HISTORY DEPENDENCE IN REPEATED BARGAINING

MARINA AGRANOV, CHRISTOPHER COTTON, AND CHLOE TERGIMAN

ABSTRACT. In models of dynamic multilateral bargaining, the literature tends to focus on stationary subgame perfect or stationary Markov perfect equilibria, which restrict attention to forward-looking, history-independent strategies. We consider the appropriateness of such refinements in the context of two simple bargaining games in which a committee repeatedly divides a budget. In one game, agenda-setting power is allocated randomly in each period; in the other game, a successful agenda setter can maintain power with support from other committee members. Although it is reasonable to believe that reputation or reciprocity may play important roles in these environments, the standard equilibrium refinements rule out such considerations and predict per-period outcomes that are identical to those in a non-repeated game. Evidence from a lab experiment shows that the standard equilibrium refinements consistently fail to match the data. Reputation and relationships matter, allocations tend to be more equitable and power is more persistent than standard theory predicts. Our results call into question the validity of restricting attention to static, history-independent strategies in dynamic bargaining games. We show that allowing for even limited history-dependence is sufficient to explain patterns in the data.

Date: **Current version: May 2018**.

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 Ryan Oprea, Guillaume Frechette, Alessandra Casella, 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, Michigan State University, University of British Columbia, Quebec Political Economy Conference, Wilfrid Laurier, 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 experiment.

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.[1] 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 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 (SSPE) or stationary Markov perfect equilibria (MPE)), which restrict attention to memoryless strategies. Limiting attention to such equilibria allows one to rule out many of the subgame perfect equilibria which exist in these games, and focus attention on what is sometimes a unique outcome. Doing so 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 also does not allow players to develop reputations and maintain relationships, things that may be important when groups must repeatedly come to agreements. Second, focusing attention on a single equilibrium imposes homogeneity in outcomes. This ignores the possibility that under a fixed set of bargaining procedures or bargaining rules, different groups may coordinate on substantially different equilibria.

In this paper, we develop simple models of repeated bargaining and design an experimental test to illustrate some issues that arise with a focus on stationarity. We have two main results. First, allowing players to condition on past actions is essential for understanding the observed behavior in our experiments. Therefore, the stationary equilibrium refinements are not suitable for predicting behavior in dynamic multilateral bargaining environments. Second, allowing for multiplicity of equilibria is important for understanding how very different levels 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 legislative bargaining in the literature. A committee

[1]For example, budget committees meet every year to bargain over the allocation of scarce resources and standing committees in legislatures repeatedly interact to determine policy and regulation.

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. Then, the committee votes on the proposal. If a majority of committee members 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 rather 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 the minimum number of other committee members with minimum allocations such that a majority is willing to pass the proposal rather than reject. In equilibrium, proposals always pass right away, the largest share of the budget goes to the agenda setter, smaller shares go to each member of a minimum winning coalition, and the remaining committee members receive nothing. 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 end once they reach an agreement.

Building upon the Baron and Ferejohn framework, our paper presents two models of repeated multilateral bargaining. In both of our models, the game does not end after the players reach an agreement. Rather, the committee must meet to agree on a budget at regular intervals, such as every budget cycle or fiscal year. This means that after a budget passes, committee members do not go their separate ways but instead begin to work on the next cycle's budget. Using the language introduced in Anesi and Siedmann (2015), we use the term repeated bargaining to refer to standing committees, members of which serve over several cycle and reach a sequence of policy agreements, while we use the term one-shot bargaining to refer to ad-hoc committees, which meet only once and adjourn for good after reaching a decision.

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. Our two models differ only in whether the 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 Random Power Model is the simplest possible version of repeated Baron and Ferejohn bargaining. The random assignment of agenda setter power each period is the standard assumption in the literature. Our second version assumes that an agenda setter can hold onto power so long as she obtains the explicit support of a majority of committee members.[2] This Majority Support Model is more consistent with real-world legislatures in which legislative and committee leaders can maintain power over many periods.[3]

Intuitively, one may expect that our two repeated bargaining models will lead to different outcomes, both compared with the one-shot game and compared 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, however, the SSPE outcome (or equivalently the stationary Markov perfect equilibrium (MPE) outcome) in each budget cycle of the two repeated games is identical to the outcome in the one-shot game. Stationarity assumes away history dependent strategies, which means that in the Random Power game is analytically equivalent to an infinite series of independent one-shot games. It also means that in Majority Support 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 Majority Support game is also one in which there is a random selection of agenda setter at the beginning of each cycle, and the outcome in each cycle is the same as the one-shot game.

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

[2]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. Our results do not necessitate that a formal vote takes place every round, only that a lack of majority support leads to a new agenda setter being selected.

[3]That an agenda setter can persist over budget cycles is a realistic assumption. For example, in the United States, we estimate that the chairman of the House appropriations committee stays in power 5.5 years on average.

In both the Random Power and Majority Support games, proposed allocations tend to be more inclusive and more generous than those predicted by the SSPE. We frequently observe an equal division of resources within winning coalitions (whether minimum winning or not), as well as 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 the Majority Support 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 of the committee, and with those members expecting the generosity to continue if they allow the agenda setter to remain in 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. Furthermore, the majority of these allocations give the agenda setter a far lower share of the allocation than the SSPE predicts and that Agranov and Tergiman (2014) have observed in a one-shot bargaining environment.

This raises a natural question: how can we reconcile observed laboratory outcomes with theoretical predictions? We show that extending the theory to allow for asymmetric strategies, risk aversion or fairness concerns does little to eliminate the disconnect between 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.[4]

The experimental data illustrate that focusing on a single subgame perfect equilibrium could not explain the majority of outcomes across the two treatments. Multiplicity itself is a defining characteristic of the repeated multilateral bargaining experiments.

In the Majority Support treatment, for example, outcomes are divided into the following categories: equal division among a minimum winning coalition (31 percent),

[4]We show that complicated strategies conditioning on complex histories of a game are not necessary. Indeed, players do not need to remember much about the history of play to have subgame perfect equilibria in which a successful agenda setter holds power in the Majority Support game. So long as players can remember the most recent player to propose an unexpected allocation or cast an unexpected vote, players may expect to be excluded from future allocations if they are seen as not doing their 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.

equal division among all players (27 percent), unequal division within a minimum winning coalition (26 percent), and unequal divisions inclusive of everyone (16 percent). This means that with the same set of bargaining rules, outcomes differ both regarding inclusiveness and regarding 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. Introducing a new equilibrium refinement that maintains uniqueness would, therefore, be a mistake, as would limiting the attention to an equilibrium refinement that assumes away reputation and relationships since neither of those would match the observed data.

Overall, our paper highlights the fundamental difference in bargaining behavior and bargaining outcomes that arise when committee members repeatedly interact 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, the vast majority of outcomes observed in one-shot bargaining settings are stationary equilibria outcomes which are memoryless. The difference between repeated and one-shot settings is especially noticeable given that the two bargaining games we consider here admit the same unique stationary subgame perfect equilibrium.

The remainder of the paper is organized as follows. Section 2 discusses the relevant literature. Section 3 presents a series of simple multilateral bargaining games, including the predictions of stationary subgame-perfect equilibrium in the environments. Section 4 describes the design of the lab experiments used to assess the theory. Section 5 presents the results of the experiments, showing a substantial gap between observed behavior and the theoretical predictions. It also considers several theoretical extensions, showing how they fail to reconcile theory and observed outcomes while maintaining focus on stationary equilibria. Section 6 documents the empirical patterns of strategies used by our experimental subjects, and makes the case that limited history dependence is required to support equilibrium outcomes observed in our experiments. It also makes the case that multiplicity of equilibria are important in these games. Section 7 concludes.

2. Related Literature

In the last few decades, legislative bargaining has received a great deal of attention both in 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 the 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 focuses on *dynamic* bargaining, we will focus our review 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) 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.[5] 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. 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, many of 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).

Our analysis is also related to papers that endogenize legislative rules within the context of a repeated bargaining game. McKelvey and Riezman (1992) and Eguia and Shepsle (2015) consider dynamic legislative bargaining setting where legislatures can choose the probability with which different members serve as proposer each period. For example, they consider the outcomes when more-senior legislators may be likely randomly selected as Proposer in any given period. The selection of an agenda setter in each period remains random, even if the probability is not the same for all committee members. Our majority-support model takes a different approach to endogenous proposal power in a repeated bargaining game by allowing a single committee member to retain power indefinitely, as long as a majority of other committee members agree to let her do so.

[5]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.

None of these models consider the possibility that an agenda-setter can, through her actions, hold onto power in the long term. Thus they are unable to study the effects of long-term persistence of power. Given how long committee chairmen and women in the United States legislature serve for, understanding how holding onto power affects bargaining outcomes is a natural and important topic to study. We thus depart from the existing literature and consider situations where 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 and mostly focused on one-cycle bargaining games (see the survey by Palfrey 2016) and has only recently moved on to dynamic bargaining experiments. Some of the experimental papers on dynamic bargaining papers focus on the evolution of status-quo policies 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 an environment. Baron, Bowen and Nunnari (2016) extend this setup by considering the 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. 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. To the best of our knowledge, this is the first paper which looks at the dynamic bargaining settings without the status-quo structure, and, instead, focuses on the agenda-setting rules used by committees.

