

PERSISTENCE OF POWER: REPEATED MULTILATERAL BARGAINING

MARINA AGRANOV, CHRISTOPHER COTTON, AND CHLOE TERGIMAN

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

Date: **Current version: November 2016.**

Key words and phrases. dynamic bargaining, political institutions, laboratory experiment

JEL: C78, D02, C92.

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

1. INTRODUCTION

Many bargaining situations involve repeated interactions. Budget committees with the same members meet every year to bargain over allocation of scarce resources. Standing committees within legislatures determine policies and regulatory measures repeatedly. The dynamic nature of these interactions creates links between bargaining decisions across cycles. This necessitates taking into account inter-temporal considerations, as opposed to analyzing each decision in isolation.

There are two categories of legislative dynamic bargaining decisions: the first one includes decisions that remain in effect in the absence of new legislation, such as social security and welfare expenditures; the second one includes decisions that require approval by voting in every cycle. This is the case for most budgetary decisions. The existing literature on dynamic bargaining has primarily focused on the first type of decisions, in which policies adopted in the current cycle become the status quo policy for the next cycle (see our review in Section 2). We take a different approach, and study, both theoretically and experimentally, environments that do not have an evolving status quo, but instead link budget cycles via the identity of the agenda setter.¹

We consider a committee with n members, which meets repeatedly and in each cycle allocates a given budget between its members using a majority voting rule. Within each cycle, the committee uses a standard Baron-Ferejohn bargaining protocol, according to which an agenda setter proposes a budget division and all committee members vote on it. If a majority of legislators support the proposal, the budget passes and the cycle ends. Otherwise, a new agenda setter is randomly selected from the committee to make a new proposal. The process continues until a proposed allocation passes, at which point the cycle ends. Our framework is repeated, in that after one budget cycle ends, a new one begins. The cycles are linked via institutional rules that determine the identity of the agenda setter based on what has occurred in the previous cycle.

Specifically, we focus on two rules that allow successful agenda setters to hold onto power, and which approximate alternative practices used by real-world committees. Under the first rule, an agenda setter who successfully passes a proposal in cycle c automatically serves as the agenda setter in cycle $c + 1$. This is akin to a framework in which a vote of confidence is attached to each policy vote. The vote of confidence procedure is a common feature of many parliamentary democracies, where failure to pass a budget proposal leads to dissolution of the current government due to "loss of supply" and to the formation of a new government. We refer to this model as the *Vote of Confidence* model. Under the second rule, the committee votes on whether to keep or replace the

¹That the agenda setter would persist over budget cycles is a realistic assumption: for example, in the US, the chairman of the House appropriations committee stays in power 5.5 years on average.

current agenda setter following the passage of each proposal. That is, the successful agenda setter in cycle c must maintain the support of a majority of committee members to serve as agenda setter in cycle $c + 1$. This is akin to the US Congress and other congressional systems where an agenda setter (e.g., Speaker of the House, committee chair) maintains power as long as a legislative majority supports him/her remaining in power. Passing a proposal is not enough to stay in power.² We refer to the second model as the *Majority Support* model. We compare both models with a *Baseline* model, in which a new legislator is randomly selected to serve as agenda setter in each cycle, independently of past success: even a successful agenda setter with the support of a majority of other legislators cannot hold onto power. The baseline model is consistent with most other theoretical models of repeated bargaining in that it assumes the random assignment of agenda-setting power in each cycle.

Our environment permits a multiplicity of subgame perfect equilibria, as is common in dynamic bargaining frameworks. Under reasonable parameters, there exists a broad range of allocations that can be supported in equilibrium. This significantly limits the predictive power of the models. The literature has responded to this issue by focusing on either Stationary, history independent strategies, whether in the form of Stationary Subgame Perfect Equilibria in stationary or cyclical environments such as in the original paper by Baron and Ferejohn (1989), or Markov Perfect Equilibria in environments with an evolving status quo (e.g. Kalandrakis 2004, Duggan and Kalandrakis 2012, Kalandrakis 2010, Anesi 2010, Baron and Bowen 2016).³ These solution concepts both assume that strategies are independent of history. In our framework, the two concepts are equivalent, except for some technical differences that do not affect the predictions of the refinements.⁴ Assuming stationary/history-independent strategies allows us to limit our attention to a much smaller subset of the potential equilibria. Typically, and in our games in particular, the refinement results in a unique equilibrium outcome. The benefits of the refinement are substantial, as it takes us from a situation in which the models have essentially no predictive power (any outcome, essentially, is consistent with subgame perfection), to a situation in which the model predicts a single outcome.

The literature has used this refinement to make predictions about bargaining outcomes in a variety of situations. For example, McKelvey and Riezman (1992) use a repeated legislative bargaining game with a stationarity refinement to consider why legislatures may

²Although we assume that such a vote takes place after the passage of each proposal, it is equivalent to a situation in which a formal vote only takes place in situations where informal communication has revealed that legislative leadership does not have the support of a majority of legislators.

³In both environments, the papers typically focus on the subset of symmetric equilibria.

⁴See the discussion about when analyses should use stationary subgame perfection versus markov perfection in Maskin and Tirole (2001).

endogenously adopt systems that award seniority. Baron (1996) uses such a framework to consider the implementation of entitlement programs. Battaglini and Coate (2007, 2008) incorporate a related framework into a model of fiscal policy, in which the legislature raises revenues via either distortionary taxes or by borrowing and uses these revenues to finance a national public good and district-specific transfers.

Following the literature, we first characterize the Stationary Subgame Perfect Equilibria in our three models. The predictions are as follows. The Vote of Confidence is the only model that predicts the persistence of agenda-setter power. This is because on the equilibrium path, proposed allocations pass without delay, which automatically means that the first agenda setter holds onto power indefinitely. In contrast, the Majority Support and the Baseline models predict frequent turnover of the agenda setter. More generally, the stationary equilibrium outcomes in the Majority Support model are identical to the ones in the Baseline model. This is because the history independence requirement prevents legislators from rewarding generous agenda setters by keeping them in power, which leads to agenda setters always being removed from power, and the game degenerating to a situation in which a new agenda setter is randomly selected each cycle. This suggests that tying the vote of confidence to the budget allocation helps solidify agenda-setter power and leads to the majority of legislators being worse off compared to a situation in which the decision to keep or replace the legislator is not tied to the policy vote.⁵

While stationary equilibrium predictions might be interesting per se, how much we learn from this analysis depends largely on whether the focus on history-independent strategies is appropriate, meaning that it is reasonable given real-world behavior. A number of papers make the argument that it is. For example, Baron and Kalai (1993) argue that a stationary subgame perfect equilibrium is the simplest and therefore most likely subgame perfect equilibrium. More recently Agranov and Tergiman (2014) and Baranski and Kagel (2015) show that the stationary equilibrium outcome often arises in one-cycle multilateral bargaining experiments. However, arguments in support of stationary equilibria are associated with one-time bargaining, where the interactions between players ends after they reach an agreement. The restriction to history-independent strategies has also been extensively utilized in repeated bargaining environments. Therefore, whether this restriction is appropriate in repeated bargaining environments is still an open and very important question (see our review in Section 2 where we describe the debate in the literature), given its prevalence in the theoretical domain. Ultimately, this

⁵In the long run, the agenda setter is better off and the other legislators are worse off as the agenda setter holds onto power indefinitely. In the short run (within a cycle), the agenda setter is worse off under the Vote of Confidence rule as he/she shares a larger portion of the budget with other legislators.

is an empirical question as data are needed to judge the circumstances under which its use is justified.

In light of the above argument, we proceed by conducting a series of laboratory experiments, in which we stage the three models described above. In the experiments, we observe both bargaining outcomes and the bargaining process that leads to these outcomes. Beyond testing equilibrium predictions, our experiment is the first one that documents the evolution of coalitions and bargaining outcomes in a dynamic multilateral framework without an evolving status quo and compares observed behavior under different legislative systems.

Our results show that the stationary/history-independent equilibrium refinement does a poor job at predicting behavior and outcomes observed in our repeated environment. This is in contrast to results documented in one-time bargaining situations. In particular, we observe high persistence of power in both the Vote of Confidence and Majority Support models. In addition, in all three games, coalitions tend to be stable across bargaining cycles in terms of size, identity of coalition partners, and shares allocated to coalition partners. Furthermore, while theory predicts that only minimum winning coalitions should be formed, we observe a substantial fraction of grand coalitions that include all members. Finally, we document a high frequency of partnerships and punishments that carry from one cycle to the next.

The failure of the stationary refinement to accommodate the observed outcomes is true regardless of the legislative structure that we consider. This raises a natural question of how to reconcile observed outcomes with theoretical predictions. To reach this goal, we consider asymmetric strategies, risk aversion and fairness concerns, all while keeping the stationary refinement intact. None of these extensions generate the behavior observed in our experiments. This calls into question the validity of restricting attention to history-independent strategies in repeated bargaining games, as is the typical practice in the literature.

