CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

Discussion Paper No. 2011-10

Vincent Anesi
and
Daniel J. Seidmann

September 2011

Bargaining over an Endogenous Agenda

CeDEx Discussion Paper Series
ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Sue Berry
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 15469
Fax: +44 (0) 115 95 14159
sue.berry@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers

Bargaining over an Endogenous Agenda

Vincent Anesi[*]
University of Nottingham

Daniel J. Seidmann[†]
University of Nottingham

September 20, 2011

Abstract

We present a model of bargaining in which a committee searches over the policy space, successively amending the default by voting over proposals. Bargaining ends when proposers are unable or unwilling to amend the existing default, which is then implemented. We characterize the policies which can be implemented from any initial default in a pure strategy stationary Markov perfect equilibrium for an interesting class of environments including multi-dimensional and infinite policy spaces. Minimum-winning coalitions may not form, and the set of equilibrium policies may be unaffected by a change in the set of proposers. The set of stable policies (which are implemented, once reached as default) forms a weakly stable set; and conversely, any weakly stable set is supported by some equilibrium. If the policy space is well ordered then the committee implements the ideal policy of the last proposer in a subset of a weakly stable set. However, this result does not generalize to other cases, allowing us to explore the effects of protocol manipulation. Variations in the quota and in the number of proposers may have surprising effects on the set of stable decisions. We also show that equilibria of our model are contemporaneous perfect ε-equilibria of a related model of repeated implementation with an evolving default; and that stable decisions in semi-Markovian equilibria form the largest consistent set.

JEL classification: C78, D71, D72.

Keywords: bargaining, evolving default, stable set.

1 Introduction

The task of a committee is to select a policy to implement from some policy space. As Compte and Jehiel (2010) note, committees in effect search over the policy space by endoge-

[*]Address: School of Economics, Room B18, The Sir Clive Granger Building, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom. Email: vincent.anesi@nottingham.ac.uk.

[†]Address: School of Economics, Room B34, The Sir Clive Granger Building, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom. Email: daniel.seidmann@nottingham.ac.uk.

nously drawing policies/proposals, and implement a proposal according to a stopping rule. Thus, in conventional noncooperative bargaining models such as Rubinstein (1982) and Baron and Ferejohn (1989), players propose in sequence, and a policy is implemented once a winning coalition of players agree to a proposal. In agenda voting models, by contrast, a default policy is pitted against alternative policies in a sequence of binary votes. This model is usually justified as a naturalistic representation of procedure in Congressional committees. However, it may be relevant elsewhere as a model in which the committee searches by taking a single policy (the current default) seriously at a time; and a policy is implemented when it can no longer be displaced by an alternative.

We analyze a model of bargaining which marries these two perspectives: the game starts with an initial default. Players have the opportunity to propose amendments to the default in a fixed sequence (the protocol). If a winning coalition of voters accepts the proposal then the default is amended: the committee takes a new policy seriously. The new default may then in turn be amended. A default is implemented when all of the proposers have failed to amend it: either because they have chosen not to propose an alternative or because their proposals have not secured sufficient support from voters. Payoffs in the game only depend on the policy implemented. The agenda is endogenous in two senses here: chosen proposals determine both the policies on the agenda and the order in which they are considered.

This model in fact describes the way that we (a committee of two with a unanimity quota) have written this paper: we have worked with a running draft (the default), which we have only changed when we agreed that a new version improved on the default; and we have only circulated the paper when we have been unable to find any revision which improves on the current version.

We use an algorithmic technique to characterize the policies which can be reached in a pure strategy stationary Markov equilibrium from any initial default and for any protocol. An equilibrium strategy combination defines an outcome function, which determines the policy implemented from any default, on or off the equilibrium path. We show that the range of this mapping — i.e. the image of the policy space (the set of possible initial defaults) — is a weakly stable set in a related simple game. We obtain the related simple game by restricting the set of winning coalitions to those which contain a proposer; and a *weakly* stable set of policies satisfies the same strict internal stability conditions as a (von Neumann-Morgenstern) stable set, but external stability is weakened to allow for weak social preference. Conversely, for any closed weakly stable set, we construct equilibria, the image of whose mapping is exactly that weakly stable set. These observations imply that the policies which can be implemented is the union of weakly stable sets.

Equilibria in our model have some striking properties: a winning coalition may amend a default to a policy which is implemented, leaving all coalition members worse off than at the initial default; the size principle — according to which only minimum-winning coalitions should form — may fail; and a player who does not propose may earn all of the surplus from agreement.

We exploit these results to consider how a committee chair can affect the policy reached from a given default by changing the protocol for a fixed set of proposers. Changing protocols may affect the policy reached from a given default. However, varying the order in which a given set of proposers move does not affect the set of winning coalitions or of weakly stable sets; so the image of an equilibrium is unchanged. Accordingly, fix an equilibrium whose image is a given weakly stable set. If the policy space is well ordered (no player is indifferent between any two policies) then a chair who proposes cannot improve on a protocol in which she proposes last; and the protocol does not affect the policy reached from any default if and only if all proposers top-rank the same policy out of those in the weakly stable set which are socially preferred to the default. These results do not generalize to games in which a player may be indifferent between policies where, as we demonstrate, the chair may be best off proposing first.

Our set theoretic approach also allows us to assess the effect of variations in the quota. According to a natural conjecture, at least as many initial defaults are implemented, the larger is the quota. Indeed, more initial defaults might be implemented because coalitions which could destabilize policies are no longer winning with a larger quota. We provide conditions for this conjecture to be true; but we also show that the conjecture may be false because changes in the set of winning coalitions have potentially conflicting effects on the internal and external stability conditions for a set of policies to be weakly stable.

We end the paper by extending our analysis in two directions:

According to our model, players only receive (undiscounted) payoffs when a policy is implemented. However, our model has essentially the same game tree as a model without a stopping rule in which either the current default or an agreed policy is implemented each round and becomes the new default; and players earn the net present value of the stream of utilities earned from the implemented policies. One might therefore conjecture that equilibria in our model are the limit of equilibria in the alternative model with repeated implementation as players become more patient. This conjecture is true if the policy space is finite and well ordered. Indeed, equilibrium strategy combinations in our model are then also equilibria of the related model when players are patient enough. More generally, we show that, for every $\varepsilon > 0$, an equilibrium strategy combination in our model is a contemporaneous perfect ε-equilibrium of the model with repeated implementation when players are patient enough.

Any weakly stable set is contained in the largest consistent set. We provide weaker conditions on the stationarity of strategies under which the image of any equilibrium mapping is a consistent set, and the union of stable policies is contained in the largest consistent set.

After reviewing the related literature in the next subsection, we present the model in Section 2. We characterize equilibria in Section 3, and explore how the policy implemented varies with the protocol and with the set of winning coalitions in Section 4. In Section 5, we provide micro-foundations for the largest consistent set, and construct contemporaneous perfect ε-equilibria in games with repeated implementation. We conclude in Section 6, and briefly discuss variants on our model with bargaining round the table; random protocols; non-singleton proposals; refinements; and mixed strategy equilibria. We relegate longer proofs to an Appendix.

Related literature

The literature contains various related models of bargaining with an evolving default in which a policy is only implemented once negotiations end:

In Bernheim et al (2006), the policy space is finite and well ordered. The default is amended over a finite number of rounds, and the default at the end of the last round is implemented. Any Condorcet winner of the original game is implemented if there are enough proposers or at least one proposer top ranks the Condorcet winner. Bernheim et al also show that the last proposer's ideal policy (her own project alone) is implemented in a pork barrel example without a Condorcet winner. We allow for an infinite number of rounds, but equilibria in our model with a well ordered policy space also exhibit the power of the last word *for a given weakly stable set*. If preferences are generated by the pork barrel example then there is a unique, singleton weakly stable set, which consists of a bare majority of projects and may exclude the last proposer's project; and this policy is implemented in our model. The analogy between our results relies on our use of backward induction arguments which, in Bernheim et al's model, start with the exogenously fixed last proposal in the tree. Our argument, by contrast, relies on our stopping rule: a default which is not amended by any proposer is implemented, ending the game. In further contrast to Bernheim et al, and to the rest of the literature surveyed below, we allow for an infinite policy space without requiring that it be well ordered (no indifference).

Harsanyi (1974) provides micro-foundations for stable sets by presenting a bargaining model in which a policy is only implemented when a default is not amended. Each equilibrium of this model supports a weakly stable set, as in our model. However, in contrast to

Bernheim et al (2006) and this paper, a chair selects coalitions which simultaneously propose policies, and her payoff depends on the number of times that the default is amended. Harsanyi's model therefore typically allows several policies to be implemented in equilibrium because players and the chair respectively only care about the implemented policy and the number of amendments.[1] By contrast, we are primarily interested in the policy implemented from a given initial default. Our approach yields much tighter predictions about the implemented policy, and also allows us to address issues of protocol manipulation. We compare Harsanyi's model with a variant on our model with a dynamic protocol in the Appendix.[2]

Harsanyi argues that stability does not adequately capture social dominance in non-simple games, where a policy might be indirectly but not directly dominated. Chwe (1994) picks up this theme, arguing that only policies outside the largest consistent set can be excluded when players are far-sighted. Chwe also sketches a view of committees akin to our interpretation, with the important difference that he treats bargaining itself cooperatively. Our results provide noncooperative foundations for the largest consistent set in simple games. The contrast to weakly stable sets turns on the stationarity of strategies, rather than on far-sightedness.