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.

A small set of recent papers have considered whether Markov perfection is consistent with behavior in experiments involving dynamic games. This literature has not reached a consensus. Several papers document that the comparative static predictions implied by Markov perfect equilibria organize experimental data well. This is the case for example in Battaglini, Nunnari and Palfrey (2012, 2016) who make this point in a dynamic legislative bargaining game with durable public goods; Salz and Vespa (2016) who study an infinite-horizon entry/exit game of oligopolistic competition; Vespa (2016) who studies a dynamic common pool game; and Agranov and Elliott (2018) who investigate decentralized bargaining games with heterogeneous trade opportunities and irreversible exit. 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). Also, Vespa and Wilson (2016) study an extension of an infinitely-repeated prisoner's dilemma game with two states and 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 to reach "better" outcomes. 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

We consider 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. **The one-shot bargaining benchmark.** There are $n \geq 3$ identical players in 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 agenda setter is randomly selected, with each member of the committee being selected with probability $1/n$. The process repeats with the new agenda setter proposing an allocation and the other players voting on whether to pass the proposal or select a new agenda setter.

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 subgame perfect equilibrium (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's strategies depends only on the forward-looking 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), The agenda setter 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 interact repeatedly over many budget cycles. After a committee passes a proposal, the game does not end. Instead, the bargaining process continues, and the committee works 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

[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.

[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.

first round agenda setter 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 a delay in passing one year's budget does not impose a delay in the following year's bargaining.

We consider two alternative models of repeated bargaining. The first model assumes that the agenda setter is randomly determined 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 model assumes that a successful agenda setter holds onto power if at least m other players vote to let her maintain power. This is more similar to a real-world setting, in which committee chairs or legislative leaders can hold onto power as long as they can maintain majority support. In this model, following the passage of a proposal, the committee votes on whether to keep or replace the successful agenda setter from the previous cycle. If at least m other players vote in favor of the agenda setter, she serves as the initial agenda setter 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" (Rand) and "Majority Support" (MS), respectively. We refer to the non-repeated Baron and Ferejohn model as either the One-Shot or Baseline model.

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 in the sense that fewer assumptions are needed for an allocation to be supported as part of an SPE.

---

[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).

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 some SPE of the 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 first proposition 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.

Maintaining any allocation as part of a SPE does not require that players utilize unfeasibly complicated history-dependent strategies. It is enough that players can condition their strategy on the identity of the most recent player to choose an unexpected strategy by deviating from equilibrium. As long as players care enough about future outcomes, simply being able to remember or infer the identity of the most-recent deviant is enough to eliminate the incentive for any player's unilateral deviation from a given equilibrium strategy. This does not require players to remember anything beyond the immediately preceding point in time; they only need remember if someone was excluded or deviated in the previous in the most-recent point in time. In this sense, players having limited, one-period memory is sufficient to support any allocation in a SPE.

To formalize this additional insight, we introduce some additional notation. Let Λ denote an arbitrary equilibrium, including the allocation and voting strategies of all players. Let $\mu_\Lambda^t \in \{\emptyset, 1, ..., n\}$ represent a history-dependent, period-specific state variable determined by the observed behavior or state variable in the immediately preceding period, $t-1$. In period t, $\mu_\Lambda^t = i$ if player i unilaterally deviated from Λ in period $t-1$, and $\mu_\Lambda^t = \mu_\Lambda^{t-1}$ if no player deviated from Λ in period $t-1$. If multiple players deviated in the previous period, μ_Λ^t identifies the player that deviated most recently, i.e., in the later stage of that period.[11] At the initiation of the game, $\mu_\Lambda^0 = \emptyset$.

To know μ_Λ^t is to know which player deviated most recently from the specified strategies of Λ. It does not provide information about when or in which way the player deviated. The following proposition shows that being able to condition strategies on this piece of information is a sufficient degree of history dependence for any allocation to be maintained as part of an equilibrium.

[11] If multiple players deviated in the same stage, then μ_Λ^t is randomly assigned to one of them.

Proposition 2. *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 some SPE,* Λ*, of the Random Power and Majority Support models that generates* $\mathbf{a}^*$ *along the equilibrium path with probability 1, and in which each player's strategy in any period* t *conditions only on* μ_Λ^t *and (when voting)* x^t*. When* $m \geq 2$*, such an equilibrium exists for every* $\gamma > 0$*.*

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 strategy. This does not require that players have a long memory; it is enough for players to remember who, if anyone, was systematically excluded or "blacklisted" in the most-recent past, without requiring anyone to remember any details about when or why that person was first blacklisted. The intuition behind this is as follows. As long as players can condition their strategies on the identity of the most recent player to deviate from some given strategy, 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 address 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 affect outcomes.[12] In the remainder of this section, we derive the SSPE of our three games noting that the same results could be obtained instead by characterizing 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.

[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).

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 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 when we revisit the theory in later sections. 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 proposed share is high enough that he 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 himself) 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 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 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 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.

[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.

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 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 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 several 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, and are equal to outcomes in the one-shot game.*

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

[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 $1 - m\bar{a} > \bar{a}m/(n-1)$.

more than he needs to with no more people than he needs to 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 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 match, subjects were randomly divided into groups of three and assigned an ID number. 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 $\gamma = 0.7$. Each cycle consists of potentially many stages, depending on whether a budget proposal failed or not. 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, i.e., the total earnings over all the cycles 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 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 agenda setter. The agenda setter 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

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

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 agenda setter. The agenda setter 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 in the first stage of every cycle the committee has 200 tokens, in the second stage, the available budget is 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 agenda setter.

In the *Random Power* treatment, each cycle of a game is identical to the first one: the agenda setter 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 holds a second vote in which all members vote on whether to retain the current agenda setter for the next cycle. To retain power, the current agenda setter needs to obtain a majority of votes in the second vote. If the current AS is voted out, the agenda setter in the next cycle is randomly chosen. Importantly, the difference in how the agenda setter changes from one cycle to the next is the only difference between treatments.

In each cycle, after the ID of the agenda setter for the current cycle was announced but before the AS submitted her proposal, members of the committee can 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 agenda setter submitted her proposal and was then disabled during the voting stage.

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 that are be payoff-relevant only if a match was not terminated before that cycle. In other words, subjects play four cycles without knowing whether or not their decisions will matter for payment. At the end of a block, subjects learn whether the match ended within that block and, if so, in which cycle. If the match was not terminated, subjects proceed to play a new block of four cycles. Subjects were paid only for the cycles that occurred before the match was 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 to obtain distinct enough predictions of the stationary subgame perfect equilibrium between treatments.

Table 1 summarizes the details of all our experimental sessions.

Table 1. Experimental Sessions

Treatment	Number of Sessions	Number of Subjects	Number of Matches	Mean number of Cycles per 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)

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

5. Experimental Results

We present our results in the following order. We start with a quick summary of the 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, the typical behavior of a subject in the one-shot bargaining environment is very close to the one predicted by the symmetric SSPE. We then compare the 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 the dynamics present in our repeated games and analyze how coalitions evolve over bargaining cycles. We conclude by analyzing the conversations 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 can consider the behavior of our experimental subjects after they have familiarized themselves with the game and interface. Also, restricting ourselves to the first block of the cycles, which all groups 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 term these *coalition types*. A *non-trivial share* is defined as share that is larger than five tokens. If only one group member receives more than five tokens, the coalition is a *dictator coalition*. If exactly two members receive non-trivial shares, the coalition is a *minimum wining coalition*. Finally, if all three members receive non-trivial shares, the coalition is a *grand coalition*. Members with non-trivial shares are *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 five tokens.

To compare the outcomes between two treatments we use regression analysis. Specifically, when we compare our two treatments (whether the fraction of a particular coalition type or the share received by the agenda setter), 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 coalitions within a treatment. 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 that we summarized in 5.2. Morton (2012) and Palfrey (2013) provide an excellent review of this literature. Despite the variations in parameters and the specifics of the experimental protocol used in different studies, the behavior of subjects and resulting bargaining outcomes are quite stable. Moreover, when subjects are allowed to communicate with each other, observed behavior is very closely aligned with the predictions of the symmetric SSPE.

Specifically, experimental play is consistent with the three main predictions 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).[17] Finally, the distribution of resources within a coalition is unequal with proposers appropriating a much higher share of resources than the coalition partners.

[17]These results hold true irrespectively of whether one defines a minimum winning coalition in its strict sense, i.e., the non-included members get exactly zero shares, or in its weak sense, i.e., non-included members are allocated tiny non-zero shares.

5.3. **Bargaining Outcomes within a Cycle in Repeated Games.** We now turn our attention to our two repeated bargaining games: the Random Power game and the Majority Support game. Recall that both of these repeated games 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.