We then document the empirical patterns of strategies used by our experimental subjects. Our data clearly show that in all three games, subjects use strategies that involve punishments, reciprocity and history dependence - all properties that contradict the stationarity refinement. Based on this evidence we argue that the disconnect between the theory and the experiments is because the theory ignores the fact that in repeated environments players may condition their current actions on their own and others' past behavior. We proceed by showing theoretically that players need not observe or remember the entire history of the game in order to generate equilibria which are consistent with the data; they just need to remember which player, if any, most recently deviated

from expected behavior. Remembering this one piece of information is enough to generate equilibria which feature both high persistence of power in the Vote of Confidence and in the Majority Support games and stable coalitions, which are either minimum winning or grand, in all three games.

The remainder of the paper is organized as follows. Section 2 discusses the relevant literature. Section 3 presents our environment and the predictions of the stationary subgame-perfect equilibrium. Section 4 outlines the experimental design. Section 5 presents the results of the experiments. Section 6 revisits the theory, considering a number of theoretical extensions, i.e., asymmetric stationary SPE, social preferences and risk aversion, in order to try to reconcile theory and observed outcomes. That section then proceeds to document the empirical patterns of strategies used by our experimental subjects and evaluate theoretically the minimal limited history dependence required to support equilibrium outcomes observed in our experiments. Section 7 comments on the features of political institutions studied in this paper and concludes.

2. RELATED LITERATURE

In the last few decades, legislative bargaining has received a great deal of attention both in the theoretical and experimental domains. The seminal paper of Baron and Ferejohn (1989) studies the legislative bargaining process, when a committee is charged with one-time allocation of a budget using a majority voting rule. Many articles extend Baron and Ferejohn's theoretical analysis to study effects of various political institutions (e.g. Baron 1996, Banks and Duggan 2000, Jackson and Moselle 2002, Merlo and Wilson 1995, Banks and Duggan 2006, Bowen and Zahran 2012, Eraslan 2002, Snyder, Ting and Ansolabehere 2005). Given that the current paper deals with *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) in order to limit the feasible proposals available to other agenda setters in the future. Kalandrakis (2004, 2010), and Duggan and Kalandrakis (2012) generalize Baron's results, allowing for multidimensional policy spaces.⁶ 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

⁶See 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.

of legislative bargaining in which an agenda setter has monopoly power over proposals, the status quo is determined by the most-recently implemented proposal, and the legislative process repeats with positive probability. Finally, some papers endogenize legislative rules within the context of a repeated bargaining game. McKelvey and Riezman (1992) and Eguia and Shepsle (2015) consider dynamic legislative bargaining when legislators must stand for reelection after each period, and shows that a legislature will endogenously adopt rules that reward more senior legislators.

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

We take a starkly different approach from the existing literature, with our analysis focusing on the rules governing how agenda-setter power changes over time, and how this affects equilibrium outcomes. To isolate the effects of our assumptions concerning agenda-setter power, we simplify the other aspects of the problem by assuming a stable, exogenous degenerate status quo policy. We are aware of no other article that focuses on the agenda-setter-authority aspect of the dynamic environment.

The experimental literature has followed the steps of theoretical research focusing first on one-cycle bargaining games (see the survey by Palfrey 2016) and recently moving on to dynamic bargaining experiments. Some of the dynamic bargaining papers focus on the evolution of status-quo policy in dynamic models of pure redistribution and consider a setting in which the status-quo policy is determined by the distribution of resources agreed upon in the previous bargaining cycle. Battaglini and Palfrey (2012) is the first paper that experimentally investigates such environment. Baron, Bowen and Nunnari (2016) extend this setup by considering effects of various communication channels available to committee members. Nunnari (2016) and Sethi and Verriest (2016) incorporate veto power and analyzes consequences of its presence. Other papers study dynamic models of public good accumulation. Battaglini, Nunnari and Palfrey (2012, 2016) consider an infinite-horizon legislative bargaining model of durable public good provision, in which status-quo policy distributes the available budget among committee members in equal private shares. Agranov et al. (2016) look at a two-period version of a similar game and decompose the inefficiency embedded in the legislative bargaining solution relative to the efficient solution into its static and dynamic components. None of these models consider the linkage of budget cycles via the agenda-setter identity.

Finally, our paper contributes to the newly emerging experimental literature that evaluates the relevance of Markov perfection in various dynamic settings, which absent such refinement generally feature large sets of subgame perfect equilibria. This literature, however, is still small and, more importantly, far from reaching a consensus regarding this question. Several papers document that comparative static predictions implied by Markov perfect equilibria organize experimental data well. Battaglini, Nunnari and Palfrey (2012, 2016) make this point in the dynamic legislative bargaining game with durable public goods. Salz and Vespa (2016) study an infinite-horizon entry/exit game of oligopolistic competition and reach the same conclusion. Vespa (2016) studies a dynamic common pool game and finds that modal behavior of subjects is consistent with Markov perfection. Finally, Agranov and Elliott (2016) investigate decentralized bargaining games with heterogeneous trade opportunities and irreversible exit and also conclude that market outcomes match MPE predictions across treatments. On the other hand, there is a large experimental literature on infinite-horizon prisoner's dilemma games, which documents that a majority of subjects use efficient, history-dependent strategies contrary to the MPE prediction of always defecting (see survey by Bó and Fréchette forthcoming). Vespa and Wilson (2016) study an extension of an infinitely-repeated prisoner's dilemma game with two states and provide evidence that suggests when the selection of MPE is more likely to occur.⁷ This debate on the validity of the stationary refinement justifies using it as a first benchmark against which to test our data.

3. REPEATED MULTILATERAL BARGAINING

Our analysis considers three models of multilateral bargaining, each of which is a repeated version of the classic closed-rule multilateral bargaining game of Baron and Ferejohn (1989). In Baron and Ferejohn, the game ends as soon as players reach an agreement on how to divide a resource between themselves. In contrast, our games do not end when the players reach an agreement. Rather, players enter a new cycle of bargaining where they again must choose how to divide a new budget.

The common features of our three models include the following. There are $n \geq 3$ identical legislative districts, each represented by a legislator. Within each cycle ($c = 1, 2, \dots$) of bargaining, the n -member legislature is responsible for splitting a budget of total size 1 between the n districts. Denote by a_i^c the share of the total budget that is allocated to player i in cycle c , where $\mathbf{a}^c = (a_1^c, \dots, a_n^c)$. An allocation \mathbf{a}^c is feasible if

⁷The authors construct an index that captures attractiveness of efficient outcomes relative to MPE outcomes, and show that this index tracks when subjects are ready to abandon MPE strategies in favor of history-dependent strategies in order to reach 'better' outcomes.

$0 \leq a_i^c \leq 1$ for each i , and $\sum_i a_i^c \leq 1$. We denote the allocation outcome across all cycles by $\mathbf{a} = (\mathbf{a}^c)_{c=1}^\infty$.

Each cycle c lasts until the legislators agree on an allocation a^c . A cycle can be divided into stages ($s = 1, 2, \dots$), with each stage involving an agenda setter (AS) proposing an allocation, and the other $n - 1$ players observing the proposal and then simultaneously casting votes in favor of or against it. The identity of the time $t = (c, s)$ AS is denoted AS^t , and his/her proposal in period t is given by $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$, where x_i^t is the proposed share for player i . A proposal is feasible if it corresponds to a feasible allocation.

If m members in addition to the AS vote in favor of the proposal, the proposal is implemented, and the game moves on to the next cycle. That is, if the proposal at time $t = (c, s)$ passes, then the cycle c allocation is $\mathbf{a}^c = \mathbf{x}^t$, and the game advances to $t' = (c + 1, 1)$. The identity of the new AS in the first stage of the new cycle depends on which version of the game is being played; we go into detail regarding the transition of proposer-power between cycles in Section 3.1. We assume that $m \in \{1, \dots, n - 2\}$, which assures that the AS needs the support of at least one other player to pass a proposal, and that unanimous support is not necessary.

If fewer than m other players vote in favor of the proposal at time $t = (c, s)$, then the proposal \mathbf{x}^t fails. Following the failure of a proposal, the game advances to the next stage within the same cycle, $t' = (c, s + 1)$. At this point, the identity of $AS^{t'}$ is randomly determined, with each of the n legislators having an equal probability of being selected as the next AS. The new AS then makes a proposal. The process repeats itself until a proposal passes. Given this, each cycle lasts at least one stage, and can potentially last infinitely many stages.

The discount factor $\delta \in (0, 1)$ applies between stages within a budget cycle. The discount factor $\gamma \in (0, \delta)$ applies between budget cycles. We assume that within-cycle delays do not make future cycles less valuable, which means that γ may be interpreted as either the between-cycle discount factor, 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 delay in passing one year's budget typically does not impose a delay upon the following year's bargaining.

3.1. Three models of repeated bargaining. Our three models of repeated bargaining differ only in the rule governing the transition of AS power from one cycle to the next. In each of the games, the AS is randomly selected in the first stage of the first cycle, as

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

well as after the failure of any proposal. The difference between the games comes only when a proposal passes and the game advances to the next cycle.

