



Persistence of power: Repeated multilateral bargaining with endogenous agenda setting authority[☆]

Marina Agranov^{a,*}, Christopher Cotton^b, Chloe Tergiman^c

^a Division of Humanities and Social Sciences, California Institute of Technology, United States of America

^b Department of Economics, Queen's University, Canada

^c Smeal College of Business, Penn State University, United States of America

ARTICLE INFO

Article history:

Received 21 March 2019

Received in revised form 4 December 2019

Accepted 5 December 2019

Available online xxxx

JEL classification:

C78

D02

C92

Keywords:

Stationary equilibrium

Dynamic bargaining

Repeated games

Laboratory experiment

ABSTRACT

We extend a simple repeated, multilateral bargaining model to allow successful agenda setters to hold on to power as long as they maintain the support of a majority of other committee members. Theoretically and experimentally, we compare this Endogenous Power environment with a standard Random Power environment in which agenda setters are appointed randomly each period. Although the theoretical analysis predicts that the two environments are outcome equivalent, the experimental analysis shows substantial differences in behavior and outcomes across the games. The Endogenous Power environment results in the formation of more stable coalitions, less-equitable budget allocations, the persistence of power across periods, and higher long-run inequality than the Random Power environment. We present evidence that the stationary equilibrium refinements traditionally used in the literature fail to predict behavior in either game.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Many bargaining situations involve repeated interactions. This is true in personal relationships and is also the case in formal decision-making bodies such as committees, legislatures, and corporate boards. 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. A growing literature on repeated multilateral bargaining explores such interactions, often focusing on how decision-making or policy outcomes evolve over time.

We extend this literature by endogenizing the turnover of proposer power (or 'agenda-setter' power) within committees. The existing

literature generally assumes that power within committees is exogenously determined, typically randomly assigned each period. Even a successful leader with support from other committee members is unable to hold onto power, meaning that frequent turnover of agenda setting power is built into the models. In contrast, we allow successful agenda setters to hold onto power as long as they maintain the support of a majority of other committee members. We explore the implications of such an assumption in an otherwise simple model of multilateral bargaining.

Our endogenous-power model offers two advantages over alternative bargaining models with random power. First, our assumption allows for outcomes that are more consistent with observations about the ability of leaders to hold onto power within prominent real world committees. For example, agenda setting authority in the US Congress rarely changes hands. Since the first US Congress in 1789, for example, there have been only 63 changes in the Speaker of the House. No more than 25 of these changes may be attributed to the speaker losing support among his party.¹

¹ Of the changes, 30 corresponded to a change in the House majority party, four were due to the death of the sitting speaker, one was due to Schuyler Colfax stepping down to become vice president, and two were due to Henry Clay temporarily stepping down to dabble in international relations. Our estimate of 25 speakers who may have lost the support of their party as being the changes in speaker that took place within or between sessions of congress in which the same party maintained power.

[☆] We appreciate feedback from Emanuel Vespa, Ryan Oprea, Guillaume Frechette, Marco Battaglini, 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, Cornell Political Economy Conference, 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-Deutch Chair in Economic and Financial Policy at Queen's University. Tergiman is grateful for generous funding from the Social Sciences and Humanities Research Council of Canada (SSHRC), which provided funding for the experiment.

* Corresponding author.

E-mail addresses: magranov@hss.caltech.edu (M. Agranov), cotton@econ.queensu.ca (C. Cotton), cjt16@psu.edu (C. Tergiman).

Committee chairmanships show similar consistency. In the US Senate, for example, there have been 35 different chairmen of the appropriations committee since its establishment in 1867.² The average Speaker of the House held the position for 4.3 years.³ Second, endogenizing agenda setting power allows us to more completely explore some of the dynamics that arise in a repeated game involving the incentives players may have to develop reputations and sustain long-term relationships or implicit agreements with one another. Such considerations may be particularly salient when maintaining support of other committee members may allow an agenda setter to hold on to power.

In exploring the endogenous-power bargaining environment, we first consider a game-theoretic analysis of a multilateral divide-the-dollar decision incorporating the revised assumptions that allows agenda setters to hold onto power across rounds of bargaining as long as they maintain the support of a majority of other committee members. If a successful agenda setter does not maintain the support from a majority, then the next round agenda setter is randomly determined. Following the theoretical analysis, we present outcomes from a laboratory experiment designed to compare bargaining outcomes in endogenous-power and random-power environments.

The theoretical analysis first establishes a folk theorem result, showing that any allocation may be maintained as part of a subgame-perfect equilibrium. We then apply the standard equilibrium refinement used in the analysis of bargaining games, focusing on stationary subgame-perfect equilibria (SSPE) (or, equivalently in our environment, stationary Markov perfect equilibria (MPE)). In the endogenous-power game, applying such a refinement restricts attention to memoryless, history-independent strategies, and rules out equilibria in which agenda setters establish coalitions or reputations that allow them to hold onto power across multiple periods. Under the standard stationary refinement, outcomes in our endogenous-power environment are equivalent to outcomes in a random-power environment where players simply cannot hold onto power.

In summary, although agenda setters maintaining power is consistent with subgame-perfect equilibrium, it is not consistent with the standard equilibrium refinements used in the multilateral bargaining literature, which predict that outcomes will be the same whether a model assumes that proposer power is randomly assigned each period, or assumes agenda setters can hold onto power with majority support. Under the standard theoretical refinements, endogenous power only complicates the analysis without changing the predicted equilibrium outcome.

In contrast to the theoretical analysis, however, the experimental analysis identifies substantial differences in committee member behavior and bargaining outcomes in the random-power and endogenous-power environments. First, we observe that agenda setters are able to exploit the institutional rules in the endogenous-power environment and hold on to power for long stretches of time; this occurs for >90% of all first-round agenda setters. On the contrary, in the random-power environment, such behavior is not feasible by design, and we observe high turnover of agenda setters. As a result, in the long run, payoffs in the random-power environment tend to be relatively equitable, reflecting the fact that no individual or group can maintain indefinite control over proposals. In the endogenous-power environment, in contrast, a single agenda setter tends to maintain power with the support of a minimum winning coalition over many periods. In such cases, those excluded from the

winning coalition receive nothing and tend to remain excluded over many periods. Long-run inequality tends to be higher in the endogenous-power environment.

There are also substantial differences between the two environments in the types of allocations we observe. In both environments, 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. However, in the endogenous-power environment, agenda setters are more likely to provide consistent and substantial allocations to a minimum winning coalition of other committee members, who in turn tend to support the agenda setter staying in power. In other words, coalitions tend to be stable across bargaining cycles both in terms of the coalition partners as well as shares allocated to coalition partners. This is in contrast to the random power environment, where agenda setters are more likely to propose allocations that divide the budget equally across all committee members, and, by institutional design, are not able to hold onto power.

Interestingly, our experimental data illustrate that different groups converged to very different outcomes after multiple repetitions of the game. In the endogenous-power environment, for example, outcomes can be classified into the following categories: equal division among a minimum winning coalition (31%), equal division among all players (27%), unequal division within a minimum winning coalition (26%), and unequal divisions inclusive of everyone (16%). This means that with the same set of bargaining rules, outcomes differ both regarding inclusiveness and regarding equality. Finally, our data also clearly show that in both environments, subjects use strategies that involve punishments, reciprocity and history dependence. These are the tools that subjects use to achieve stability of coalitions across bargaining cycles in the endogenous-power environment. Interestingly, some agenda setters also use these tools in the random-power environment, in an attempt to counter the institutional rotation of power.