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

	Random Power	Majority Support
Coalition type		
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 (% among MWC)	80.8%	56.0%
Equal split (% among Grand coalitions)	83.1%	65.0%

In Table 2 we present the distribution of coalition types for proposals that passed without delay. The results differ from the one-shot environment. While in the one-shot games minimum winning coalitions are the most common structure of coalitions by far, in both versions of the repeated game, the fraction of grand coalitions is substantial. In fact, in the Random Power treatment, which is a mere repetition of the one-shot game, grand coalitions are formed more than 70% of the time. In the Majority Support treatment, the fraction of grand coalitions reaches 41.8%. Regression analysis confirms that the proportion of three-person coalitions is higher in the Random Power than in the Majority Support treatment ($p = 0.088$).[18]

Further, while in the one-shot setting proposers appropriate a higher share of resources compared to to the shares of their coalition partners, in the repeated game this is not necessarily the case. While in both treatments, on average, agenda setters receive higher shares than their coalition partners, the last two rows of Table 2 show that in both treatments, for both coalition types, allocations between coalition members are, in their

[18]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.

majority, equal splits.[19] Naturally, coalition size affects the share that the agenda setter can appropriate for herself. Figure 1 shows the histograms of shares received by agenda setters conditional on coalition type in each of our treatments. For each coalition type, the vertical lines indicate the average share of agenda setters. 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).[20] Comparing across treatments, we find that the shares of agenda setters in the Random Power treatment are significantly lower than in the Majority Support treatment ($p = 0.025$ for MWCs and $p = 0.085$ for Grand coalitions.)[21]

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

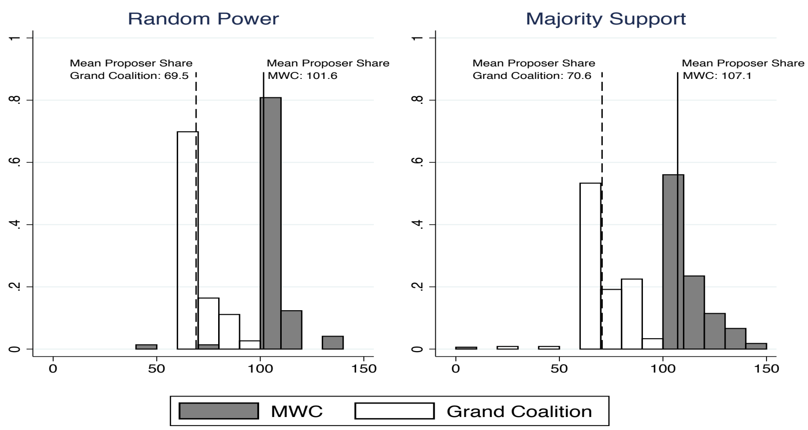

5.4. **Dynamics: behavior and bargaining outcomes across cycles in repeated games.** We now turn towards analyzing behavior across cycles. We no longer restrict ourselves to proposals that pass right away, but instead look at behavior dynamics both in groups that had proposals rejected and those that didn't.

[19]Since agenda setters almost never give themselves smaller shares than others, it is not surprising that on average they would receive more, as shown by a series of panel OLS regression with clustering at the session level, using Share as the dependent variable and whether someone was an agenda setter as the independent one. The p-values on agenda setters are at most 0.006. The coefficients are all positive.

[20]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.

[21]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.

5.4.1. *Persistence of power.* As we have seen in the previous section, agenda setters obtain shares that are on average larger than the shares obtained by any other committee member. Thus, holding the agenda setter seat has obvious benefits. The two repeated games that 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 game, it is largely hindered by design in the Random Power game, in which the identity of the agenda setter in each cycle is determined randomly and independently of past proposals and behavior.

In the Random Power treatment, the event in which the *same* agenda setter serves in all four cycles of the first block is quite rare, as it only happens 7% of the time. In contrast, 91.7% of Majority Support committees operate with the same agenda setter in all four cycles of the first block.

The number of cycles in which the same agenda setter holds onto power directly affects his/her long-run payoff in the game. In the Majority Support treatment, the first agenda setter in a match earns, on average, 355 tokens compared with 265 tokens for the first agenda setter in the Random Power treatment (these are significantly different, $p < 0.001$).[22]

5.4.2. *Evolution of coalitions.* We begin the analysis of the evolution of coalitions types by considering the frequency with which these change. Table 3 shows the likelihood of a coalition type being proposed, conditional on the type of coalition that passed in the previous cycle. As evident from the transition matrix, coalition types are highly persistent: in 87% or more cases, the next cycle proposal has the same coalition type as the one passed in the previous cycle. This is also the case in the Random Power treatment where proposers change across cycles.

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

[22] This is the p-value 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 -88.8. We cluster at the session level.

Next, we consider the stability of coalition partners across cycles. To do this, we focus on the persistence of the minimum winning coalition partner in all instances where the agenda setter was the same in two consecutive cycles. [23] Our data show that when an agenda setter retains her seat in two consecutive cycles, the chances that she will re-invite the same non-proposer in her coalition are 80.0% 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 agenda setters who are forming minimum winning coalitions are not choosing their coalition partners randomly.[24] That is, minimum winning coalitions tend to be stable across cycles. Additionally, 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 Power and Majority Support treatments, respectively. Thus, not only are coalitions stable regarding the identity of coalition members, but when that is the case, the shares given to the coalition partners also are largely constant. In other words, agenda setters seek stability.

We conclude this section by documenting the long-run inequality in group members' payoffs. 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 cumulative distribution functions of the Gini coefficient in each committee.

FIGURE 2. Empirical cummulative distribution functions of the GINI coefficients by treatment

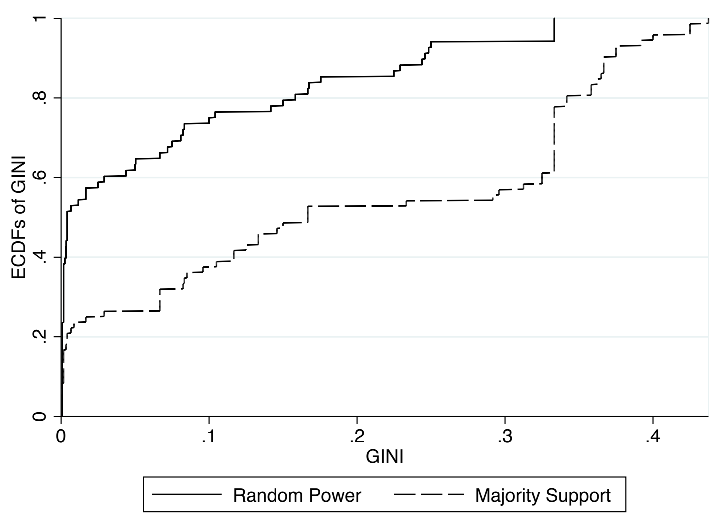

As evident from Figure 2, the Random Power treatment features a much more equal distribution of long-run payoffs compared to the Majority Support treatment ($p < 0.001$

[23] This is the only non-trivial case, since in grand coalitions all members are coalition partners by definition.
[24] In the Random Power treatment we obtain $p = 0.020$, while in Majority Support $p < 0.001$.

for a Kolmogorov-Smirnov test). Two forces contribute to this result. First, the frequent turnover of the agenda setter, which is a built-in feature of the Random Power treatment, increases the chances that different members serve as agenda setters in different cycles. Consequently, when the agenda setter 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, in turn, lead to a more equal distribution of resources within a committee compared with two-person coalitions.

5.5. **Summary of our experimental results in relation to the SSPE.** In this section, we summarize the results of our experiments and compare them with the predictions of the stationary equilibrium refinement.

Bargaining outcomes within a cycle are efficient in that we observe very few delays. This is the case in both environments and is consistent with the theoretical predictions of the symmetric SSPE. 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 very 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 of whether the coalitions are minimum winning or grand. In particular, agenda setters share the surplus equally with their coalition partner in 81% of Random Power minimum winning coalitions, and 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 an agenda setter should appropriate a strictly higher share of the resources compared with the coalition partner. These per cycle outcomes are also in sharp contrast to the behavior in the symmetric one-shot bargaining games, which closely aligns with the predictions of SSPE.

Turning to bargaining outcomes across cycles, our data reveal high a high level of 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 regarding their size, the identity of coalition partners and the shares of coalition partners. All of these observations are at odds with the predictions of the symmetric SSPE. Further, the long-run payoffs of agenda setters are higher in the Majority Support treatment compared with the Baseline treatment, even though the theory stipulates that these payoffs should be identical. Finally, we document that among our two treatments, the Random Power treatment features the lowest inequality regarding long-run payoffs between committee members.

Overall, the symmetric SSPE predictions fail to accommodate the data. 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 regarding the structure of successful proposals or the comparative static predictions of dynamic outcomes across treatments fail to be supported by the data.

