PERSISTENCE OF POWER: REPEATED MULTILATERAL BARGAINING

MARINA AGRANOV, CHRISTOPHER COTTON, AND CHLOE TERGIMAN

Abstract. OLD ABSTRACT. In a variety of settings, budgets are set by a committee that interacts repeatedly over many budget cycles. To capture this, we study a model of repeated multilateral bargaining by a budget committee. Our focus is on the transition of agenda setting power from one cycle to the next, and how such considerations affect bargaining and coalition formation over time. Specifically, we compare a rule that approximates the budget process in many parliamentary democracies in which a vote of confidence is traditionally attached to each budget proposal, and a rule that approximates the budget process in congressional systems where party leadership must maintain the support of a majority of other legislators to hold onto power. As is standard in the literature, we use stationary equilibrium refinements to make predictions about behavior in our environments. In a controlled laboratory experiment, we find no support for the standard equilibrium refinements used in the literature. In sharp contrast to the theoretical predictions, in the experiment, both rules give rise to stable and persistent coalitions in terms of coalition size, identity, and shares of coalition partners and feature high persistence of agenda-setter power. Our results call into question the validity of restricting attention to history independent strategies in dynamic bargaining games. We conclude by showing that weakening the standard equilibria concepts to allow players to condition on one piece of history (the most recent deviator) is enough to generate equilibria which are consistent with outcomes and behavior observed in the experiments.

Date: **Current version: August 15, 2017**.

Key words and phrases. stationary equilibrium, dynamic bargaining, repeated games, laboratory experiment

JEL: C78, D02, C92.

Agranov: magranov@hss.caltech.edu, Division of Humanities and Social Sciences, California Institute of Technology. Cotton: cotton@econ.queensu.ca, Department of Economics, Queen's University. Tergiman: cjt16@psu.edu, Smeal College of Business, Penn State University. We appreciate feedback from Alessandra Cassela, Amanda Friedenberg, Alex Hirsch, Navin Kartik, Justin Leroux, Salvatore Nunnari, Pietro Ortoleva, Thomas Palfrey and participants at seminars and workshops at Columbia University, UCSD Rady, Caltech, University of British Columbia, the Warwick-Princeton Political Economy workshop in Venice, and the Canadian Public Economic Group meetings. Cotton is grateful for financial support provided through his position as the Jarislowsky-Deutsch Chair in Economic and Financial Policy at Queen's. Tergiman is grateful for generous funding from the Social Sciences and Humanities Research Council of Canada (SSHRC), which provided funding for the project.

1. Introduction

Many bargaining situations involve repeated interactions. This is true in personal relationships, and is also the case in more formal decision making bodies such as committees, legislatures, and corporate boards. Budget committees meet every year to bargain over the allocation of scarce resources. Standing committees in legislatures repeatedly interact to determine policy and regulation. In repeated interactions such as these, members of the decision-making body may be able to develop reputations and sustain long-term relationships or implicit agreements with one another that may not be possible in groups that only ever interact once and agree on a single decision.

Despite ample evidence suggesting that reputation and relationships play an important role in organizational decision making, theoretical models of repeated bargaining tend to limit attention to stationary equilibria (depending on the environment, either stationary subgame perfect equilibria or stationary Markov perfect equilibria), which restrict attention to memoryless strategies. Limiting attention to such equilibria allows analyses to rule out many of the subgame perfect equilibria which exist in these games, and focus attention on what is sometimes a unique outcome. This simplifies analysis and increases prediction power.

However, a focus on stationarity comes with two substantial costs. First, stationarity assumes that players are only forward looking, preventing them from conditioning their strategies on past behavior. This rules out the potential to reward or punish based on past generosity or the lack there of. It doesn't allow players to develop reputations and maintain relationships, things that may very well be important when groups must repeatedly come to agreements. Second, focusing attention on a single equilibrium imposes homogeneity in outcomes. Further, the possibility that under a fixed set of bargaining procedures or bargaining rules different groups may coordinate on wildly different equilibria is ignored.

In this paper, we develop a simple model of repeated bargaining and an experimental test to illustrate some issues with focusing on stationarity. We show two main results. First, allowing players to condition on past actions is essential for understanding the observed behavior in our experiments. Second, allowing for multiplicity of equilibria is important for understanding how very different outcomes in terms of cooperation and equality may arise under similar committee rules.

Throughout the analysis, the simplest version of Baron and Ferejohn's (1989) multilateral bargaining model serves as a benchmark. This is the well-studied framework that serves as a workhorse model of bargaining in the literature. A committee must choose how to divide a fixed budget among its members. One committee member is selected as the initial agenda setter. This player proposes an allocation. The committee votes on the proposal. If a majority support it, then the budget is allocated accordingly. Otherwise, the process starts over with a new agenda setter being randomly selected to serve as the next agenda setter. This continues until a proposal passes. Once a proposal passes, the game ends. Players find delay costly, and all else equal prefer to reach an agreement sooner than later.

Baron and Ferejohn (1989) focus on the symmetric stationary subgame perfect equilibrium (SSPE) of their game. The SSPE involves the initial agenda setter proposing an allocation that provides just-enough other committee members with just-large-enough budget shares that they are willing to pass the proposal rather than reject it with the hopes of themselves being randomly selected as the next agenda setter. Proposals always pass right away with equilibrium allocation giving the majority of the budget to the agenda setter, smaller shares to a minimum winning coalition of other committee members, and nothing to others. In the one shot game, there are reasons to think that the SSPE is reasonable. Baron and Kalai (1993) argue that a SSPE is the simplest and therefore most likely of the many subgame perfect equilibria that exist in such games. Agranov and Tergiman (2014) and Baranski and Kagel (2015) show that the stationary equilibrium outcome often arises in one-cycle multilateral bargaining experiments. However, these arguments and evidence in support of stationary equilibria are associated with one-time bargaining, where the interactions between players ends after they reach an agreement.

Building upon Baron and Ferejohn framework, in this paper we present two models of repeated multilateral bargaining. We assume that the game does not end after the players reach an agreement in a one-time budget. Rather, the committee must meet to agree on a budget at regular intervals, for example every budget period or fiscal year. This means that after a budget passes, the committee doesn't go their separate ways but rather begins work on the next cycle's budget.

We present two alternative versions of this repeated bargaining framework. Both versions assume that bargaining within each budget cycle follows the standard Baron and Ferejohn model. They differ only in whether successful agenda setters can hold onto power from one budget cycle to the next. Our first version of repeated bargaining assumes that agenda setting power is randomly assigned at the beginning of each budget cycle. This is the standard assumption in the literature. Our second version assumes that an agenda setter can hold onto power with the explicit support of a majority of the committee.[1] This is consistent with real world legislative processes where legislative and committee leaders can maintain power over many periods.[2]

Intuitively, we may expect different outcomes to arise in our repeated bargaining models compared both with the one-shot game and with each other. The two dynamic environments make it feasible for players to develop reputations and maintain relationships with other players that are not possible in the one-shot game. These considerations may be even more important in the second dynamic model where a favorable reputation may enable a player to hold onto power over time.

In contrast to intuition, the SSPE outcome in each period of the two repeated games is identical to the outcome in the one-shot game. Stationarity assumes away history dependent strategies, which means that the first repeated game is analytically equivalent to an infinite series of independent one-shot games. It also means that in the second game, committee members cannot condition their vote to keep the sitting agenda setter on how generous he has been, which means they will never keeping an agenda setter in power. Because of this, the second repeated game is also one in which there is a random selection of agenda setter at the beginning of each cycle, and the outcome each cycle is the same as the one-shot game.

To assess whether our concerns regarding the stationary equilibrium refinement are reasonable, we conduct a series of laboratory experiments. We build on the experiment conducted in Agranov and Tergiman (2014) for the one-shot bargaining game, to consider our two repeated environments. From the earlier work, we know that SSPE is a reasonable predictor of outcomes in the one-shot bargaining experiments. This is not the case for the two repeated bargaining games, where outcomes differ widely from the predictions of SSPE.

In both of our repeated games, proposed allocations tend to be more inclusive and more generous than predicted by the SSPE. We frequently observe equal division within winning coalitions, and a substantial fraction of grand coalitions that include all members. Our data also clearly show that in both games, subjects use strategies that involve punishments, reciprocity and history dependence - all properties that contradict the stationarity refinement. There are also notable differences between the two repeated treatments, with the most notable difference being that in our second repeated game, the agenda setter holds onto power from one cycle to the next almost all of the time. This is consistent with the agenda setter behaving generously-enough with at least a majority

[1]We assume that the committee first votes on whether to pass a proposal, and then if the proposal passes, the committee votes again on whether to keep the current agenda setter or randomly select a new one. But, the results do not formal vote take place every round, only that a lack of majority support leads to a new AS being selected.

[2]That the agenda setter would persist over budget cycles is a realistic assumption: for example, in the US, the chairman of the House appropriations committee stays in power 5.5 years on average.

of the committee, and this coalition expecting the generosity to continue if they allow the agenda setter to keep power.

Only 5 percent of proposals in the random power treatment and 26 percent of proposals in the majority support treatment involve splitting allocations unequally within a minimum winning coalition, even though stationarity predicts that such proposals will be made 100 percent of the time. And the majority of these allocations give the agenda setter a lower share of the allocation than SSPE predicts.

