V)$. Similarly, we can define the probabilities the proposal passes or fails if you vote *no* $P_n^i(V)$ and $P_n^i(V)$, which in the N $= 3$ case are $P_n^i(V) = \Pi_{l \neq i}\sigma_k^{jl}(V)$ and $F_n^i(V) = 1 - \Pi_{l \neq i}\sigma_k^{jl}(V)$. The expected utility from voting no is $(x_{ki} + \delta v_{ki}) P_n^i(V) + \left(x_{ji} + \delta v_{ji}\right) F_n^i(V)$. Since we are using Logit equilibrium, we use the logit quantal response function, which yields the following collection of MLE conditions characterizing the voting stage:

[25] For a theoretical analysis of some of these issues, see Battaglini and Coate (2006, 2007, 2008) and Battaglini et al. (2009).

[26] For an extensive discussion on the concept of quantal response equilibrium in normal and extensive form games see McKelvey and Palfrey (1995, 1998).

$$\sigma_k^{ji}(V) = \frac{\exp\lambda\left\{\begin{array}{c}(x_{ki}+\delta v_{ki})\,P_y(V)\\+\left(x_{ji}+\delta v_{ji}\right)F_y(V)\end{array}\right\}}{\exp\lambda\left\{\begin{array}{c}(x_{ki}+\delta v_{ki})\,P_y(V)\\+\left(x_{ji}+\delta v_{ji}\right)F_y(V)\end{array}\right\}+\exp\lambda\left\{\begin{array}{c}(x_{ki}+\delta v_{ki})\,P_n(V)\\+\left(x_{ji}+\delta v_{ji}\right)F_n(V)\end{array}\right\}}$$

where $\sigma_k^{ji}(V)$ is the probability that committee member i would vote for proposal x_k if the status quo is x_j.

7.1.2 Proposal equilibrium conditions

We define $u_k^{ji}(V)$ the expected utility of agent i in state x_j when x_k is proposed (before the vote):

$$\begin{aligned}u_k^{ji}(V) = (x_{ki}+\delta v_{ki})&\left[\sigma_k^{ji}(V)\,P_y^i(V)+\left(1-\sigma_k^{ji}(V)\right)P_n^i(V)\right]\\&+\left(x_{ji}+\delta v_{ji}\right)\left[\sigma_k^{ji}(V)\,F_y^i(V)+\left(1-\sigma_k^{ji}(V)\right)F_n^i(V)\right].\end{aligned}$$

Hence, the equilibrium conditions for the proposal stage are given by $\rho_k^{ji}(V) = \frac{\exp\left(\lambda u_k^{ji}(V)\right)}{\sum_{k'=1}^{J}\exp\left(\lambda u_{k'}^{ji}(V)\right)}$, where $\rho_k^{ji}(V)$ is the probability that committee member i would propose x_k in state x_j.

7.1.3 The fixed-point problem

The following expression defines a $N \times J$ equations that map expected utilities V to expected utilities V.

$$v_{ji} = \sum_{l=1}^{N}\alpha_l\left[\sum_{k=1}^{J}\rho_k^{jl}(V)\,u_k^{ji}(V)\right] \tag{7}$$

where α_l is the probability that agent l is selected as proposer. The fixed points of this equation are MLE of the bargaining game. The fixed point of (7) at a given λ° is computed by homotopy methods. We know the fixed point of (7) at $\lambda = 0$. Though (7) is not a contraction, it behaves as a contraction in a neighborhood of a fixed point. We can therefore find the fixed point at λ° by tracing the fixed points of (7) as λ gradually increases. Obviously one cannot compute the equilibrium for all positive values of λ, and one must cut off the computation at some maximum λ°. In the analysis of an approximate Nash equilibrium, in Sect. 3, we compute the equilibrium up to $\lambda^\circ = 20$. Changes in MLE strategies beyond $\lambda = 10$ were negligible.