5.6. **Robustness of the failure of stationarity.** In the Online Appendix, we consider several alternative assumptions, while continuing to focus on stationary equilibrium refinements. We consider whether reasonable alternative assumptions regarding equilibrium structure or preferences can lead to theoretical predictions that are more in line with the experimental evidence. We show that although the alternative assumptions improve the ability of the theory to match some dimensions of observed behavior, their ability to do so is limited, and can lead to a worse fit with observed behavior on other dimensions. None of the alternative assumptions eliminate our concerns about the stationary equilibrium refinements.

5.6.1. *Asymmetric stationary equilibria.* The main analysis follows Baron and Ferejohn (1989) and much of the literature by focusing on the symmetric SSPE. However, one can alternatively consider the asymmetric SSPE in which players still use stationary strategies but can treat other players asymmetrically, particularly when the agenda setter chooses which player to include in her minimum-winning coalition in each period. 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 cases: the asymmetric SSPE with high persistence of power and the asymmetric SSPE with 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. The asymmetric SSPE does no better than the symmetric SSPE in explaining the data.

5.6.2. *Risk-averse preferences.* Another natural way to extend the theory is to consider outcomes that emerge when bargainers are risk-averse. The introduction of risk-averse preferences leads to a more-unequal split of resources in favor of the agenda setter compared with the risk-neutral case. This pattern is the opposite of what we observe in our data. Intuitively, as risk aversion increases, a coalition partner becomes willing to accept a lower share rather than reject a proposal and risk not being included in the next minimum-winning coalition. Moreover, there is no symmetric SSPE in which there is persistence of power in the Majority Support game. We show that combining risk-averse bargainers doesn't help reconcile theory and data either, as persistence of power in the Majority Support game with asymmetric stationary strategies. Therefore, incorporating risk aversion moves the stationary equilibrium predictions even further away from observed behavior.

5.6.3. *Fairness concerns.* Finally, another possibility is that players care about fairness.[25] To allow for this, we incorporate other-regarding preferences in line with the model of Fehr and Schmidt (1999). As one may expect, if fairness concerns are large and players find it sufficiently costly to provide unequal allocations, then there exists a SSPE of the game in which all players receive an equal share of the allocation in each cycle. Alternatively, when other-regarding preferences are weak, the SSPE allocations resemble those with standard utility functions except that a minimum-winning coalition member needs to be offered a higher allocation to offset the costs of inequality. However, incorporating fairness concerns does not explain other observed behavior. Specifically, we show that in our experimental game, we should never observe equal division within a minimum-winning coalition. This is because any agenda setter who prefers to split equally with her minimum-winning coalition partner will believe it is even better to split equally with *all* committee members. That is, an agenda setter who would consider an even division within a minimum-winning coalition would instead deviate to proposing an equal split in a grand coalition instead. This is the case in all of our games, given the parameter values of our experiments. Thus, fairness concerns may explain some, but not all, of our data. 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 that is very common in both our games as shown in Section 5.3.

6. Reaching stable outcomes in repeated bargaining

Allowing players to condition actions on past events leads to a folk theorem type of result, as we have shown in Propositions 1 and 2. In particular, essentially any outcome can be supported by a subgame perfect equilibrium, including outcomes featuring high persistence of power in the Majority Support game, and allocations that are equal among coalition partners in both the Majority Support and Random Power games. Our data in both treatments indicates that different groups have converged to very different outcomes, regarding both the coalition size and the distribution of resources within a

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

coalition (see Table 2). This section considers how different groups reached such different outcomes.

While all outcomes that we observe in both games are consistent with predictions of Proposition 2, the pre-requisite for implementing outcomes which are *not* the ones predicted by SSPE is the use of history-dependent strategies.[26] In Section 6.1, we explore whether subjects indeed condition their actions on past events. In Section 6.2, we analyze data on the communication between the subjects in our experiment to shed light on how groups achieved such different outcomes. Lastly, in Section 6.3, we discuss the broader layer of behavioral literature, which suggests why many subjects may choose to coordinate on "fair" (equal) allocations irrespectively of the chosen coalition size.

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.

We start by documenting strategies that include punishment. In the Random Power treatment, if a previously excluded member becomes the agenda setter, she excludes the previous agenda setter from a minimum winning coalition 83.3% of the time. A one-sided test of proportions shows that this fraction is significantly larger than 50% ($p < 0.01$).[27] Given the very high persistence of power observed in the Majority Support treatment, to obtain a reasonable number of observations related to punishment behavior, we look at all cases in which there was turnover in agenda setter identity and no longer restrict the data to the last four matches. In these situations, the agenda setter who failed to pass the proposal in the previous cycle is excluded from the new agenda setter's minimum winning coalition in 83.3% of the cases, just as in the Random Power treatment.[28]

Additionally, in the Random Power treatment, we observe reciprocity-type of behavior between former coalition partners. This happens when a minimum-winning coalition partner from cycle $c - 1$ is selected to serve as the agenda setter in cycle c. In this case, the former minimum-winning coalition partners invite the previous agenda setter into their coalitions 81.8% of the time, a fraction that is significantly greater than 50% according to a one-sided test of proportions ($p = 0.017$). Thus, committee members attempt to

[26]Recall that about 94% of all allocations observed in the Random Power treatment and 75% of all allocations observed in the Majority Support treatment are not minimum winning coalitions with unequal shares among coalition partners.

[27]In almost 75% of cases this new agenda setter proposes a minimum winning coalition.

[28]We only have 6 observations of this type in the Majority Support treatment, and the one-sided test of proportions is significant with p=0.051.

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 is more likely to invite the former proposer in his/her minimum winning coalition in the future.

Next, we analyze communication within groups to explore the processes by which different groups reach different equilibria.

6.2. **Effects of communication on equilibrium selection.** Establishing partnerships and long-term relationships may require committee members to agree on specific terms. This is where communication plays an important role: our analysis below shows that communication allows group members to coordinate on playing a particular equilibrium.

We start by noting that our subjects use the communication tool very often: in the Random Power treatment, 79% of groups (54 out of 68) engage in conversations with each other before budget proposals are submitted during the first block of interactions in experienced cycles. In the Majority Support treatment, this fraction is 97% (70 out of 72 groups).[29]

What do subjects discuss? To investigate the effects of communication on equilibrium selection, we 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 different aspects that were relevant to the game (game rules, strategies, proposals, threats, etc..). Most relevant messages fell into one of two broadly defined content categories: "Fair" or "Selfish".[30] The category "Fair" 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 "Selfish" category includes messages that contain information about one's own share and lobbying for the interest of oneself or a subgroup, potentially at the expense of another subject, 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 how the coders classified statements from each individual within a cycle, agreements are very high and range between 87.1% and 89.2%.[31]

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%.[32]

[29]Frequency of meaningful conversations is the highest in the first cycle of a block in both treatments: more than 60% of groups talk before the first budget proposal is submitted. The later cycles feature smaller number of meaningful conversations.

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

[31]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.

[32]The Kappa scores were 0.85, 0.69, 0.81, 0.78 for fair statements in the Random and Majority support treatments and self statements in those same treatments.

FIGURE 3. Fraction of conversations with various content.

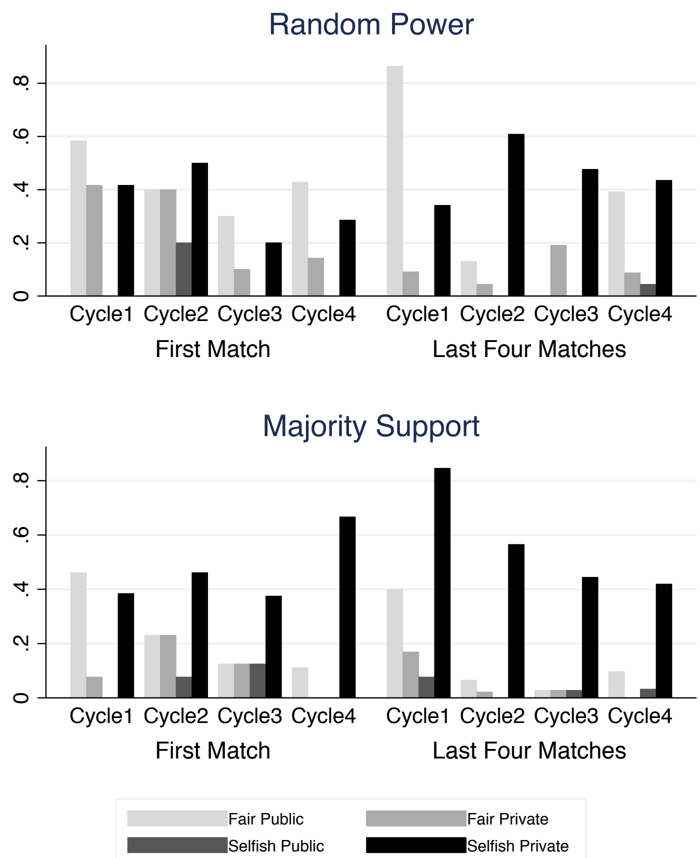

Figure 3 shows the content of conversations that occur in each cycle. For each conversation, we classify relevant messages according to whether message content is about equality and fairness, or, about lobbying for one's own interest. We separate messages that are sent privately to one other member of the group and those sent to the entire group. Thus, we have with four categories: private/public messages that contain fairness and equality statements (Fair Private/Fair Public), and private/public messages that contain statements with requests for own share of resources (Selfish Private/Selfish Public). Interestingly, even in the last 4 matches (the experienced matches), we observe great variety in types of messages that subjects send to each other and communication channels used to transmit these messages. Indeed, in Random Power treatment, some subjects lobby for equality and fairness using public messages, while others talk in private and lobby for own shares. Similarly, in the Majority Support treatment, many discussions are done in private and involve subjects lobbying for their own interests; however, there is also a significant fraction of conversations, in which subjects publicly lobby for fairness (about 40% in the first cycle of the experienced games).