– **Baseline model**

Following the passage of a proposal in any cycle c , there is a new selection of an AS, with each of the n legislators having a $1/n$ probability of being randomly selected to serve as AS in the first stage of cycle $c + 1$. This is the standard assumption in the literature on repeated legislative bargaining.⁹

– **Vote of Confidence model**

Following the passage of a proposal in any cycle c , the most recent AS (the one who proposed the successful allocation) automatically serves as the AS in the first stage of the next cycle $c + 1$. There is an up or down *vote of confidence* attached to each budget proposal: the legislature votes once at each stage and the adoption of a proposal implies retaining the same AS in the first stage of the next cycle.

– **Majority Support model**

Following the passage of a proposal in any cycle c , the legislature holds another vote that determines whether to retain the most recent AS (the one who proposed the successful allocation) to serve as the AS in the first stage of the next cycle $c + 1$. If at least m other (non-AS) legislators vote to retain the AS, then the AS serves as the first-stage AS during cycle $c + 1$. If fewer than m others vote in favor of the AS, then there is a new random selection of a legislator to serve as AS in stage 1 of cycle $c + 1$, with each of the n legislators having an equal probability of being selected. Here, an AS who successfully passes a proposal must maintain support of m other members to retain power.

The three models of legislative bargaining described above, while obviously stylized, capture essential characteristics of different bargaining procedures used in legislative politics around the world. The Vote of Confidence model includes features common to parliamentary democracies in which the failure to pass a major piece of legislation (budget allocations included) is considered a vote of "no confidence," and leads to new elections. In more than 30 countries with parliamentary systems, a budget bill is seen as a default pseudo confidence vote.¹⁰ The Majority Support model approximates the US and other congressional systems where an AS (e.g. Speaker of the House, or committee chair) maintains power as long as a legislative majority supports him/her remaining in

⁹See 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).

¹⁰Traditionally, in the Westminster system, the defeat of a budget bill is followed by the resignation of the government or dissolution of parliament, since a government that cannot pass budget bill has no money to continue functioning. This event is termed 'loss of supply'. See https://en.wikipedia.org/wiki/Westminster_system#Current_countries.

power. Our Majority Support game incorporates this idea into the legislative bargaining framework by assuming that after the legislature passes a budget allocation, it holds a second vote in which the legislature votes on whether to keep or replace the current AS. Although we assume that such a vote takes place after the passage of each proposal, it is equivalent to a situation in which a formal vote only takes place in situations where informal communication has revealed that the legislative leadership does not have the support of a majority of legislators. In other words, passing a proposal is not enough to stay in power. Finally, in the Baseline model, the link between bargaining cycles is completely removed, as there is no institutional procedure that allows an AS to keep his/her power between cycles. This model features the automatic re-shuffling of AS power and will be mostly used in the analysis as a control environment.

3.2. Subgame Perfect Equilibrium. In our environment, when players care enough about the future, any feasible allocation can be maintained as part of a SPE.

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

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

3.3. Stationary Equilibria: Subgame and Markov Perfection. In each of our games, many SPE exist. To deal with the multiplicity of equilibria in multilateral bargaining games, the literature typically follows Baron and Ferejohn (1989) and focuses on stationary refinements of SPE, whether focusing on *Stationary Subgame Perfect Equilibria* (SSPE) as Baron and Ferejohn did, or *Stationary Markov Perfect equilibrium* (MPE). In our environments, the two concepts are equivalent. In the remainder of this section, we derive the SSPE of our three games noting that the same results could be obtained by characterizing instead the MPE in each of our three games.

The SSPE concept requires that players choose the same strategies in every structurally equivalent subgame.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, a SSPE requires that each player follows the same proposal strategy every time he/she serves as AS, and has the same voting strategy every time he/she does not serve as AS. Equilibrium strategies cannot condition on the history of play, although a player's vote in favor of or against a proposal will depend on his/her proposed share of the allocation. In what follows, we make two additional assumptions that are common in this literature: First, we initially focus on symmetric SSPE implying that the strategies are symmetric across all players. We consider asymmetric SSPE in Section 6.1. Second, we restrict attention to equilibria strategies that are not weakly dominated, implying that players who are indifferent between voting in favor of or against a proposal (or sitting AS) will choose the alternative that they would choose if they were certain to cast the deciding vote.¹²

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

3.3.1. *Stationary equilibria in the Baseline model.* From an ex ante perspective, in any period of play (be it a new stage within a cycle or the first stage in a new cycle), each player is

¹¹Two 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).

¹²This 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.

selected as AS with probability $1/n$ and is included in the MWC with probability m/n . Thus, just like in the one-cycle bargaining game in which the game ends after the first allocation passes, the expected payoff from rejecting the proposal is $\frac{1}{n}$. When a non-AS supports proposal x^t , he/she expects to get

$$x_i^t + \frac{\gamma}{1-\gamma} \cdot \left(\frac{1}{n}(1-m\bar{a}) + \frac{m}{n}\bar{a} \right) = x_i^t + \frac{\gamma}{1-\gamma} \cdot \frac{1}{n} \quad (1)$$

If he/she votes against x^t , he/she expects to get

$$\frac{1}{n}\delta + \frac{\gamma}{1-\gamma} \cdot \frac{1}{n} \quad (2)$$

A player will vote in favor of a proposal when (1) is greater than (2). This means that a proposal x^t passes if and only if at least m players (besides the AS) get shares at least as large as $\frac{\delta}{n}$. In the symmetric SSPE, the AS randomly selects m other legislators and allocates the smallest acceptable share of $\frac{\delta}{n}$ to each of them and keeps the remaining $1 - m\frac{\delta}{n}$ for herself. That is, the SSPE of the Baseline model is fully characterized by the equilibrium allocation provided to each MWC member:¹³

$$a_{\text{Baseline}}^{\text{SSPE}} = \frac{\delta}{n}$$

3.3.2. *Stationary equilibria in the Vote of Confidence model.* Here, the players recognize that a successful AS automatically holds onto power. This changes the incentives players have to vote in favor of another AS's proposals. Player i who votes in favor of the proposal x^t expects a net present value of current and future payoffs equal to

$$x_i^t + \frac{\gamma}{1-\gamma} \cdot \frac{m}{n-1}\bar{a} \quad (3)$$

If he/she votes against the proposal, he/she receives a net present value of expected payoff that is again given by Eq. (2). Setting (3) equal to (2) and solving for \bar{a} gives us the equilibrium allocation to a MWC member:

$$a_{\text{Vote of Confidence}}^{\text{SSPE}} = \frac{1}{n} \frac{(n-1)(\gamma + \delta - \gamma\delta)}{(n-1)(1-\gamma) + \gamma m}$$

Further, the MWC member share in the Vote of Confidence model, $a_{\text{Vote of Confidence}}^{\text{SSPE}}$, is strictly increasing in how intensely players care about the future (i.e. in both δ and γ), and converging to δ/n as players stop caring about future cycles (i.e. as $\gamma \rightarrow 0$).

¹³Anesi and Deidmann (2015), Baron and Bowen (2016) and Anesi and Duggan (2016) highlight the multiplicity of stationary Markov perfect equilibria in multilateral bargaining environments with an evolving status quo that equals the most recently passed policy. With a fixed status quo, the stationary subgame perfect equilibrium is unique.

3.3.3. *Stationary equilibria in the Majority Support model.* Our third game is complicated by an additional vote that takes place after the passage of a proposal. When a proposal passes, the players then cast a second vote to determine whether or not the current AS continues to serve as AS in the first stage of the next cycle. 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 themselves will be selected as AS in the next cycle.

Because of this, under SSPE, the Majority Support model is equivalent to the Baseline model, with a new AS being randomly selected at the start of each cycle.¹⁴ Thus,

$$a_{\text{Majority Support}}^{\text{SSPE}} = \frac{\delta}{n}$$

3.3.4. *Testable predictions of stationary equilibrium.* The model generates a number of testable comparative static predictions between the three models of repeated bargaining in terms of distribution of resources within a cycle and the evolution of coalitions across cycles. These predictions are summarized in Proposition 2.

Proposition 2. *In the unique symmetric stationary SPE of the three models of repeated bargaining described in Section 3.1, at each time $t = (c, s)$*

- (i) *Proposals pass without delay.*
- (ii) *There is low persistence of AS power in the Baseline and Majority Support models, while there is high persistence of AS power in the Vote of Confidence model.*
- (iii) *The proposal provides a positive allocation to exactly m other players.*
- (iv) *The identity of the MWC partners are randomly determined at each time, even if the AS remains the same.*
- (v) *The shares of MWC partners are the same in the Baseline and Majority Support models and are higher in the Vote of Confidence model.*
- (vi) *The expected payoff of the AS is higher in the Vote of Confidence model than in the two other models.*

Under the symmetric stationary equilibrium refinement, outcomes are identical in environments where an AS is randomly selected at the start of each cycle (Baseline model) and where a successful AS maintains power with majority support (Majority Support model). This is because the stationary refinement does not allow legislators to vote in support of legislators based upon past generosity. Compared to the Baseline and

¹⁴One can verify that the legislators do prefer to vote to replace the AS in this situation. The expected benefit of being the AS is $1 - m\bar{a}$ each stage, and the expected benefit of not being the AS is $\bar{a}m/(n - 1)$ each stage. Thus, the non-ASs vote to replace the AS since $1 - m\bar{a} > \bar{a}m/(n - 1)$.