Our model is also related to Baron and Ferejohn's (1989) open rule game, where proposers can amend the existing default. In contrast to Bernheim et al (2006), this game can last indefinitely; but, in contrast to our model, the game only ends when a player proposes moving the previous question. The difference in stopping rules is crucial, as many of our results rely on backward induction arguments which do not apply to open rule bargaining.

Following Baron (1996), a recent literature has studied equilibria of games with repeated implementation (as described in the last subsection). The most closely related paper is Acemoglu et al (forthcoming), which essentially shares our game tree, but allows the set of winning coalitions to depend on the default.[3] Acemoglu et al prove existence when social preferences are acyclic (their Theorem 2).[4] We focus on characterization, rather than existence results, but explore a much larger class of policy spaces which includes non-acyclic social preferences. To see why we eschew existence results, consider games with a well ordered policy space. Any weakly stable set is then stable; so existence of an equilibrium in our model is equivalent to existence of a stable set in simple games. It is

[1] Consider, for example, cases in which several policies in a stable set socially dominate the initial default.

[2] In contrast to Harsanyi (1974) and this paper, Hortala-Vallve (forthcoming) studies play in a related model without a stable set.

[3] Acemoglu et al. allow for some exogenous proposals (and exploit this possibility in their proofs).

[4] See Duggan and Kalandrakis (2011) and Kalandrakis (2004) and (2010) for existence results in related models.

well known that stable sets may not exist, even if the policy space is finite (e.g. when there are Condorcet cycles), but stable sets have been characterized for some games with infinite policy spaces, such as three player divide the dollar games.[5] More general conditions for existence remain an open question. We sidestep this issue by characterizing those equilibria which exist; and this approach allows us to study a much wider class of policy spaces.

Given this approach, Anesi (2010) is related most closely to this paper. Anesi demonstrates that that any stable set is the absorbing set of some Markov perfect equilibria in a legislative bargaining game with a finite, well ordered policy space, random proposers and repeated implementation.[6] We extend Anesi (2010) in two respects. First, the model provides bargaining foundations for (weakly) stable sets in a larger class of environments, allowing for infinite policy spaces that are not well ordered. Second, we obtain a complete equivalence between the class of stable sets and the class of absorbing sets of Markov perfect equilibria when the policy space is finite and well ordered. Anesi only proved that the former is a subset of the latter, demonstrating by example that the legislature may choose policies outside stable sets.[7] The latter result turns on the supposition that proposers are selected randomly. We address the implications of random proposers in our model in the Conclusion.

Our assumption that the default can be amended recalls a literature (surveyed by Austen-Smith and Banks (2005)) in which players vote successively over a finite, well ordered agenda. (Our algorithmic approach highlights the similarities.) This literature has largely focused on successive elimination and amendment agendas, in which a default is implemented when it (respectively) beats the next contender and all subsequent contenders. Duggan (2006) is related most closely to our paper. He assumes that players first add policies to an amendment agenda according to some protocol, and the committee then votes over the agenda; so the agenda is endogenous in our sense. In contrast, our model integrates proposing and voting; a given policy may be repeatedly placed on the agenda, which need not be finite; and the default is implemented when it has not been amended.[8] We follow the literature by considering how a chair could manipulate the agenda - though in our model, the chair directly manipulates the protocol (the order in which proposers are recognized) because the agenda itself is endogenous.

[5] Ordeshook (1986) discusses these issues; Hirai (2009) provides some recent results. More generally, Lucas (1992) surveys the literature on stable sets.

[6] Acemoglu et al show (in an online Appendix) that equilibria in their model support the unique stable set, which coincides with the largest consistent set.

[7] Anesi (2006) obtained the equivalence between stable sets and absorbing sets of equilibrium processes of coalition formation (cf. Konishi and Ray (2003)) in a cooperative model of committee voting over a finite well ordered policy space. Our model provides noncooperative bargaining foundations for stable sets.

[8] Our model therefore integrates features of successive elimination and amendment agendas.

2 The model

We consider a finite committee consisting of $m \geq 1$ *proposers*, $M \equiv \{1, \ldots, m\}$, and $n \geq 2$ *voters*, $N \equiv \{1, \ldots, n\}$. The set of committee members, or *players*, is thus $C \equiv M \cup N$. A player may be both a proposer and a voter, but we also allow for the possibility that $M \cap N = \emptyset$. Our model therefore encompasses agencies (such as the EPA and the FDA) which consult stakeholders, the FOMC (whose meetings are attended by the nonvoting Federal Reserve Bank Presidents) and the Supreme Court (where opinions are notoriously drafted by clerks), as well as committees which restrict deliberation to voters, such as juries.

Let X be a compact metric space of policies, which may be finite or a subspace of finite-dimensional Euclidean space. The preferences of each player $i \in C$ on X are represented by a weak order $\succeq_i$. Let $\succ_i$ and $\sim_i$ denote the asymmetric and symmetric parts of $\succeq_i$, respectively. We will say that the policy space is well ordered if every player has a linear order over X. We assume that preferences are continuous. Specifically:

Assumption A0. *Continuous Preferences:* For all $i \in C$, and all $x \in X$, the upper and lower contour sets of x associated with $\succeq_i$ are closed.

The committee has to reach a collective choice from X, with initial default policy $x^0 \in X$. Decision making takes place as follows. Each of a (possibly) infinite number of discrete rounds, indexed by $t = 1, 2, \ldots$, starts in the shadow of an ongoing default policy x^{t-1}. For each possible default $x \in X$, there is a fixed protocol $\pi_x : \{1, \ldots, m_x\} \to M$, $m_x \in \mathbb{N}$, that determines the order in which the proposers (i.e. the players in M) are given the opportunity to propose policies to amend $x = x^{t-1}$. That is, when $x \in X$ is the current default, protocol π_x allows for m_x opportunities to amend x and, for each $k \in \{1, \ldots, m_x\}$, the kth opportunity is given to proposer $\pi_x(k) \in M$. Each proposer $i \in M$ has at least one opportunity to amend the default in every round: $\left|\pi_x^{-1}(i)\right| \geq 1$ for all $i \in M$ and all $x \in X$. We denote the collection of protocols by $\pi \equiv \{\pi_x\}_{x \in X}$.

The outcome of a vote depends on the set of winning coalitions of voters $\mathbf{W} \subseteq 2^N \setminus \{\emptyset\}$. Throughout, we make the following assumption:

Assumption A1. $\mathbf{W}$ is

(i) *monotonic:* $S \in \mathbf{W}$ and $N \supseteq S' \supseteq S$ implies $S' \in \mathbf{W}$; and

(ii) *proper:* $S \in \mathbf{W}$ implies $(N \setminus S) \notin \mathbf{W}$.

In words: (i) every superset of a winning coalition is winning, and (ii) a coalition and its complement cannot both be winning.

Bargaining is then represented as follows:[9]

1. If the kth proposer, $\pi_{x^{t-1}}(k)$, is given the opportunity to make a proposal, she proposes $y_k^t \in X$.

2. a) If $y_k^t \neq x^{t-1}$ then y_k^t is put to an immediate vote against x^{t-1}. Members of N sequentially vote 'yes' or 'no' (in an arbitrary order). If the set of players who voted 'yes' is an element of $\mathbf{W}$ then y_k^t is accepted; otherwise it is rejected and x^{t-1} remains the default.

b) If $y_k^t = x^{t-1}$ (i.e. the proposer 'passes') then there is no voting and x^{t-1} remains the default.

3. a) If $y_k^t \neq x^{t-1}$ is accepted then it displaces x^{t-1} as the default policy and the round ends.

b) If $y_k^t \neq x^{t-1}$ is rejected or if there is no voting because $y_k^t = x^{t-1}$ and $k < m_x$, then the game moves to step 2 with k increased by 1; if $k = m_{x^{t-1}}$, x^{t-1} is implemented and the game ends.

Players only care about the policy which is eventually implemented, rather than the route from the initial default to the implemented policy. When comparing two different paths, each player $i \in C$ thus prefers the one yielding the best final policy outcome with respect to $\succeq_i$. We assume that bargaining indefinitely makes all players worse off than if any policy is implemented after a finite number of rounds.[10] Let $\Gamma\left(\pi, x^0\right)$ be the bargaining game defined by this process.

Following the lead of the previous literature, our main focus will be on subgame perfect equilibria of $\Gamma(\pi, x^0)$ in which players use pure stationary Markov strategies. A strategy consists of two components, one specifying a player's choice when given the opportunity to propose, the other specifying a voter's choice after a proposal is made. In proposal stages, strategies only depend on the default and the identity of the remaining proposers in the current round; in voting stages, strategies only depend on the current default, the proposal just made, votes already cast, and the remaining proposers in the current round. Unless otherwise stated, we will refer to stationary Markov pure strategy equilibria as 'equilibria.'[11]

[9] A similar bargaining process is used in Acemoglu et al. (forthcoming).