We note that the content of conversations and the use of communication channels are very different in the repeated bargaining games studied here and in the one-shot bargaining games documented in Agranov and Tergiman (2014) and in Baranski and Kagel (2015). Indeed, in the one-shot bargaining games, vast majority of conversations happen in private between the proposer and potential coalition partners and most of these conversations include lobbying for one's own interest. The distinctive feature of the bargaining negotiations in the repeated setting depicted in Figure 3 is the variety of channels and types of messages that subjects continue to use even after gaining experience with the game (last 4 matches).

Conversations between members of a group affect the size of the coalition that the proposer forms, as the random-effects GLS regression analyses in Table 4 highlight. The dependent variable is an indicator of proposing a 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 both treatments, proposers are less likely to form minimum winning coalitions when some group members talk about fairness and equality using a public chat message.

TABLE 4. Effect of Conversations on Coalition Size

	Random Power	Majority Support
Indicator for Fair Public message this Cycle	$-.11^{**}$ (0.06)	-0.23^{***} (0.08)
Indicator for Fair Private message this Cycle	0.04 (0.25)	-0.15^{**} (0.07)
Indicator for Selfish Public message this Cycle	omitted	-0.07 (0.38)
Indicator for Selfish Private message this Cycle	0.68^{***} (0.06)	0.48^{***} (0.08)
Match	-0.07^{***} (0.02)	-0.00 (0.02)
Constant	0.54^{***} (0.17)	0.34 (0.23)
# of observations	68	72
# of subjects	36	45
R-square overall	0.557	0.258

Notes: Errors are clustered at the session level. ***, **, * show significance at 1%, 5% and 10% levels.

Conversations between members of a group also affect the likelihood of proposing a coalition with equal shares to coalition members conditional on the coalition size. In particular, in the Random Power game, we observe a 20% increase in the fraction of coalitions in which resources are divided equally between all members of the coalition (be that MWC or grand coalition) in response to group conversations that discuss fairness and equality. Similarly, this increase is equal to 22% in the Majority Support game.[33]

Overall, analyses of the chats suggest that communication serves as a coordination device for equilibrium selection between group members. Proposers take these conversations seriously (despite chats being cheap talk) and respond to them regarding both coalition size and the division of resources within a coalition.

6.3. **Coordination on fair outcomes.** One fascinating feature of our experimental data is the prevalence of equal distribution of resources between coalition partners both when grand coalitions are formed as well as within minimum winning coalitions. This feature is present in both the Random Power and Majority Support treatments, which suggests that outcomes involving equal division of resources among coalition members might be a more general feature of allocations in repeated (rather than one-shot) bargaining games. From the communications data, we know that subjects often focus on the fairness of allocations during discussions. Why is coordinating on equal allocations so appealing to subjects? To answer this question, we turn to the broader behavioral and theoretical literatures and discuss several reasons for the emergence of equal allocations in the repeated bargaining setup.

The literature on equilibrium selection in games provides 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 a subset) may serve as focal points, and help facilitate coordination on a particular equilibrium. 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).[34]

[33]Specifically, in Random Power game, the probability of proposing allocation with equal shares to all coalition members is 92.3% when group conversations involved discussing fairness and equality, while such fraction is only 72.4% absent such discussions. This difference is significant at the 5% level ($p = 0.0276$). Similarly, in Majority Support game, the probability of proposing allocation with equal shares to all coalition members is 66.7% when group conversations involved discussing fairness and equality, while such fraction is only 45.2% absent such discussions. These differences are marginally significant ($p = 0.0720$).

[34]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-shot bargaining framework.

Further, recent work by Andreoni et al. (2016) introduces the notion of myopic fairness to support the idea of equal division of resources within a minimum-winning coalition. Instead of evaluating proposed allocations in terms of the 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, one might view the equal division equilibrium in a repeated environment as the simplest and most intuitive one as compared with other allocations including 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 allocations that will be offered each period. Equilibria involving an equal division among members of a winning coalition on the other hand, involve little complex reasoning, with the agenda-setter each period splitting the allocation equally with at least m coalition partners, who in turn vote in favor of the allocation (and vote in favor of the agenda setter in the Majority Support game). 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.

7. Conclusion

In bargaining situations, repeated interactions are ubiquitous. With repeated interactions, participants may develop reputations and form relationships that may allow them to sustain a level of cooperation that would not be possible with a one-time interaction. Intuition suggests that such considerations may matter for the outcomes of repeated bargaining games, just as they matter in other settings. Despite this, the theoretical work on dynamic bargaining typically relies on a strong stationarity assumption in order to obtain a unique equilibrium predictions. By assumption, these theories rule out any history-dependent strategies, thereby ruling out strategies that involve reciprocity or punishment, and the possibility that committee members are concerned about developing relationships and reputations.

To clearly illustrate our concerns, we extend Baron-Ferejohn (1989)'s one-shot bargaining model to develop two simple models of repeated bargaining. The first is the simplest-possible model of repeated bargaining, with the one-shot model played over and over again. The second model allows a successful agenda setter in one period to hold onto power into the next period if she has the support of a majority of other committee members. In both repeated games, there is opportunity for reciprocity and punishment for past behavior, and the possibility that relationships and reputation matter. However, when we apply that standard stationarity equilibrium refinement to our games, none of these possibilities matter. In both games, the per period SSPE outcome is the same as the outcome in the non-repeated bargaining game. In settings where we intuitively think that dynamics should matter, they do not; at least not under the standard equilibrium refinement used in the literature.

We design an experiment to study the behavior of subjects when placed in both those settings. We show that stationarity refinements are at odds with laboratory data, which itself aligns with the intuition that in repeated interactions, people condition on the past. Indeed, our laboratory results show that subjects largely condition behavior on past events. What's more, we show that outcomes differ widely across homogeneous groups, suggesting that the multiplicity of reasonable equilibrium outcomes seems to be a inherent characteristic of repeated multilateral bargaining. In particular, agenda setters are often awarded for fair behavior by being allowed by others to hold onto power. Even in the case when agenda setting power is randomly assigned each period, an agenda setter's behavior is often rewarded or punished in subsequent rounds. Regarding allocations, we observe both frequent grand coalitions and minimum-winning coalitions. Within these coalitions types, divisions sometimes favor the AS, but more often involve an equally split across coalition members. This suggests that in repeated interactions, coordinating on "fair" outcomes is a normative behavior. These empirical findings are in sharp contrast with equilibrium predictions under the standard stationarity assumption used in the literature.

We show that the stationary equilibrium assumption included in SSPE or MPE leads to predictions in dynamic multilateral bargaining games that are inconsistent with the way that most groups play such games. At the same time, however, our results suggest that any particular equilibrium refinement will fail to explain the majority of the data. The multiplicity of outcomes is important, as outcomes differ widely across groups. Which equilibrium outcome is most common in any given environment may be less of a theoretical question, and more of an empirical one. In any case, our work provides a clear case for broadening the equilibrium concepts used in repeated settings.

References

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

Agranov, Marina, and Chloe Tergiman. 2017. "Communication in Bargaining Games with Unanimity." *Working Paper*.

Agranov, Marina, and Matt Elliott. 2018. "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.

Anesi, Vincent, and Daniel Siedmann. 2015. "Bargaining in Standing Committees with an Endogenous Agenda." *Review of Economic Studies*, 82: 825–867.

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.

Appendix A. Proof for Main Theoretical Results

Proposition 1 follows immediately from Proposition 2. We prove Proposition 2 below. The online appendix walks through the analysis of the alternative models considered in Section 5.6, showing how they fail to reconcile theory and observation while maintaining the stationarity assumption.

Proof to Proposition 2. Let $\bar{\mathbf{a}}_{\setminus i} = (\bar{a}_1, ..., \bar{a}_n)$ be some a time-invariant allocation in which $\bar{a}_i = 0$, and for all other players $\bar{a}_j \in [0,1]$ and $\sum \bar{a}_j = 1$. Let K^j define a subset of m agents which does not include j. Let $\lambda_{\setminus i}$ denote the following strategy profile:

- The player serving as AS_t proposes allocation $\mathbf{x}^t = \bar{\mathbf{a}}_{\setminus i}$.
- Each player $j \in K^{AS_t}$ votes in favor of the period t proposal (and in Majority Support votes in favor of AS_t) if and only if $\mathbf{x}^t = \bar{\mathbf{a}}_{\setminus i}$.
- Each player $j \notin K^{AS_t}$ votes against the period t proposal (and in Majority Support votes against AS_t).