Majority Support models, tying a vote of confidence to each budget proposal (as in the Vote of Confidence model) leads to low equilibrium turnover of AS power. In this case, players have a greater incentive to vote against any proposal with the hopes of becoming the AS themselves in the next stage. As a result, an AS must offer a larger share to other legislators to secure their support. In the short run, this can make the AS worse off under the Vote of Confidence model, as the AS shares a larger portion of the budget with other legislators each cycle in equilibrium. In the long run, however, the AS is better off and other legislators worse off from the vote of confidence rule, as the AS holds onto power for the long term.

4. EXPERIMENTAL DESIGN

All our experiments were conducted at the Center for Experimental Social Sciences at New York University using Multistage software.¹⁵ Subjects were recruited from the general undergraduate population and each subject participated in only one session. A total of 156 subjects participated in our experimental sessions.

We ran three treatments: (*Baseline*, *Vote of Confidence* and *Majority Support*). Those correspond to the three models of repeated bargaining described in Section 3.1. In what follows we describe the details of the experimental protocol used in each treatment and refer the reader to the Online Appendix for the full instructions received by subjects.

In each session subjects played the repeated bargaining game eight times. We refer to each of those as a match.¹⁶ In each cycle, subjects had 200 tokens to divide. At the end of a session one match was selected at random for payment, and earnings in that match were converted into USD (10 tokens = \$1). These earnings, together with the participation fee are what a subject earned in this experiment. The sessions lasted about two hours and on average subjects earned \$20, including a participation fee of \$7.

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

In all three treatments, at the beginning of the first stage of the first cycle of a match, one committee member was randomly chosen to serve as the AS. The AS was asked to propose how to distribute the 200 tokens between the three committee members and this

¹⁵The Multistage package is available for download at <http://software.ssel.caltech.edu/>.

¹⁶In one session of the *Vote of Confidence* treatment, subjects played only 7 matches, as the experiment lasted longer than expected. We observe no other significant differences in behavior in this session relative to the other sessions.

proposal was presented to all group members for a vote. If the proposal was accepted by a majority of votes (at least two out of three members), then the cycle ended. With probability 70%, the group moved on to the second cycle of the match and with probability 30% the match was terminated. If, however, the proposal was rejected, then the group remained in the first cycle and the second bargaining stage started. At the beginning of the second bargaining stage one member was randomly selected to serve as the new AS. The AS was asked to submit a budget proposal, which was then voted on by all committee members. However, the rejection of a proposal triggered a 20% reduction in the budget (that is, the within-cycle discount factor is $\delta = 0.8$). In other words, while in the first stage of every cycle the committee had 200 tokens, in the second stage, the available budget was reduced to 160 tokens, and if a committee reaches the third stage it was further reduced to 128 tokens, etc. This procedure continued until a majority of committee members voted in favor of the budget proposed by the AS.

In the *Baseline* treatment, each cycle of a game is identical to the first one: the AS in the first stage of every cycle is chosen randomly among the three committee members. In the *Vote of Confidence* treatment, the AS that successfully passed a proposal in cycle c remained the AS in cycle $c + 1$. In the *Majority Support* treatment, following the successful passage of a proposed budget, the committee held a second vote in which all members voted on whether to retain the current AS for the next cycle. To retain power, the current AS needed to obtain a majority of votes in the second vote. If the current AS was voted out, the AS in the next cycle was randomly chosen. *The difference in how the AS changes from one cycle to the next is the only difference between treatments.*

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

Finally, we implemented the Random Block Termination design developed and tested by Frechette and Yuksel (2013), in which subjects receive feedback about the termination of a match in blocks of cycles. In our implementation, each block consisted of four cycles. Within each block, subjects receive no feedback about whether the match has ended or not and they make choices which will be payoff-relevant conditional on a match actually reaching this cycle. At the end of a block, subjects learned whether the match ended within that block and, if so, in which cycle. If the match was not terminated, subjects

proceeded to play a new block of four cycles. Subjects were paid only for the cycles that occurred before match was actually terminated. The advantage of using the Block design is that it allows for the collection of long strings of data (at least four cycles) even with a relatively small discount factor of $\gamma = 0.7$. This small discount factor was chosen in order to obtain distinct enough predictions of the stationary subgame perfect equilibrium between treatments.

Table 1 summarizes the details of all our experimental sessions.

TABLE 1. Experimental Design

Treatment	# of Sessions	# of Subjects	# of Matches	Mean # of Cycles/Match
<i>Baseline</i>	3 sessions	(18,18,15)	(8,8,8)	(4,7,6)
<i>Vote of Confidence</i>	3 sessions	(15,18,18)	(7,8,8)	(6,6,5)
<i>Majority Support</i>	3 sessions	(21,15,18)	(8,8,8)	(4,6,6)

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

$$a_{\text{Baseline}}^{\text{SSPE}} = a_{\text{Majority Support}}^{\text{SSPE}} = 53 \text{ tokens} < a_{\text{Vote of Confidence}}^{\text{SSPE}} = 96 \text{ tokens}$$

5. EXPERIMENTAL RESULTS

We begin the empirical analysis by comparing bargaining outcomes within each cycle across our three treatments. We then shift our attention to the dynamic outcomes and analyze how coalitions evolve across bargaining cycles. At the end of this section, we compare patterns of behavior with our initial theoretical analysis, and consider how well the symmetric stationary equilibrium refinement predicts behavior in our environments.

5.1. Approach to 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 the final four matches of each session, we are able to consider the behavior of our experimental subjects after they have familiarized themselves with the game. Focusing on the first four cycles, which all groups certainly play, allows us to have a balanced dataset with identical amounts of experience within a match across all treatments.

We classify proposals in terms of the number of members that receive non-trivial shares and refer to these as *coalition types*. A *non-trivial share* is defined as share that is larger than 5 tokens. A proposal in which only one group member receives more than 5 tokens is a *dictator coalition*. A proposal in which exactly two members receive

non-trivial shares is a *minimum winning coalition*. Finally, a proposal, in which all three members receive non-trivial shares is a *grand coalition*. We call members with non-trivial shares *coalition partners*. Finally, we refer to some proposals as *equal split* proposals. Equal split proposals are defined as the ones in which the difference between the shares of any two coalition partners is at most 5 tokens.

To compare the outcomes between two treatments we use regression analysis. Specifically, to compare outcomes between two treatments (whether the fraction of a particular coalition type or the share received by the AS), we run random-effects GLS 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 cluster standard errors by sessions, recognizing the interdependencies between observations that come from the same session, since subjects are randomly rematched between matches.

5.2. Bargaining Outcomes within a Cycle. In all three treatments, almost all proposals pass in the first stage of each cycle. This is the case in 96.3%, 94.8% and 99.7% of experienced cycles in the Baseline, Vote of Confidence and Majority Support treatments, respectively. In the remainder of this subsection we concentrate on those proposals that passed without delay.

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

	Baseline	Vote of Confidence	Majority Support
<i>Coalition size</i>			
Dictator (1-person coalition)	0.0%	0.0%	0.3%
MWC (2-person coalition)	27.9%	48.1%	57.8%
Grand (3-person coalition)	72.1%	51.9%	41.8%
<i>Allocations within coalitions</i>			
Equal split MWC	80.8%	50.4%	56.0%
Equal split Grand coalition	83.1%	58.9%	65.0%

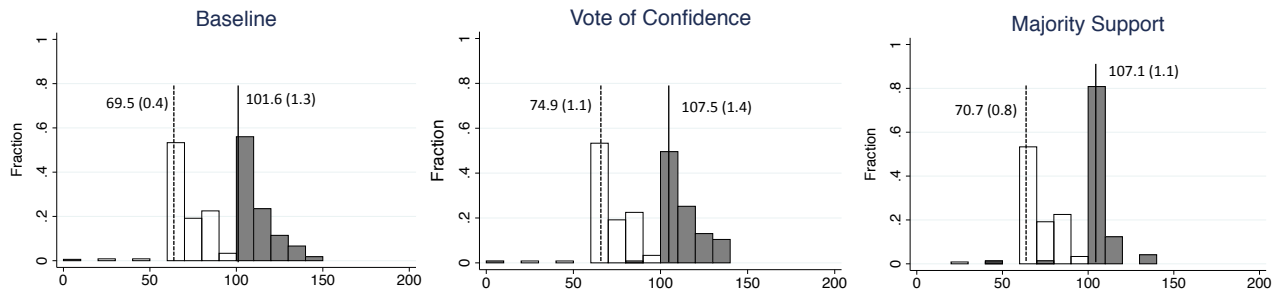
Notes: In the last two rows, we report the fraction of equal splits conditional on the coalition size being a two-person coalition (fourth row) and a three-person coalition (fifth row).

In Table 2 we present the distribution of proposals that passed without delay in terms of coalition size. Both two-person and three-person coalitions are common in all three treatments, with the largest fraction of grand coalitions (over 70%) observed in the Baseline treatment. Regression analysis confirms that the proportion of three-person coalitions is higher in the Baseline than in the two other treatments ($p = 0.041$ for Baseline versus Vote of confidence and $p = 0.088$ for Baseline versus Majority Support). At the same time, there is no significant difference between the fraction of grand coalitions

passed in the Vote of Confidence and Majority Support treatments ($p = 0.896$): under both of these institutional rules, roughly one half of all passed proposals are grand coalitions with the remaining half being two-person coalitions. We also note that, conditional on coalition type, allocations across treatments are, in their majority, equal splits.