7.1.4 Proof of Proposition 1

We first show that the strategies described in the proposition are an equilibrium for any strictly increasing utility U. We can normalize utility by $U(0) = 0$, $U(30) = 1$, and

$U(20) = a \in (0, 1)$. Then, the value functions for player i can be written as follows:

$$u_i(x_0) = a + \delta v_i(x_0), \quad u_i(x_i) = \delta v_i(x_i), \quad u_i(x_{-i}) = 1 + \delta v_i(x_{-i}), \tag{8}$$

where $v_i(x)$ is the continuation expected utility of being in state x. The continuation values are determined by the recognition probabilities (1/3 each) and the proposal strategies, and the voting strategies. By hypothesis, these strategies are as stated in Proposition 1, and therefore we get

$$\begin{aligned} v_i(x_0) &= \frac{1}{3}u_i(x_i) + \frac{2}{3}u_i(x_{-i}), \\ v_i(x_i) &= \frac{2}{3}u_i(x_i) + \frac{1}{3}u_i(x_{-i}), \\ v_i(x_{-i}) &= \frac{1}{6}u_i(x_i) + \frac{5}{6}u_i(x_{-i}). \end{aligned} \tag{9}$$

From here, it is straightforward to show that, for all $a \in (0, 1)$, the value functions are ordered

$$u_i(x_{-i}) > u_i(x_0) > u_i(x_i), \tag{10}$$

which is sufficient to show that the strategies described in Proposition 1 form an equilibrium. There are two steps. First, we show that $u_i(x_{-i}) > u_i(x_i)$. To see this, subtract the second equation in 8 from the third equation, to get

$$u_i(x_{-i}) - u_i(x_i) = 1 + \delta \left[v_i(x_{-i}) - v_i(x_i)\right] = 1 + \frac{1}{2}\delta \left[u_i(x_{-i}) - u_i(x_i)\right]$$

with the second step resulting from the substitution of the last two equations of 9. Hence, $u_i(x_{-i}) - u_i(x_i) = \frac{2}{2-\delta} > 0$. To see that $u_i(x_{-i}) - u_i(x_0) > 0$, observe that from 8 we get:

$$u_i(x_{-i}) - u_i(x_0) = (1 - a) + \delta \left[v_i(x_{-i}) - v_i(x_0)\right].$$

Using 9, the result follows immediately. It is also easy to show, using exactly the same line of argument, that the same result (rotating minimum winning coalitions) extends to any odd number N of committee members, where the set of alternatives includes all equal-split minimum winning coalitions, plus the universal outcome.

Note that in any equilibrium in which (10) is true, strategies must be as described in the proposition. So to establish uniqueness when utilities are linear, we only need to show that (10) must be satisfied in an equilibrium (as defined in Sect. 2). It is easy to see that in any (symmetric) equilibrium, the utility of each agent is identical in state x_0: $u_i(x_0) = u_j(x_0)\ \forall i, j$. Moreover, by symmetry, we also have that $u_i(x_j) = u_k(x_j)$ for $i, k \neq j$, and $u_i(x_i) = u_j(x_j)$ for $\forall j, i$. Finally note that, with linear utilities, for any state x we must have $\sum_{i=1}^{3} u_i(x) = \frac{60}{1-\delta}$ for any x. Together with symmetry, these conditions imply that $u_i(x_0) = \frac{20}{1-\delta}\ \forall i$.

We first show that $u_i(x_j) \geq u_i(x_0)$ in any equilibriium. Suppose to the contrary that in some equilibrium $u_i(x_j) < u_i(x_0)$ for $j \neq i$. This would imply that $u_i(x_i) < u_i(x_j)$. For, if instead we had $u_i(x_i) \geq u_i(x_j)$ then the worst continuation utility for agent i in state x_j would be $u_i(x_j)$. This would imply $u_i(x_j) \geq 30 + \delta u_i(x_j)$, i.e. $u_i(x_j) \geq \frac{30}{1-\delta} > \frac{20}{1-\delta} = u_i(x_0)$, a contradiction. Now $u_i(x_i) < u_i(x_j)$ implies

$$60 = (1-\delta) \sum_{i=1}^{3} u_i(x_j) = (1-\delta)\left[2u_i(x_j) + u_j(x_j)\right] < 3\,(1-\delta)\,u_i(x_j)$$

since $u_j(x_j) = u_i(x_i)$ by symmetry, which in turn would imply $u_i(x_j) > u_i(x_0)$, a contradiction. We conclude that $u_i(x_j) \geq u_i(x_0)$.