Suppose that Λ defines a strategy profile such that (i.) in any period t in which $\mu^t_\Lambda = i$ (for any i), players choose strategies according to $\lambda_{\setminus i}$, and (ii.) in any period $t = (c, r)$ in which $\mu^t_\Lambda = \emptyset$, players choose strategies as follows:

- The player serving as AS_t proposes allocation $\mathbf{x}^t = \mathbf{a}^{r*}$.
- Each player $j \in K^{AS_t}$ votes in favor of the period t proposal (and in Majority Support votes in favor of AS_t) if and only if $\mathbf{x}^t = \mathbf{a}^{r*}$.
- Each player $j \notin K^{AS_t}$ votes against the period t proposal (and in Majority Support votes against AS_t).

The strategy profile defined by Λ involves the AS conditioning only on μ^t_Λ, and player voting strategies conditioning only on μ^t_Λ and $\mathbf{x}^t$. Therefore, Λ meets the requirements imposed on it by the conditions in the Proposition.

We derive conditions under which Λ is an SPE of each game, Random Power and Majority Support. In equilibrium Λ, any player that deviates from Λ expects to be excluded from all future allocations. We will show that this threat of exclusion is sufficient to maintain any allocation profile $\mathbf{a}^{r*}$ on the path of play in Λ. Since this result holds for any $\mathbf{a}^{r*}$ including time-invariant allocations in which a given player is excluded (e.g.

including $\bar{\mathbf{a}}_{\setminus i}$), it implies that the exclusion of a deviating player from all future coalitions is sequentially rational off the equilibrium path of play; therefore exclusion of a deviating player is a credible threat, sustainable as part of a SPE.

Let $\mathbf{a}^t_\Lambda$ denote the equilibrium allocation given by Λ in period t (accounting for the history of the game up to t). Depending on game history, $\mathbf{a}^t_\Lambda$ is either $\bar{\mathbf{a}}_{\setminus i}$ for some i or $\mathbf{a}^{r*}$.

In both games $(RandomPower, MajoritySupport)$, the AS each period has no incentive to deviate from offering $\mathbf{x}^t = \mathbf{a}^t_\Lambda$. If she deviates to offer something besides this, then other players vote against the proposal and it fails, and the AS is excluded in future periods, returning a NPV of current and future periods equal to 0. Even if the AS receives nothing herself from $\mathbf{a}^{r*}$, she has no incentive to deviate as she receives 0 given any deviation.

In both games $(RandomPower, MajoritySupport)$, no player who is expected to vote in favor of a proposal has an incentive to vote against it. For every $j \in K^{AS_t}$ in period t, voting against $\mathbf{a}^t_\Lambda$ causes the proposal to fail, and causes the game to enter a subgame equilibrium in which j is excluded from all future allocations. This leads to a NPV of current and future payoffs equal to 0. Even a player who is excluded in the current equilibrium has no incentive to deviate, as she also receives 0 given any deviation.

In both games $(RandomPower, MajoritySupport)$, no player who is expected to vote against the equilibrium allocation has an incentive to deviate, as voting in favor of the equilibrium allocation will not change the current period outcome and will lead to the deviating player being excluded for the duration of the game.

It remains to show that no player would choose to vote in favor of a proposal other than $\mathbf{a}^t_\Lambda$ in any period t. When $m \geq 2$, no player has an incentive to vote in favor of $\mathbf{x}^t \neq \mathbf{a}^t_\Lambda$, as a single player cannot pass a proposal on his own. Therefore, when $m \geq 2$, deviating to vote for an off equilibrium path proposal does not change the current outcome, but leads to the deviating player being excluded in future periods. Thus, when $m \geq 2$, no player will ever deviate from the strategy of voting against any $\mathbf{x}^t \neq \mathbf{a}^t_\Lambda$.

When $m = 1$, however, a single player can pass a proposal on his own. We must rule out that possibility that any player is willing to vote in favor of an off equilibrium path proposal giving him any share $x_j \leq 1$ where $\mathbf{x} \neq \mathbf{a}^t_\Lambda$. It is sufficient to determine when players are unwilling to accept an off equilibrium proposal offering them $x_j = 1$, as it will imply that j is also unwilling to accept any deviant offer giving him $x_j < 1$. Accepting a proposal with $x_j = 1$ leads to a payoff to player j of 1 in the current period, and to j being excluded in future periods. (Remember μ^t_Λ identifies the most recent deviant, and will therefore be the deviant during the voting stage rather than the proposal stage in the event that both deviate during a given period.) Thus, the NPV of accepting for j of current and future period payoffs is simply equal to the current period payoff of 1.

Rejecting such a proposal in period t leads to an equilibrium in which the current AS is excluded rather than player j, as the current AS had deviated from the equilibrium to make the deviant proposal. This leads to a NPV future payoffs equal to

$$\bar{a}^{j}_{\setminus AS_t}\left(\delta + \frac{\gamma}{1-\gamma}\right).$$

We require that in any period t, no player $j \neq AS_t$ has an incentive to deviate. Therefore, it must be that for each j

$$\bar{a}^{j}_{\setminus AS_t}\left(\delta + \frac{\gamma}{1-\gamma}\right) \geq 1.$$

This requires that for each j, $\bar{a}^{j}_{\setminus AS_t}$ is positive, and γ is sufficiently large. The range of γ for which such a condition is met is determined by the minimum value of $\bar{a}^{j}_{\setminus AS_t}$ across all players $j \neq AS_t$. Thus, the least restriction is placed on γ when $\bar{a}^{j}_{\setminus AS_t} = 1/(n-1)$ for all $j \neq AS_t$. Therefore, we can rewrite the equilibrium restriction on γ as

$$\frac{1}{n-1}\left(\delta + \frac{\gamma}{1-\gamma}\right) \geq 1.$$

Which holds as long as $\gamma \geq \bar{\gamma}$, where

$$\bar{\gamma} \equiv \frac{n-1-\delta}{n-\delta} \in (0,1).$$

As long as $\gamma \geq \bar{\gamma}$, there exists a SPE of any subgame in which any player i is excluded indefinitely.

Appendix B. ONLINE MATHEMATICAL APPENDIX

INTENDED FOR ONLINE PUBLICATION.

The following subsections walk through the analysis of the alternative theoretical models discussed in Section 5.6 of the paper. These alternative models are unable to reconcile the disconnect between theory and observed behavior while maintaining the stationarity assumption.

B.1. **SSPE without symmetry.**

We begin by relaxing the symmetry requirement of SSPE. Rather than require that the players' strategies are independent of other player's identities, as is the standard assumption, we allow for stationary strategies which treat other players asymmetrically, specifically when the AS each period chooses which players to include in her MWC. We focus on pure strategy equilibria in this environment.

Let $\mathbf{x}^j = (x_1^j, ..., x_n^j)$ denote player j's equilibrium proposal strategy, which she makes in every period that she serves as AS. Let $\bar{a}_i^j$ denote player i's voting strategy, where i votes for a proposal made by player j in any period t that i serves as AS if and only if $x_i^t \geq \bar{a}_i^j$.

Consider the following stationary, but asymmetric, strategy profile:

- Each player j chooses a MWC K_j made up on m other players. Player j's proposal gives $x_i^j = X$ for each $i \in K_j$, and $x_i^j = 0$ for each $i \notin \{K_j, j\}$.
- Each player i is included in the MWC of exactly m other players.
- Each player i votes in favor of proposal $\mathbf{x}^t$ if and only if $x_i^t \geq X$ when $i \in K_{AS_t}$ and if and only if $x_i^t \geq Y$ when $i \notin K_{AS_t}$.

We determine the values of X and Y such that the above constitutes an asymmetric SSPE.

First, consider such strategies in the context of Random Power. Here, the switch from symmetric to asymmetric strategies does not change the incentives that players have to accept or reject proposals each period. A player who is offered $\hat{x}_j$ can accept the proposal and expect a NPV of

$$\hat{x}_j + \left(\frac{1}{n}(1 - mX) + \frac{m}{n}X\right)\frac{\gamma}{1-\gamma} = \hat{x} + \frac{1}{n}\frac{\gamma}{1-\gamma}$$

or he can reject the proposal and expect a NPV of

$$\left(\frac{1}{n}(1 - mX) + \frac{m}{n}X\right)\left(\delta + \frac{\gamma}{1-\gamma}\right) = \frac{1}{n}\left(\delta + \frac{\gamma}{1-\gamma}\right).$$

In equilibrium, $\hat{x}_j = X = Y$, and such an offer leaves a MWC member indifferent between accepting and rejecting the proposal each period. Thus, $X = Y = \frac{\delta}{n}$.