The comparison of observed outcomes, bargaining dynamics, and subjects' strategies with those predicted by the stationary SSPE shows that this refinement generally fails to predict the majority of observed outcomes in either the random-power or endogenous-power experiments. More importantly, it fails to predict the observed differences in behavior between the two environments. The one dimension on which the standard stationary equilibrium refinement tends to perform well is in predicting long-run payoffs in the random power environment. However, this is not because the model reasonably predicts behavior or allocation sizes in this environment. Rather, it reflects the fact that the environment is constructed in a way that assures an equal distribution of power in the long run. On all other notable dimensions, the stationary equilibrium refinement is a poor predictor of behavior or allocations for a majority of the groups in the experiment. These results highlight the fundamental difference between the one-shot and the repeated bargaining environments: stationary SSPE organizes the data well in experiments involving ad hoc committees that bargain only once, while it fails to do so when the members bargain over and over across bargaining cycles. This is the case even in the random power environment, despite the fact that it is a mere repetition of the one-shot bargaining game.

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, we show that allowing for even a very limited form of history dependence results in any allocation being consistent

² The mode time in the position is 2 years, or one term of congress. However, some chairs held the position for more than a decade, including Clarence Cannon (19 years), George H. Mahon (15 years), and Jamie L. Whitten (15 years).

³ There have been 53 individuals to hold the position since 1789, representing an average time in the position in excess of 4.3 years. However, many have served substantially longer, including Sam Rayburn (17 years).

with equilibrium, suggesting that reconciling the disconnect between the theory and evidence is not straightforward.⁴

Rather than suggesting that there exists a single equilibrium refinement that predicts the majority of behavior, our data show that multiplicity of equilibrium itself is a defining characteristic of the repeated multilateral bargaining experiments. This suggests that the level of partisanship and cooperation observed within a committee or legislature may have less to do with specific institutional rules and more to do with equilibrium selection. Partisanship and inclusive cooperation both occur very frequently. Therefore, it is likely a mistake to introduce an equilibrium refinement that limits attention to one outcome given that similar groups tend to coordinate on different outcomes in the data. The remainder of the paper is organized as follows. Section 2 discusses the relevant literature. Section 3 presents the two simple multilateral bargaining games (random-power and endogenous-power) and derives predictions of stationary subgame-perfect equilibrium in the two games. 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 is important in these games. Section 7 concludes.

2. Related literature

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 et al., 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 pre-determined and 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 settings 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 proposer 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.

None of these other 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 et al. (2016) extend this setup by considering the effects of various communication channels available to committee members. Nunnari (2019) incorporates veto power and analyzes experimentally the consequences of its presence.⁶ Other papers study dynamic models of public-good accumulation. Battaglini et al. (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

⁴ Section 3.2 summarizes the detailed analysis from Appendix A, in which 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 Endogenous Power 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.

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.

⁶ For the theoretical treatment of the effects of veto power in dynamic bargaining settings see Sethi and Verriest (2016) and Nunnari (2018).

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. In contrast, Battaglini et al. (2016) embed a model of bargaining into a dynamic model of public investment and show that outcomes of an experiment are generally consistent with the theoretical predictions under MPE. Ours is the first to assess the suitability of history-independence in a dynamic bargaining environment with endogenous agenda setting authority. 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 et al. (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 de-centralized 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 Fréchet, 2020). 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. We contribute to this debate by comparing a treatment in which the incentives for long-term collaboration between players is generally minimized with an alternative treatment where there is greater incentives for building coalitions or developing relationships in order to show that the appropriateness of history-independent refinements depends largely on the environment being studied.

3. Repeated multilateral bargaining framework

We consider two models of repeated multilateral bargaining. The first model (Random Power) is a simple version of a repeated multilateral bargaining game in which proposer power is randomly determined each period, regardless of outcomes in previous periods. The second model (Endogenous Power) extends the Random Power game to allow an agenda setter who passes a proposal in one period to continue as the agenda setter the next period so long as the majority of other committee members agree to it. Where the Random Power game may be viewed as a simplified version of the repeated bargaining environments common in the literature (which generally assume the random allocation of proposer power each period), the Endogenous Power game represents a theoretically novel environment.

3.1. Model

Within both the Random Power and Endogenous Power games, a committee of $n \geq 3$ identical players must repeatedly agree how to divide a fixed budget of size 1 between its members. The games take place across many budget cycles, with each cycle requiring agreement over a budget allocation before the game moves on to the next cycle. The game within each budget cycle is a classic closed-rule multilateral bargaining game from Baron and Ferejohn (1989). Each budget cycle

may be made up of multiple rounds of bargaining before a specific budget allocation is agreed upon. Denote any period by $t = (c, r)$, where $c = 1, 2, \dots$ denotes the budget cycle, and $r = 1, 2, \dots$ the proposal round within any cycle.

We first describe the game that takes place in each cycle, which is a common component in both the Random Power and Endogenous Power games.

3.1.1. Bargaining within one budget cycle

Within each budget cycle, the game takes place as follows. One of the n players will serve as the initial 'agenda setter' (AS). The identity of this initial AS will differ across the Random Power and Endogenous Power environments. This initial AS will propose how to divide the budget across the n players. 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 cycle ends. If fewer than m other players vote in favor of the proposal, the proposal fails and the game moves on to the next round of bargaining within the same budget cycle. In the new round, a committee member is randomly selected (each with probability $1/n$) to serve as the new AS. The process repeats with the new AS proposing an allocation and the other players voting on whether to pass the proposal or again reject the proposal and the start of another round.

The game can potentially last many rounds if proposals consistently fail to gain majority support. AS^{*t*} is the AS with proposer power in period t . Let $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$ denote the proposal made by the AS in period $t = (c, r)$, and let $\mathbf{a}^c = (a_1^c, \dots, a_n^c)$ denote the allocation that is eventually implemented in cycle c , with each player i earning $u_i^c(a^c) = a_i^c$ from that budget cycle's allocation. An allocation \mathbf{a}^c is feasible if $0 \leq a_i^c \leq 1$ for each i , and $\sum_i a_i^c \leq 1$; 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.⁷ Delay is costly, with within-cycle discount factor $\delta \in (0, 1)$ applies between rounds within a budget cycle.

3.1.2. Repeated bargaining

Within each budget cycle, our repeated bargaining games each take place as described above. A cycle starts with the 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 entire cycle game repeats.

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, or the probability that the game enters another cycle.⁸ 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 (**Random Power**) 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.⁹

The second model (**Endogenous Power**) assumes that a successful agenda setter holds onto power if at least m other players vote to let her maintain power following the successful passage of a proposal.

⁷ Almost all of the theoretical results continue to hold as long as $m \in \{1, \dots, n - 2\}$, which assures that the AS cannot pass a proposal unilaterally, and that unanimity is not required.

⁸ That is, the next cycle is discounted at γ , and not $\delta^r \gamma$ when the current cycle lasts r stages. The alternative formulations of discounting lead to qualitatively similar results.

⁹ 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).

This is 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.