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

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



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

In all three treatments, ASs that form grand coalitions appropriate a smaller share of resources than those that form minimum winning coalitions ($p < 0.05$ within each treatment). Comparing across treatments, we find that the shares of ASs in the Baseline treatment are significantly lower than in the Majority Support and Vote of Confidence treatments.^{17,18}

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

¹⁷We obtain $p = 0.085$ for Baseline versus Majority Support among three-person coalitions, and $p < 0.05$ for all other pairwise comparisons.

¹⁸Statistically, the Vote of Confidence AS shares are higher than those in Majority support in three-person coalitions ($p = 0.025$), though the magnitude of the difference (107.1 versus 107.4 is quite small). All other comparisons of AS shares are not statistically different ($p > 0.10$).

5.3.1. *Persistence of Power.* As we have seen in the previous section, ASs appropriate larger shares of the budget compared to other committee members. Thus, holding the AS seat has obvious benefits within each cycle. We now turn to how AS power evolves across cycles and then ask whether institutional rules that might temper the persistence of such power over time do that in practice. This question is only meaningful for the Vote of Confidence and Majority Support treatments since by design the Baseline treatment prevents persistence of AS power.¹⁹

Our data indicate that both the Vote of Confidence and Majority Support treatments feature high persistence of power in that ASs keep their seat for two consecutive cycles in more than in 94% of the cases. Regression analysis confirms that the likelihood of observing the same AS in two consecutive cycles is the same between the Vote of Confidence and Majority Support treatments ($p = 0.254$) and both are significantly larger than the likelihood of such an event in the Baseline treatment, which happens in 32.9% of all cases (this is very close to the theoretical likelihood of 33.3%).²⁰

Another related way of documenting the persistence of power is to observe how often the *same* AS served in all four cycles of the first block. In the Baseline treatment this event is rare and happens only 7% of the time. In the other two treatments this characterizes the overwhelming majority of cases: 84% of Vote of Confidence committees and 92% of Majority Support committees operate with the same sitting AS in power in all four cycles of the first block. These last two fractions are not statistically different ($p = 0.336$).

The number of cycles in which the same AS holds onto power directly affects his/her long-run payoff in the game as measured by the total payoff that the AS first selected obtains over the course of an entire block of four cycles. Our data indicate that this long-run payoff increases from the Baseline to the Vote of Confidence and Majority Support treatment, with average payoffs of 273 tokens, 318 tokens and 335 tokens, respectively. Statistical analysis confirms that first-cycle ASs obtain lower long-run payoffs in the Baseline treatment than in the Vote of Confidence and Majority Support games with no difference between the latter two treatments ($p = 0.087$ for Baseline versus Vote of Confidence, $p = 0.033$ for Baseline versus Majority Support, and $p = 0.618$ for Vote of Confidence versus Majority Support).

5.3.2. *Evolution of Coalitions.* We begin this analysis by considering the frequency with which coalition types change. Table 3 shows the likelihood of proposed coalition types conditional on the type of coalition that passed in the previous cycle. As evident from the transition matrix, we observe a high level of persistence of coalition types between

¹⁹Despite the inability of ASs to keep power between bargaining cycles in the Baseline game, in Section ?? we discuss and document the existence of partnerships used by committee members in order to circumvent the inherent turnover that exists in this treatment.

²⁰ $p < 0.01$ for both pairwise comparisons.

cycles in all three treatments: in 87% or more cases, the next cycle proposal has the same coalition size as the one passed in the previous cycle. Interestingly, the degree of persistence is very similar across the three legislative environments. This is surprising given that in the Baseline treatment, ASs are very likely to be changing across cycles.

TABLE 3. Transition of coalition types across cycles

	Cycle $c + 1$					
	Baseline		Vote of Confidence		Majority Support	
	MWC	Grand	MWC	Grand	MWC	Grand
Cycle c						
MWC	0.87	0.12	0.89	0.09	0.94	0.06
Grand	0.11	0.89	0.09	0.91	0.07	0.93

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

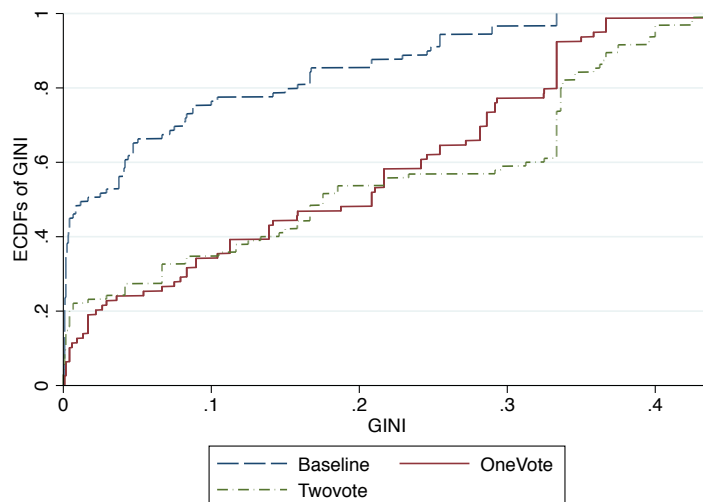
Next, we consider the persistence of coalition members across cycles. To do this, we focus on the persistence of the minimum winning coalition partner in all instances where the AS was the same for two consecutive cycles.²¹ Our data show that when a non-proposer is invited into a minimum winning coalition in one cycle, the probability that he/she will be re-invited into a minimum winning coalition in the following cycle is 78.6%, 72.9% and 89.7% in the Baseline, Vote of Confidence and Majority Support treatments. A series of two-sided tests of probability show that these percentages are significantly different than 50%, which means that proposers who are forming minimum winning coalitions are not choosing their coalition partners randomly.²² That is, minimum winning coalitions tend to be stable across cycles. In addition, our data indicate that the shares of those coalition partners stay the same across cycles in 91%, 79% and 84% of the cases in the Baseline, Vote of Confidence and Majority Support treatments, respectively. Thus, we conclude that not only are coalitions stable in terms of the identity of coalition members, but, in addition, when that is the case, the shares given to the coalition partners also are largely constant. In other words, ASs seek stability.

5.3.3. *Long-run Inequality.* We conclude our data analysis by documenting the long-run inequality in group members' payoffs induced by three institutional rules we consider in our paper. As before, the long-run payoffs are measured by the members' total payoffs over the course of an entire block of four cycles. Figure 2 presents the empirical CDFs of the Gini coefficient for each committee.

²¹This is the only non-trivial case, since in grand coalitions all members are coalition partners by definition.

²²In the Baseline treatment we obtain $p = 0.033$, while in the remaining two treatments $p < 0.01$.

FIGURE 2. Empirical CDFs of GINI coefficients, by treatment



As evident from Figure 2, the Vote of Confidence and Majority Support treatments result in similar and statistically indistinguishable distributions of long-run payoffs among committee members (a Ranksum test yields a p value of 0.268). The Baseline treatment, however, features a much more equal distribution of long-run payoffs ($p < 0.01$ for both pairwise Ranksum tests). Two forces contribute to this result. First, the frequent turnover of AS, which is a built-in feature of the Baseline treatment, increases the chances that different members serve as ASs in different cycles. Consequently, committee members "take turns" in obtaining the higher shares appropriated by the AS. Second, in the Baseline treatment grand coalitions are more common than in the other treatments. These grand coalitions naturally produce a more equal distribution of resources within a committee compared with two-person coalitions.

5.4. Summary of Results and Symmetric SSPE. In this section we summarize the results of our experiments and compare them with the predictions of the stationary equilibrium refinements. We focus on symmetric SSPE, which coincide with the predictions of Markov Perfect equilibria in our three games (see discussion in Section 3.3 and summary in Proposition 2).

Bargaining outcomes within a cycle are efficient (no delays) in all three bargaining games, consistent with theoretical predictions. However, while the symmetric SSPE predicts that all passed proposals should feature two-person minimum winning coalitions, our data show a different pattern. We observe that both minimum winning and grand coalitions are common in all three settings. The highest fraction of grand coalitions is in the Baseline treatment, in which over 70% of all passed proposals include non-trivial shares to all three group members. Finally, conditional on coalition size, at least 50%

of passed proposals feature an equal division of the surplus between coalition partners, irrespective of the treatment. In particular, ASs share the surplus equally with their coalition partner in 81% of Baseline minimum winning coalitions, while they do so in 50% and in 56% of such proposals in the Vote of Confidence and Majority Support treatments. This is in sharp contrast with the symmetric SSPE prediction, according to which ASs appropriate strictly higher shares of resources than their coalition partner.