This raises a natural question of how to reconcile observed outcomes with theoretical predictions. We show that extending the theory to consider asymmetric strategies, risk aversion and fairness concerns does little to reconcile the disconnect between the theory and observed behavior if we maintain focus on stationary strategies. The disconnect arises because the theory ignores the fact that in repeated interactions, players may condition their current actions on their own and others' past behavior. However, when we allow for even a very limited form of history dependence, any allocation is consistent with subgame perfect equilibria, suggesting that reconciling the disconnect between the theory and evidence is not so straightforward.[3]

We do not introduce a new equilibrium refinement that maintains uniqueness while better matching the data. Doing so would be a mistake, as the experiment illustrates that focusing on a single equilibrium could not explain the majority of outcomes across the two treatments. Multiplicity itself is a defining characteristics of the simple, repeated multilateral bargaining experiments. In the majority support treatment, for example, outcomes can be categorized as follows: equal division among a minimum winning coalition (31 percent), equal division among all players (27 percent), unequal division within a minimum winning coalition (26 percent), unequal divisions but inclusive of everyone (16 percent). Given the same bargaining rules, outcomes differ in terms of inclusiveness and equality. This suggests that the level of partisanship and cooperation observed within a legislature may be less the result of specific rules, and more to do with equilibrium selection. High partisanship and high cooperation both occur a substantial portion of the time even in very simple models of repeated bargaining. We miss this if we limit attention to an equilibrium refinement that assumes away reputation and relationships.

[3]We show that players do not need to remember much about the history of play to have subgame perfect equilibria in which successful agenda setters hold power in the Majority Support game. Complicated strategies conditioning on complex histories of a game are not necessary. As long as players can remember the most recent player to propose an unexpected allocation or cast an unexpected vote, a player may expect to be excluded from future allocations if she is seen to not do her part in the current period. In that case, as long as players care enough about future cycles, any allocation and the persistence of agenda setter power are consistent with a subgame perfect equilibrium.

Overall, our paper highlights the fundamental difference in bargaining behavior and bargaining outcomes that arise when committee members interact repeatedly with each other rather than only once. Bargaining in repeated settings presents a variety of different outcomes and often features long-term relationships between its members who use history-dependent strategies to facilitate such relationships. At the same time, vast majority of outcomes observed in the one-shot bargaining settings are stationary equilibria outcomes which are memoryless. The difference between repeated and one-shot settings is especially noticeable given that three bargaining games we consider here admit the same unique stationary subgame perfect equilbrium.

The remainder of the paper is organized as follows. Section 2 discusses the relevant literature. Section 3 presents our environment and the predictions of the stationary subgame-perfect equilibrium. Section 4 outlines the experimental design. Section 5 presents the results of the experiments. Section 7 revisits the theory, considering a number of theoretical extensions, i.e., asymmetric stationary SPE, social preferences and risk aversion, in order to try to reconcile theory and observed outcomes. That section then proceeds to document the empirical patterns of strategies used by our experimental subjects and evaluate theoretically the minimal limited history dependence required to support equilibrium outcomes observed in our experiments. Section **??** comments on the features of political institutions studied in this paper and concludes.

2. Related Literature

In the last few decades, legislative bargaining has received a great deal of attention both in the theoretical and experimental domains. The seminal paper of Baron and Ferejohn (1989) studies the legislative bargaining process, when a committee is charged with one-time allocation of a budget using a majority voting rule. Many articles extend Baron and Ferejohn's theoretical analysis to study effects of various political institutions (e.g. Baron 1996, Banks and Duggan 2000, Jackson and Moselle 2002, Merlo and Wilson 1995, Banks and Duggan 2006, Bowen and Zahran 2012, Eraslan 2002, Snyder, Ting and Ansolabehere 2005).

Given that our paper deals with *dynamic* bargaining, we will focus our review of the theoretical work on the subset of this literature that studies legislative bargaining in a dynamic setting. Baron (1996) develops a model of dynamic bargaining in which the status quo in any period is the previous policy that the legislature implemented. In equilibrium, agenda setters strategically propose policies (and manipulate the status quo) in order to limit the feasible proposals available to other agenda setters in the future. Kalandrakis (2004, 2010), and Duggan and Kalandrakis (2012) generalize Baron's results, allowing for multidimensional policy spaces.[4] Battaglini and Coate (2007, 2008) allow the legislature to choose policies that affect government spending, taxes, and debt, considering how these variables fluctuate over time. Diermeier and Fong (2011) develop an alternative model of legislative bargaining in which an agenda setter has monopoly power over proposals, the status quo is determined by the most-recently implemented proposal, and the legislative process repeats with positive probability. Finally, some papers endogenize legislative rules within the context of a repeated bargaining game. McKelvey and Riezman (1992) and Eguia and Shepsle (2015) consider dynamic legislative bargaining when legislators must stand for reelection after each period, and shows that a legislature will endogenously adopt rules that reward more senior legislators.

Each of these dynamic applications of legislative bargaining assumes that the status quo policy evolves over time, determined by past-period bargaining outcomes. To focus on how the status quo evolves over time, these articles make the simplifying assumption that agenda-setter power is exogenous, independent of past policy outcomes. This is the case when an agenda setter is randomly selected each period (e.g. Duggan and Kalandrakis 2012, Bowen and Zahran 2012), or when the identity of a future agenda setter is common knowledge (e.g. Diermeier and Fong 2011).

We take a starkly different approach from the existing literature, considering the possibility that the identity of the agenda setter, rather than the status quo policy, is endogenous. To construct our argument as clearly as possible, we abstract from other aspects of the bargaining environment, including assuming a stable, exogenous degenerate status quo policy. We are aware of no other article that focuses on the agenda-setter authority aspect of the dynamic environment.

The experimental literature has followed the steps of theoretical research focusing first on one-cycle bargaining games (see the survey by Palfrey 2016) and recently moving on to dynamic bargaining experiments. Some of the experimental papers on dynamic bargaining papers focus on the evolution of status-quo policy in dynamic models of pure redistribution and consider a setting in which the status-quo policy is determined by the distribution of resources agreed upon in the previous bargaining cycle. Battaglini and Palfrey (2012) is the first paper that experimentally investigates such environment. Baron, Bowen and Nunnari (2016) extend this setup by considering effects of various communication channels available to committee members. Nunnari (2016) and Sethi and Verriest (2016) incorporate veto power and analyzes consequences of its presence.

[4]See also Gomes and Jehiel (2005) who develop a model of dynamic bargaining between coalitions which allows for fully transferable utility between agents. Additionally, Dahm and Glazer (2015) consider a game in which the bargaining process is repeated only once, to consider how an agenda setter may promise future benefits to legislators who support him in the first period.

Other papers study dynamic models of public good accumulation. Battaglini, Nunnari and Palfrey (2012, 2016) consider an infinite-horizon legislative bargaining model of durable public good provision, in which status-quo policy distributes the available budget among committee members in equal private shares. Agranov et al. (2016) look at a two-period version of a similar game and decompose the inefficiency embedded in the legislative bargaining solution relative to the efficient solution into its static and dynamic components. None of these models consider the linkage of budget cycles via the agenda-setter identity.

Finally, our paper contributes to the literature that evaluates the relevance of stationary equilibrium refinements including Markov perfection. Baron and Kalai (1993) argue that a stationary subgame perfect equilibrium is the simplest and therefore most likely subgame perfect equilibrium. More recently, Agranov and Tergiman (2014) and Baranski and Kagel (2015) show that the stationary equilibrium outcome often arises in one-cycle multilateral bargaining experiments. However, these arguments and evidence in support of stationary equilibria are associated with one-time bargaining, where the interactions between players ends after they reach an agreement. Ours is the first to assess the suitability of stationarity in a repeated bargaining environment, where it is feasible for players to develop long term reputations and relationships. To do this, we extend the experimental design from Agranov and Tergiman (2014) to match the dynamic theory.

Other papers have considered whether Markov perfection is consistent with behavior in experiments involving dynamic games. This literature, however, is still small and far from reaching a consensus regarding this question. Several papers document that comparative static predictions implied by Markov perfect equilibria organize experimental data well. Battaglini, Nunnari and Palfrey (2012, 2016) make this point in the dynamic legislative bargaining game with durable public goods. Salz and Vespa (2016) study an infinite-horizon entry/exit game of oligopolistic competition and reach the same conclusion. Vespa (2016) studies a dynamic common pool game and finds that modal behavior of subjects is consistent with Markov perfection. Finally, Agranov and Elliott (2016) investigate decentralized bargaining games with heterogeneous trade opportunities and irreversible exit and also conclude that market outcomes match MPE predictions across treatments. On the other hand, there is a large experimental literature on infinite-horizon prisoner's dilemma games, which documents that a majority of subjects use efficient, history-dependent strategies contrary to the MPE prediction of always defecting (see survey by Bó and Fréchette forthcoming). Vespa and Wilson (2016) study an extension of an infinitely-repeated prisoner's dilemma game with two states and provide evidence that suggests when the selection of MPE is more likely to occur.[5] This debate on the validity of the stationary refinement justifies using it as a first benchmark against which to test our data.

3. Repeated Multilateral Bargaining

Our analysis considers two models of repeated multilateral bargaining, both of which are repeated versions of the classic closed-rule multilateral bargaining game of Baron and Ferejohn (1989).

3.1. **One-shot bargaining benchmark.** There are $n \geq 3$ identical players on a committee, which must decide how to divide a fixed budget of size 1 between its members. One of the n players serves as the initial 'agenda setter' (AS) who proposes an allocation of the budget. Then, the $n-1$ other players vote on whether or not to implement the proposal. If m other players vote in favor of the proposal, the budget is divided accordingly and the game ends. Otherwise, the proposal fails and the game continues. In this case, a new AS is randomly selected, with each member of the committee being selected with probability $1/n$. The process repeats with the new AS making a proposal and the other players voting on whether to pass the proposal or select a new AS.

The game can potentially last many rounds, if proposals consistently fail to gain majority support. Let $\mathbf{x}^r = (x_1^r, ..., x_n^r)$ denote the proposal made by the AS in round r, and let $\mathbf{a} = (a_1, ..., a_n)$ denote the allocation that is eventually implemented, with each player i earning $u_i(a) = a_i$. An allocation $\mathbf{a}$ is feasible if $0 \leq a_i \leq 1$ for each i, and $\sum_i a_i \leq 1$; a proposal is feasible if it corresponds to a feasible allocation. For simplicity, we assume that n is odd, and that $m = (n-1)/2$, meaning that only a simple majority of votes is needed for a proposal to pass.[6] Delay is costly, with discount factor δ applying each time a proposal fails.