3.2. Equilibrium

In Appendix A, we establish a folk-theorem result for the repeated bargaining environment, showing that as long as discount factors are not too small, any allocation is consistent with subgame-perfect equilibrium in both the Random Power and Endogenous Power games. To address the multiplicity of equilibria inherent in multilateral bargaining games, the literature typically follows Baron and Ferejohn (1989) and focuses on stationary refinements of subgame-perfect equilibrium, 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.¹⁰ In the remainder of this section, we follow the literature and derive the SSPE of our repeated bargaining 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.¹¹ This means that strategies can only condition on payoff-relevant information, and must ignore payoff irrelevant information about the history of the game.

Applied to our framework, any 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. The analysis makes 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 in later sections and the Online Appendix. Second, we restrict attention to equilibrium 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.¹²

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 ($x_{AS_i}^t = 1 - m\bar{a}$) for herself, and nothing ($x_i^t = 0$) to everyone else. The m players receiving share \bar{a} voting in favor of the proposal. In this group of m players, m 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 budget cycles have a present discounted value of expected payoffs equal to $v \equiv \frac{1}{n} \frac{\gamma}{1 - \gamma}$. A player prefers a proposal that provides him with current-cycle allocation x_i to pass than to fail as long as $x_i + v \geq \frac{1}{n} + v$, or equivalently $x_i \geq \frac{\delta}{n}$.

Note that when players choose symmetric stationary strategies, the incentives that any player has to vote in favor of a proposal are identical in the repeated game with Random Power and a non-repeated, one-budget game. 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.

The equilibrium of the Endogenous Power game is 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 Endogenous Power model collapses to the Random Vote model, with a new AS being randomly selected at the start of each cycle.¹³ This, in turn, implies that the per cycle SSPE outcome in the Endogenous Power game is also identical to the outcomes in a non-repeated game.

In each period of the only SSPE of both the Random Power and Endogenous Power games, the equilibrium budget allocation is identical to the budget allocation in a non-repeated game. A new AS is randomly selected, she proposes to allocate δ/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.

3.2.1. Testable equilibrium predictions

The models generate several testable predictions, summarized in Proposition 1.

Proposition 1. *The unique symmetric SSPE of the Endogenous Power game is outcome equivalent to that of the Random Power game. 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 is randomly determined, independent of the past actions of others.*

¹⁰ See the discussion about when analyses should use SSPE versus MPE in Maskin and Tirole (2001).

¹¹ 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).

¹² 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.

¹³ 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-AS's vote to replace the AS since $1 - m\bar{a} > \bar{a}m/(n - 1)$.

- (iii) *Proposals pass without delay.*
- (iv) *There is low persistence of AS power.*

These predictions have implications for the long-term outcomes of the game, which are based on the observation that in both environments, players are equally likely to serve as the agenda setter and equally likely to be selected for any MWC. Therefore, proposer power will rotate between players and the overall share of the budget received by each player will be equal in the long run.

Corollary 1. *In the unique symmetric SSPE of the Random Power and Endogenous Power games, in the long run each committee member will*

- (i) *serve as agenda setter in the same share of periods, and*
- (ii) *receive an equal share of the total budget across all periods.*

4. Experimental design

All experiments were conducted at the Center for Experimental Social Sciences at New York University using Multistage software.¹⁴ 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. 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 experimental session, subjects played a multiple-cycle repeated bargaining game eight times. We refer to each of the multiple-cycle games 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. This means that the experimental setup is theoretically equivalent to an infinitely repeated game with a between-cycle discount factor of $\gamma = 0.7$. 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 2 h 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 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 δ

= 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 *Endogenous Power* 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 4 cycles. Within each block, subjects receive no feedback about whether the match has ended or not and they make choices that are payoff-relevant only if a match was not terminated before that cycle. In other words, subjects play 4 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 4 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 4 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.¹⁵

Table 1 summarizes the details of all our experimental sessions.

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

¹⁵ We recognize that subjects who choose to participate in an experiment realize that there is no chance that the experiment will last far beyond the time they were called for, and thus the lab conditions differ from the precise conditions of the theory. However, experimental evidence has shown that subjects do behave consistently with the theory when they are faced with indefinitely repeated games and that they behave differently when faced with finitely repeated games of the same expected length. This shows that the experimental implementation of indefinitely repeated games may indeed lead subjects to consider the indefinite nature of the game (see Dal Bó (2005) for example). In terms of our specific implementation, we note that the instructions did not clarify what would happen in the event that the experiment continued beyond the time it was called for (as opposed to clearly stating that it would end, no matter what, by a certain time or round). In short, our experiment implements an infinite-time horizon in a laboratory by using a finite-time experiment with an indeterminate horizon. An open and worthwhile future question, which we leave to the future research, is under what conditions the finite settings can approximate fully infinitely-repeated games.

¹⁴ The Multistage package is available for download at <http://software.ssel.caltech.edu/>.

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)
Endogenous power	3 sessions	(21,15,18)	(8,8,8)	(4,6,6)

5. Experimental results

We present our results in the following order. First, we document the general characteristics of bargaining outcomes and behavior starting with the dynamics observed *across* bargaining cycles. We then shift to outcomes reached *within* a cycle. Following that, we compare outcomes in our repeated bargaining games with those observed in the one-shot bargaining games and comment on the suitability of the stationarity refinement to account for the observed outcomes in our settings. We finish by looking into the types of strategies used by our experimental subjects and use conversations between subjects as the window into the mechanism through which subjects form coalitions.

Throughout the analysis, we focus on the Endogenous Power game as our main treatment, and contrast outcomes observed in this treatment with those in the Random Power treatment, which serves as a benchmark. This comparison is informative as it highlights the difference between endogenous and exogenous agenda-setter power in repeated bargaining.

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 data set with identical amounts of experience within a match across all treatments.¹⁶

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 winning 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.¹⁷

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 inter-dependencies between observations that come from the same session since subjects are randomly re-matched between matches.

¹⁶ We conduct a similar analysis using the first block of 4 cycles in all matches, as well as using the first block of 4 cycles in the first four matches to investigate learning over the course of the experiment. Our message is unchanged. These results are in the Online Appendix.

¹⁷ Our analyses are robust to defining coalition types in a strict sense. However, since the 200 tokens can never be equally split among three coalition members, in our stricter definitions we reduce the allowed difference across members to 2 tokens but cannot make it zero. Our results are unchanged and presented in Online Appendix.

5.2. Behavior and bargaining outcomes across cycles

5.2.1. Persistence of power

Our study of repeated bargaining is motivated in large part by the desire to understand whether agenda setters are able to hold on to power for long periods of time. The Endogenous Power treatment enables agenda setters to do that if they can secure the support of a majority of the committee. At the same time, the model can also accommodate the rotation of power if agenda setters are voted down at the end of each cycle. In the Random Power treatment on the other hand, the rotation of power is built-in to the institution itself since the identity of the agenda setter in each cycle is determined randomly and independently of past proposals and behavior.