Turning to the examination of bargaining outcomes across cycles, we note that our data reveal high persistence of power in the Vote of Confidence and in Majority Support treatments, while such persistence of power is theoretically predicted only for the Vote of Confidence game. In all three games, the observed coalitions are stable across cycles in terms of their size, the identify of coalition partners and the shares of coalition partners, which is also at odds with predictions. Further, long-run payoffs of ASs are higher in the Majority Support compared with the Baseline treatment, while the theory stipulates that these payoffs should be the same. Finally, we document that among our three treatments, the Baseline treatment features the lowest inequality in terms of long-run payoffs between committee members.

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

6. RECONSIDERING THE THEORY

In this section, we seek to reconcile observed bargaining outcomes with theoretical predictions. Rather than rationalizing each treatment in isolation, we seek a unified approach that can account for behavior observed in all three legislative environments. To recapitulate, the three main features of our experimental data are (1) the high persistence of power in the Vote of Confidence and Majority Support games, (2) the existence of stable two-person coalitions observed in all three games, and (3) the fact that grand coalitions are observed in all three games.

We first consider whether the behavior we observe in our experiment can be accounted for by relaxing a number of the simplifying assumptions imposed in the original theory, while maintaining the focus on history-independent strategies. To this end, we separately allow for asymmetric strategies within the SSPE, and incorporate non-linear preferences, considering risk aversion and fairness concerns. As we will see, none of these extensions lead to the behavior observed in the experiments.

We then turn our attention to the assumption that strategies are *history-independent*. We start by documenting empirical patterns of strategies used by our experimental subjects. This analysis reveals that our subjects clearly use history-dependent strategies that feature punishments and rewards. Among other findings, in all three games, we observe that the ASs who fail to pass proposals are very likely to be excluded from future coalitions, i.e., punished. Based on this evidence we argue that it is the assumption that players ignore past behavior that leads to the disconnect between the theory and experiments. We then proceed to show theoretically that players need not observe (or remember) the entire history of the game in order to generate equilibria that are consistent with the theory; they just need to remember which player, if any, most-recently deviated from expected behavior. The observed behavior in the experiments is consistent with subgame perfect equilibria in which actions are history-independent along the equilibrium path of play, and the most-recent person who deviates from the equilibrium is remembered to have deviated and believes that they may be (credibly) punished by being excluded by the other players in future allocations. Simply allowing players to remember one piece of information about past behavior (who deviated most-recently) is enough to generate theoretical equilibria consistent with the theory.

This section concludes with a discussion of the multiplicity of equilibria under the relaxed equilibrium refinement. Allowing players to condition actions on past behavior leads to equilibrium predictions that are consistent with experimental observation. Multiplicity of equilibria is again a concern, just as it was with the unrefined subgame perfect equilibrium concept. Given this, we argue that we should not expect the standard stationary equilibrium refinements—which require history independence—to generate predictions that are consistent with observed behavior. Rather, we should consider which allocations are likely to serve as focal points and guide player behavior. The literature offers guidance along these lines.

6.1. Asymmetric SSPE. Here we relax the symmetry requirement used in Section 3.3. Rather than require that the players' strategies are independent of other player's identities, in this section we allow for stationary strategies that treat other players asymmetrically, in particular when the AS each period chooses which player to include in his/her MWC. We focus on pure strategy equilibria in this environment.

Let $x^j = \{x_1^j, \dots, x_n^j\}$ denote player j 's equilibrium proposal strategy, which he/she makes in every period that he/she serves as AS. Let \bar{a}_i^j denote player i 's voting strategy, where i votes for a proposal made by player j in any period where j serves as AS if and only if $x_i^t \geq \bar{a}_i^j$. Consider the following stationary, but asymmetric, strategy profile:

- Each player j chooses a MWC K_j made up of m other players. Player j 's proposal gives $x_i^j = X$ for each $i \in K_j$, and $x_i^j = 0$ for each $i \notin \{K_j, j\}$.

- Each player i is included in the MWC of exactly m other players.
- Each player i votes in favor of proposal x^t if and only if $x_i^t \geq X$ when $i \in K_{AS^t}$, and if and only if $x_i^t \geq Y$ when $i \notin K_{AS^t}$.

For each game, we determine the values of X and Y such that the above constitutes an asymmetric SSPE. We present here the main intuition of asymmetric SSPE and refer the reader to the Online Appendix for a detailed derivation.

In the Baseline game, the switch from symmetric to asymmetric strategies does not change the incentives players have to accept or reject proposals in each cycle, and, thus,

$$\text{Baseline: } X = Y = \frac{\delta}{n}$$

In the Vote of Confidence game, however, the switch to asymmetric strategies changes things, as players who are included in the current MWC expect to continue being included as long as the AS stays in power (which is forever on the equilibrium path). Thus, coalition partners are willing to accept lower shares than those required in support of a proposal in the symmetric SSPE. Specifically,

$$\text{Vote of Confidence: } X = \frac{\delta + \gamma - \delta\gamma}{n} \text{ and } Y = \frac{1}{n} \frac{\delta + \gamma - \delta\gamma}{1 - \gamma}$$

Finally, to determine the asymmetric SSPE in the Majority Support game, one needs to consider two possibilities: either the member of K_j will reelect j when he/she is AS in each cycle, or they will not. Those not in K_j have no incentive to reelect player j as AS. The restriction that $\gamma \in (0, 1)$ rules out the possibility of asymmetric SSPE with high persistence of AS power, which leaves the low persistence equilibrium as the only viable option. Suppose that we are in an equilibrium with low persistence of AS power. Thus, players vote against the AS in every cycle. In this case, the incentives to vote for or against a given proposal are the same as in the Baseline game, as there is a new AS draw in each cycle. As such,

$$\text{Majority Support: } X = Y = \frac{\delta}{n}$$

It remains to verify that members of K_j prefer to draw a new AS the next cycle, rather than reelect the current one, which is true given that $\delta \leq 1$. Thus, in the asymmetric SSPE of the Majority Support game, the equilibrium resembles that of the Baseline game with low persistence of AS power, and, as a consequence, non-stable coalitions.

From this analysis, we see that the asymmetric SSPE do no better, and worse in the case of the Vote of Confidence game, than the symmetric SSPE in explaining the observed behavior in the experiments.

6.2. SSPE with concave utilities. Another natural avenue for extending the results of SSPE is to consider outcomes that emerge when bargainers are risk-averse. Specifically, assume that the overall utility of member i is given by

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

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

In the symmetric SSPE of all three games, a player i that votes against a proposed allocation obtains an expected net present value of

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

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

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

in the Baseline and Majority Support games, and gets

$$u(x_i^t) + \frac{\gamma}{1 - \gamma} \cdot \frac{m}{n - 1} u(a^{\text{Game}})$$

in the Vote of Confidence game.

In the Online Appendix, we solve for the symmetric SSPE allocations when players have identical CRRA utility functions:

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

In equilibrium, the share of the MWC partner is strictly decreasing in α , which means that introducing risk-aversion leads to a more-unequal split of resources in favor of the AS compared with the risk-neutral case. This pattern is opposite of what we observe in our data. Intuitively, as α increases, coalition partners are willing to accept a lower share rather than reject such a proposal and risk not being included in the next MWC, since MWC partners are chosen randomly. Moreover, there is no symmetric SSPE in which there is persistence of power in the Majority Support game. It turns out that combining risk-averse bargainers with asymmetric SSPE does not help either, as one cannot obtain a high persistence of power in the Majority Support game with asymmetric stationary strategies.

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

6.3. Fairness concerns. Another possibility is that players care about fairness. To allow for this, we incorporate other-regarding preferences in line with the model of Fehr and Schmidt (1999). Player i 's utility from allocation a in any given period is:

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

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

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

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

Thus, fairness concerns may explain some, but not all, of the observed allocations during the experiments. The main feature that the SSPE coupled with fairness concerns cannot explain is the equal splits among coalition partners within minimum winning coalitions, a behavior which is common in all our three games as shown in Section 5.2.

6.4. Limited history dependence.

6.4.1. Empirical evidence of history-dependent strategies. The dynamic nature of our bargaining environment creates potential links between cycles and allows subjects to form and execute history-dependent strategies. In this section we investigate whether our

subjects use history-dependent strategies, and, if so, what the main features of these strategies are.

We start by documenting strategies that include punishment. In the Baseline treatment, if a previously excluded member becomes the AS, he/she excludes the previous AS from a minimum winning coalition 81.8% of the time. A two-sided test of proportions shows that this fraction is significantly different from 50% ($p < 0.01$).²³ Given the very high persistence of power observed in the remaining two treatments, in order to obtain a reasonable number of observations related to punishment behavior, we look at all cases in which there was turnover in AS power (no longer restricting the data to the last four matches). For the cases where we observe such turnover, the AS who failed to pass the proposal in the previous cycle is excluded from the new AS's MWC in 100% of the cases in the Vote of Confidence treatment and in 80% of the cases in the Majority Support treatment.^{24,25}

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

In summary, in all three treatments, history-dependent strategies cannot be overlooked.

6.4.2. Theoretical treatment of limited history dependence. Here, we relax the focus on history independence that is assumed as part of the stationary equilibrium refinements. In the context of a SSPE or MPE, this section allows for an expanded definition of the state space, where the state may now include payoff-irrelevant information about the history of the game. When the state space includes the entire history of past behavior, the requirement that players choose stationary strategies, playing the same within-period strategy in all identical states, has no impact, since no two nodes of the game tree will have the same state. Putting no restriction on which aspects of the game's history a

²³In almost 75% of cases this new AS proposes a minimum winning coalition.