[10] We make this assumption in order to ensure that the one-shot deviation principle applies even though the game is not continuous at infinity. We could dispense with this assumption in games with a finite policy space by supposing that payoffs are discounted by the number of rounds, and that players are patient enough: such games would be continuous at infinity.

[11] In Section 5, we will also consider strategies that are measurable with respect to other elements in the history of the game.

Our restriction to pure strategies precludes existence in some well known cases, such as the Condorcet Paradox; in other cases, there may be multiple equilibria. We will discuss the implications of allowing for mixed strategy equilibria and of refining our solution concept in the Conclusion.

Any stationary Markov strategy σ generates an outcome function f^σ, which assigns to every $x \in X$ and every $k \in \{1, \ldots, m_x\}$ the unique final outcome $f^\sigma(x, k)$ eventually implemented (given σ) when x is the ongoing default and the kth proposer is about to move (in any round t). Of particular interest is $f^\sigma\left(x^0, 1\right)$ as it gives the final policy outcome of the game from any initial default $x^0 \in X$ when players act according to σ. As we will often refer to it in what follows, we will sometimes abuse notation and write $f^\sigma\left(x^0\right)$ instead of $f^\sigma\left(x^0, 1\right)$. The characterization of this function for all possible equilibria of $\Gamma\left(\pi, x^0\right)$ is the subject matter of the next section.

3 Equilibrium characterization

3.1 Computation

There are two principal sorts of questions we want to address: the first concerns the determination of equilibrium behavior and policy outcomes from any initial default; the second concerns how institutional details affect the set of policy outcomes. We address the former in this section, and postpone the latter to Section 4.

First of all, we need to modify the collection of winning coalitions, $\mathbf{W}$, in order to obtain a collection of coalitions that better accounts for the distribution of power among committee members. Let $\mathcal{W} \equiv \{S \subseteq C : (S \cap N) \in \mathbf{W} \ \& \ (S \cap M) \neq \emptyset\}$. That is, a coalition S belongs to $\mathcal{W}$ if the voters in S constitute a winning coalition and S includes at least one proposer. Note that $\mathcal{W}$ inherits monotonicity and properness from $\mathbf{W}$.

We define two social preference relations, which we call strict and weak dominance relations respectively, as follows: for all $x, y \in X$,

$$xPy \Longleftrightarrow \exists S \in \mathcal{W} : x \succ_i y \ , \forall i \in S \ ,$$

$$xRy \Longleftrightarrow \exists S \in \mathcal{W} : x \succeq_i y \ , \forall i \in S \ .$$

A subset of policies $V \subseteq X$ is said to be P-internally stable if and only if it satisfies

(IS$_P$) $\forall x, y \in V$: $\neg(xPy)$.

Furthermore, Y is said to be R-externally stable if and only if it satisfies

(ES$_R$) $\forall x \in X \setminus V$, $\exists y \in V$: yRx.

We say that V is a *weakly stable set* if and only if it is both P-internally stable and R-externally stable. The collection of weakly stable sets is denoted by $\mathcal{V}$.

Weakly stable sets will play a central role in the analysis to follow. Before we proceed any further, it is therefore worth discussing some of their properties. First of all, a (von Neumann-Morgenstern) stable set is a weakly stable set which is P-externally stable (which is defined by replacing R with P in (ES_R)). Conversely, when the policy space X is well ordered (i.e., when all the $\succeq_i$'s are linear orders), $\mathcal{V}$ corresponds to the collection of stable sets. This is not true when X is not well ordered: there may be policy sets that are weakly stable but not stable, as the following example illustrates:

Example 3.1. Let $M = N = \{1, 2, 3\}$, $X = \{x, y, z\}$ and every pair of players is winning, with preference orderings $z \succ_1 x \succ_1 y$, $x \sim_2 y \succ_2 z$, and $y \succ_3 x \sim_3 z$. It is easy to confirm that yPz, and that $\{x, z\}$ is weakly stable, but is not stable. ($\{x, y\}$ is stable.)

$\square$

The predictive power of weakly stability, like stability, depends on other parameters of the model: there may be a unique and small weakly stable set (e.g. any Condorcet solution); there may be a unique but large weakly stable set (e.g. every division of the pie in two-player bargaining: see Example 3.2 below); there may be several weakly stable sets (e.g. in three-player divide the pie bargaining: cf. Ordeshook (1986) Ch 9.2); and no weakly stable set need exist (e.g. in the Condorcet Paradox example).

Finally, a weakly stable set may contain Pareto dominated policies; and the closure of a weakly stable set is itself weakly stable. To see the latter, suppose there is some $x \in \partial V$ that does not belong to V. Evidently, $V \cup \{x\}$ satisfies (ES_R). If $V \cup \{x\}$ does not satisfy (IS_P) then there exists $v \in V$ such that either vPx or xPv. In the former case, by continuity of individual preference relations, there exists a neighborhood of x, N_x, such that v P-dominates all the elements of N_x. As $x \in \partial V$, $N_x \cap V \neq \emptyset$ and, therefore, v P-dominates some members of V: a contradiction of (IS_P). In the latter case, by continuity of individual preference relations, there exists a neighborhood of x, N_x, such that any $x' \in N_x$ P-dominates v. As $x \in \partial V$, $N_x \cap V \neq \emptyset$ and, therefore, v is P-dominated by some member of V: a contradiction of (IS_P).

We now return to our main purpose in this section, which is to describe an algorithmic procedure capable of finding the set of possible equilibrium policy outcomes from any initial default $x^0 \in X$. We first need some more notation. For any binary relation Q on X, $x \in X$ and any subset $Y \subseteq X$, we use the notation $Q(x) \equiv \{y \in X : yQx\}$, $Q_Y(x) \equiv \{y \in Y : yQx\}$, and $\mathbf{M}(Q, Y) \equiv \{y \in Y : \forall y' \in Y \setminus \{y\}\ , y'Qy \text{ implies } yQy'\}$. The elements of the latter set will often be referred to as the Q-maximal policies in Y.

Our procedure starts with a weakly stable set $V \in \mathcal{V}$. It then constructs a tree $\mathfrak{T}^{\pi}(V, x)$ — whose nodes are elements of $V \cup \{x\}$ — as follows. The initial node of $\mathfrak{T}^{\pi}(V, x)$ is x. If $x \notin V$ then the successors of x in the tree are obtained in m_x steps $k = m_x, m_x - 1, \ldots, 1$:

- **$k = m_x$:** The set of immediate successors of x is

$$\mathfrak{s}^{\pi}_{m_x}(V, x) \equiv \bigcup_{Y \subseteq R_V(x)} \mathbf{M}\left(\succeq_{\pi_x(m_x)}, P_V(x) \cup \{x\} \cup Y\right) .$$

- **$1 \leq k \leq m_x - 1$:** For each $y_{k+1} \in \mathfrak{s}^{\pi}_{k+1}(V, x)$, the set of immediate successors of y_{k+1} is

$$\mathfrak{s}^{\pi}_{k}(V, y_{k+1}) \equiv \bigcup_{Y \subseteq R_V(y_{k+1})} \mathbf{M}\left(\succeq_{\pi_x(k)}, P_V(y_{k+1}) \cup \{y_{k+1}\} \cup Y\right) .$$

If $x \in V$ then the tree has a single path in which all nodes are equal to x: $\mathfrak{s}^{\pi}_{k}(V, x) = \{x\}$ for each $k = 1, \ldots, m_x$.

Having constructed the tree $\mathfrak{T}^{\pi}(V, x)$ with the above procedure, we obtain a (possibly empty) set of terminal nodes of paths of length m_x. Let $F^{\pi}(V, x)$ be the set of terminal nodes that belong to V: that is, $y \in F^{\pi}(V, x)$ if and only if there exists a sequence $(y_1, \ldots, y_{m_x+1})$ such that $y_1 = y \in V$, $y_{m_x+1} = x$, and $y_k \in \mathfrak{s}^{\pi}_{k}(V, y_{k+1})$ for each $k = 1, \ldots, m_x$.