There are many subgame perfect equilibria (SPE) of this game. In fact, when δ and n are sufficiently large, any allocation can be maintained as part of a SPE (Baron and Ferejohn 1989). For more predictive power, Baron and Ferejohn focus on the unique symmetric stationary subgame perfect equilibrium (SSPE) of the game.[7] The SSPE limits attention to equilibria in which a player strategies depend only on the forward-looking

[5]The authors construct an index that captures attractiveness of efficient outcomes relative to MPE outcomes, and show that this index tracks when subjects are ready to abandon MPE strategies in favor of history-dependent strategies in order to reach 'better' outcomes.

[6]Almost all of the theoretical results continue to hold as long as $m \in \{1, ..., n-2\}$, which assures that the AS cannot pass a proposal unilaterally, and that unanimity is not required.

[7]The common application of the refinement also assumes that strategies are weakly undominated, the importance of which we discuss in more detail in the context of repeated games.

game tree, and are therefore independent of history that has not affected payoff-relevant information in the state variable.[8] Strategies must be memoryless.

In the SSPE, the agenda setter in any round randomly selects m other players to include in a minimum winning coalition (MWC), and then proposes an allocation $\mathbf{x}^*$ which gives $x^*_m = \delta/n$ to the m members of the MWC, $x^*_{AS} = 1 - \delta m/n$ to himself, and 0 to all other players. The allocation x^*_m is just enough to entice the members of the MWC to vote in favor of the proposal, rather than to vote against the proposal and potentially be selected as AS in the next round. The proposal passes in round 1, ending the game.

3.2. **Repeated bargaining games.** We extend the benchmark model to allow committees to repeated interact over many budget cycles. After a committee passes a proposal, the game does not end. Rather, the bargaining process continues with the committee beginning work on a new budget cycle.

Denote any period by $t = (c, r)$, where $c = 1, 2, ...$ denotes the budget cycle, and $r = 1, 2, ...$ the proposal round within any cycle. AS^t is the AS with proposal power in period t, $\mathbf{x}^t$ is the proposal made by AS^t in that period, and $\mathbf{a}^c$ is the implemented budget allocation in cycle c. Within each budget cycle, our repeated bargaining games are identical to the one-shot benchmark model described above. A cycle starts with the first round AS in that cycle proposing an allocation, and ends once a proposal passes. Then the game transitions to the next budget cycle and the process repeats.

The discount factor $\delta \in (0, 1)$ applies between stages within a budget cycle. The discount factor $\gamma \in (0, \delta)$ applies between budget cycles. We assume that within-cycle delays do not make future cycles less valuable, which means that γ may be interpreted as either the between-cycle discount factor, of the probability that the game enters another cycle.[9] This interpretation of γ leads to a more straightforward experimental design and does not drive our theoretical results. It is also justified given our focus on budget decisions, where delay in passing one year's budget typically does not impose a delay upon the following year's bargaining.

We consider two alternative versions of repeated bargaining. The first version assumes that agenda setting power is randomly assigned at the beginning of each cycle. This is the standard assumption in the literature on repeated bargaining games, albeit one that is typically made in more complicated models with more moving parts such as an evolving status quo.[10] The second version assumes that a successful AS can hold onto

[8]Strategies can condition on the one's share in the allocation in which they are currently voting, but not on past allocations or voting behavior.

[9]That is, the next cycle is discounted at γ, and not $\delta^s\gamma$ when the current cycle lasts s stages. The alternative formulations of discounting lead to qualitatively similar results.

[10]See for example, Baron (1996), Kalandrakis (2004, 2010), Baron and Herron (2003), Battaglini and Coate (2007, 2008), Bowen and Zahran (2012) and Duggan and Kalandrakis (2012).

power with the support of at least m other players. This is more similar to a real world setting, in which committee chairs or legislative leaders can hold onto power as long as they maintain majority support. Between cycles, following the passage of a proposal, the committee votes on whether to keep or replace the successful AS from the previous cycle. If at least m other players vote in favor of the AS, he serves as the initial AS in the following budget cycle. If fewer than m other players vote in favor of the AS, then a new AS is randomly drawn with each player having a $1/n$ probability of making the first proposal in the following cycle.

We refer to the two repeated versions of the model as Repeated Bargaining with "Random Power" (RP) and "Majority Support" (MS), respectively. We refer to the non-repeated Baron and Ferejohn model as either the One-Shot or Baseline model (B).

3.2.1. *Subgame Perfect Equilibrium.* Unsurprisingly, we can show that the concerns about multiplicity of subgame perfect equilibria extend to the environment with repeated bargaining. We present the result formally to show that the anything-is-possible result is even stronger in the repeated environment.

Proposition 1. *Consider any feasible allocation profile $\mathbf{a}^* = \{\mathbf{a}^{r*}\}_{r=1}^{\infty}$, such that for every r, $a_i^{r*} \in [0,1]$ for each i and $\sum_i a_i^{r*} = 1$. As long as γ is sufficiently large, there exists SPE of Random Power and Majority Support models that generates $\mathbf{a}^*$ along the equilibrium path with probability 1. When $m \geq 2$, such an equilibrium exists for every $\gamma > 0$.*

Our proposition 1 may be viewed as a repeated-game version of Proposition 2 from Baron and Ferejohn (1989), which asserted that any allocation could occur as part of a SPE in a one cycle bargaining game, as long as the (within-cycle) discount factor δ and the number of players n are sufficiently large. Neither δ nor n appear in the repeated game result, however. When we extend the result to our repeated environment, the key parameter for determining whether any allocation can be sustained as part of equilibrium is the between-cycle discount factor γ. This is because an off equilibrium path threat of being excluded from future cycle allocations provides a stronger incentive for cooperation than any within-cycle concerns.

We can restate Proposition 1 even more strongly, showing that maintaining any allocation as part of a SPE does not require that players have a complete memory of past behavior and outcomes. It is enough that players are able to remember the identity of the most recent player to choose an unexpected strategy. We say that a strategy exhibits "limited history dependence" if the actions depend only on the forward looking game tree and the identity of the player who most-recently deviated from expected equilibrium behavior.[11]

[11]CHRIS - IS THIS THE ONLY PIECE OF INFORMATION THAT LEAD TO SUCH RESULT? CAN WE SAY SOMETHING GENERAL THAT THE NEXT PROPOSITION WOULD MOST LIKELY HOLD

Proposition 2. *Proposition 1 holds even when we restrict attention to strategies with limited history dependence.*

Any allocation can be maintained as part of a SPE, even when we require that strategies condition on only payoff-relevant information (i.e. the forward-looking game tree), and the identity of the most-recent player to deviate from an expected set of strategies. The intuition behind this is as follows. Players are able to condition their strategies on the most-recent player to deviate from some given strategy, and this is enough to permit punishment strategies that exclude any player who deviates from the equilibrium strategies from future allocations. When the across-cycle discount factor is sufficiently high, this threat of future exclusion is substantial enough to prevent players from deviating from the equilibrium strategies, and to ensure that the punishment strategies are credible.

3.2.2. *Stationary Equilibria: Subgame and Markov Perfection.* To deal with the multiplicity of equilibria in multilateral bargaining games, the literature typically follows Baron and Ferejohn (1989) and focuses on stationary refinements of SPE, whether focusing on *Stationary Subgame Perfect Equilibria* (SSPE) in stationary or cyclical environments such as the current paper, or *Stationary Markov Perfect equilibrium* (MPE) in environments with an evolving status quo (e.g. Kalandrakis 2004, 2010, Duggan and Kalandrakis 2012, Anesi 2010, Baron and Bowen 2016). These solution concepts both assume that strategies are independent of history. In our environments, the two concepts are equivalent, except for some technical differences that do not effect outcomes.[12] In the remainder of this section, we derive the SSPE of our three games noting that the same results could be obtained by characterizing instead the MPE in each of our three games.

The SSPE concept requires that players choose the same strategies in every structurally equivalent subgame.[13] This means that strategies can only condition on payoff-relevant information, and must ignore payoff irrelevant information about the history of the game.

Applied to our framework, a SSPE requires that each player follows the same proposal strategy every time he/she serves as AS, and has the same voting strategy every time he/she does not serve as AS. Equilibrium strategies cannot condition on the history of play, although a player's vote in favor of or against a proposal will depend on his/her proposed share of the allocation. In what follows, we make two additional assumptions

FOR OTHER SIMILAR VARIATIONS REGARDING MEMORY OF THE PLAYERS? MAYBE FOOTNOTE HERE?

[12]See the discussion about when analyses should use SSPE versus MPE in Maskin and Tirole (2001).

[13]Two subgames are structurally equivalent if and only if the sequence of moves is the same, the action sets are the same at each corresponding node, and the preferences of the players are the same in each period. See Baron (1998) and Baron and Ferejohn (1989).

that are common in this literature: First, we initially focus on symmetric SSPE implying that the strategies are symmetric across all players. We consider asymmetric SSPE in Section **??**. Second, we restrict attention to equilibria strategies that are not weakly dominated, implying that players who are indifferent between voting in favor of or against a proposal (or sitting AS) will choose the alternative that they would choose if they were certain to cast the deciding vote.[14]

In the SSPE of our games, a player votes in favor of a proposal when his/her proposed share is high enough that he/she prefers the proposal to pass and for the game to move on to the next cycle rather than for the proposal to fail and for a new AS (possibly him/herself) to be selected and continue with the current cycle. This means that the voting strategy is defined by an allocation threshold $\bar{a}$, where each player votes in favor of a proposal if and only if it offers him/her an allocation of at least $\bar{a}$. Anticipating this, the AS at any time t proposes an allocation offering the minimum acceptable share ($x_i^t = \bar{a}$) to exactly m other players, a higher share ($x_{AS_t}^t = 1 - m\bar{a}$) for him/herself, and nothing ($x_i^t = 0$) to everyone else. The m players receiving share $\bar{a}$ voting in favor of the proposal. This group of m players is collectively referred to as the *Minimum Winning Coalition* (MWC), and we denote their allocation by x_m^t. The $n - m - 1$ players receiving nothing vote against the proposal. In the SSPE, each player's proposal strategy randomly chooses which other players to include in the MWC and which to exclude each period that he/she serves as AS. On the path of play, proposals always pass, and each cycle lasts only one stage.