Our data show that the vast majority of committees in the Endogenous Power treatment operate with the same agenda setter in all 4 cycles of the first block: this happens in 91.7% of all cases. In contrast, in the Random Power treatment, the event in which the same agenda setter serves in all 4 cycles of the first block is quite rare, as it only happens 8.8% of the time. 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 Endogenous Power 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, $p < 0.001$).¹⁸

5.2.2. Evolution of coalitions

How do agenda setters in the Endogenous Power treatment manage to remain in power for such long periods? To answer this question, we first analyze the evolution of coalition types by considering the frequency with which these change. Table 2 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 the Endogenous Power treatment: in 93% or more cases, the next cycle proposal has the same coalition type as the one passed in the previous cycle.¹⁹

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.²⁰ Our data show that when an agenda setter retains her seat in two consecutive cycles, the probability that she will re-invite the same non-proposer in her coalition is 89.6% in Endogenous Power treatment. Test of proportions shows that this percentage is significantly higher than 50% ($p < 0.001$), which means that agenda setters who are forming minimum winning coalitions are not choosing their coalition partners randomly. 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 85.4% of the cases. Thus, not only are the coalitions stable with respect to the

¹⁸ 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.

¹⁹ Coalitions also persist across all four blocks. For example, if we compare the types of proposals that are first proposed in each cycle, conditional on a grand coalition being proposed in the first cycle, 77% and 86% of all remaining cycles are also see a Grand coalition proposed in the Random and Endogenous Power treatments, respectively. In terms of minimum winning coalitions, those numbers are 57% (we note a relatively small number of observations here) and 86%.

²⁰ This is the only non-trivial case, since in grand coalitions all members are coalition partners by definition.

Table 2
Transition of coalition types across cycles.

Cycle c	Cycle c + 1			
	Random power		Endogenous power	
	MWC	Grand	MWC	Grand
MWC	0.87	0.12	0.94	0.06
Grand	0.11	0.89	0.07	0.93

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 in all dimensions in the Endogenous Power game.

Interestingly, despite the exogenously imposed rotation of agenda setter power in the Random Power treatment, the evolution of coalitions in this treatment looks very similar to that observed in the Endogenous Power treatment, though we have a much smaller number of cases in which the same agenda setter was randomly selected to serve in two consecutive cycles. Indeed, when this happens, the agenda setter tends to form the same types of coalitions in the two consecutive cycles (this happens in 87% of cases), and tends to invite the same non-proposer into their minimum winning coalition (this happens in 80% of cases). Further, they always offer the same share to their long-term coalition partners (this happens in 100% of cases).

5.3. Bargaining outcomes within a cycle

We now turn our attention to outcomes observed within a cycle in our repeated bargaining games. In both treatments, vast majority of 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 Endogenous Power treatments, respectively. Therefore, in the remainder of this subsection we concentrate on those proposals that passed without delay.

In Table 3 we present the distribution of coalition types for proposals that passed without delay. In both versions of the repeated game, the fraction of grand coalitions is substantial, reaching 42% in the Endogenous Power treatment and 70% in the Random Power treatment. Regression analysis confirms that the proportion of three-person coalitions is higher in the Random Power than in the Endogenous Power treatment ($p = 0.088$).²¹ Further, while in both treatments, agenda setters receive higher shares than their coalition partners on average, the last two rows of Table 3 show that in both treatments, for both coalition types, allocations between coalition members are, in their majority, equal splits.²²

Naturally, coalition size affects the share that agenda setters can appropriate for themselves. Fig. 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).²³ Comparing across treatments, we find that the shares of agenda setters in the

²¹ This is the p -value on the treatment coefficient in a panel probit regression using Grand Coalition as the dependent variable and the Endogenous Power treatment as the explanatory one. The coefficient itself is -1.667 . We cluster at the session level.

²² 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.

²³ 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 Endogenous Power, respectively. We cluster at the session level.

Table 3
Coalition types for proposals that passed without delay, by treatment.

	Random power	Endogenous power
<i>Coalition type</i>		
Dictator (1-person coalition)	0.0%	0.3%
MWC (2-person coalition)	27.9%	57.8%
Grand (3-person coalition)	72.1%	41.8%
<i>Allocations within coalitions</i>		
Equal split (% among MWC)	80.8%	56.0%
Equal split (% among Grand coalitions)	83.1%	65.0%

Random Power treatment are significantly lower than in the Endogenous Power treatment ($p = 0.025$ for MWCs and $p = 0.085$ for Grand coalitions).²⁴

5.4. Long-run payoffs

Finally, we turn to the long-run payoffs of committee members. Our results above suggest that one should see a greater dispersion in members' long-run payoffs in the Endogenous Power treatment compared to the Random Power treatment. We measure the long-run payoffs of subjects by their average-per-cycle payoffs over the course of an entire block of 4 cycles and present the distribution of these long-run payoffs of subjects in each game separately in Fig. 2.

As evident from Fig. 2, the average-per-cycle payoffs of subjects in the Random Power treatment is centered around 66.7, which is the third of the total budget of 200 tokens. At the same time the distribution of payoffs in the Endogenous Power treatment looks much different. It is in fact tri-modal with significant masses at zero, 66.7, and 100. Kolmogorov-Smirnov test of equality of the distributions rejects that the average-per-cycle payoffs observed in two games come from the same distribution ($p = 0.021$).

Fig. 3 corroborates these findings by presenting the empirical cumulative distribution functions of the Gini coefficients in each committee, which we use to measure the long-run inequality in group members' payoffs. As evident from Fig. 3, the Random Power treatment features a much more equal distribution of long-run payoffs compared to the Endogenous Power treatment ($p < 0.001$ for a Kolmogorov-Smirnov test).

5.5. Summary of experimental results

In this section, we summarize the results of our experiments and compare them with (a) the bargaining outcomes documented in one-shot bargaining games (to highlight the effect of dynamics) and (b) the predictions of the stationary equilibrium refinement.

Overall, our experimental results show that when agenda setter power is determined endogenously as in the Endogenous Power treatment, agenda setters use this institutional feature to their advantage and remain in power for long stretches of time. This, coupled with the fact that agenda setters obtain on average higher shares than other members creates a relatively high level of inequality in the long-run payoffs between committee members. A majority of coalitions (about 60%) formed in this treatment are minimum winning coalitions, which include the smallest subset of members that can successfully pass proposals and vote to keep the same agenda setter in power in the future cycle. However, the remaining 40% of coalitions are grand coalitions, which include all committee members. In general, the evolution of coalitions across cycles features stability across several dimensions: coalition size, the identity of coalition partners and well as their shares. In particular, agenda setters that form minimum

²⁴ 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.