The idea behind this construction is as follows. Suppose that all players believe that the bargaining process 'converges' to V, in the sense that, starting from any default, bargaining must lead to a policy in V which will never be amended. Thus, when considering whether and how to amend the initial default $x \notin V$, they only consider policies in V. Suppose the m_xth proposer, $\pi_x(m_x)$, is given the opportunity to make a proposal in the first round. The set of policies she can induce includes the default x (if she passes, the unamended default will be implemented) and the set of policies in V that winning coalitions are willing to accept. The latter set must include $P_V(x)$. Indeed, if an offer y in V is accepted then it will be implemented; if it is rejected then x will be implemented. Consequently, voters who strictly prefer y to x must vote 'yes'. Voters who are indifferent between x and y may vote either 'yes' or 'no'. Thus, the set of policies that the last proposer can induce is of the form $P_V(x) \cup \{x\} \cup Y$, where $Y \subseteq R_V(x)$. The set Y is determined by the voting behavior of indifferent voters. For instance, a situation where indifferent voters always vote 'no' can be described by setting $Y = \emptyset$, while a situation where indifferent voters always vote 'yes' can be described by setting $Y = R_V(x)$. Thus, each path of the tree corresponds to a different assumption about voting behavior. The m_xth proposer must optimally choose a policy, y_{m_x}, which is $R_{\pi_x(m_x)}$-maximal in that 'feasible set' and, therefore, $y_{m_x} \in \mathfrak{s}^{\pi}_{m_x}(V, x)$. Now, consider the $(m_x - 1)$th proposer's choice. She faces the same problem as the m_xth proposer, except that x must be replaced by y_{m_x}: players anticipate that if the $(m_x - 1)$th proposer's proposal is rejected then y_{m_x} will be the final policy outcome. Hence, the policy y_{m_x-1} she chooses to induce belongs to $\mathfrak{s}^{\pi}_{m_x-1}(V, y_{m_x})$. Moving backward, we can repeatedly apply the same reasoning to all proposers until the first, $\pi_x(1)$, whose choice must thus be in $F^{\pi}(V, x)$.

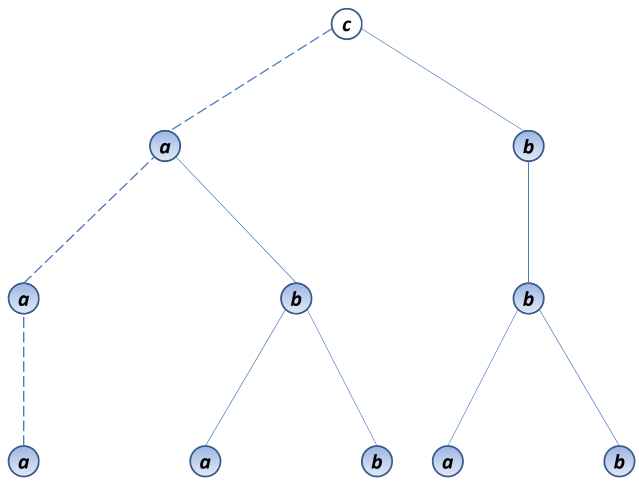

Figure 1: Tree $\mathfrak{T}^{\pi}(\{a, b\}, c)$

Before we proceed any further, it may be helpful to illustrate this intuition with a simple example. Suppose there are three proposers/voters — i.e.: $M = N = \{1, 2, 3\}$ — and $X = \{a, b, c\}$. The set of winning coalitions $\mathcal{W} = \mathbf{W}$ is the collection of majority coalitions (i.e. those coalition that include at least two players). Players' preferences are: $a \succ_1 b \succ_1 c$, $b \succ_2 a \succ_2 c$, $a \sim_3 b \succ_3 c$. Assume further that the initial default is $x^0 = c$, and that the protocol is defined as $\pi_x(i) = i$ for all $x \in X$ and all $i \in N$, thus completing the description of game $\Gamma\left(\pi, x^0\right)$. The tree $\mathfrak{T}^{\pi}(\{a, b\}, c)$ is depicted in Figure 1 — it is readily checked that $\{a, b\}$ is a weakly stable set. The set of immediate successors of the initial node, c, is $\mathfrak{s}^{\pi}_3(\{a, b\}, c) = \{a, b\}$. Indeed, the last proposer is player 3 and her ideal policies in $P_{\{a,b\}}(c) \cup \{c\}$ are a and b (in this example, $R_{\{a,b\}}(c) = P_{\{a,b\}}(c) = \{a, b\}$). Intuitively, proposer 3 anticipates that either a or b would be voted up while c would be the final outcome if she passed. Hence, she optimally proposes either a or b. Suppose she proposes the former. The set of immediate successors of node $a \in \mathfrak{s}^{\pi}_3(\{a, b\}, c)$ is $\mathfrak{s}^{\pi}_2(\{a, b\}, a) = \{a, b\}$. To see this, note first that $P_{\{a,b\}}(a) = \emptyset$ and $R_{\{a,b\}}(c) = \{a, b\}$. Thus, a is $\succeq_2$-maximal in $\{a\}$, and b is $\succeq_2$-maximal in $\{a, b\}$. The reason why policy a

may be the only option to proposer 2 is that voter 3, who is indifferent between a and b, may vote 'no' when offered b. If she does then player 1, anticipating that default c will be amended to her ideal policy a if she makes no proposal, optimally passes. Formally, $\{a\} = \mathbf{M}(\succeq_1, \{a,b\}) = \mathfrak{s}_1^\pi(\{a,b\},a)$, so that a is a final node of tree $\mathfrak{T}^\pi(\{a,b\},c)$. This completes the description of the dotted path in Figure 1. One could apply the same procedure and intuition to the other paths of $\mathfrak{T}^\pi(\{a,b\},c)$, so as to obtain $F^\pi(\{a,b\},c) = \{a,b\}$.

The intuition above is confirmed by our first two results:

Proposition 1. *Suppose that V is the closure of a weakly stable set, and let $f \in V^X$ be a selection of $F^\pi(V,\cdot)$: $f(x) \in F^\pi(V,x)$ for all $x \in X$. There exists an equilibrium σ such that $f^\sigma(x) = f(x)$ for all $x \in X$. Hence, $\bigcup_{x\in X} f^\sigma(x) = V$.*

This proposition says that, if V is the closure of a weakly stable set (and is therefore a weakly stable set itself) then *any* selection $f(\cdot)$ of $F^\pi(V,\cdot)$ can be supported by an equilibrium of $\Gamma(\pi,x^0)$. Put differently, all final nodes of length-m_x paths in tree $\mathfrak{T}^\pi(V,x)$ are equilibrium policy outcomes of continuation games starting with x as the initial default. In particular, all policies in $F^\pi(V,x^0)$ are equilibrium outcomes of $\Gamma(\pi,x^0)$. We assume that V is the closure of a weakly stable set to ensure that $F^\pi(V,x)$ is nonempty: if V is not closed then the set $\mathbf{M}\left(\succeq_{\pi_x(k)}, P_V(y_{k+1}) \cup \{y_{k+1}\} \cup Y\right)$ may be empty.

We will say that an equilibrium *supports* a weakly stable set V when exactly the initial defaults in V are not amended in that equilibrium. (Recall that we use the term 'implementation' to refer to the policy reached in $\Gamma(\pi,x^0)$.

This result prompts the following question: Can there be equilibria of $\Gamma(\pi,x^0)$, whose outcomes do not belong to $F^\pi(V,x^0)$? The next proposition answers this question in the negative.

Proposition 2. *If σ is an equilibrium of $\Gamma(\pi,x^0)$ then there exists $V \in \mathcal{V}$ such that $f^\sigma(x) \in F^\pi(V,x)$ for $x \in X$. Hence, $V = \bigcup_{x\in X} f^\sigma(x)$.*

To prove Propositions 1 and 2, we establish stronger results. First, weak stability of V implies that, for every $x \in X$ and every length-m_x path $(x, y^m(x), \ldots, y_1(x))$ of tree $\mathfrak{T}^\pi(V,x)$ with $y_1(x) \in V$, there exists an equilibrium σ such that $f^\sigma(x,k) = y_k(x)$ for each $k \in \{1,\ldots,m_x\}$. Second, for every equilibrium σ of $\Gamma(\pi,x^0)$ and every $x \in X$, there exists a weakly stable set V and a length-m_x path $(x, y^m(x), \ldots, y_1(x))$ of $\mathfrak{T}^\pi(V,x)$, with $y_1(x) \in V$, such that $y_k(x) = f^\sigma(x,k)$ for each $k \in \{1,\ldots,m_x\}$. Thus, the construction of trees associated with weakly stable sets also provides a complete characterization of equilibrium behavior both on and off equilibrium paths.

The policies which can be implemented in $\Gamma(\pi,x^0)$ depend on the initial default unless some policy is a weak Condorcet winner: viz. it weakly dominates every other policy.

Propositions 1 and 2 jointly yield a complete characterization of the set of policy outcomes that can be reached from any particular default policy $x^0 \in X$.

Corollary 1. *Let $\Sigma^* \left(\pi, x^0\right)$ be the set of equilibria of $\Gamma \left(\pi, x^0\right)$. The set of equilibrium policy outcomes in $\Gamma \left(\pi, x^0\right)$ is given by*

$$\bigcup_{\sigma \in \Sigma^*(\pi, x^0)} f^\sigma \left(x^0\right) = \bigcup_{V \in \mathcal{V}} F^\pi \left(V, x^0\right) \ .$$

An immediate implication of this result is that the set of policy outcomes that can result from all equilibria and from all initial defaults is the union of all weakly stable sets. Put differently, a policy in X can be obtained as the policy outcome of the bargaining game from some initial default if and only if it belongs to some weakly stable set. Thus, as a byproduct of our analysis, we obtain a new bargaining interpretation for (weakly) stable sets in voting games; one which, in contrast to the existing literature (e.g., Anesi (2006)), extends to situations with an infinite policy space.