Consider first the Random Power game. Regardless of if a proposals passes, the ex ante expected payoff to any player in period t is

$$\frac{1}{n}(1 - m\bar{a}) + \frac{m}{n}\bar{a} = \frac{1}{n}.$$

This means that future buget cycles have a present discounted value of expected payoffs equal to $\nu \equiv \frac{1}{n}\frac{\gamma}{1-\gamma}$. A player prefers a proposal that provides him with current-cycle allocation x_i to pass than to fail as long as $x_i + \nu \geq \delta\frac{1}{n} + \nu$, or equivalently $x \geq \frac{\delta}{n}$.

That is, when players choose symmetric stationary strategies, the incentives that any player has to vote in favor of a proposal are identical in the one-shot and RV models. The Random Power game is equivalent to a series of many independent one-shot games. Within each period, the outcomes of the two environments are equivalent.

Now, consider the Majority Support game, which is potentially complicated by the additional vote that takes place after each passage of a proposal. The SSPE refinement greatly simplifies this analysis. It rules out proposal strategies in which an AS conditions

[14]This standard assumption rules out equilibria in which a player not included in the minimum winning coalition votes in favor of the proposal and has no incentive to deviate because the proposal passes with or without that legislator's support. We assume that a player who remains indifferent votes in favor.

allocations on who supported him/her in the past, which eliminates any incentives that players may have to keep an AS in power. Instead, the other players vote against the current AS hoping that they themselves will be selected as AS in the next cycle. Because of this, under SSPE, the Majority Support model collapses to the Random Vote model, with a new AS being randomly selected at the start of each cycle.[15] This in turn implies that the per cycle SSPE outcome in the Majority Support game is also identical to the one shot game.

In each period of the only SSPE of both the Random Power and Majority Support games, the outcome is identical to the one shot game. A new AS is randomly selected, she proposes to allocate $\frac{\delta}{n}$ to each of m randomly selected other players, $1 - \frac{\delta m}{n}$ to herself, and 0 to everyone else. This implies that the AS receives more than half the allocation herself. Non-agenda setters vote in favor of any allocation that gives them at least δ/n. Thus, the proposals pass, but the agenda setter is not reelected. The SSPE outcomes are identical.

3.2.3. *Testable predictions of stationary equilibrium.* The models generate a number of testable predictions, summarized in Proposition 3.

Proposition 3. *In the unique symmetric SSPE of the Random Power and Majority Support models, in each cycle*

(i) *Proposals assign a majority share to the AS, and a positive share to a MWC of exactly m other players. Other players get nothing.*
(ii) *The identity of the MWC partners are randomly determined, independent of the past actions of others.*
(iii) *Proposals pass without delay.*
(iv) *There is low persistence of AS power.*
(v) *Outcomes are independent of whether we are in the Random Power or Majority Support environment.*

Some of these predictions sound reasonable. It is perfectly reasonable that proposals pass without delay. And, it is intuitive that a strategic agenda setter would share no more than he needs to with no more people than he needs to in order to pass a proposal. But, other predictions are less intuitive. In the Majority Support game, where it is feasible for an AS to maintain power, one may imagine that the AS will be more generous when making proposals in an effort to maintain power. But, such behavior is assumed away by the stationary equilibrium refinement. Next, we conduct an experiment to test these

[15]One can verify that the legislators do prefer to vote to replace the AS in this situation. The expected benefit of being the AS is $1 - m\bar{a}$ each stage, and the expected benefit of not being the AS is $\bar{a}m/(n-1)$ each stage. Thus, the non-ASs vote to replace the AS since $1 - m\bar{a} > \bar{a}m/(n-1)$.

predictions. As we will see, even some of the more-intuitive predictions fail to find empirical support.

4. Experimental Design

All our experiments were conducted at the Center for Experimental Social Sciences at New York University using Multistage software.[16] Subjects were recruited from the general undergraduate population and each subject participated in only one session. A total of 105 subjects participated in our experimental sessions.

We ran treatments that correspond to the two models of repeated bargaining described in Section 3.2. In what follows we describe the details of the experimental protocol used in each treatment and refer the reader to the Online Appendix for the full instructions received by subjects.

In each session subjects played the repeated bargaining game eight times. We refer to each of those as a match. In each cycle, subjects had 200 tokens to divide. At the end of a session one match was selected at random for payment, and earnings in that match were converted into USD (10 tokens = $1). These earnings, together with the participation fee are what a subject earned in this experiment. The sessions lasted about two hours and on average subjects earned $20, including a participation fee of $7.

In each match subjects were randomly divided into groups of three and assigned an ID number. Each match consisted of many cycles and each cycle consisted of potentially many stages. Subjects kept the same ID within all cycles of a given match. The number of cycles in a match was uncertain and determined by a random draw: with probability 30% each cycle was the last cycle of the game. That is, the between-cycle discount factor is $\gamma = 0.7$.

In both treatments, at the beginning of the first stage of the first cycle of a match, one committee member was randomly chosen to serve as the AS. The AS was asked to propose how to distribute the 200 tokens between the three committee members and this proposal was presented to all group members for a vote. If the proposal was accepted by a majority of votes (at least two out of three members), then the cycle ended. With probability 70%, the group moved on to the second cycle of the match and with probability 30% the match was terminated. If, however, the proposal was rejected, then the group remained in the first cycle and the second bargaining stage started. At the beginning of the second bargaining stage one member was randomly selected to serve as the new AS. The AS was asked to submit a budget proposal, which was then voted on by all committee members. However, the rejection of a proposal triggered a 20% reduction in the budget (that is, the within-cycle discount factor is $\delta = 0.8$). In other words, while

[16] The Multistage package is available for download at at http://software.ssel.caltech.edu/.

in the first stage of every cycle the committee had 200 tokens, in the second stage, the available budget was reduced to 160 tokens, and if a committee reaches the third stage it was further reduced to 128 tokens, etc. This procedure continued until a majority of committee members voted in favor of the budget proposed by the AS.

In the *Random Power* treatment, each cycle of a game is identical to the first one: the AS in the first stage of every cycle is chosen randomly among the three committee members. In the *Majority Support* treatment, following the successful passage of a proposed budget, the committee held a second vote in which all members voted on whether to retain the current AS for the next cycle. To retain power, the current AS needed to obtain a majority of votes in the second vote. If the current AS was voted out, the AS in the next cycle was randomly chosen. *The difference in how the AS changes from one cycle to the next is the only difference between treatments.*

In each cycle, after the ID of the AS for the current cycle was announced but before the AS submitted his/her proposal, members of the committee could communicate with each other using a chat box. We implemented the unrestricted communication protocol used in Agranov and Tergiman (2014). Subjects could send any message to any subset of members; in particular, subjects could send a private message to a specific member of the committee, or send a public message that would be delivered to all members of the group. The chat option was available until the AS submitted his proposal and was then disabled during the voting stage. Our software recorded all the chats.

Finally, we implemented the Random Block Termination design developed and tested by Frechette and Yuksel (2013), in which subjects receive feedback about the termination of a match in blocks of cycles. In our implementation, each block consisted of four cycles. Within each block, subjects receive no feedback about whether the match has ended or not and they make choices which will be payoff-relevant conditional on a match actually reaching this cycle. At the end of a block, subjects learned whether the match ended within that block and, if so, in which cycle. If the match was not terminated, subjects proceeded to play a new block of four cycles. Subjects were paid only for the cycles that occurred before match was actually terminated. The advantage of using the Block design is that it allows for the collection of long strings of data (at least four cycles) even with a relatively small discount factor of $\gamma = 0.7$. This small discount factor was chosen in order to obtain distinct enough predictions of the stationary subgame perfect equilibrium between treatments.

Table 1 summarizes the details of all our experimental sessions.

Given our parameterization ($n = 3$, $m = 1$, $\delta = 0.8$, $\gamma = 0.7$, and a budget of 200 tokens), in both games any feasible allocation profile $\mathbf{a}^*$ can be maintained as part of a SPE. Further, the unique symmetric stationary SPE predicts that per-cycle allocation

Table 1. Experimental Design

Treatment	# of Sessions	# of Subjects	# of Matches	Mean # of Cycles/Match
Random Power	3 sessions	(18,18,15)	(8,8,8)	(4,7,6)
Majority Support	3 sessions	(21,15,18)	(8,8,8)	(4,6,6)

to coalition partners of $a^{SSPE} = 53$ tokens, and a per-cycle allocation to the AS of 147 tokens.

5. Experimental Results

We present results in the following order. We start with a quick summary of behavior that has been documented in one-shot bargaining experiments. The one-shot bargaining setup serves as a natural benchmark as it admits the same unique symmetric SSPE as the two repeated games considered here. Moreover, typical behavior of subjects in the one-shot bargaining environment is very close to the one predicted by the symmetric SSPE. We then compare bargaining outcomes within a cycle in our two repeated games with the outcomes documented in the one-shot game. Following that, we shift our attention to dynamics in repeated games and analyze how coalitions evolve over bargaining cycles. We conclude by analyzing the recorded chats between subjects.

5.1. **Approach to the data analysis.** Most of the analysis is performed using the first block of four cycles in the last four matches of each session. We refer to these as *experienced cycles*. By focusing on behavior in these experienced cycles, we are able to consider the behavior of our experimental subjects after they have familiarized themselves with the game. In addiont, restricting ourselves to the first block of the cycles, which all groups certainly play, allows us to have a balanced dataset with identical amounts of experience within a match across all treatments.