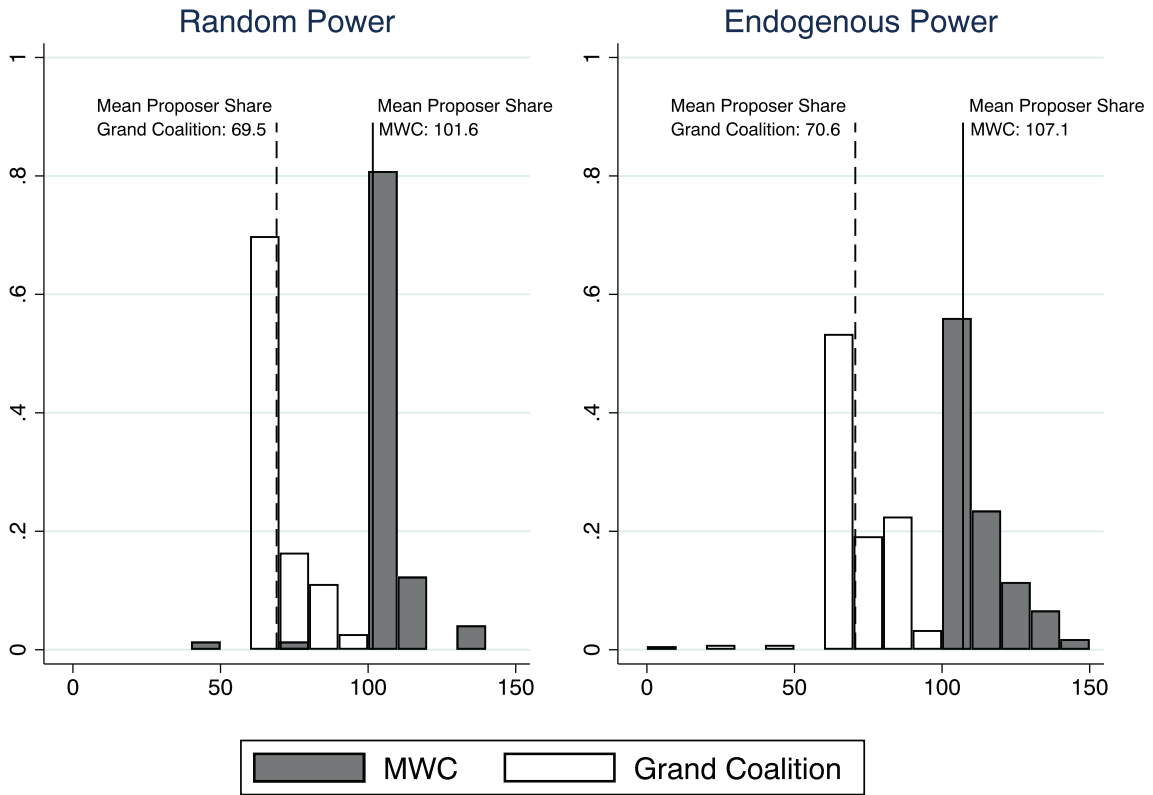


Fig. 1. Agenda setters' shares in proposals that passed without delay.

winning coalitions tend to invite the same partners and offer them the similar shares across cycles.

On the contrary, when rotation in agenda setters' power is institutionalized as in the Random Power treatment, persistence of power is hindered by design, which reduces the inequality in the long-run payoffs between committee members. This also affects which types of

coalitions are formed and passed within each cycle. In contrast with the Endogenous Power game in which most of the coalitions are minimum winning, in the Random Power treatment over 70% of all successful coalitions are grand coalitions which mostly divide the budget equally among the committee members. However, even in this treatment we observe some traits of stability. In the rare cases in which the

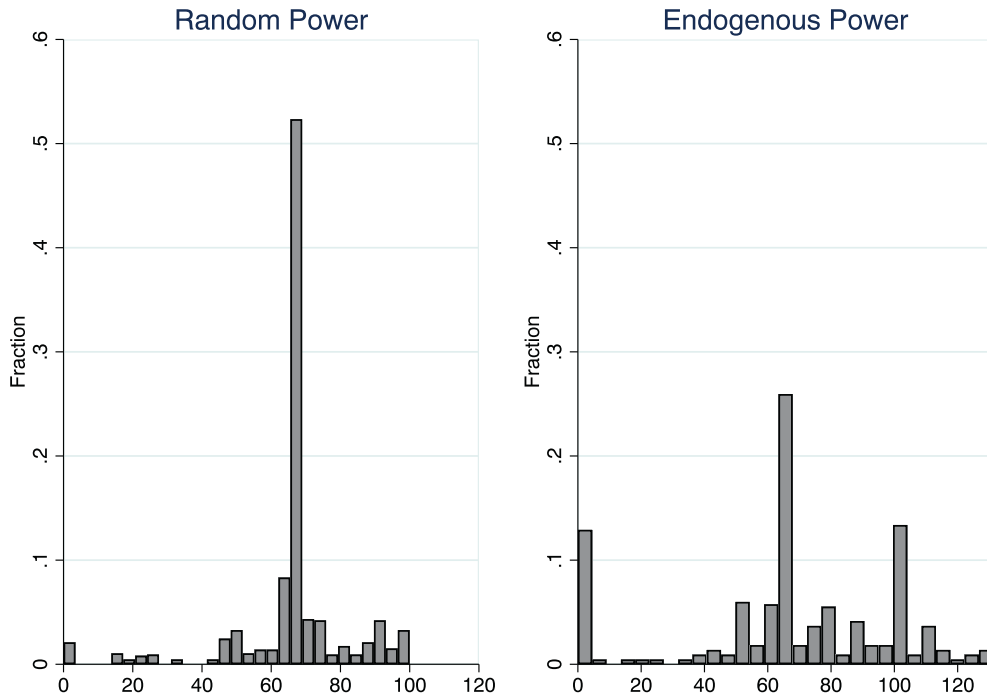


Fig. 2. Long-run payoffs of committee members, by treatment.




Fig. 3. Empirical cumulative distribution functions of the GINI coefficients by treatment.

same agenda setter is randomly selected to serve the two consecutive cycles and she forms the minimum winning coalition in the both cycles, she tends to invite the same coalition partner and allocate to him the same share.

5.5.1. Comparison of results to one-shot bargaining games

An interesting and meaningful benchmark for our results are the outcomes documented in one-shot bargaining games. In particular, we note that our Random Power treatment is a mere repetition of the one-shot bargaining game and so one would naturally expect it to deliver similar outcomes. This is, however, not the case.

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 Section 3. Morton (2012) and Palfrey (2013) provide a thorough 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 and very different from the results we document here for both the Endogenous Power and Random Power games. Specifically, experimental play in one-shot bargaining games features mostly minimum winning coalitions. Depending on the study this fraction is between 65% as in Frechette et al. (2003) and 90% as in Agranov and Tergiman (2014).²⁵ Moreover, the distribution of resources within a coalition is unequal with proposers appropriating a much higher share of resources than coalition partners. Both of these observations are in contrast to what we observe in the repeated games, especially in the Random Power treatment indicating that dynamic bargaining environments are truly different from the static ones.²⁶

5.5.2. Comparison of the results to predictions of symmetric stationary SPE

Further, we compare our experimental results with the theoretical predictions described in Section 3.2, which uses the stationarity refinement, and find that these predictions fail to organize the observed outcomes in both repeated bargaining games. The only two outcomes that match the stationary predictions are that bargaining outcomes are efficient in that we observe very few delays, and that long-run outcomes are on average in line with predictions. The remaining outcomes are not aligned with those derived by the symmetric stationary SPE, which predicts identical outcomes in both treatments. In particular,

²⁵ These results hold true irrespective 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.

²⁶ The experimental literature on one-shot bargaining games has documented that when subjects are allowed to communicate with each other, observed behavior is very closely aligned with the predictions of the symmetric stationary SPE (Agranov and Tergiman (2014) and Baranski and Kagel (2015)).

while the symmetric stationary SPE predicts that all passed proposals should feature two-person minimum winning coalitions, our data shows that both minimum winning and grand coalitions are very common in both games. In addition, conditional on coalition size, at least 50% of passed proposals feature an equal division of the surplus between coalition partners, which is at odds with the symmetric stationary SPE, which predicts that an agenda setter should appropriate a strictly higher share of the resources compared with the coalition partner.