The weakly stable sets in a game only depend on the protocol via M, the set of proposers. Propositions 1 and 2 imply that variations in the protocol do not affect the set of policies which can be implemented across initial defaults. However, as we will see in the next section, variations in the protocol may affect the policies which can be implemented from a given initial default.

3.2 Properties of the equilibrium correspondence

In this subsection, we illustrate some interesting properties of the equilibria of $\Gamma \left(\pi, x^0\right)$ via some examples which will prove useful in subsequent sections.

Example 3.2. Suppose that two players (1 and 2) can divide a pie, earning their share of the pie, if and only if they both agree; that player 1 proposes before player 2 in each round; and that both players earn 0 at the initial default (x^0). If x_1 denotes player 1's share then the policy space consists of x^0 and every $x_1 \in [0, 1]$. This policy space is not well ordered because each player is indifferent between x^0 and a division which yields her none of the pie. There is a unique weakly stable set, consisting of every division of the pie. $\Gamma \left(\pi, x^0\right)$ then has a unique equilibrium outcome in which player 2 takes the whole pie.

□

Example 3.3. Consider the preferences in Bernheim et al's (2006) benchmark pork barrel model, where each player i earns b_i if her project is implemented, and pays $c_j > 0$ for every project j implemented; the policy space is well ordered; and any bare majority (viz. $(n+1)/2$, n odd) of voters is winning. Suppose, in addition, that every player is both a proposer and a voter. We will now argue that this game has a unique weakly stable set, consisting of a single policy: the $(n+1)/2$ cheapest projects are implemented:

Well ordering and a majority quota imply that the game is strong; so any weakly stable set must be a singleton, say y. Bernheim et al's additional assumption A3 implies that a bare majority of the projects is implemented in our model.[12] To see this, consider a policy x which differs from y by removing some of the projects implemented at y. Every player whose project is implemented at x and every player whose project is not implemented at y prefers x over y. Hence, y would fail external stability if more than a bare majority of projects were implemented. On the other hand, if fewer than a bare majority of projects were implemented at y then A3 implies that a policy x which adds a bare majority of projects to y would be preferred over the latter by all of the beneficiaries.

In light of these arguments, a weakly stable set must implement the cheapest projects (policy y^*). To see this, suppose that another policy (y) with a bare majority of beneficiaries formed a weakly stable set. By construction, some player must be a member of both coalitions. This player must prefer y^* over y, as must the bare minority of players who are beneficiaries at y^*; so y^* would strictly dominate y, contrary to external stability. Hence, $\{y^*\}$ is the unique weakly stable set and, by Propositions 1 and 2, y^* is the unique equilibrium policy.

□

Example 3.4. Suppose that any two out of three players can agree to any division of a dollar. It is well known that the union of stable sets for this game is the entire triangle: viz. every division of the dollar (cf. Ordeshook (1986) Ch 9.2). If at least two players can propose then $\mathcal{W}$ is the set of pairs of players. As every stable set is weakly stable, the set of weakly stable policies is also the entire triangle. This observation implies that a player who cannot propose may nevertheless earn the entire dollar in some equilibrium of a game whose initial default is no agreement. By contrast, a player who cannot propose earns 0 in Baron and Ferejohn's (1989) closed rule model. Furthermore, policies in the interior of the triangle may be implemented in an equilibrium, contrary to the size principle, which also holds in Baron and Ferejohn's (1989) closed rule model.[13]

□

We argued above that $\{x, z\}$ is a weakly stable set in Example 3.1; so there is an equilibrium in which these policies alone can be implemented from any initial default.

[12] This assumption states that a mutually beneficial policy (relative to the default) exists for all coalitions of $(n+1)/2$ or fewer individuals.

[13] See also Wilson (1971) on main simple stable sets. On the other hand, the size principle fails in Baron and Ferejohn's open rule model when players are impatient enough.

Consider a protocol in which players propose in the order 1, 3, 2 in each round and suppose that y is the initial default. It is readily checked that path (y, x, z, z) is a path of tree $\mathfrak{T}^{\pi}(\{x, z\}, y)$: $x \in \mathfrak{s}_3^{\pi}(\{x, z\}, y)$, $z \in \mathfrak{s}_2^{\pi}(\{x, z\}, x)$, and $z \in \mathfrak{s}_1^{\pi}(\{x, z\}, z)$. Consequently, there is an equilibrium in which player $1 = \pi_y(1)$ passes, player $3 = \pi_y(2)$ amends the default to z, which is then implemented. This equilibrium demonstrates that a committee can implement a policy which is strictly dominated by the initial default. This property must hold in Bernheim et al's (2006) benchmark pork barrel model (Example 3.3. above) when there is no Condorcet winner: the last proposer's ideal policy is then implemented; but it violates Acemoglu et al's (forthcoming) Desirability Axiom.

We record the arguments in this subsection as

Observation 1. *a) A player who does not propose may nevertheless earn all of the surplus from agreement;*

b) The size principle may fail in an equilibrium;

c) The members of some winning coalition may all strictly prefer the initial default x^0 to the final policy outcome.

4 Comparative statics

In this section, we consider how variations in the model's parameters affect the policies that are implemented from any initial default. In Section 4.1, we explore the effect of changing the protocol on the policies implemented in a given weakly stable set. In Section 4.2, we focus on the implications of changes in the set of weakly stable sets.

4.1 The protocol

Thus far, we have studied play in games with a fixed protocol. In this subsection, we study situations in which a player, the *chair,* chooses a protocol $\pi \in \Pi$ after observing the initial default x^0; and the game $\Gamma\left(\pi, x^0\right)$ is then played. In order to describe the chair's decision problem in such situations, we must make an assumption about her preferences over protocols. To do so, we need to address the implications of multiple equilibria.

Indeed, the analysis of equilibria in the previous section revealed that there may be equilibrium multiplicity at two levels in the bargaining game (for a given protocol π). First, Proposition 1 says that any weakly stable set can be supported by an equilibrium. The possible multiplicity of weakly stable sets may thus be a source of equilibrium multiplicity. Second, Proposition 1 also implies that, for a given weakly stable set $V \in \mathcal{V}$, any terminal node of tree $\mathfrak{T}^{\pi}\left(V, x^0\right)$ is the policy outcome of some equilibrium of $\Gamma\left(\pi, x^0\right)$. Hence, each weakly stable set may contain several equilibrium policy outcomes. As the chair chooses protocols π that will apply in the bargaining game $\Gamma\left(\pi, x^0\right)$, her choice will depend on her anticipations about future behavior in that game, and therefore we must make an assumption about her predictions of equilibrium behavior in $\Gamma\left(\pi, x^0\right)$ for every protocol π.

One extreme assumption would be that the chair prefers a protocol π_1 to another protocol π_2 if there are equilibria σ and σ' of $\Gamma\left(\pi_1, x^0\right)$ and $\Gamma\left(\pi_2, x^0\right)$, respectively, such that she prefers the equilibrium policy $f^{\sigma_1}\left(x^0\right)$ in $\Gamma\left(\pi_1, x^0\right)$ to the equilibrium policy $f^{\sigma_2}\left(x^0\right)$ in $\Gamma\left(\pi_2, x^0\right)$. The difficulty with this assumption is that, for any protocols π_1 and π_2, one will typically find pairs of equilibria in $\Gamma\left(\pi_1, x^0\right)$ and $\Gamma\left(\pi_2, x^0\right)$ such that the chair both prefers π_1 to π_2 and π_2 to π_1, leaving her indifferent between all protocols. At the other extreme, one could assume that the chair prefers π_1 to π_2 if she prefers all the equilibrium policies in $\Gamma\left(\pi_1, x^0\right)$ to all those in $\Gamma\left(\pi_2, x^0\right)$. With this assumption, however, it will typically be the case that the chair neither prefers π_1 to π_2 nor π_2 to π_1. Based on our characterization results in Section 3, our assumption — as we formally elaborate on below — will instead be that the chair prefers π_1 to π_2 if, for every weakly stable set V, she prefers all the policies implemented in equilibria supporting V in $\Gamma\left(\pi_1, x^0\right)$ to all the policies implemented in equilibria supporting V in $\Gamma\left(\pi_2, x^0\right)$.

Accordingly, we will exploit Proposition 2 by focussing on the implications of protocol variations when equilibria which support a given weakly stable set are played, allowing us to use the structure of weakly stable sets.[14]

We divide this subsection's results into two parts. First, we define conditions under which the policy implemented from any initial default is independent of the protocol — thus leaving the chair indifferent over all protocols — and provide necessary and sufficient conditions for cases in which the policy space is both finite and well ordered. We then characterize conditions under which the chair cannot improve on proposing last each round: the analog of Bernheim et al's (2006) power of the last word.

4.1.1 Order independence

Propositions 1 and 2 simplify the set of equilibrium policies we must consider when evaluating the impact of the protocol on equilibrium outcomes. However, as explained above, there still remains a potential multiplicity of equilibria to be evaluated. This evaluation can be simplified by grouping equilibria into classes, where each element of a class supports an identical weakly stable set. From Proposition 2, we thus obtain a partition of the set of equilibria. We will say that the bargaining game exhibits 'order independence' if the set

[14]We could, alternatively, focus on the set of policies implemented across equilibria, and therefore across weakly stable sets; but the union of weakly stable sets is typically not weakly stable.

of equilibrium policies *within each partition element* is independent of proposal orders in the protocol.