Second, consider such strategies in the context of Majority Support. For this game, we must also describe the voting strategies for the players when deciding whether to keep or replace the current period AS. There are two possibilities: either the members of K_j will reelect j, or they will not. Those not in K_j have no incentive to reelect player j as AS.

Suppose that we are in an equilibrium of Majority Support with high persistence of AS power. Thus, for every AS j, players in K_j vote in favor of player j retaining power whenever j is AS.

In this case, we consider the incentives to vote for or against a given proposal. Here, an asymmetric proposal strategy means that players expect to continue to be included in the MWC of an AS who includes them in her proposal strategy. This means that if player j votes in favor of a proposal giving him $\hat{x}_j$ that is made by an AS such that $j \in K_{AS}$, then j expects a NPV of

$$\hat{x}_j + X\frac{\gamma}{1-\gamma}.$$

Accepting the same proposal made by an AS such that $j \notin K_{AS}$ returns a NPV of only $\hat{x}_j$ to player j, as j does not expect to be included in the future MWCs of that AS. In either case, if j votes against the proposal, he again expects

$$\frac{1}{n}\left(\delta + \frac{\gamma}{1-\gamma}\right).$$

In equilibrium, for proposals made by an AS such that $j \in K_{AS}$, $\hat{x}_j = X$ and this leaves player j indifferent between accepting and rejecting. Thus,

$$X + X\frac{\gamma}{1-\gamma} = \frac{1}{n}\left(\delta + \frac{\gamma}{1-\gamma}\right) \rightarrow X = \frac{1}{n}(\delta + \gamma - \delta\gamma)$$

For proposals made by an AS such that $j \notin K_{AS}$, $\hat{x}_j = Y$ and Y leaves player j indifferent between accepting and rejecting. Thus,

$$Y = \frac{1}{n}\frac{\delta + \gamma - \delta\gamma}{1-\gamma}.$$

Given the parameter values, $X < Y$. This means that it is less expensive for an AS to include a player in K_{AS} in her MWC than a non member. Therefore, the AS does not want to deviate to include others in her MWC.

For this case, it remains to determine when the members of K_j prefer to reelect the AS rather than to draw a new AS the next period. At the time the players vote for the AS, they are choosing between a favorable vote, which returns NPV of expected future payoffs equal to

$$X\frac{\gamma}{1-\gamma},$$

and an unfavorable vote which returns NPV of expected future payoffs equal to

$$\frac{1}{n}\frac{\gamma}{1-\gamma}.$$

Thus, players in K_j prefer to retain j as AS as long as $X \geq 1/n$. Plugging in for the value of X determined previously, this gives

$$\frac{1}{n}(\delta + \gamma - \delta\gamma) \geq \frac{1}{n} \rightarrow \gamma \geq 1.$$

This is a contradiction, as $\gamma \in (0,1)$, ruling out the possibility that such an asymmetric SSPE with high persistence of AS power in Majority Support exists.

Next, suppose that we are in an equilibrium of Majority Support with low persistence of AS power. Thus, players vote against the AS in each period. In this case, the incentives to vote for or against a given proposal are the same as in Random Power, as there is a new draw of AS power each period. As such X and Y are the same as in Random Power, with $X = Y = \delta/n$.

For this case, it remains to determine when the members of K_j prefer to draw a new AS the next period, rather than reelect the current AS. Our assumption that players ignore weakly dominated strategies means that the players vote as if they were casting the deciding vote. We need the players in K_j to each prefer to vote against the AS. The calculations are the same as in the case with AS retention, except with a reversed sign of the inequality. Thus, players in K_j prefer to replace j as AS as long as $X \leq 1/n$. Plugging in for X from Random Powergives

$$\frac{\delta}{n} \leq \frac{1}{n} \rightarrow \delta \leq 1.$$

This condition always holds. Thus, in the asymmetric SSPE of Majority Support, the equilibrium resembles that of Random Power, with low persistence of AS power.

B.2. **Risk aversion.** In the symmetric SSPE of both dynamic games, a player i that votes against a proposed allocation obtains expected net present value of

$$\left(\frac{1}{n} \cdot u_i(1 - ma^{\text{Game}}) + \frac{m}{n}u_i(a^{\text{Game}}) + \frac{n-1-m}{n}u_i(0)\right)\left(\delta + \frac{\gamma}{1-\gamma}\right)$$

where a^{Game} denotes equilibrium share of the coalition partner in a specific game. If, on the contrary, i supports the proposed allocation at time t, she gets

$$u_i(x_i^t) + \frac{\gamma}{1-\gamma}\left[\frac{1}{n}u_i(1 - ma^{\text{Game}}) + \frac{m}{n}u_i(a^{\text{Game}}) + \frac{n-1-m}{n}u_i(0)\right]$$

in the Random Power and Majority Support games.

We assume that players have identical CARA utility functions, with

$$u_i(x) = u(x) = 1 - e^{-r\cdot x} \quad \text{for all } i.$$

In the SSPE of Random Power and Majority Support, the minimum acceptable offer $\bar{a}$ solves

$$u(\bar{a}) + \tfrac{\gamma}{1-\gamma}\left[\tfrac{1}{n}u(1-m\bar{a}) + \tfrac{m}{n}u(\bar{a}) + \tfrac{n-1-m}{n}u(0)\right] =$$
$$\left(\tfrac{1}{n}\cdot u(1-m\bar{a}) + \tfrac{m}{n}u(\bar{a}) + \tfrac{n-1-m}{n}u(0)\right)\left(\delta + \tfrac{\gamma}{1-\gamma}\right)$$

This simplifies to

$$1 - e^{-r\cdot\bar{a}} = \left(\frac{1}{n}\cdot(1-e^{-r\cdot(1-m\bar{a})}) + \frac{m}{n}(1-e^{-r\cdot\bar{a}}) + \frac{n-1-m}{n}(1-e^{0})\right)\delta.$$

Plugging in the parameters from the experiment (i.e. m, n, δ) gives

$$7 = 11e^{-r\cdot\bar{a}} - 4e^{-r\cdot(1-\bar{a})}.$$

Solving for $\bar{a}$ gives

$$\bar{a} = \frac{1}{r}\ln\left(-\frac{7}{8}e^{r} + \frac{1}{8}e^{r/2}\sqrt{49e^{r}+176}\right)$$

Using a numerical analysis in Mathematica, we show that this expression for $\bar{a}$ are strictly decreasing in r. Thus, as risk aversion increases, the share allocated to the MWC player decreases. Risk aversion leads to even more inequality.

B.3. **Preferences for fair behavior.** Here, we incorporate other regarding preferences, as proposed by Fehr and Schmidt (1999). In each period, a player i's period utility is

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

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. We focus on the case from the experiment where $n = 3$ and $m = 1$.

Equal division in a grand coalition. First, we determine conditions under which there exists a SSPE in which an equal share is allocated to all players.

Suppose that allocation $\mathbf{a}$ assigns $a_i = 1/3$ for each i. In equilibrium, entering a new period of bargaining gives any player an expected payoff of $1/3$. Fairness concerns do not affect payoffs in the case of equal division.

Anticipating a payoff of $1/3$ in the next period, if the current period proposal does not pass, a player requires utility of at least $\delta/3$ to vote for the current period proposal. Therefore, if the AS in any given period deviates from equal division, she must offer a

MWC player at least $\bar{a}$ for the proposal to pass, where $\bar{a}$ solves

$$\bar{a} - \alpha\bar{a} - \beta(1 - 2\bar{a}) = \delta/3.$$

Thus,

$$\bar{a} = \frac{3\beta + \delta}{3(1 - \alpha + 2\beta)}.$$

For equal division to be an equilibrium, the AS must prefer to allocate evenly, earning 1/3 in any period, than to allocate $\bar{a}$ to a single MWC player. This will be the case if

$$1 - \bar{a} - \alpha(2(1 - \bar{a}) - \bar{a}) \leq 1/3.$$

Plugging in for $\bar{a}$ and simplifying the expression gives the required parameter condition

$$\alpha \geq 1/3.$$

Therefore, as long as α is sufficiently large, there exists an equilibrium in which the players allocate evenly each period.

B.3.1. *Equal split with MWC.* Next, we consider the possibility that there exists a SSPE in which an AS and a MWC partner split the allocation evenly each period, excluding another player.

In equilibrium, each period the AS and MWC partner receive

$$\frac{1}{2} - \alpha\frac{1}{2}\frac{1}{2} = \frac{2 - \alpha}{4},$$

and the excluded player receives

$$-\beta\frac{1}{2}\frac{1}{2} = -\frac{\beta}{4}.$$

From an ex ante perspective, the expected per period utility for each player is

$$\frac{1}{3} - \frac{2}{3}\alpha\frac{1}{2}\frac{1}{2} - \frac{1}{3}\beta\frac{1}{2}2\frac{1}{2} = \frac{2 - \alpha - \beta}{6}.$$

Consider the baseline model, where the AS is randomly selected each period. For equal division between an AS and MWC to be an equilibrium, the AS must prefer such an allocation to any alternative.