Across cycles, our data reveal high a high level of persistence of power in the Endogenous Power treatment, despite this being ruled out by the stationarity refinement in the theoretical analysis. In both games, the observed coalitions are stable across cycles in terms of their size, the identity and shares of coalition partners. All of these observations are at odds with the more refined predictions of the symmetric SSPE suggesting that while the symmetric SSPE is able to predict the long-run payoffs of committee members in the Random Power treatment, it does not capture the more refined predictions regarding coalition structures. With respect to the Endogenous Power treatment, the symmetric SSPE fails to capture both the observed long-run payoffs of committee members as well as the coalition structures and its dynamics.

Given the failure of the symmetric stationary SPE to organize observed outcomes, 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, especially in the case of endogenous proposer power. In the Online Appendix, we show that although some alternative assumptions help match specific dimensions of observed behavior, their ability to do so is limited, and can lead to a worse fit with observed behavior on other dimensions. Overall, none of the alternative assumptions eliminate our concerns regarding the stationary equilibrium refinements when proposer power is endogenous.²⁷

6. Reaching stable outcomes in repeated bargaining

In this section, we investigate how subjects reach stable coalitions and sustain persistence of power in the Endogenous Power treatment. Furthermore, given that different groups converge to very different outcomes, for instance, in terms of coalition sizes and distributions of resources within a coalition, we aim to understand how do different groups reach such different agreements and manage to preserve them throughout the play.

We proceed as follows. First, we provide empirical evidence of the use of history-dependent strategies, which is a pre-requisite for implementing outcomes that are different from those predicted by the stationary SPE. Second, we show that communication plays an important role in the establishment of partnerships and long-term relationships, as it provides a platform to discuss and agree on which particular SPE to play. This is important in this setting, as a folk-type theorem holds with respect to outcomes that can be supported as SPE. Third, we document the explicit coordination on “fair” outcomes, and argue that it is both the simplest and most intuitive one as compared with other possible allocations including the symmetric stationary SPE allocation.

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-

²⁷ Specifically, we consider three alternative models: one with asymmetric stationary equilibria, one with risk-averse preferences and one with fairness concerns. The only model that may explain some but not all of our data is the model in which agents have other-regarding preferences and, thus, by design care about fairness. However, even that extension fails to accommodate the equal division of resources within a minimum-winning coalitions, which is the prominent feature of our data.

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% ($p < 0.01$).²⁸ Given the very high persistence of power observed in the Endogenous Power 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.²⁹

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 $>50\%$ according to a one-sided test of proportions ($p = 0.017$). Thus, committee members attempt to establish stability even when, by treatment design, stability is hard to establish. Stability increases both because proposers tend to re-invite the same partner in their minimum winning coalition, and because the invited partner 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 to engage in conversations directly pertaining to the game very often³⁰: in the Random Power treatment, 79% of groups (54 out of 68) engage in pertinent conversations with each other before budget proposals are submitted during the first block of interactions in experienced cycles. In the Endogenous Power treatment, this fraction is 97% (70 out of 72 groups).

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.). Conditional on being relevant to the game, most relevant messages fell into one of two broadly defined content categories: "Fair" or "Selfish".³¹ The category "Fair" includes any message that can be interpreted as lobbying for fairness and equality, for example, "equal is nice," "let's just do equal," and "just play fair". The "Selfish"

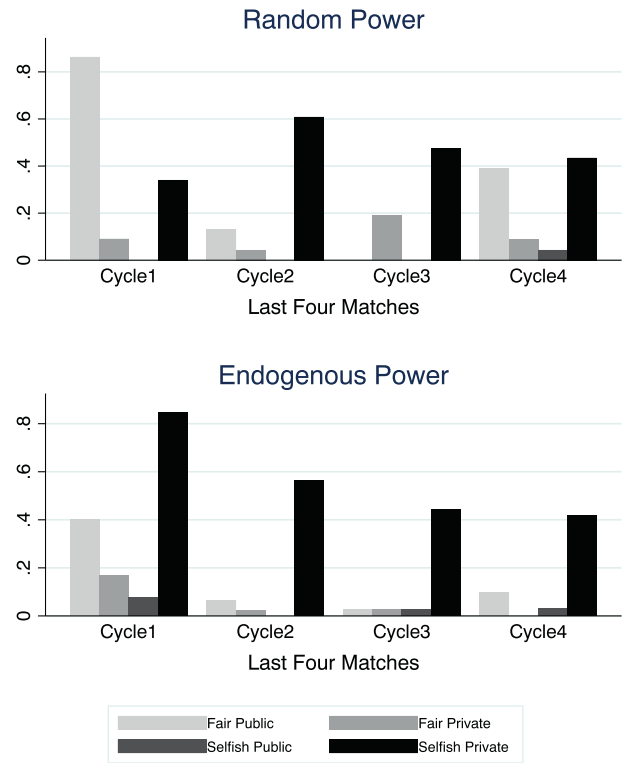


Fig. 4. Fraction of conversations with various contents.

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 will 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%.³²

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%.³³

Fig. 4 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 a great variety in the types and contents of messages that subjects send to each other. Indeed, in the Random Power treatment, some subjects lobby for equality and fairness using public messages, while others talk in private and lobby for their own shares. Similarly, in the Endogenous Power 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 compared with the one-shot bargaining games

²⁸ In addition, in almost 75% of cases this new agenda setter proposes a minimum winning coalition.

²⁹ We only have 6 observations of this type in the Endogenous Power treatment, and the one-sided test of proportions is significant with $p = 0.051$.

³⁰ A statement was counted as pertaining to the game so long as it was relevant to the game. For example, discussions about shares, how many blocks were left, history of play were all counted as pertinent. Statements that contained only salutations or random thoughts about student life were not counted as pertinent.

³¹ 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)).

³² The Kappa scores are 0.71, 0.73, 0.78, 0.68 for fair statements in the Random and Endogenous Power treatments and self-statements in those same treatments.

³³ The Kappa scores were 0.85, 0.69, 0.81, 0.78 for fair statements in the Random and Endogenous Power treatments and self-statements in those same treatments.

Table 4
Effect of conversations on coalition size.

	Random power	Endogenous power
Indicator for fair public message this cycle	-0.11 ^{***} (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. The indicator for Selfish Public message is omitted in the Random Power regression because of colinearity.

* Shows significance at 10% level.

** Show significance at 5% level.

*** Show significance at 1% level.

documented in Agranov and Tergiman (2014) and in Baranski and Kagel (2015). Indeed, in the one-shot bargaining games, the 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 repeated settings, depicted in Fig. 4, 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 for each 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.

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 coalitions) in response to group conversations that discuss fairness and equality. Similarly, this increase is equal to 22% in the Endogenous Power game.³⁴

Overall, the analyses of 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.

The effects of communication that we observe in our repeated bargaining setting are consistent with those documented by Baron et al. (2017) who study dynamic bargaining with a status quo structure. In particular, Baron et al. (2017) find that the communication channel and the content of conversations affects both the duration and the shape of the coalitions formed. Subjects who tend to exclusively communicate using private channels tend to lobby for their own interests

³⁴ 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 Endogenous Power 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 ($p = 0.0720$).

and are more likely to propose a MWC as opposed to universal coalitions compared with subjects who only send messages using the public chat. Moreover, messages related to fairness and equality were positively correlated with universal coalitions, which allocated resources between all committee members. The consistency in our results suggests that the effects of communication are robust to the specifics of bargaining protocol in the dynamic bargaining settings.