Formally, let Π be set of protocols π, as defined in Section 2 (holding M constant). The tree construction in the previous section gives us a clear description of the equilibria of $\Gamma(\pi, x^0)$ for each $\pi \in \Pi$, $\Sigma^*\left(\pi, x^0\right)$. We know from Proposition 2 that each member of $\Sigma^*\left(\pi, x^0\right)$ supports those policies in some weakly stable set V. We can therefore partition $\Sigma^*\left(\pi, x^0\right)$ into a collection $\left\{\Sigma_V^*\left(\pi, x^0\right)\right\}_{V \in \mathcal{V}^*}$, where $\mathcal{V}^* \subseteq \mathcal{V}$ is the class of weakly stable sets that are supported by some equilibrium of $\Gamma(\pi, x^0)$: $\mathcal{V}^* \equiv \{V \in \mathcal{V} : F^\pi(V, x) \neq \emptyset \ , \forall x \notin V\}$. This partition groups together those equilibria that support identical weakly stable sets. We say that the class of games $\left\{\Gamma(\pi, x^0) : \pi \in \Pi \ \& \ x^0 \in X\right\}$ satisfies *order independence* if the following statement is true for any initial default $x^0 \in X$ and any $V \in \mathcal{V}^*$: $F^{\pi_1}\left(V, x^0\right) = F^{\pi_2}\left(V, x^0\right)$ for any $\pi_1, \pi_2 \in \Pi$.[15]

Here are two cases which satisfy order independence:

1. If the game has a Condorcet winner then this policy is the only weakly stable set. Propositions 1 and 2 imply that this policy is implemented for every protocol and from every initial default.

2. There are four policies $\{a, b, c, d\}$ and four players, who are both voters and proposers (i.e.: $M = N = \{1, 2, 3, 4\}$), and whose preferences are

$$a \succ_1 b \succ_1 c \succ_1 d$$
$$b \succ_2 c \succ_2 d \succ_2 a$$
$$c \succ_3 d \succ_3 a \succ_3 b$$
$$d \succ_4 a \succ_4 b \succ_4 c$$

   Every coalition of at least three players is winning.[16] For every policy x, a unique policy weakly dominates x; and the weakly stable sets are $\{a, c\}$ and $\{b, d\}$. Consequently, the games $\{\Gamma(\pi, x^0)\}$ satisfy order independence: for example, if $\{b, d\}$ is supported in equilibrium and a is the initial default then d must be implemented, irrespective of the protocol; and any default in $\{b, d\}$ is stable, irrespective of the protocol. The same argument applies when c is the initial default; and analogous arguments apply for equilibria which support $\{a, c\}$ and any initial default.

General conditions which ensure order independence turn out, however, to be uninstructively complex. Nevertheless, we can make much more progress by focussing on cases

[15]We define order independence by fixing the weakly stable set implemented and varying the protocol, whereas Moldovanu and Winter's (1995) definition fixes a given strategy combination.

[16]This example appeared in Anesi (2006).

in which X is both finite and well ordered (so any weakly stable set is stable). The following result will be useful:

Proposition 3. *If X is finite and well ordered then, for any $x \notin f^{\sigma}(X)$ and every equilibrium σ of $\Gamma\left(\pi, x^0\right)$:*

$$f^{\sigma}(x) = \mathbf{M}\left(\succ_{\pi(k)}, R(x) \cap f^{\sigma}(X)\right) \ ,$$

where $k \equiv \max\left\{l \in \{1, \ldots, m_x\} : \mathbf{M}\left(\succ_{\pi(l)}, R(x) \cap f^{\sigma}(X)\right) \succ_{\pi(l)} x\right\}$.

In Proposition 3, the kth proposer is the last proposer among those who have an incentive to amend the ongoing default x in equilibrium σ; namely those who strictly prefer some equilibrium policy that is 'reachable' to the default. Denoting by $V = f^{\sigma}(X)$ the weakly stable set supported by σ (Proposition 2), we will refer to those 'amenders of x' as

$$M(V, x) \equiv \{i \in M : \mathbf{M}(\succ_i, R(x) \cap V) \succ_i x\} \subseteq M \ .$$

Proposition 3 thus implies that the ideal policy in $R(x) \cap V$ of the last amender according to π is implemented in every equilibrium $\sigma \in \Sigma_V\left(\pi, x^0\right)$. Note that, while the identity of the last amender of x depends on which element of Π we consider, the set of amenders of x is the same for all permutations of π.

Interestingly, while a non-amender (say, i) would pass if she were the last proposer in a round, she might otherwise amend x^0: for example, to some $v \in V$ in an equilibrium if, after i passed, the next amender would propose a policy v' which weakly dominates v and is strictly worse for i.

Using the proposition above, we can characterize conditions for order independence when the policy space is finite and well ordered.

Proposition 4. *If X is finite and well ordered then order independence is satisfied if and only if the following is true for any weakly stable set $V \in \mathcal{V}$: For each $x \notin V$ and all $i, j \in M(V, x)$,*

$$\mathbf{M}(\succ_i, R(x) \cap V) = \mathbf{M}(\succ_j, R(x) \cap V) \ . \tag{1}$$

Condition (1) offers a very simple necessary and sufficient condition for order independence in the finite, well ordered case: the bargaining game satisfies order independence if and only if, for any weakly stable set V and any possible default x, all amenders of x share the same ideal policy among those in V which dominate x. Although very simple, this condition seems to be quite demanding. One can therefore expect that in most cases there will be room for manipulation of the protocol. This is the subject matter of Section 4.1.2.

4.1.2 Manipulating the protocol

In this subsection, we explore conditions under which the chair can influence the policy implemented from a given default by varying the protocol. We will simplify exposition by supposing that the chair is a proposer.

X well ordered: 'power of the last word'

Our main results above concerned cases in which the policy space is well ordered: Proposition 4 then specifies conditions under which the chair is indifferent across all protocols (order independence). These conditions are satisfied if there is a Condorcet winner. The chair can also not affect the policy implemented in Bernheim et al's (2006) related model with an exogenous deadline (and well ordered policy space) if there is a Condorcet winner.

On the other hand, Proposition 3 implies that a chair who wants to amend the initial default can never improve on making the last proposal at the initial default when the policy space is well ordered. This feature is reminiscent of the 'power of the last word' in Bernheim et al's (2006) pork barrel model, which we included above as Example 3.3: they show that the last proposer's ideal policy (her own project) is implemented in every equilibrium.

However, there are important differences between our respective results. In particular, Proposition 3 implies that the last amender gets her best policy in some weakly stable set that weakly dominates the initial default rather than her best policy. We argued above that the only weakly stable set consists of the cheapest bare majority of projects. The chair's best protocol might therefore not entail implementation of her own project, in radical contrast to Bernheim et al's result. This difference reflects an important distinction between our respective use of backward induction arguments. In Bernheim et al, the game must end after the last proposal; in our model, any amendment must lead to the implementation of a policy in a weakly stable set. More generally, Bernheim et al's result relies on the pork barrel structure and the existence of winning coalitions which exclude some players. This allows the final proposer to play off putative members of the winning coalition. Proposition 3, by contrast, allows for cases in which only the grand coalition is winning: e.g. in variants on Example 3.2 where the pie can only be split in a finite number of proportions.

General case: 'power of the last two words'

We now turn to cases in which the policy space is not well ordered. We start with a useful result, which states that the set of policies which can be implemented in equilibrium — i.e., from Corollary 1: $\bigcup_{V\in\mathcal{V}} F^{\pi}\left(V,x^0\right)$ — only depends on the protocol at the initial default: π_{x^0}.

Proposition 5. *Let $\pi^1 \equiv \left\{\pi^1_x\right\}$ and $\pi^2 \equiv \left\{\pi^2_x\right\}$ be two protocols in Π. If $\pi^1_{x^0} = \pi^2_{x^0}$ then*

$$\bigcup_{V\in\mathcal{V}} F^{\pi^1}\left(V,x^0\right) = \bigcup_{V\in\mathcal{V}} F^{\pi^2}\left(V,x^0\right) \ .$$

Thus, the chair's best protocols only depend on the protocol at the initial default. In particular, Proposition 5 implies that, for any $\pi \in \Pi$, there is a constant protocol $\bar{\pi}$ — i.e. $\bar{\pi}_x = \bar{\pi}_y$ for all $x, y \in X$ — such that the same policies are implemented in $\Gamma\left(\pi,x^0\right)$ and $\Gamma\left(\bar{\pi},x^0\right)$. As the chair's choice of protocols only affects her payoff via π_{x^0}, the chair cannot improve on selecting some $\bar{\pi}$ which is constant across X.[17]

This result is of independent interest; but we will exploit it here to simplify our analysis of protocol manipulation by the chair: we can without loss of generality restrict attention to constant protocols. Fix a weakly stable set $V \in \mathcal{V}$, a policy $x^0 \notin V$, and some constant protocol $\pi \in \Pi$. To lighten the notation, let $x_k \equiv \mathbf{M}\left(\succeq_k, V\right)$ and $y_k \equiv \mathbf{M}\left(\succeq_k, R_V\left(x^0\right)\right)$ for each amender $k \in M\left(V,x^0\right)$: x_k and y_k are amender k's ideal policies in V and $R_V\left(x^0\right)$, respectively.