We classify proposals in terms of the number of members who receive non-trivial shares and refer to these as *coalition types*. A *non-trivial share* is defined as share that is larger than 5 tokens. A proposal in which only one group member receives more than 5 tokens is a *dictator coalition*. A proposal in which exactly two members receive non-trivial shares is a *minimum wining coalition*. Finally, a proposal, in which all three members receive non-trivial shares is a *grand coalition*. We call members with non-trivial shares *coalition partners*. Finally, we refer to some proposals as *equal split* proposals. Equal split proposals are ones in which the difference between the shares of any two coalition partners is at most 5 tokens.

To compare the outcomes between two treatments we use regression analysis. Specifically, to compare outcomes between two treatments (whether the fraction of a particular coalition type or the share received by the AS), we run random-effects regressions, in which we regress the outcome under investigation on a constant and a dummy that takes a value of 1 for one of the two considered treatments. We use the same method to compare outcomes between different types of coalition members. In both cases we cluster standard errors by session, recognizing the interdependencies between observations that come from the same session, since subjects are randomly rematched between matches.

5.2. **Bargaining Outcomes in One-shot Bargaining Games.** There is a large experimental literature that tests one-shot bargaining games using the standard bargaining protocol of Baron and Ferejohn summarized in 5.2. Morton (2012) and Palfrey (2013) provide an excellent review of this literature. Despite variations in the parameters and the specifics of experimental protocol used in different studies, the behavior of subjects and resulting bargaining outcomes are quite stable and closely admit to predictions of the symmetric SSPE especially when subjects are allowed to communicate with each other.

Specifically, experimental play is consistent with the three main identifying features of the stationary SPE. First, the vast majority of proposals pass without delay. Second, most coalitions are minimum winning. Depending on the study this fraction is between 65% as in Frechette et al (2003) and 90% as in Agranov and Tergiman (2014). Finally, the distribution of resources within a coalition is unequal with proposers appropriating a much higher share of resources than the coalition partners. **State range of numbers here.**

5.3. **Bargaining Outcomes within a Cycle in the Repeated Games.** We now turn our attention to the two repeated bargaining games (Random Power and Majority Support) which have the same unique symmetric SSPE as the one-shot bargaining game. Similarly to the one-shot game, almost all proposals pass without delay in the first stage of each cycle. This is the case in 96.3% and 99.7% of experienced cycles in the Random Power and Majority Support treatments, respectively. In the remainder of this subsection we concentrate on those proposals that passed without delay.

In Table 2 we present the distribution of proposals that passed without delay in terms of coalition size. The results differ from the one-shot environment. While in the one-shot games minimum winning coalitions were the most common structure of coalitions, in both versions of the repeated game, a different coalition structure emerges in addition to the minimum-winning one: the grand coalition, in which all three members of the committee receive positive shares. In fact, in the Random Power treatment, which is a mere repetition of the one-shot game, grand coalitions were formed more than 70% of the time and we observe fewer than 30% of MWCs. In the Majority Support treatment, there are slightly more MWCS (57.8%) but the fraction of grand coalitions remains substantial

TABLE 2. Coalition types of proposals that passed without delay, by treatment

	Random Power	Majority Support
Coalition size		
Dictator (1-person coalition)	0.0%	0.3%
MWC (2-person coalition)	27.9%	57.8%
Grand (3-person coalition)	72.1%	41.8%
Allocations within coalitions		
Equal split | MWC	80.8%	56.0%
Equal split | Grand coalition	83.1%	65.0%

Notes: In the last two rows, we report the fraction of equal splits conditional on the coalition size being a two-person coalition (fourth row) and a three-person coalition (fifth row).

(41.8%). Regression analysis confirms that the proportion of three-person coalitions is higher in the Random Power than in the Majority Support treatments ($p = 0.088$).[17]

Further, while in the one-shot setting, proposers appropriate a higher share of resources compared to to the shares of coalition partners, in the repeated game this is not necessarily the case. The last two rows of Table 2 show that conditional on the coalition type, allocations across treatments are, in their majority, equal splits.

Coalition size affects the share that the AS can appropriate for him/herself. Figure 1 depicts the histograms of the shares received by ASs conditional on coalition size in each of our treatments. For each coalition type, the vertical lines indicate the average share of the ASs.

In both treatments, ASs receive higher shares than their coalition partners, and this is true for both types of coalitions that we observe.[18] Those proposers that form grand coalitions appropriate a smaller share of resources than those that form minimum winning coalitions ($p < 0.001$ within each treatment).[19] Comparing across treatments, we find that the shares of ASs in the Random Power treatment are significantly lower than

[17]This is the p-value on the treatment coefficient in a panel probit regression using Grand Coalition as the dependent variable and the Random Power treatment as the explanatory one. The coefficient itself is -1.667. We cluster at the session level.

[18]A series of panel OLS regression with clustering at the session level, using Share as the dependent variable and whether someone was a proposer as the independent one shows that the pvalues on Proposer are at most 0.006. The coefficients are all positive.

[19]These are the p-values on the treatment coefficient in a panel OLS regression using Shares as the dependent variable and the Grand Coalition as the explanatory one. The coefficients themselves are -31.46 and -34.57 for Random Power and Majority Support, respectively. We cluster at the session level.

Figure 1. Agenda Setters' shares in proposals that passed without delay

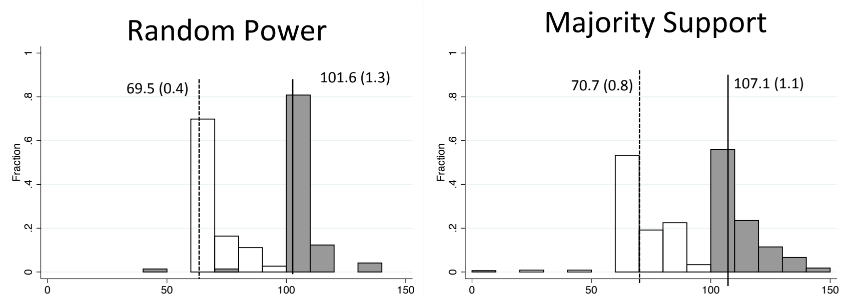

Notes: Dark bars depict the shares of ASs in two-person coalitions. White bars depict the shares of ASs in three-person coalitions. Vertical lines indicate the average share of ASs conditional on the type of the coalition: solid line for two-person coalitions and dashed line for three-person coalitions. The numbers next to the lines are the average shares of ASs conditional on the coalition type. The numbers in parenthesis are the standard error of the mean.

in the Majority Support treatment ($p = 0.025$ for MWCs and $p = 0.085$ for Grand coalitions.)[20]

5.4. **Dynamics: Behavior and Bargaining Outcomes across Cycles in Repeated Games.** We now turn towards analyzing behavior across cycles. Thus, we no longer restrict ourselves to proposals that pass right away, but instead look at behavior dynamics both in groups that had proposals rejected and those that didn't.

5.4.1. *Persistence of Power.* As we have seen in the previous section, ASs obtain shares that are on average larger than shares obtained by any other committee members. Thus, holding the AS seat has obvious benefits within each cycle. The two repeated games we consider differ in whether persistence of power is institutionalized: while it is possible for proposers to submit proposals and bargain with other group members to retain power in the Majority Support treatment, it is prevented by design in the Random Power game, in which the identity of AS in each cycle is determined randomly and independently of past proposals and behavior.

In the Random Power treatment, the event in which the *same* AS serves in all four cycles of the first block is quite rare (7% of the time). In contrast, 91.7% of Majority Support committees operate with the same sitting AS in power in all four cycles of the first block.

[20]These are the p-values on the treatment coefficient in a panel OLS regression using Shares as the dependent variable and the Random Power treatment as the explanatory one. The coefficients themselves are -4.65 and -1.76 for the Grand Coalition and Minimum Winning Coalitions, respectively. We cluster at the session level.

The number of cycles in which the same AS holds onto power directly affects his/her long-run payoff in the game as measured by the total payoff that the AS first selected obtains over the course of an entire block of four cycles. Given the persistence of power observed in the Majority Support treatment, the first AS in this treatment on average earns 355 tokens compared with 265 tokens for the first AS in the Random Power treatment (these are significantly different $p < 0.001$).[21]

5.4.2. *Evolution of Coalitions.* We begin this analysis by considering the frequency with which coalition types change. Table 3 shows the likelihood of proposed coalition types conditional on the type of coalition that passed in the previous cycle. As evident from the transition matrix, we observe a high level of persistence of coalition types between cycles in both treatments: in 87% or more cases, the next cycle proposal has the same coalition size as the one passed in the previous cycle.

TABLE 3. Transition of coalition types across cycles

	Cycle $c+1$			
	Random Power		Majority Support	
	MWC	Grand	MWC	Grand
Cycle c				
MWC	0.87	0.12	0.94	0.06
Grand	0.11	0.89	0.07	0.93

Notes: In this table, we consider only proposals that passed without delay.

Next, we consider the persistence of coalition members across cycles. To do this, we focus on the persistence of the minimum winning coalition partner in all instances where the AS was the same for two consecutive cycles.[22] Our data show that when a non-proposer is invited into a minimum winning coalition in one cycle, the probability that he/she will be re-invited into a minimum winning coalition in the following cycle is 76.9% and 89.6% in the Random Power and Majority Support treatments. A series of tests of probability show that these percentages are significantly higher than 50%, which means that proposers who are forming minimum winning coalitions are not choosing their coalition partners randomly.[23] That is, minimum winning coalitions tend to be stable across cycles. In addition, our data indicate that the shares of those coalition partners stay the same across cycles in 100% and 85.4% of the cases in the Random

[21] This is the p-values on the treatment coefficient in a panel OLS regression using total payment as the dependent variable and the Random Power treatment as the explanatory one. The coefficient itself is -97.16. We cluster at the session level.