²⁴We have very few observations for both these treatments as the overwhelming majority of proposals pass. We have 3 observations in the Vote of Confidence treatment and 5 in the Majority Support treatment.

²⁵More generally, proposers that submit proposals that fail tend to be excluded by the next proposer. This is the case in over 60% of the cases in all treatments.

strategy can condition on permits overly complicated strategies with no obvious justification for considering. In response to these concerns, we allow players to condition their strategies each period on only a single aspect of the history of play: the identity of the player who most-recently deviated from expected equilibrium behavior.

We allow the players to condition on this specific aspect of game history; however, this is not the only way we could relax history independence. We chose to allow conditioning on the most-recent deviant because this is arguably the simplest piece of information that allows for punishment strategies, but it is not the only piece of past information that would allow this. Its simplicity helps illustrate the fragility of stationarity, as allowing players to condition their strategies on this one piece of pertinent information about the history of the game is enough to ensure that any allocation is consistent with equilibrium. We formalize this result in Proposition 3. We refer to a SPE in which player strategies condition only on payoff-relevant information and the identity of the most-recent deviant as a "SPE with limited history dependence."

Proposition 3. *Consider any feasible allocation profile $\mathbf{a}^* = \{\mathbf{a}^{r*}\}_{r=1}^{\infty}$, such that for every r , $a_i^{r*} \in [0, 1]$ for each i and $\sum_i a_i^{r*} = 1$. As long as γ is sufficiently large, there exists a SPE with limited history-dependence of game $\Gamma \in \{\mathbf{Rand}, \mathbf{Auto}, \mathbf{Vote}\}$ that generates \mathbf{a}^* along the equilibrium path with probability 1. 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 set of strategies. The intuition behind this is as follows. Players are able to condition their strategies on the most-recent player to deviate from some given strategy, and this is enough to permit punishment strategies that exclude any player who deviates from the equilibrium strategies from future allocations. When the discount factor is sufficiently high, this threat of future exclusion is substantial enough to prevent players from deviating from the equilibrium strategies, and to ensure that the punishment strategies are credible.

We don't need players to remember a detailed history of play, or to formulate complex strategies in order for every allocation to be consistent with equilibrium. All we need is for players to be able to remember an aspect of the history of the game as simple as the identity of the most recent player to do something unexpected. When they can do this, any behavior can again be justified.

This result illustrates the fragility of the SSPE refinement. Combined with the fact that our experimental subjects typically do not behave in accordance with the history-independent stationary equilibria, we dismiss the standard refinements as having little predictive power in repeated multilateral bargaining games.

6.5. Justification for equal division within a coalition. Without being able to rely on the standard stationary equilibrium refinements to narrow down the set of potential outcomes, we again find ourselves in a situation in which any allocations can be justified as consistent with an equilibrium. Given this inconsistency, we look to the literature for an alternative equilibrium refinement that may be more consistent with the experimental evidence.

The literature on equilibrium selection in games provides guidance, with evidence that players tend to coordinate on equal or "fair" outcomes in games with multiple Pareto dominated equilibria (e.g. Yaari and Bar-Hillel 1984, Young 1993, 1996, Roth 2005, Janssen 2006). This suggests that equal divisions (among all players or among a subset), when they are associated with an equilibrium, may serve as focal points, and help facilitate coordination. This view is also consistent with empirical evidence concerning the division of resources in legislative decision-making. Gamson's Law highlights the empirical regularity that coalitions of legislators tend to divide resources (e.g. cabinet positions) between parties in proportion to each party's share of total votes within the coalition (Gamson 1961, Browne and Franklin 1973, Browne and Frensdreis 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).²⁶

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

Finally, we argue that equal division equilibrium in the repeated environment may be more simple and more intuitive for the players than the SSPE. Although the SSPE may involve the simplest dynamics, with players choosing the same actions regardless of past outcomes, the per period proposal requires players to engage in some degree of complex reasoning to estimate the asymmetric equilibrium allocation that will be offered each period. Equilibria involving equal division among a winning coalition, on the other hand, involve little complex reasoning, with the AS each period splitting the allocation equally with at least m coalition partners, who in turn vote in favor of an allocation (and vote in favor of the AS in the Majority Support game) as long as the AS doesn't

²⁶See Fréchet, Kagel and Morelli (2005) for a comparison of the predictions of Gamson's Law and the stationary Baron and Ferejohn (1989) bargaining outcomes in a one-cycle bargaining framework.

deviate from the equal division strategy. Even the punishment strategies played off the equilibrium path are intuitive, with players simply excluding anyone who deviated from the equilibrium strategy in the past. This suggests that Baron and Kalai (1993)'s claim that the SSPE is likely to serve as a focal point because of its simplicity may be less likely to apply in a repeated environment. Rather, we see SPE with equal division among a winning coalition (whether grand or minimum-winning) and the threat of exclusion as a potentially simpler equilibrium in a repeated environment.

7. DISCUSSION OF POLITICAL INSTITUTIONS AND CONCLUSIONS

The results have implications for the theoretical literature on multilateral bargaining, and for the real-world design of political institutions.

First, we contribute to the literature on equilibrium selection in bargaining games. Models of multilateral bargaining tend to have many subgame perfect equilibria, limiting their ability to predict outcomes. To deal with the issues surrounding the multiplicity of equilibria, the literature tends to focus on stationary refinements of subgame perfection, including SSPE and MPE. Such refinements limit attention to history-independent strategies, and tend to reduce the set of potential outcomes to a unique equilibrium. This potentially improves the predictive power of the bargaining models, but only to the extent that we believe that the refinements draw attention to the equilibrium that is actually played. A number of papers suggest that the refinement is reasonable in one-cycle bargaining environments, in which a committee must reach a decision, but only once (see for example Baron and Kalai (1993) and Agranov and Tergiman (2014)). In contrast, our analysis shows that the stationary equilibrium refinements applied so widely in the literature may be inappropriate in models of repeated bargaining.

Experimental evidence suggests that players in repeated bargaining games do condition on some past outcomes, and choose more equitable allocations than the standard theory predicts. The analysis finds no support for the standard stationary equilibrium refinements in repeated bargaining. We then explore alternative theoretical explanations for the observed behavior.

Second, this is the first paper to develop a model (and experiments) of repeated bargaining that focuses on the transition of agenda power between budget cycles. We formally consider rules that capture key features of parliamentary systems and congressional democracies. This gives insight into how legislative rules influence outcomes.

Our Vote of Confidence model includes features that are common to parliamentary democracies in which the failure to pass a budget is considered a vote of "no confidence,"

and leads to the formation of a new government.²⁷ Diermeier and Feddersen (1998) show that the vote of confidence procedure is sufficient to explain why parliamentary systems exhibit strong governing parties with majority members voting in line with their party leadership.²⁸ This is consistent with our own theoretical predictions. Theoretically, in our Vote of Confidence model, the AS should be able to extract far higher shares than in the other two models. Indeed, in order to change proposer, non-proposers face a high cost: that of voting the AS's proposal down and risk not being in the next winning coalition. Our Majority Support model removes the link between changing proposer and that cost. Indeed, in Majority Support, the passage of the bill is independent of whether the proposer remains in power. In other words, in this treatment, the proposer loses some of his/her power. In fact, the Majority Support model is isomorphic to the Baseline model.

Our experimental treatments offer additional insight. Empirically, there is little difference between the the Vote of Confidence and Majority Support treatments. Proposers in the Majority Support treatment hold on to power and ruling coalitions are very stable. In fact, our data in the Baseline treatment, where proposers are randomly selected in each cycle, is striking in that even though turnover is institutionalized, players establish partnerships in order to maintain coalition stability. In short, when it is feasible to keep the same proposer in place, even if it is costless to change him/her as in the Majority Support treatment, subjects tend to support coalition stability and vote to maintain the same power structure. When subjects have no control over who the next proposer will be (such as in the Baseline treatment), subjects labor to create alliances so that power is concentrated among a few players. Our data strongly imply that both cycle-by-cycle and long-term distribution of wealth are more equal in the Baseline treatment because it is impossible for subjects to guarantee the same coalition will remain in power throughout the game. Legislative institutions and rules often aim to strike a balance between the stability of governing coalitions, a "fair" distribution of resources, and the concentration of power.²⁹ Our paper makes a case for such principles.

²⁷In more than 30 countries with parliamentary systems, a budget bill is seen as a default pseudo confidence vote. See Huber (1996).

²⁸Tergiman (2015) shows that in the lab a Vote of Confidence procedure indeed increases proposer power. A famous example is that of John Major's government who in July 1993 submitted a bill to adopt the Maastricht Treaty. The bill was rejected, in part because members of his own party voted against it. John Major reintroduced it with a vote of confidence attached, and the bill passed.

²⁹In the US for example, executive power is restricted by term limits: the President can only serve for two terms, as is the case with governors. A number of states have similar rules for state legislators.