Write $K \in M\left(V,x^0\right)$ for the last amender according to π. We have already seen that, when X is well ordered, if the chair j is an amender then she can never improve on making the last proposal (i.e., being the last amender: $j = K$): from Proposition 3, y_K must be implemented in equilibrium. If she is not an amender then she is best off choosing π (and therefore the last amender K) such that $y_K \succeq_j y_k$ for all $k \in M\left(V,x^0\right)$.

If X is not well-ordered then there is an equilibrium in which K amends x^0 to y_K. However, we will now argue that *a chair who is an amender can improve on protocol π (in which she is the last amender K) if $x_K \neq y_K$*. To see this, take another constant protocol π_i which differs from π by adding a proposal by some player $i \in M\left(V,x^0\right)$ immediately after K, and consider $\Gamma\left(\pi_i,x^0\right)$ in which protocol π_i applies at every default. Player K is better off in some equilibrium with protocol π_i than with π if and only if the following is true:

($\mathbf{C}_i$) There is $v \in \left[V \setminus R_V\left(x^0\right)\right]$ which weakly dominates y_i (i.e.: vRy_i) such that $v \succ_K y_K$ and $v \succ_K y_i$.

It is readily checked that, when (C_i) holds, a path of tree $\mathfrak{T}^{\pi_i}\left(V,x^0\right)$ describes an equilibrium of $\Gamma\left(\pi_i,x^0\right)$ in which: any proposal by a player moving before K to amend x^0 is

[17] Proposition 5 also applies in a different 'dynamic' game where the chair selects the next proposer immediately after each vote which does not end the game (see the last section of the Appendix for a formal proof).

rejected; K successfully proposes v; any other proposal by K would be rejected; and player i would amend x^0 to y_i if K did not amend x^0.

Suppose (C_i) holds for $i = K$. In this case, the chair $K \in M\left(V, x^0\right)$ is best off with a protocol in which she is the penultimate and the last proposer if $x_K \notin R_V\left(x^0\right)$ and $x_K R y_K$: in the equilibrium of $\Gamma\left(\pi_K, x^0\right)$ described in the previous paragraph, all proposers pass till the chair (who is the last amender in π_j) makes the last two proposals, at the first of which x^0 is amended to x_K. All players anticipate that rejecting x_K would result in y_K being proposed and accepted. Since $x_K R y_K$, a winning coalition of voters is (weakly) better off accepting x_K. This result contrasts with equilibria in Diermeier and Fong's (2011) game with repeated implementation and a single proposer, who may lose and cannot gain if she proposes more than once.

Now suppose that (C_i) holds for $i \neq K$. The chair K would then be better off in an equilibrium if she were the penultimate but not the last amender. By the same argument as in the previous paragraph, the equilibrium of $\Gamma\left(\pi_i, x^0\right)$ would yield policy v ($\succ_K y_K$). This result contrasts with the equilibria in Bernheim et al (2006) and with our results for well ordered X, where each amender is at least as well off proposing last.

We summarize the discussion above in

Proposition 6. *If the policy space is not well ordered then the chair may prefer to make the last two or the penultimate proposals over only proposing last.*

4.2 The set of winning coalitions

Thus far, we have considered how varying the protocol affects play for a given set of weakly stable sets. In this subsection, we explore the effects of changing the set of winning coalitions, and thereby the weakly stable sets. We consider two reasons why the winning coalitions might change: in Section 4.2.1, we study the effects of increasing the quota; in Section 4.2.2, we consider how changing the number of proposers affects play.

4.2.1 Quotas

In conventional bargaining models with spatial preferences on the real line, an increase in the quota makes voters with more extreme preferences decisive. The committee can only amend a default if the decisive voters agree; so committees with a greater quota have a larger gridlock interval.[18] We will discuss this conclusion in the context of our model with an arbitrary policy space.

[18] See Compte and Jehiel (2010) for a collective search perspective.

We say that $\Gamma\left(\pi, x^0\right)$ is a *quota game* if the collection of winning coalitions (of voters) $\mathbf{W}$ is of the form $\mathbf{W}^s \equiv \{S \subseteq N : |S| \geq s\}$ with $s \geq \frac{n+1}{2}$. Our goal is thus to study how the set of equilibrium policies of a quota game is affected by an increase in the quota. Given the collection of winning coalitions (of voters) $\mathbf{W}^s$, we can define the corresponding social preference relations R_s and P_s, and the corresponding collection of weakly stable sets $\mathcal{V}^s$ as we did in Section 3. An immediate consequence of our characterization results is that, in our context, the conclusion above can be reformulated as: $q > r$ implies that $\bigcup \mathcal{V}^r \subseteq \bigcup \mathcal{V}^q$ (where $\bigcup \mathcal{V}^s \equiv \{v \in X : v \in V \text{ for some } V \in \mathcal{V}^s\}$). We will refer to this property as the *Conjecture,* and to "$q > r$ implies $\bigcup \mathcal{V}^r \subset \bigcup \mathcal{V}^q$" as the *strong form of the Conjecture.*

A sufficient condition for the Conjecture is that there is enough conflict of interest that $X \in \mathcal{V}^q$, for we must then have $\bigcup \mathcal{V}^r \subseteq X = \bigcup \mathcal{V}^q$ for any r. Indeed, $X \in \mathcal{V}^r$ implies $X \in \mathcal{V}^q$, but the converse is false; so the strong form of the Conjecture may hold when there is enough conflict of interest that $X \in \mathcal{V}^q$, but not enough that $X \in \mathcal{V}^r$.

At the other extreme of enough common interest, there is a Condorcet winner (which must be the only stable set) with the higher quota. The same policy must then be the Condorcet winner with the lower quota; so the Conjecture trivially holds. Indeed, the strong form of the Conjecture may hold if $xP_r y$ for every other policy y, but there is $y \neq x$ such that $\neg\left(xP_q y\right)$.

We now turn to intermediate cases, where there is no Condorcet winner for quota r and $X \notin \mathcal{V}^q$. The former condition implies that $\neg\left(xP_r y\right)$ and $\neg\left(yP_r x\right)$ for some policies $x \neq y$; the latter condition implies that there are $x, y \in X$ such that $xP_q y$; and $r < q$ implies that $xP_r y$. Hence, no $V \in \mathcal{V}^r$ or $V' \in \mathcal{V}^q$ can contain both x and y.

We approach these cases from two perspectives. We first consider whether, for any $V \in \mathcal{V}^r$, there is $V' \in \mathcal{V}^q$ such that $V \subseteq V'$ when X is well ordered; we then consider cases where $\bigcup \mathcal{V}^s = X$ for some $s \in \{q, r\}$.

If X is well ordered and V is a stable set then no set of policies which contains V can be a stable set. If $V \in \mathcal{V}^r$ then V satisfies P_q-internal stability. Define U_V as $\{x \in X \setminus V : \neg\left(yR_q x\right), \forall y \in V\}$. If U_V is empty for every $V \in \mathcal{V}^r$ then every stable set with quota r is also a stable set with quota q. Indeed, if X is well ordered then the games with quotas q and r have the same stable sets, so the Conjecture holds.

However, V would fail external stability with quota q if U_V were nonempty for some $V \in \mathcal{V}^r$. As V satisfies P_r-external stability, we cannot have $xP_q y$ for any $x \in U_V$ and $y \in V$. Now consider a game with the same players and preferences, a quota of q but possible policies of U_V (rather than X), and let T_V be a stable set of this game. We must then have $T_V \cup V \in \mathcal{V}^q$. The strong form of the Conjecture then holds if U_V is nonempty for some $V \in \mathcal{V}^q$, and some T_V exists for every nonempty U_V. More generally, the Conjecture holds if every game with possible policies $U \subseteq X$ has a stable set with quota q.

We now turn to the other perspective, demonstrating by example that the Conjecture might fail for intermediate cases because $\bigcup \mathcal{V}^q \subset X = \bigcup \mathcal{V}^r$, even if X is well ordered:

Example 4.1. Suppose that $X = \{w, x, y, z\}$, all players propose (i.e., $M = N$), and there are two quotas: $r < q$. Social preferences satisfy xP_ry, wP_rz, yP_rw, zP_rx, $\neg (xP_rw)$, $\neg (wP_rx)$, $\neg (yP_rz)$, $\neg (zP_ry)$, $\neg (xP_qy)$, wP_qz, yP_qw, and zP_qx.