[22] This is the only non-trivial case, since in grand coalitions all members are coalition partners by definition.

[23] In the Random Power treatment we obtain $p = 0.026$, while in Majority Support $p < 0.001$.

Power and Majority Support treatments, respectively. Thus, we conclude that not only are coalitions stable in terms of the identity of coalition members, but, in addition, when that is the case, the shares given to the coalition partners also are largely constant. In other words, ASs seek stability.

We conclude this section by documenting the long-run inequality in group members' payoffs induced by the two versions of the repeated bargaining. We measure the long-run inequality in group members' payoffs by the members' total payoffs over the course of an entire block of four cycles. Figure 2 presents the empirical CDFs of the Gini coefficient for each committee.

FIGURE 2. Empirical CDFs of GINI coefficients, by treatment

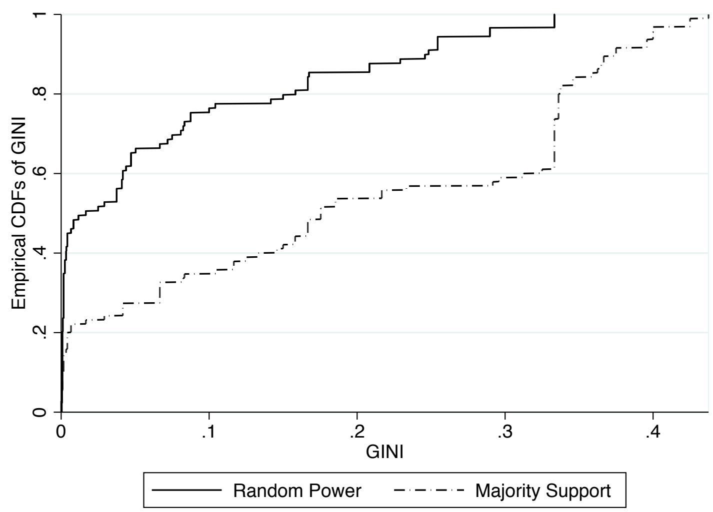

As evident from Figure 2, the Random Power treatment features a much more equal distribution of long-run payoffs compared to the Majority Support treatment ($p < 0.001$ for a kolmogorov-smirnov test). Two forces contribute to this result. First, the frequent turnover of the AS, which is a built-in feature of the Random Power treatment, increases the chances that different members serve as AS in different cycles. Consequently, when the AS does receive an above average payoff, the committee members "take turns" in obtaining the higher shares. Second, in the Random Power treatment grand coalitions are more common than in the other treatments. These grand coalitions naturally produce a more equal distribution of resources within a committee compared with two-person coalitions.

5.5. **Summary of Experimental Results and SSPE.** In this section we summarize the results of our experiments and compare them with the predictions of the stationary equilibrium refinement. We focus on symmetric SSPE, which coincide with the predictions of Markov Perfect equilibria in our two games (see discussion in Section 3.2.2 and summary in Proposition 3).

Bargaining outcomes within a cycle are efficient (no delays) in both environments, consistent with theoretical predictions. However, while the symmetric SSPE predicts that all passed proposals should feature two-person minimum winning coalitions, our data show a different pattern. We observe that both minimum winning and grand coalitions are common in both settings. The highest fraction of grand coalitions is in the Random Power treatment, in which over 70% of all passed proposals include non-trivial shares to all three group members. Finally, conditional on coalition size, at least 50% of passed proposals feature an equal division of the surplus between coalition partners, irrespective of the treatment and whether the coalitions are minimum winning or grand. In particular, ASs share the surplus equally with their coalition partner in 81% of Random Power minimum winning coalitions, while they do so in 56% of such proposals in the Majority Support treatment. This is in sharp contrast with the symmetric SSPE prediction, according to which ASs appropriate strictly higher shares of resources than their coalition partner. These per cycle outcomes are also in sharp contrast to the behavior in one-shot bargaining games documented in the previous literature, which closely admits to the predictions of SSPE.

Turning to the examination of bargaining outcomes across cycles, our data reveal high persistence of power in the Majority Support treatment, despite this being ruled out by the stationarity refinement in the theoretical analysis. In both games, the observed coalitions are stable across cycles in terms of their size, the identify of coalition partners and the shares of coalition partners, which is also at odds with predictions. Further, long-run payoffs of ASs are higher in the Majority Support compared with the Baseline treatment, while the theory stipulates that these payoffs should be the same. Finally, we document that among our two treatments, the Random Power treatment features the lowest inequality in terms of long-run payoffs between committee members.

Overall, the symmetric SSPE predictions clearly fail to accommodate the observed outcomes. In fact, the symmetric SSPE only correctly predicts: (a) efficient outcomes in both treatments, and (b) the existence of minimum winning coalitions. All the remaining predictions, whether in terms of the structure of passed proposals, or the comparative static predictions of dynamic outcomes across treatments fail to be supported by the data.

5.5.1. *Relaxing the Symmetry Assumption.* One may wonder whether relaxing the symmetry assumption leads to a reconciliation of experimental findings and theoretical predictions. Indeed, one can characterize an asymmetric SSPE in which players still use stationary strategies but can treat other players asymmetrically, particularly when the AS each period chooses which player to include in her MWC. Here we present the main intuition of such an asymmetric SSPE and refer the reader to Appendix ?? for a detailed derivation. In the Random Power game, the switch from symmetric to asymmetric strategies does not change players' incentives to accept or reject proposals in each cycle. Thus the situation is similar to the symmetric case. In the Majority Support game, one needs to consider two possibilities: the asymmetric SSPE with high and low persistence of power. The restriction that $\gamma \in (0,1)$ rules out the possibility of an asymmetric SSPE with high persistence of AS power, which leaves the low persistence equilibrium as the only viable option. In this case the incentives to vote for and against a given proposal are the same as in the Random Power game, and so we are back to the same prediction of non-stable coalitions. Clearly, the asymmetric SSPE does no better than the symmetric SSPE in explaining the data.

5.5.2. *SSPE with Risk-Averse Members.* Another natural avenue for extending the results of the SSPE is to consider outcomes that emerge when bargainers are risk-averse. Specifically, we assume that the overall utility of member i is given by

$$U_i = \sum_{c=1}^{\infty} \gamma^c \cdot u_i \left(\delta^{s_c} x_i^c\right)$$

, where x_i^c denotes the allocation in cycle c that passed in stage s_c and $u_i(\cdot)$ is the per cycle concave and well-behaved Bernoulli utility function of member i.

In the Online Appendix, we solve for the symmetric SSPE allocations when players have identical CRRA utility functions: $u_i(x) = 1 - e^{-\alpha \cdot x}$ for all i. In equilibrium, the share of the MWC partner is strictly decreasing in α, which means that introducing risk-aversion leads to a more-unequal split of resources in favor of the AS compared with the risk-neutral case. This pattern is the opposite of what we observe in our data. Intuitively, as α increases, coalition partners are willing to accept a lower share rather than reject a proposal and risk not being included in the next MWC, since MWC partners are chosen randomly. Moreover, there is no symmetric SSPE in which there is persistence of power in the Majority Support game. It turns out that combining risk-averse bargainers with asymmetric SSPE does not help either, as one cannot obtain a high persistence of power in the Majority Support game with asymmetric stationary strategies.

Therefore, incorporating risk aversion moves the SSPE predictions even further away from observed behavior.

5.5.3. *Fairness concerns.* Finally, another possibility is that players care about fairness.[24] To allow for this, we incorporate other-regarding preferences in line with the model of

[24]For the study investigating fairness concerns in the one-shot bargaining games see Montero (2007).

Fehr and Schmidt (1999). Player i's utility from allocation a in any given period is:

$$u_i(a) = a_i - \alpha \frac{1}{n-1} \sum_{j \neq i} \max\{x_i - x_j, 0\} - \beta \frac{1}{n-1} \sum_{j \neq i} \min\{x_j - x_i, 0\},$$

where $\alpha \in (0, 1)$ is a cost incurred from others being treated "unfairly" relative to oneself, and $\beta \in [\alpha, 1)$ is a cost incurred by being treated "unfairly" oneself. To simplify the analysis, we focus on the three-player case, with $n = 3$ and $m = 1$.

In the Online Appendix, we solve for the symmetric SSPE of both games after incorporating such fairness concerns. As one may expect, when players find it sufficiently costly to provide unequal allocations to others (i.e. when α is high), there exists a SSPE of the game in which each player receives an equal share of the allocation in each cycle. Additionally, when other regarding preferences are weak (i.e. when α and β are sufficiently small), the SSPE allocations resemble those with standard utility functions, except that a MWC member needs to be offered a higher allocation in order to offset the costs of inequality.

Less obvious is whether or not such fairness concerns can lead to equilibria which are consistent with other behavior that we observe during the experiments. Specifically, we look at whether they can result in SSPE in which the AS splits the allocation evenly within a MWC each period. In doing so, we focus on the parameter values from our experiment (i.e. $\delta = 0.8$ and $\gamma = 0.7$), and show that no such equilibria exist. Intuitively, if α is sufficiently high that an AS prefers to split the allocation evenly with a MWC rather than offering the MWC a lower acceptable allocation, then the AS will receive an even higher payoff from splitting the allocation evenly among all players rather than just a MWC. An AS that would consider an even division within a MWC would deviate to even division in a grand coalition instead. This is the case in all three of our games, given the parameter values of our experiments.

Thus, fairness concerns may explain some, but not all, of the observed allocations during the experiments. The main feature that the SSPE coupled with fairness concerns cannot explain is the equal splits among coalition partners within minimum winning coalitions, a behavior which is very common in both our games as shown in Section 5.3.