We next show that $u_i(x_j) > u_i(x_0)$. Suppose to the contrary that $u_i(x_j) = u_i(x_0)$. There are two possible cases. Case 1 arises if $u_i(x_i) \geq u_i(x_0)$. In this case, the continuation utility for agent agent i in state x_j is never below $u_i(x_0)$, so $u_i(x_j) \geq 30+\delta u_i(x_0) = 10+u_i(x_0) > u_i(x_0)$, a contradiction. Case 2 arises if $u_i(x_i) < u_i(x_0)$. Then, for $j \neq 0$, we have

$$\begin{aligned} 60 &= (1-\delta) \sum_{i=1}^{3} u_i(x_j) = (1-\delta)\left[2u_i(x_j) + u_j(x_j)\right] \\ &< (1-\delta)\left[2u_i(x_j) + u_j(x_0)\right] = 3\,(1-\delta)\,u_i(x_0) \end{aligned}$$

which implies that $u_i(x_0) > \frac{20}{1-\delta}$, again a contradiction. Hence, we have shown that $u_i(x_j) > u_i(x_0)$.

Finally, we show that $u_i(x_i) < u_i(x_0)$. Suppose to the contrary that $u_i(x_i) \geq u_i(x_0)$. Then, using the same logic as before, we have that for $j \neq 0$: $60 = (1-\delta)\left[2u_i(x_j) + u_j(x_j)\right] > (1-\delta)\sum_{i=1}^{3} u_i(x_0)$, a contradiction. We conclude that $u_i(x_j) > u_i(x_0) > u_i(x_i)$, which proves the result. □

7.1.5 Proof of Proposition 2

If strategies are myopic as defined in Sect. 3.1.1, the value function is described by (3) and (6). It is easy to verify that $v_i(x_i) > v_i(x_0) > v_i(x_j)$. This implies that the voting behavior described in a myopic strategy profile is optimal. Given this, the weakly stage undominated strategy for an agent i is always to propose x_i in state x_i or x_0. Consider now the case of an agent i in state x_j. It is easy to verify that it is never optimal to propose x_k for $k \neq i, 0$: This choice would certainly yield a lower utility than proposing, for example, x_0, which would always win. When the state is x_j, it is optimal for agent i to propose x_i if it yields higher expected utility than x_0, that is if $\left(\frac{45}{2} + \frac{\delta}{2}(v_i(x_i) + v_i(x_j))\right) \geq \frac{20}{(1-\delta)}$. The right hand side of this inequality is the utility if x_0 is proposed (given that it wins with probability one), and the left hand side is the expected utility of proposing x_i, which wins only with probability $\frac{1}{2}$. Given (6), it can be verified that this inequality is true for all $\delta \in [0, 1]$. We conclude that the myopic strategies described in Sect. 3.1.1 are an equilibrium.

We now prove that they must be the unique equilibrium. Proceeding as in Proposition 1, we can show that a symmetric equilibrium must be monotonic: so $u_i(x_i) > u_i(x_0) > u_i(x_j)$. We therefore only need to prove that any monotonic equilibrium must adopt myopic strategies. It is easy to see that any monotonic equilibrium must have the same voting behavior as described in the myopic strategies. Given this, it is also immediate to see that proposal behavior must be as described in myopic strategies for each agent i in states x_i and x_0. Consider now proposal behavior of agent i in state x_j. As before we can rule out the case in which x_k for $k \neq i, 0$ is proposed. Assume that the agent proposes x_0 with positive probability, say with probability a. In this case, it must be that

$$\left(\frac{45}{2} + \frac{\delta}{2}(v_i(x_i) + v_i(x_j))\right) \leq \frac{20}{(1-\delta)}. \tag{11}$$

Moreover, it must be that in any state the sum of payoffs sums up to $\frac{60}{(1-\delta)}$, so by symmetry:

$$v_i(x_i) + 2v_i(x_j) = \frac{60}{(1-\delta)}. \tag{12}$$

From (11)–(12), we obtain $v_i(x_j) \geq \frac{15\delta+5}{\delta(1-\delta)}$. Using this fact, we note that $u_i(x_j) = 15 + \delta v_i(x_j) \geq \frac{20}{(1-\delta)} = u_i(x_0)$, a contradiction. We conclude that a symmetric equilibrium must be in myopic strategies.