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 Endogenous Power 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).³⁵

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 Endogenous Power 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's (1993) claim that the

³⁵ See Fréchet et al. (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.

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

Most theoretical and experimental work involving multilateral bargaining limits attention to settings in which agenda setting authority or proposer power is exogenous. This simplifies the environment on one dimension, allowing these papers to focus attention on other considerations such as endogenizing status quo outcomes and the evolution of policy or budgets over time. However, in many real world bargaining environments, such as those involving legislatures and boards, agenda setting power is endogenous. Those in leadership positions are typically able to hold onto power if they maintain the support of other committee members.

We consider a simple model of multilateral bargaining with endogenous agenda setting authority both theoretically and experimentally. From a theoretical perspective, introducing endogenous agenda setting authority does not change predicted outcomes under the standard equilibrium refinements. This would suggest that endogenizing proposer power may represent an unnecessary complication to the bargaining environment that adds no additional insight. However, the experimental analysis shows that this conclusion is inconsistent with actual behavior.

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 when proposer power is endogenous. In contrast, standard equilibrium refinements predict quite well long-run payoffs of committee members when proposer power is exogenously determined.

Indeed, the introduction of endogenous agenda setting authority into a bargaining environment leads to substantial changes in observed behavior, the sharing of power, and the equality of outcomes that are not predicted by a theoretical analysis that takes the standard approach and limits attention to SSPE or MPE. In a setting where agenda setters can hold onto power with majority support, players have a greater incentive to develop relationships across periods in an attempt to hold onto power. This means any equilibrium refinement that rules out such behavior is less appropriate in such an environment than it is when agenda setting power is exogenously determined.

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 an 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.

In any case, our work provides a clear case for broadening the equilibrium concepts used to study dynamic bargaining games. At the same time, however, we are unable to identify any single equilibrium refinement that predicts the majority of the outcomes in our games. For example, in the Appendix A, we relax the stationarity refinement only slightly, allowing only the conditioning of current actions on the most recent period, and show that the folk-theorem results hold. In other

words, allowing for even a minimum amount of history dependence in strategies leads to anything being possible as part of an equilibrium. Our results suggest that any single equilibrium refinement will fail to explain the majority of the observed data. Rather, multiplicity of outcomes is itself important, as outcomes differ widely across otherwise similar groups. Which equilibrium outcome is most common in any given environment may be less of a theoretical question, and more of an empirical one.

Appendix A. Theoretical appendix

In this section, we formally establish a folk-theorem result for the repeated bargaining games in our paper. We show that any feasible outcome may be maintained as part of a subgame-perfect equilibrium. Furthermore, we establish that this does not require that players utilize unfeasibly complicated history-dependent strategies. It is enough that players can choose strategies to effectively punish the player who deviated most recently from an equilibrium strategy. 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 for establishing the folk-theorem result.

Let Λ denote an arbitrary equilibrium, including the allocation and voting strategies of all players. Let $\mu_\Lambda^t \in \{\emptyset, 1, \dots, n\}$ 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.³⁶ 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.

Proposition 2. *Consider any feasible allocation profile $\mathbf{a}^* = \{\mathbf{a}_i^*\}_{i=1}^n$, such that for every c , $a_i^c \in [0, 1]$ for each i and $\sum_i a_i^c = 1$. As long as γ is sufficiently large, there exists some SPE, Λ , of the Random Power and Endogenous Power 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 and is not weakly dominated in equilibrium. 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

³⁶ If multiple players deviated in the same stage, then μ_Λ^t is randomly assigned to one of them.

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.

A.1. Proof to Proposition 2

Let $\bar{a}_{\setminus i} = (\bar{a}_1, \dots, \bar{a}_n)$ be some 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^i 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 Endogenous Power votes in favor of AS_t) if and only if $\mathbf{x}^t = \bar{\mathbf{a}}_{\setminus i}$
- Each player $j \in K^{AS_t}$ votes against the period t proposal (and in Endogenous Power votes against AS_t).

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

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

The strategy profile defined by Λ involves the AS conditioning only on $\mu_{\setminus i}^t$ and player voting strategies conditioning only on $\mu_{\setminus i}^t$ 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 Endogenous Power. 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}^{c*} on the path of play in Λ . Since this result holds for any \mathbf{a}^{c*} 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 $\mathbf{a}_{\setminus i}^t$ denote the equilibrium allocation given by Λ in period t (accounting for the history of the game up to t). Depending on game history, $\mathbf{a}_{\setminus i}^t$ is either $\bar{\mathbf{a}}_{\setminus i}$ for some i or \mathbf{a}^{c*} .

In both games (*RandomPower*, *EndogenousPower*), the AS each period has no incentive to deviate from offering $\mathbf{x}^t = \mathbf{a}_{\setminus i}^t$. 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}^{c*} , she has no incentive to deviate as she receives 0 given any deviation.

In both games (*RandomPower*, *EndogenousPower*), 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 $\mathbf{a}_{\setminus i}^t$ 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*, *EndogenousPower*), 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 $\mathbf{a}_{\setminus i}^t$ in any period t . When $m \geq 2$, no player has an incentive to vote in favor of $\mathbf{x}^t \neq \mathbf{a}_{\setminus i}^t$, 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 $\mathbf{x}^t \neq \mathbf{a}_{\setminus i}^t$.

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 $\mathbf{x} \neq \mathbf{a}_{\setminus i}^t$. 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 $\mu_{\setminus i}^t$ 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

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

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

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

This requires that for each j , $\bar{a}_{\setminus AS_t}^j$ 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 $\bar{a}_{\setminus AS_t}^j$ across all players $j \neq AS_t$. Thus, the least restriction is placed on γ when $\bar{a}_{\setminus AS_t}^j = 1/(n-1)$ for all $j \neq AS_t$. Therefore, we can rewrite the equilibrium restriction on γ as

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

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. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jpubeco.2019.104126>.

References

Agranov, Marina, Elliott, Matt, 2018. Commitment and (In)Efficiency: A Bargaining Experiment (working paper).
 Agranov, Marina, Tergiman, Chloe, 2014. Communication in multilateral bargaining. *J. Public Econ.* 118, 75–85.
 Agranov, Marina, Tergiman, Chloe, 2017. Communication in Bargaining Games with Unanimity (Working Paper).