6. Forces Behind Reaching Stable Outcomes in Repeated Bargaining

Our previous analyses uncover interesting differences between how people bargain in one-shot versus repeated settings. While behavior in the one-shot setting is consistent with predictions of the SSPE, this is not the case with repeated bargaining. Given this discrepancy, in this section we investigate the important forces that govern bargainers' behavior in the repeated environment: how do subjects get to the outcomes documented above?

6.1. **Empirical evidence of history-dependent strategies.** The dynamic nature of our bargaining environment creates potential links between cycles and allows subjects to form and execute history-dependent strategies. As we show below, in both treatments, subjects rely extensively on the history of play and use both punishments and rewards to enforce partnerships and long-term relationships documented in the previous section.

We start by documenting strategies that include punishment. In the Random Power treatment, if a previously excluded member becomes the AS, he/she excludes the previous AS from a minimum winning coalition 81.8% of the time. A two-sided test of proportions shows that this fraction is significantly different from 50% ($p < 0.01$).[25] Given the very high persistence of power observed in the the second treatment (Majority Support), in order to obtain a reasonable number of observations related to punishment behavior, we look at all cases in which there was turnover in AS power (no longer restricting the data to the last four matches). For the cases where we observe such turnover, the AS who failed to pass the proposal in the previous cycle is excluded from the new AS's MWC in 80% of the cases.[26]

In addition, in the Baseline treatment, we observe reciprocity-type of behavior between former coalition partners. This happens when a MWC partner from cycle $c-1$ is selected to serve as the AS in cycle c. In this case, the former MWC partners invite the previous AS into their coalitions 81.8% of the time, a fraction that is significantly different from 50% according to two-sided test of proportions ($p = 0.035$). Thus, committee members attempt to establish stability even when, by treatment design, stability is hard to establish. Stability increases both because proposers tend to re-invite the same partner in their minimum winning coalition, and because the invited partner, if selected to be the next AS, invites the former proposer in his/her minimum winning coalition.

Overall, in both repeated games history-dependent strategies are commonly observed.

6.2. **Effects of Communication on Establishing Partnerships.** Establishing partnerships and long-term relationships may require committee members to agree on particular terms. Communication serves this purpose. Indeed, we observe that subjects often engage in conversations with each other before budget proposals are put on the table for vote.

To analyze communication, we have hired two independent research assistants who were not privy to the purpose of this experiment. Both research assistants classified the chats into several categories based on whether or not conversations included discussion of things relevant to the game (game rules, strategies, proposals, threats, etc..). Most of the relevant messages fall into one of the two broadly defined content categories: "Nice"

[25] In almost 75% of cases this new AS proposes a minimum winning coalition.

[26] We have only 5 observations of this kind in the Majority Support treatment.

or "Self".[27] The category "Nice" includes any message that can be interpreted as lobbying for fairness and equality, i.e., "equal is nice," "let's just do equal," and "just play fair". The category "Self" includes messages that contain information about one's own share and lobbying for ones interest, i.e., "lets do half half ill do the same with you", "ok wanna 100/100 every time?", and "Wanna collaborate? 101/99?". Looking at whether the coders classified statements from each individual within a cycle, agreements are very high and range between 87.1% and 89.2%.[28]

In the remainder of this section, we use group-level conversations in a cycle as the unit of observation, where agreement ranged from 84.3% to 96.3%.[29] Figure 3 displays the frequency of conversations that contained relevant messages in the first block of either the first or the last 4 matches. As evident from Figure 3 subjects often discuss things relevant to the experiment. Even in the experienced games more than 70% of groups have meaningful conversations at the beginning of a block. In both treatments, majority of groups discuss relevant things more often at the beginning of the first cycle than later on in the block.

FIGURE 3. Frequency of Conversations

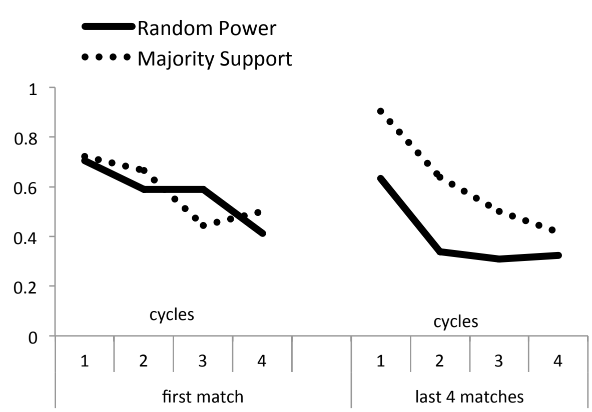

What do subjects talk about? Figure 4 displays the features of conversations that occur in each cycle. For each conversation, we code relevant messages within this conversations according to the two dimensions: whether message content is about equality and fairness or, on the contrary, about lobbying for one's own interest, and whether the message is sent privately to one other member of the group or to the entire group. Thus, we end up with four categories: private/public messages that contain fairness and equality

[27]These types of messages were also documented to affect play in the one-shot bargaining games (see Agranov and Tergiman (2014) and Agranov and Tergiman (2017).

[28]The Kappa scores are 0.71, 0.73, 0.78, 0.68 for fair statements in the Random and Majority support treatments and self statements in those same treatments.

[29]The Kappa scores were 0.85, 0.69, 0.81, 0.78.

statements (Nice Private/Nice Public), and private/public messages that contain statements with requests for own share of resources (Self Private/Self Public). In the Random Power treatment, both in the very first match as well as in the last 4 matches, subjects often lobby for equality and fairness using public messages that are delivered to the entire group. On the contrary, in the Majority Support treatment, most discussions are done in private and contain lobbying for one's own interests. The common feature between two treatments is the fact that messages about one's own interest are frequent and observed in all cycles of the last 4 matches and not just in the first cycle. These messages about one's own interest serve as reinforcements of the agreements that members reach at the beginning of the first cycle.

FIGURE 4. Content of Conversations

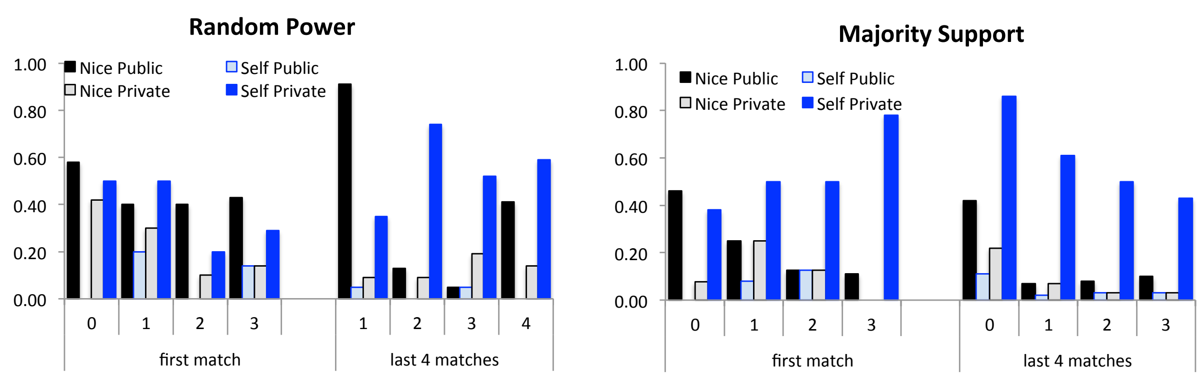

The conversations between members affect the size of the coalition that the proposer forms. Table 4 shows the results of regression analyses to show this point. The dependent variable is the indicator of proposing minimum winning coalition in the first cycle of the first block. The right-hand side variables include the match number to capture learning effects as well as indicators of the four types of messages described above. The likelihood of forming a minimum winning coalition increases substantially in both games when proposers receive private communication from one of the members with a message containing a "selfish" motive. Moreover, in the Majority Support treatment, proposers are less likely to form minimum winning coalitions when some group members talk about fairness and equality in the public chat. Such effect is not present in the Random Power treatment, in which essentially all conversations preceding first cycle of a block contains at least one public message with plead for fairness.

Table 4. Effect of Conversations on Coalition Size

	Random Power	Majority Support
Indicator for Nice Public message this Cycle	-0.03 (0.08)	-0.35^{**} (0.10)
Indicator for Nice Private message this Cycle	-0.17 (0.17)	0.07 (0.11)
Indicator for Self Private message this Cycle	0.53^{**} (0.09)	0.41^{**} (0.10)
Indicator for Self Public message this Cycle	-0.29 (0.22)	-0.27 (0.21)
Match	-0.05^{*} (0.03)	-0.04 (0.03)
Constant	0.44^{**} (0.18)	0.62^{**} (0.20)
# of observations	68	72
# of subjects	36	45
R-square overall	0.4655	0.2089

Notes: Results of Random-effects GLS regression are reported. Dependent variable is Indicator of Proposing minimum winning coalition in the first cycle of the first block in the experienced games (last 4 matches). We focus on proposals that pass right away without delay. Errors are clustered at the subject level. ** (*) indicates significance at 5% level (10% level).

7. Discussion

Without being able to rely on the standard stationary equilibrium refinements to narrow down the set of potential outcomes, we again find ourselves in a situation in which any allocation can be justified as consistent with an equilibrium (as we have shown in the Theory section). Given this inconsistency, we look to the literature for an alternative equilibrium refinement that may be more consistent with the experimental evidence.

The literature on equilibrium selection in games provides guidance, with evidence that players tend to coordinate on equal or "fair" outcomes in games with multiple Pareto dominated equilibria (e.g. Yaari and Bar-Hillel 1984, Young 1993, 1996, Roth 2005, Janssen 2006). This suggests that equal divisions (among all players or among a subset), when they are associated with an equilibrium, may serve as focal points, and help facilitate coordination. This view is also consistent with empirical evidence concerning the division of resources in legislative decision-making. Gamson's Law highlights the empirical regularity that coalitions of legislators tend to divide resources (e.g. cabinet positions) between parties in proportion to each party's share of total votes within the coalition (Gamson 1961, Browne and Franklin 1973, Browne and Frendreis 1980). Applied to our games, where each player has equal voting weight in any coalition, Gamson's Law suggests that legislators are likely to divide resources evenly among a winning coalition of players each period (whether minimum-winning or grand).[30]

[30]See Fréchette, Kagel and Morelli (2005) for a comparison of the predictions of Gamson's Law and the stationary Baron and Ferejohn (1989) bargaining outcomes in a one-cycle bargaining framework.