REFERENCES

- Agranov, Marina, and Chloe Tergiman.** 2014. "Communication in Multilateral Bargaining." *Journal of Public Economics*, 118: 75–85.
- Agranov, Marina, and Matt Elliott.** 2016. "Commitment and (In)Efficiency: a Bargaining Experiment." working paper.
- Agranov, Marina, Guillaume Fréchet, Thomas R. Palfrey, and Emanuel Vespa.** 2016. "Static and Dynamic Underinvestment: An Experimental Investigation." *Journal of Public Economics*, 143: 125–141.
- Andreoni, James, Deniz Aydn, Blake Barton, B. Douglas Bernheim, Deniz Aydn, Blake Barton, B. Douglas Bernheim, and Jeffrey Naecker.** 2016. "When Fair Isn't Fair: Sophisticated Time Inconsistency in Social Preferences." *working paper*.
- Anesi, Vincent.** 2010. "Noncooperative foundations of stable sets in voting games." *Games and Economic Behavior*, 70(2): 488–493.
- Anesi, Vincent, and Daniel J. Deidmann.** 2015. "Bartaining in Standing Committees with an Endogenous Default." *Review of Economic Studies*, 82: 825–867.
- Anesi, Vincent, and John Duggan.** 2016. "Existence and Indeterminacy of Markovian Equilibria in Dynamic Bargaining Games." working paper.
- Banks, Jeffrey S., and John Duggan.** 2000. "A Bargaining Model of Collective Choice." *American Political Science Review*, 94(1): 73–88.
- Banks, Jeffrey S., and John Duggan.** 2006. "A General Bargaining Model of Legislative Policy-Making." *Quarterly Journal of Political Science*, 1: 49–85.
- Baranski, Andrzej, and John H. Kagel.** 2015. "Communication in Legislative Bargaining." *Journal of the Economic Science Association*, 1(1): 59–71.
- Baron, David P.** 1996. "A Dynamic Theory of Collective Goods Programs." *American Political Science Review*, 90(2): 316–330.
- Baron, David P.** 1998. "Comparative Dynamics of Parliamentary Governments." *American Political Science Review*, 92(3): 593–609.
- Baron, David P., and John A. Ferejohn.** 1989. "Bargaining in Legislatures." *American Political Science Review*, 83(4): 1181–1206.
- Baron, David P., and Michael Herron.** 2003. "A Dynamic Model of Multidimensional Collective Choice." In *Computational Models in Political Economy.*, ed. K Kollman, J Miller and S Page. Cambridge, MA: MIT Press.
- Baron, David P., and T. Renee Bowen.** 2016. "Dynamic coalitions." working paper.
- Baron, David P., Renee Bowen, and Salvatore Nunnari.** 2016. "Dynamic Coalitions and Communication: Public versus Private Negotiations." working paper.
- Baron, Davin, and Ehud Kalai.** 1993. "The Simplest Equilibrium of a Majority-Rule Division Game." *Journal of Economic Theory*, 61: 290–301.

- Battaglini, Marco, and Stephen Coate.** 2007. "Inefficiency in Legislative Policymaking: A Dynamic Analysis." *American Economic Review*, 97(1): 118–149.
- Battaglini, Marco, and Stephen Coate.** 2008. "A Dynamic Theory of Public Spending, Taxation, and Debt." *American Economic Review*, 98(1): 201–236.
- Battaglini, Marco, and Thomas R. Palfrey.** 2012. "The Dynamics of Redistributive Politics." *Economic Theory*, 49(3): 739–777.
- Battaglini, Marco, Salvatore Nunnari, and Thomas R. Palfrey.** 2012. "Legislative Bargaining and the Dynamics of Public Investment." *American Political Science Review*, 106(2): 407–429.
- Battaglini, Marco, Salvatore Nunnari, and Thomas R. Palfrey.** 2016. "The Political Economy of Public Debt: A Laboratory Study." working paper.
- Bó, Pedro Dal, and Guillaume R. Fréchette.** forthcoming. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature*.
- Bowen, T. Renee, and Zaki Zahran.** 2012. "On Dynamic Compromise." *Games and Economic Behavior*, 76: 391–419.
- Browne, E. C., and J. P. Frendreis.** 1980. "Allocating coalition payoffs by conventional norm: Assessment of the evidence for cabinet coalition situations." *American Journal of Political Science*, 24(4): 753–768.
- Browne, E. C., and M. Franklin.** 1973. "Aspects of coalition payoffs in European parliamentary democracies." *American Political Science Review*, 67(2): 453–469.
- Dahm, Matthias, and Amihai Glazer.** 2015. "A carrot and stick approach to agenda setting." *Journal of Economic Behavior and Organization*, 116: 465–480.
- Diermeier, Daniel, and Pohan Fong.** 2011. "Legislative Bargaining with Reconsideration." *Quarterly Journal of Economics*, 126(2): 947–985.
- Diermeier, Daniel, and Timothy J. Feddersen.** 1998. "Cohesion in Legislatures and the Vote of Confidence Procedure." *American Political Science Review*, 92(3): 611–621.
- Duggan, John, and Tasos Kalandrakis.** 2012. "Dynamic Legislative Policy Making." *Journal of Economic Theory*, 147(5): 1653–1688.
- Eguia, Jon X., and Kenneth A. Shepsle.** 2015. "Legislative Bargaining with Endogenous Rules." *Journal of Politics*, 77(4): 1079–1088.
- Eraslan, Hulya K. K.** 2002. "Uniqueness of Stationary Equilibrium Payoffs in the Baron-Ferejohn Model." *Journal of Economic Theory*, 103: 11–30.
- Fehr, Ernst, and Klaus M. Schmidt.** 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics*, 114(3): 817–868.
- Fréchette, Guillaume R., and Sevgi Yuksel.** 2013. "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination." working paper.

- Fréchette, Guillaume R., John H. Kagel, and Massimo Morelli.** 2005. "Gamson's Law versus non-cooperative bargaining theory." *Games and Economic Behavior*, 51: 365–390.
- Gamson, William A.** 1961. "A Theory of Coalition Formation." *American Sociological Review*, 26: 373–382.
- Gomes, Armando, and Philippe Jehiel.** 2005. "Dynamic Processes of Social and Economic Interactions: On the Persistence of Inefficiencies." *Journal of Political Economy*, 113(3).
- Huber, John D.** 1996. "The Vote of Confidence in Parliamentary Democracies." *American Political Science Review*, 90(2): 269–282.
- Jackson, Matthew O., and Boaz Moselle.** 2002. "Coalition and Party Formation in a Legislative Voting Game." *Journal of Economic Theory*, 103(1): 49–87.
- Janssen, Maarten C.W.** 2006. "On the strategic use of focal points in bargaining situations." *Journal of Economic Psychology*, 27: 622–634.
- Kalandrakis, Anastassios.** 2004. "A Three-Player Dynamic Majoritarian Bargaining Game." *Journal of Economic Theory*, 116: 294–322.
- Kalandrakis, Tasos.** 2010. "Minimum Winning Coalitions With Endogenous Status Quo." *International Journal of Game Theory*, 39: 617–643.
- Maskin, Eric, and Jean Tirole.** 2001. "Markov Perfect Equilibrium: I. Observable Actions." *Journal of Economic Theory*, 100: 191–219.
- McKelvey, Richard D., and Raymond Riezman.** 1992. "Seniority in legislatures." *The American Political Science Review*, 86(4): 951–965.
- Merlo, Antonio, and Charles A. Wilson.** 1995. "A Stochastic Model of Sequential Bargaining with Complete Information." *Econometrica*, 35(2): 371–399.
- Nunnari, Salvatore.** 2016. "Dynamic Legislative Bargaining with Veto Power: Theory and Experiments." working paper.
- Palfrey, Thomas R.** 2016. "Experiments in Political Economy." 347–434.
- Roth, Alvin.** 2005. "Towards a focal-point theory of bargaining." 259–268.
- Salz, Tobias, and Emanuel Vespa.** 2016. "Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation." working paper.
- Sethi, Ravideep, and Ewout Verriest.** 2016. "The Power of the Agenda Setter: A Dynamic Legislative Bargaining Model." working paper.
- Snyder, James M., Michael M. Ting, and Stephen Ansolabehere.** 2005. "Legislative Bargaining under Weighted Voting." *American Economic Review*, 95(4): 981–1004.
- Tergiman, Chloe.** 2015. "Institution design and public good provision: an experimental study of the vote of confidence procedure." *Experimental Economics*, 18(4): 697–717.
- Vespa, Emanuel.** 2016. "An Experimental Investigation of Strategies in the Dynamic Common Pool Game." working paper.

- Vespa, Emanuel, and Alistair J. Wilson.** 2016. "Communication With Multiple Senders: An Experiment." *Quantitative Economics*, 7(1): 1–36.
- Yaari, M. E., and M. Bar-Hillel.** 1984. "On dividing justly." *Social Choice and Welfare*, 1: 1–24.
- Young, H. Peyton.** 1993. "An evolutionary model of bargaining." *Journal of Economic Theory*, 59: 145–168.
- Young, H. Peyton.** 1996. "The economics of convention." *Journal of Economic Perspectives*, 110: 105–122.