References

Baron, D.P.: A dynamic theory of collective goods procedures. Am Polit Sci Rev **90**, 316–330 (1996)

Baron, D.P., Diermeier, D., Fong, P.: Dynamic theory of parliamentary democracy. Economic Theory (2011, this issue)

Baron, D.P., Ferejohn, J.A.: Bargaining in legislatures. Am Polit Sci Rev **83**, 1181–1206 (1989)

Baron, D.P., Herron, M.: A dynamic model of multidimensional collective choice. In: Kollman, K., Miller, J., Page, S. (eds.) Computational Models of Political Economy, pp. 13–47. Cambridge: MIT Press (2003)

Battaglini, M., Coate, S.: Inefficiency in legislative policymaking: a dynamic analysis. Am Econ Rev **97**, 118–149 (2006)

Battaglini, M., Coate, S.: Fiscal policy over the real business cycle: a positive theory. NBER Working Paper No. 14047 (2007)

Battaglini, M., Coate, S.: A dynamic theory of public spending, taxation and debt. Am Econ Rev **98**(1), 201–236 (2008)

Battaglini, M., Nunnari, S., Palfrey, T.: Political institutions and the dynamics of public investment. Social Science Working Paper #1318, California Institute of Technology (2009)

Bowen, R., Zahran, Z.: On dynamic compromise. Research Paper No. 2020, Stanford Graduate School of Business (2009)

Campo, S., Guerre, E., Perrigne, I., Vuong, Q.: Semiparametric estimation of first-price auctions with risk averse bidders. Rev Econ Stud **78**, 112–147 (2011)

Diermeier, D., Fong, P.: Dynamic legislative bargaining with reconsideration. mimeo (2009)

Diermeier, D., Morton, R.: Experiments in majoritarian bargaining. In: Austen-Smith, D., Duggan, J. (eds.) Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks, pp. 201–226. Heidelberg: Springer (2006)

Diermeier, D., Gailmard, S.: Self-interest, inequality, and entitlement in majoritarian decision-making. Q J Polit Sci **1**, 327–350 (2006)

Duggan, J., Kalandrakis, T.: Dynamic legislative policy making. Working paper, University of Rochester (2010)
Epple, D., Riordan, M.: Cooperation and punishment under repeated majority rule. public Choice **55**(1–2), 41–73 (1987)
Frechette, G., Kagel, J.H., Lehrer, S.F.: Bargaining in legislatures: an experimental investigation of open versus closed amendment rules. Am Polit Sci Rev **97**, 221–232 (2003)
Frechette, G., Kagel, J.H., Morelli, M.: Gamson's law versus non-cooperative bargaining theory. Games Econ Behav **51**, 365–390 (2005a)
Frechette, G., Kagel, J.H., Morelli, M.: Nominal bargaining power, selection protocol, and discounting in legislative bargaining. J. Public Econ **89**, 1497–1517 (2005b)
Frechette, G., Kagel, J.H., Morelli, M.: Behavioral identification in coalitional bargaining: an experimental analysis of demand bargaining and alternating offers. Econometrica **73**, 1893–1938 (2005)
Frechette, G., Kagel, J.H., Morelli, M.: Pork versus public goods: an experimental study of public good provision within a legislative bargaining framework. Econ Theory (2011, this issue)
Goeree, J.K., Palfrey, T.R., Holt, C.A.: Quantal response equilibrium and overbidding in first price auctions. J Econ Theory **104**(1), 247–272 (2002)
Goeree, J.K., Palfrey, T.R., Holt, C.A.: Risk averse behavior in generalized matching pennies games. Games Econ Behav **45**(1), 97–113 (2003)
Holt, C.A., Laury, S.: Risk aversion and incentive effects. Am Econ Rev **92**(5), 1644–1655 (2002)
Kalandrakis, T.: A three player dynamic majoritarian bargaining game. J Econ Theory **16**(2), 294–322 (2004)
Kreps, D. M., Wilson, R. B.: Sequencial equilibria. Econometrica **50**, 863–894 (1982)
McKelvey, R.D: An experimental test of a stochastic game model of committee bargaining. In: Palfrey, T. (ed.) Laboratory Research in Political Economy, pp. 139–169. Ann Arbor: University of Michigan Press (1991)
McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. Games Econ Behav **10**, 6–38 (1995)
McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for extensive form games. Exp Econ **1**, 9–41 (1998)
Penn, E.M.: A model of farsighted voting. Am J Polit Sci **53**(1), 36–54 (2009)