Recent work by Andreoni et al. (2016) corroborates the idea of equal division of resources within a coalition (be that grand coalition or minimum-winning) based on the notion of myopic fairness. Instead of evaluating proposed allocations in terms of overall inequality between all committee members, bargainers might focus somewhat narrowly on the subset of people involved in the deal directly. This narrowly framed fairness notion takes as given the coalition size and ignores parties that are excluded from the deal.

Finally, we argue that equal division equilibrium in the repeated environment may be more simple and more intuitive for the players than the SSPE. Although the SSPE may involve the simplest dynamics, with players choosing the same actions regardless of past outcomes, the per period proposal requires players to engage in some degree of complex reasoning to estimate the asymmetric equilibrium allocation that will be offered each period. Equilibria involving equal division among a winning coalition, on the other hand, involve little complex reasoning, with the AS each period splitting the allocation equally with at least m coalition partners, who in turn vote in favor of an allocation (and vote in favor of the AS in the Majority Support game) as long as the AS doesn't deviate from the equal division strategy. Even the punishment strategies played off the equilibrium path are intuitive, with players simply excluding anyone who deviated from the equilibrium strategy in the past. This suggests that Baron and Kalai (1993)'s claim that the SSPE is likely to serve as a focal point because of its simplicity may be less likely to apply in a repeated environment. Rather, we see SPE with equal division among a winning coalition (whether grand or minimum-winning) and the threat of exclusion as a potentially simpler equilibrium in a repeated environment.

References

Agranov, Marina, and Chloe Tergiman. 2014. "Communication in Multilateral Bargaining." *Journal of Public Economics*, 118: 75–85.

Agranov, Marina, and Matt Elliott. 2016. "Commitment and (In)Efficiency: a Bargaining Experiment." working paper.

Agranov, Marina, Guillaume Fréchette, Thomas R. Palfrey, and Emanuel Vespa. 2016. "Static and Dynamic Underinvestment: An Experimental Investigation." *Journal of Public Economics*, 143: 125–141.

Andreoni, James, Deniz Aydn, Blake Barton, B. Douglas Bernheim, Deniz Aydn, Blake Barton, B. Douglas Bernheim, and Jeffrey Naecker. 2016. "When Fair Isn't Fair: Sophisticated Time Inconsistency in Social Preferences." *working paper*.

Anesi, Vincent. 2010. "Noncooperative foundations of stable sets in voting games." *Games and Economic Behavior*, 70(2): 488–493.

Banks, Jeffrey S., and John Duggan. 2000. "A Bargaining Model of Collective Choice." *American Political Science Review*, 94(1): 73–88.

Banks, Jeffrey S., and John Duggan. 2006. "A General Bargaining Model of Legislative Policy-Making." *Quarterly Journal of Political Science*, 1: 49–85.

Baranski, Andrzej, and John H. Kagel. 2015. "Communication in Legislative Bargaining." *Journal of the Economic Science Association*, 1(1): 59–71.

Baron, David P. 1996. "A Dynamic Theory of Collective Goods Programs." *American Political Science Review*, 90(2): 316–330.

Baron, David P. 1998. "Comparative Dynamics of Parliamentary Governments." *American Political Science Review*, 92(3): 593–609.

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

Baron, David P., and Michael Herron. 2003. "A Dynamic Model of Multidimensional Collective Choice." In *Computational Models in Political Economy*. , ed. K Kollman, J Miller and S Page. Cambridge, MA: MIT Press.

Baron, David P., and T. Renee Bowen. 2016. "Dynamic coalitions." working paper.

Baron, David P., Renee Bowen, and Salvatore Nunnari. 2016. "Dynamic Coalitions and Communication: Public versus Private Negotiations." working paper.

Baron, Davin, and Ehud Kalai. 1993. "The Simplest Equilibrium of a Majority-Rule Division Game." *Journal of Economic Theory*, 61: 290–301.

Battaglini, Marco, and Stephen Coate. 2007. "Inefficiency in Legislative Policymaking: A Dynamic Analysis." *American Economic Review*, 97(1): 118–149.

Battaglini, Marco, and Stephen Coate. 2008. "A Dynamic Theory of Public Spending, Taxation, and Debt." *American Economic Review*, 98(1): 201–236.

Battaglini, Marco, and Thomas R. Palfrey. 2012. "The Dynamics of Redistributive Politics." *Economic Theory*, 49(3): 739–777.

Battaglini, Marco, Salvatore Nunnari, and Thomas R. Palfrey. 2012. "Legislative Bargaining and the Dynamics of Public Investment." *American Political Science Review*, 106(2): 407–429.

Battaglini, Marco, Salvatore Nunnari, and Thomas R. Palfrey. 2016. "The Political Economy of Public Debt: A Laboratory Study." working paper.

Bó, Pedro Dal, and Guillaume R. Fréchette. forthcoming. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature*.

Bowen, T. Renee, and Zaki Zahran. 2012. "On Dynamic Compromise." *Games and Economic Behavior*, 76: 391–419.

Browne, E. C., and J. P. Frendreis. 1980. "Allocating coalition payoffs by conventional norm: Assessment of the evidence for cabinet coalition situations." *American Journal of*

Political Science, 24(4): 753–768.

Browne, E. C., and M. Franklin. 1973. "Aspects of coalition payoffs in European parliamentary democracies." *American Political Science Review*, 67(2): 453–469.

Dahm, Matthias, and Amihai Glazer. 2015. "A carrot and stick approach to agenda setting." *Journal of Economic Behavior and Organization*, 116: 465–480.

Diermeier, Daniel, and Pohan Fong. 2011. "Legislative Bargaining with Reconsideration." *Quarterly Journal of Economics*, 126(2): 947–985.

Duggan, John, and Tasos Kalandrakis. 2012. "Dynamic Legislative Policy Making." *Journal of Economic Theory*, 147(5): 1653–1688.

Eguia, Jon X., and Kenneth A. Shepsle. 2015. "Legislative Bargaining with Endogenous Rules." *Journal of Politics*, 77(4): 1079–1088.

Eraslan, Hulya K. K. 2002. "Uniqueness of Stationary Equilibrium Payoffs in the Baron-Ferejohn Model." *Journal of Economic Theory*, 103: 11–30.

Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 114(3): 817–868.

Frechette, Guillaume R., and Sevgi Yuksel. 2013. "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination." working paper.

Fréchette, Guillaume R., John H. Kagel, and Massimo Morelli. 2005. "Gamson's Law versus non-cooperative bargaining theory." *Games and Economic Behavior*, 51: 365–390.

Gamson, William A. 1961. "A Theory of Coalition Formation." *Americal Sociological Review*, 26: 373–382.

Gomes, Armando, and Philippe Jehiel. 2005. "Dynamic Processes of Social and Economic Interactions: On the Persistence of Inefficiencies." *Journal of Political Economy*, 113(3): 626–667.

Jackson, Matthew O., and Boaz Moselle. 2002. "Coalition and Party Formation in a Legislative Voting Game." *Journal of Economic Theory*, 103(1): 49–87.

Janssen, Maarten C.W. 2006. "On the strategic use of focal points in bargaining situations." *Journal of Economic Psychology*, 27: 622–634.

Kalandrakis, Anastassios. 2004. "A Three-Player Dynamic Majoritarian Bargaining Game." *Journal of Economic Theory*, 116: 294–322.

Kalandrakis, Tasos. 2010. "Minimum Winning Coalitions With Endogenous Status Quo." *International Journal of Game Theory*, 39: 617–643.

Maskin, Eric, and Jean Tirole. 2001. "Markov Perfect Equilibrium: I. Observable Actions." *Journal of Economic Theory*, 100: 191–219.

McKelvey, Richard D., and Raymond Riezman. 1992. "Seniority in legislatures." *The American Political Science Review*, 86(4): 951–965.

Merlo, Antonio, and Charles A. Wilson. 1995. "A Stochastic Model of Sequential Bargaining with Complete Information." *Econometrica*, 35(2): 371–399.

Montero, Maria. 2007. "Inequity Aversion May Increase Inequality." *Economic Journal*, 117(519): 192–204.

Nunnari, Salvatore. 2016. "Dynamic Legislative Bargaining with Veto Power: Theory and Experiments." working paper.

Palfrey, Thomas R. 2016. "Experiments in Political Economy." *The Handbook of Experimental Economics, Vol. 2*, , ed. John Kagel and Alvin Roth, 347–434. Princeton University Press.

Roth, Alvin. 2005. "Towards a focal-point theory of bargaining." *Game-theoretic models of bargaining*, , ed. Alvin Roth, 259–268. Cambridge University Press.

Salz, Tobias, and Emanuel Vespa. 2016. "Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation." working paper.

Sethi, Ravideep, and Ewout Verriest. 2016. "The Power of the Agenda Setter: A Dynamic Legislative Bargaining Model." working paper.

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

Vespa, Emanuel. 2016. "An Experimental Investigation of Strategies in the Dynamic Common Pool Game." working paper.

Vespa, Emanuel, and Alistair J. Wilson. 2016. "Communication With Multiple Senders: An Experiment." *Quantitative Economics*, 7(1): 1–36.

Yaari, M. E., and M. Bar-Hillel. 1984. "On dividing justly." *Social Choice and Welfare*, 1: 1–24.

Young, H. Peyton. 1993. "An evolutionary model of bargaining." *Journal of Economic Theory*, 59: 145–168.

Young, H. Peyton. 1996. "The economics of convention." *Journal of Economic Perspectives*, 110: 105–122.