It is straightforward to show that an AS prefers equal division with a MWC to any success allocation that gives more than 1/2 to a MWC:

$$\frac{2 - \alpha}{4} + V_{AS} \geq 1 - a_m - \frac{1}{2}\alpha(1 - a_m) - \frac{1}{2}\beta(a_m - (1 - a_m)) + V_{AS},$$

where V_{AS} is the expected payoff to the current period AS from future periods, if the current period proposal passes. V_{AS} depends on which one of the games is being played.

This inequality simplifies to

$$2a_m(2-\alpha+2\beta) \geq 2-\alpha+2\beta.$$

Given that $\alpha, \beta < 1$, this further simplifies to

$$a_m \geq 1/2.$$

Thus, the AS always prefers $a_m = 1/2$ to $a_m > 1/2$ when splitting only with a MWC.

She must also prefer such an allocation to any allocation that gives $a_m < 1/2$ to a MWC partner if

$$\frac{1}{2} - \frac{1}{2}\alpha\frac{1}{2} + V_{AS} \geq 1 - a_m - \frac{1}{2}\alpha(2(1-a_m) - a_m) + V_{AS}.$$

This condition simplifies to

$$2a_m(2-3\alpha) \geq 2-3\alpha,$$

and given that $a_m < 1/2$, it further simplifies to the required condition that

$$\alpha \geq 2/3.$$

The AS must also prefer to allocate evenly with only a MWC rather than to allocate evenly amongst the grand coalition. This is the case if

$$\frac{1}{2} - \frac{1}{2}\alpha\frac{1}{2} + V_{AS} \geq 1/3 + V_{AS} \rightarrow \alpha \leq 2/3.$$

This implies that, except for a knife edge case where α is exactly 2/3, the two conditions cannot be simultaneously satisfied. It is unreasonable to believe that the knife edge condition is satisfied.[35] Therefore, we conclude that incorporating other regarding preferences a la Fehr and Schmidt (1999) cannot lead to equal division with a MWC being consistent with SSPE in the baseline game.

Finally, we must establish that the other players would accept an allocation of equal division amongst a grand coalition, if the AS were to deviate from equal division with a MWC to make such a proposal. (Otherwise the AS's preference for such an allocation over equal division with a MWC is not an acceptable deviation.)

A player votes in favor of equal division if

$$1/3 + V_i \geq \left(\frac{2-\alpha-\beta}{6}\right)\left(\delta + \frac{1}{1-\gamma}\right),$$

where V_i is the player's expected future payoff from the proposal passing. V_i depends on the game, and whether we are considering symmetric or asymmetric SSPE.

[35] Assuming that the common α is the realization of any continuous distribution with no mass points implies that $\alpha = 2/3$ is a zero probability event.

In the baseline game and the symmetric SSPE of the majority support game (where reelection does not occur as part of equilibrium for the same reasons it did not occur originally), $V_i = (2 - \alpha - \beta)/6$ and the required inequality simplifies to

$$1/3 \geq \frac{2 - \alpha - \beta}{6}\delta \rightarrow 2(1 - \delta) + (\alpha + \beta)\delta \geq 0\delta,$$

which is clearly satisfied given $0 < \alpha, \beta, \delta < 1$.

In the asymmetric SSPE of the majority support game (where the equilibria are of the structure considered in the earlier subsection on asymmetric equilibria), $V_i = (1/2)\gamma/(1 - \gamma)$ for the player that is included in the AS's MWC strategy, and $V_i = 0$ for the player that is excluded. The included player will clearly support the equal division within a grand coalition deviation, rather than risk a player that excludes him being selected as AS in the future.

The above analysis rules out SSPE with equal shares to the AS and a MWC for the Random Power and Majority Support games.

Appendix C. Instructions for the Majority Support treatment

This is an experiment in the economics of decision making. The instructions are simple. If you follow them carefully and make good decisions you may earn a considerable amount of money which will be paid to you at the end of the experiment. The currency in this experiment is called tokens. The total amount of tokens you earn in the experiment will be converted into US dollars: 10 Tokens = $1. You will also get a participation fee upon completion of the experiment.

General Instructions

(1) In this experiment you will be playing 8 Matches. During each Match, you will be randomly assigned an ID and you will be asked to make decisions over a sequence of Rounds.

(2) The number of Rounds in a Match is randomly determined as follows:

You will play every Match in blocks of 4 Rounds. Even though you will complete all 4 Rounds in each block you play, not all Rounds in a block will necessarily count towards your earnings for the Match.

The *first* Round in a Match will always count towards your earnings for that Match. Whether any of the following ones will count will be randomly determined according to the **"70% rule:"** after each Round that counts towards your earnings in a match, there is a 70% chance that the next Round will also count towards your earnings in a Match. The computer will determine this by randomly choosing a number between 1 and 100. If the number is less or equal to 70 then the next Round will also count towards your earnings for this Match.

Note however, that this random draw is done **"silently."** That is, you will play all four Rounds in a block but you will only find out at the end of the block which Rounds actually count towards your earnings for this Match. If each random draw the computer makes in a block is less or equal to 70, then you will move to the next block of 4 Rounds and so on. **Your earnings for a Match consist of the sum of all your earning over <u>all the Rounds</u> up until the computer drew a number above 70 for the <u>first time in the Match</u>.** The Match ends after the last Round of the block in which the computer drew a number above 70 for the first time.

(3) Once a Match ends, you will be randomly and anonymously rematched with two other people in this room to start a new Match. Each member in the group will again be randomly assigned an ID number. Thus, while your ID remains the same over Rounds *within* a Match, it is very likely to vary from Match to Match

and you will not be able to identify who you've interacted with in previous or future Matches.

(4) **What Happens in Each Match**

- In each Match you will be randomly matched into groups of three members. Each member in the group is randomly assigned an ID number. Thus, while your ID remains the same over Rounds *within* a Match, it is very likely to vary from Match to Match.
- At the start of each Match, one of the three members in your group will be randomly chosen to be the Proposer.
- Step 1: The Proposer's task is to propose how to split a budget of 200 tokens between himself and the two other members of his/her group.
- Step 2: Once the Proposer has submitted a budget proposal, all members of your group will observe the budget proposal and will vote on it.
  - (a) If a proposal receives a **simple majority of votes** (i.e. two or more members in your group vote in favor of the proposal), then the proposal passes and for this Round the earnings for each of you in the group will correspond to the number of tokens offered to them in that proposal.
  - (b) If a proposal receives **fewer than 2 votes** then it is defeated. If a proposal is defeated, you will remain in the same Round, but the computer will then randomly choose one of the three members of your group to be the "new" Proposer. Each member of your group (including the previous proposer) has the same chance of being chosen (1 in 3). Whoever is chosen will submit a new proposal. However, the number of tokens to be divided will be reduced by 20% relative to the preceding proposal and rounded to the nearest integer. Thus, if the first proposal is rejected, then after a "new" Proposer is randomly selected, his/her proposal will involve splitting 160 tokens. If this proposal is rejected, again a "new" proposer will be chosen and his/her proposal will involve splitting 128 tokens, etc... This goes on until a proposed allocation gets 2 or more votes and passes.

Once a proposal receives two or more votes (whether right away or after a delay), you remain in the same group and will move onto the next Round. The Proposer who submitted the successful proposal remains in place, the budget then restarts at 200 tokens and you return to Step 1. This process repeats itself until a Match ends, which is determined by the 70% rule described above. Once a Match ends, you will start a new Match and will be randomly re-matched to form new groups of three. Remember: while your ID remains the same over Rounds *within* a Match, it is very likely to vary from Match to Match.

(5) **Communication:** In each Round, before the Proposer submits his/her proposal, members of your group will have the opportunity to communicate with each other using a chat box. The communication is structured as follows. On the top of the screen, each member of the group will be told her ID number. You will also know the ID number of the Proposer. Below you will see a box, in which you will see all messages sent to either all members of your group or to you personally. You will not see the chat messages that are sent privately to other members of your group. You can type your own message and send it to one or both members of your group, and only the person(s) you select as recipient(s) will receive your message. The chat option will be available until the Proposer submits his/her proposal. At this moment the chat option will be disabled.

(6) Remember that in each Match subjects are randomly matched into groups and the ID numbers of the group-members are randomly assigned. Thus, while your ID remains the same over Rounds *within* a Match, it is very likely to vary from Match to Match.

(7) **Your Payment:** You will each receive a show-up fee. In addition, at the end of the experiment, the computer will randomly choose one out of 8 Matches that you played. You will be paid for ***all the Rounds that actually counted towards your payment within that Match*** (determined according to the 70% rule).

(8) **Screenshots:** We will now slowly go through different screenshots so you can familiarize yourself with the types of screens you'll be seeing. The examples we are about to go through are not meant to show you what you ought to do in this experiment but are just there to show you on screen the different possible stages of a Match. Please raise your hand if you have any questions about the experiment and/or interface.