- Agranov, Marina, Fréchette, Guillaume, Palfrey, Thomas R., Vespa, Emanuel, 2016. Static and dynamic underinvestment: an experimental investigation. *J. Public Econ.* 143, 125–141.
- Andreoni, James, Aydn, Deniz, Barton, Blake, Douglas Bernheim, B., Aydn, Deniz, Barton, Blake, Douglas Bernheim, B., Naecker, Jeffrey, 2016. When Fair Isn't Fair: Sophisticated Time Inconsistency in Social Preferences (working paper).
- Anesi, Vincent, 2010. Noncooperative foundations of stable sets in voting games. *Games and Economic Behavior* 70 (2), 488–493.
- Banks, Jeffrey S., Duggan, John, 2000. A bargaining model of collective choice. *Am. Polit. Sci. Rev.* 94 (1), 73–88.
- Banks, Jeffrey S., Duggan, John, 2006. A general bargaining model of legislative policy-making. *Quart. J. Polit. Sci.* 1, 49–85.
- Baranski, Andrzej, Kagel, John H., 2015. Communication in legislative bargaining. *JESA* 1 (1), 59–71.
- Baron, David P., 1996. A dynamic theory of collective goods programs. *Am. Polit. Sci. Rev.* 90 (2), 316–330.
- Baron, David P., 1998. Comparative dynamics of parliamentary governments. *Am. Polit. Sci. Rev.* 92 (3), 593–609.
- Baron, David P., Bowen, T. Renee, 2016. Dynamic coalitions (working paper).
- Baron, David P., Ferejohn, John A., 1989. Bargaining in legislatures. *Am. Polit. Sci. Rev.* 83 (4), 1181–1206.
- Baron, David P., Herron, Michael, 2003. A dynamic model of multidimensional collective choice. In: Kollman, K., Miller, J., Page, S. (Eds.), *Computational Models in Political Economy*. MIT Press, Cambridge, MA.
- Baron, David P., Bowen, Renee, Nunnari, Salvatore, 2016. Dynamic Coalitions and Communication: Public Versus Private Negotiations (working paper).
- Baron, Davin, Kalai, Ehud, 1993. The simplest equilibrium of a majority-rule division game. *J. Econ. Theory* 61, 290–301.
- Battaglini, Marco, Coate, Stephen, 2007. Inefficiency in legislative policymaking: a dynamic analysis. *Am. Econ. Rev.* 97 (1), 118–149.
- Battaglini, Marco, Coate, Stephen, 2008. A dynamic theory of public spending, taxation, and debt. *Am. Econ. Rev.* 98 (1), 201–236.
- Battaglini, Marco, Palfrey, Thomas R., 2012. The dynamics of redistributive politics. *Economic Theory* 49 (3), 739–777.
- Battaglini, Marco, Nunnari, Salvatore, Palfrey, Thomas R., 2012. Legislative bargaining and the dynamics of public investment. *Am. Polit. Sci. Rev.* 106 (2), 407–429.
- Battaglini, Marco, Nunnari, Salvatore, Palfrey, Thomas R., 2016. The Political Economy of Public Debt: A Laboratory Study (working paper).
- Bó, Pedro Dal, Fréchette, Guillaume R., 2020. On the determinants of cooperation in infinitely repeated games: a survey. *J. Econ. Lit.* (forthcoming).
- Bowen, T. Renee, Zahran, Zaki, 2012. On dynamic compromise. *Games and Economic Behavior* 76, 391–419.
- Browne, E.C., Franklin, M., 1973. Aspects of coalition payoffs in European parliamentary democracies. *Am. Polit. Sci. Rev.* 67 (2), 453–469.
- Browne, E.C., Frendreis, J.P., 1980. Allocating coalition payoffs by conventional norm: assessment of the evidence for cabinet coalition situations. *Am. J. Polit. Sci.* 24 (4), 753–768.
- Dahm, Matthias, Glazer, Amihai, 2015. A carrot and stick approach to agenda setting. *J. Econ. Behav. Organ.* 116, 465–480.
- Dal Bó, Pedro, 2005. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *Am. Econ. Rev.* 95, 1591–1604.
- Diermeier, Daniel, Fong, Pohan, 2011. Legislative bargaining with reconsideration. *Q. J. Econ.* 126 (2), 947–985.
- Duggan, John, Kalandrakis, Tasos, 2012. Dynamic legislative policy making. *J. Econ. Theory* 147 (5), 1653–1688.
- Eguia, Jon X., Shepsle, Kenneth A., 2015. Legislative bargaining with endogenous rules. *J. Polit.* 77 (4), 1079–1088.
- Eraslan, Hulya K.K., 2002. Uniqueness of stationary equilibrium payoffs in the Baron-Ferejohn model. *J. Econ. Theory* 103, 11–30.
- Fréchette, Guillaume R., Yuksel, Sevgi, 2013. Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination (working paper).
- Fréchette, Guillaume R., Kagel, John H., Morelli, Massimo, 2005. Gamson's law versus non-cooperative bargaining theory. *GEB* 51, 365–390.
- Gamson, William A., 1961. A theory of coalition formation. *Am. Sociol. Rev.* 26, 373–382.
- Gomes, Armando, Jehiel, Philippe, 2005. Dynamic processes of social and economic interactions: on the persistence of inefficiencies. *J. Polit. Econ.* 113 (3), 626–667.
- Jackson, Matthew O., Moselle, Boaz, 2002. Coalition and party formation in a legislative voting game. *J. Econ. Theory* 103 (1), 49–87.
- Janssen, Maarten C.W., 2006. On the strategic use of focal points in bargaining situations. *J. Econ. Psychol.* 27, 622–634.
- Kalandrakis, Anastassios, 2004. A three-player dynamic majoritarian bargaining game. *J. Econ. Theory* 116, 294–322.
- Kalandrakis, Tasos, 2010. Minimum winning coalitions with endogenous status quo. *Int. J. Game Theory* 39, 617–643.
- Maskin, Eric, Tirole, Jean, 2001. Markov perfect equilibrium: I. Observable actions. *J. Econ. Theory* 100, 191–219.
- McKelvey, Richard D., Riezman, Raymond, 1992. Seniority in legislatures. *Am. Polit. Sci. Rev.* 86 (4), 951–965.
- Merlo, Antonio, Wilson, Charles A., 1995. A stochastic model of sequential bargaining with complete information. *Econometrica* 63 (2), 371–399.
- Nunnari, Salvatore, 2018. Dynamic Legislative Bargaining With Veto Power (CEPR Discussion Paper, DP12938).
- Nunnari, Salvatore, 2019. Veto Power in Standing Committees: An Experimental Study. Working Papers 649. IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University.
- Palfrey, Thomas R., 2016. Experiments in political economy. In: Kagel, John, Roth, Alvin (Eds.), *The Handbook of Experimental Economics*. Vol. 2. Princeton University Press, pp. 347–434.
- Roth, Alvin, 2005. Towards a focal-point theory of bargaining. In: Roth, Alvin (Ed.), *Game-Theoretic Models of Bargaining*. Cambridge University Press, pp. 259–268.
- Salz, Tobias, Vespa, Emanuel, 2016. Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation (working paper).
- Sethi, Ravideep, Verriest, Ewout, 2016. The Power of the Agenda Setter: A Dynamic Legislative Bargaining Model (working paper).
- Snyder, James M., Ting, Michael M., Ansolabehere, Stephen, 2005. Legislative bargaining under weighted voting. *Am. Econ. Rev.* 95 (4), 981–1004.
- Vespa, Emanuel, 2016. An Experimental Investigation of Strategies in the Dynamic Common Pool Game (working paper).
- Vespa, Emanuel, Wilson, Alistair J., 2016. Communication with multiple senders: an experiment. *Quant. Econ.* 7 (1), 1–36.
- Yaari, M.E., Bar-Hillel, M., 1984. On dividing justly. *Soc. Choice Welf.* 1, 1–24.
- Young, H. Peyton, 1993. An evolutionary model of bargaining. *J. Econ. Theory* 59, 145–168.
- Young, H. Peyton, 1996. The economics of convention. *J. Econ. Perspect.* 10, 105–122.