The strict social preferences and $r < q$ imply that this game is intermediate in the sense above. If the quota is r then there are two stable sets: $\{w, x\}$ and $\{y, z\}$; so $\bigcup \mathcal{V}^r = X$. As $r < q$, we must also have: $\neg (xP_qw)$, $\neg (wP_qx)$, $\neg (yP_qz)$, and $\neg (zP_qy)$. These conditions imply that $\{y, z\} \in \mathcal{V}^q$. We will now argue that x cannot be in any stable set when the quota is q. Policy x alone cannot be a stable set because zP_qx, which also implies that z cannot be in any stable set containing x. No stable set can contain x and y because it would have to include w to dominate z; and this would violate (IS_{P_q}) because yP_qw. The latter also excludes any stable set which contains x and w, but not y. In sum, x is in no stable set, so $\bigcup \mathcal{V}^q \subset \bigcup \mathcal{V}^r = X$, contrary to the Conjecture.

$\square$

Note that Example 4.1 is structured such that the game restricted to $\{w, x, y\}$ does not have a stable set with quota q. On the other hand, it is easy to construct intermediate case examples in which the strong form of the Conjecture holds because $\bigcup \mathcal{V}^r \subset X = \bigcup \mathcal{V}^q$. To see this, modify the example above by adding v to X, and consider social preferences which satisfy $xP_{qr}y$, wP_qz, yP_qw, zP_qx, $\neg (xP_qw)$, $\neg (wP_qx)$, $\neg (yP_qz)$, $\neg (zP_qy)$, uP_rv and $\neg (uP_qv)$ $\forall u \neq v$. No set in $\mathcal{V}^r$ can contain v, which is contained in both sets in $\mathcal{V}^q$, whose union is X.

We summarize the arguments above in

Proposition 7. *Suppose* $\Gamma\left(\pi, x^0\right)$ *is a quota game. An increase in the quota weakly expands the set of equilibrium decisions if, for the higher quota, there is a Condorcet winner or if* X *is a stable set or if the game restricted to each subset of decisions has a stable set. However, an increase in the quota may contract the set of equilibrium decisions in other intermediate cases.*

4.2.2 Adding proposers

It is widely believed that players can never lose if they are given the opportunity to propose: for a proposer could always make an offer which will be rejected. This argument has been influential, for example, in the design of regulatory agencies, which are required to include stakeholders in their decision making process; and the argument is correct in our model *for any fixed weakly stable set.* However, adding a proposer can change the set of coalitions in $\mathcal{W}$, and thereby the weakly stable sets. Consequently, as we argue below, *a player may be worse off if she is given the opportunity to propose.*

We will henceforth focus on the special case where there is initially a single proposer (say, player 1): both for expositional convenience and in order to compare our results with Diermeier and Fong (2011), who study a model with repeated implementation. They show that the single proposer may be worse off when she has the opportunity to propose in several rounds than when she proposes once in the game. This property is clearly impossible in our model: on the one hand, adding another proposal by player 1 does not change the set of weakly stable sets; on the other hand, player 1 could pass at her first opportunity to propose. The same argument implies that the set of policies which can be implemented in some equilibrium is unchanged by adding another proposer (say, player 2) with the same preferences as player 1.

Adding a proposer with different preferences from player 1 may affect play for two reasons. We have argued above that, for given weakly stable sets, player 1 may be better off if player 2 proposes before her, provided that X is not well ordered. We will now demonstrate by example that adding a proposal by player 2 may make player 1 better off and player 2 worse off because of changes in the weakly stable sets, even if X is well ordered.

Example 4.2. There are four policies: $X = \{w, x, y, z\}$, and X is well ordered. If player 1 can alone propose — so that all winning coalitions in $\mathcal{W}$ include her — then zPw and wPx; and no other pair of policies can be socially ranked. The only weakly stable set is then $\{x, y, z\}$. In particular, an initial default of x is then implemented.

Player 2's preferences are $x \succ_2 y \succ_2 z \succ_2 w$. If player 2 can propose then we have $yP'z$, $zP'w$ and $wP'x$; and no other pair of policies can be socially ranked. In this case, the only weakly stable set is $\{w, y\}$. Player 1 is then better off when the initial default is x because w is then implemented.

Our results in Section 3 then imply that player 2 is strictly better off *not* proposing when the initial default is either w or x, and only gains from proposing if z is the initial default.

$\square$

5 Extensions

5.1 Implementation

According to the model analyzed above, payoffs only depend on the policy (if any) that is eventually implemented, at which point the game ends. In a variant on our model, bargaining continues indefinitely; but payoffs are determined by the policy implemented. Equilibrium outcomes in this related model clearly correspond to equilibrium outcomes in our model because play after implementation is payoff-irrelevant. The extensive form in this variant is exactly that studied in the literature on repeated implementation, where players earn a per-round utility which depends on the ongoing default, and payoffs are the net present value of the utilities earned each round. Consequently, for any fixed strategy combination, each player's payoff in a repeated implementation model with a discount factor $\delta \simeq 1$ is close to that in the variant on our model. This observation suggests that, for every $\varepsilon > 0$, there is $\delta < 1$ such that an equilibrium strategy combination in our model (or, more precisely, in the related model) might be an ε-equilibrium in the repeated implementation model. We explore such an intuition in this subsection.

More specifically, we consider a variant of $\Gamma\left(\pi, x^0\right)$ in which the bargaining process continues *ad infinitum.* At the end of each round t, $t = 1, \ldots, \infty$, the new default policy x^t is implemented and each player i receives an instantaneous payoff $(1-\delta)u_i\left(x^t\right)$, where $\delta \in (0,1)$ is the common discount factor and $u_i \in \mathbb{R}^X$ is a continuous utility function which represents $\succeq_i$ — Assumption A0 guarantees that such a utility function exists. Thus, player i's payoff from a sequence of defaults $\left\{x^t\right\}_{t=1}^{\infty}$ is $(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} u_i\left(x^t\right)$. We will refer to the game thus obtained as $\Gamma^\delta\left(\pi, x^0\right)$. We will say that an equilibrium of $\Gamma^\delta\left(\pi, x^0\right)$ is absorbing if there is a round T such that $x^t = x^T$ for every subsequent round: $t > T$.[19]

Our next result confirms the intuition above: If we weaken the equilibrium concept by only requiring approximate best responses, we can obtain a similar result as Proposition 1 in the repeated implementation model $\Gamma^\delta\left(\pi, x^0\right)$, when δ is large enough. Indeed, Proposition 8 below states that some equilibria of $\Gamma\left(\pi, x^0\right)$ are also contemporaneous perfect ε-equilibria (Mailath et al., 2005) of $\Gamma^\delta\left(\pi, x^0\right)$ when δ is close enough to 1.

We abuse terminology in the next two results by identifying equilibria in our model — $\sigma \in \Sigma^*\left(\pi, x^0\right)$ — with equilibria in the related model with continued (but payoff-irrelevant) bargaining:

Proposition 8. *Let V be the closure of a weakly stable set. There exists $\sigma \in \Sigma_V^*\left(\pi, x^0\right)$*

[19] Existence and characterization of absorbing equilibria in legislative bargaining games with repeated implementation are discussed in Acemoglu et al (forthcoming), Anesi (2010) and Diermeier and Fong (2011).

such that the following is true: for any $\varepsilon > 0$, there exists $\delta_\varepsilon \in (0,1)$ such that σ is a contemporaneous perfect ε-equilibrium of $\Gamma^\delta\left(\pi, x^0\right)$ for all $\delta > \delta_\varepsilon$.

Proposition 8 implies that every policy which can be implemented in our game, the union of weakly stable sets, is a possible policy outcome in a contemporaneous perfect ε-equilibrium of the game with repeated implementation. Our last result in this subsection strengthens this result for the case of finite, well ordered X:

Proposition 9. *If X is finite and well ordered then there exists $\bar{\delta} \in (0,1)$ such that the following statement is true whenever $\delta > \bar{\delta}$: σ is an equilibrium of $\Gamma\left(\pi, x^0\right)$ if and only if it is an absorbing stationary Markov equilibrium of $\Gamma^\delta\left(\pi, x^0\right)$.*

As δ becomes arbitrarily close to 1, player i's discounted payoff from a (converging) sequence of defaults $\left\{x^t\right\}$ becomes arbitrarily close to her instantaneous payoff from the limit policy, say x^T:

$$\sum_{t=1}^{\infty} \delta^{t-1} u_i\left(x^t\right) \to u_i\left(x^T\right) \text{ as } \delta \to 1 \ .$$

The assumption that X is finite and well ordered thus guarantees that there exists a sufficiently large $\delta < 1$ ($\bar{\delta}$) such that players evaluate sequences of defaults similarly in $\Gamma^\delta\left(\pi, x^0\right)$ and $\Gamma\left(\pi, x^0\right)$: only final (or limit) policies matter. Put differently, xP_iy if and only if player i strictly prefers any sequence of defaults converging to x to any sequence converging to y.

5.2 The largest consistent set

In this subsection, we study the relation between our framework and that in Chwe (1994). Although the latter's approach to farsighted coalitional stability is cooperative, it is closely related to ours: as in our model, when a coalition S contemplates a deviation from the ongoing default, its members anticipate (and only take into account) the final outcome that will result from the sequence of deviations triggered by S's initial deviation. Chwe argues that the set of stable outcomes — i.e., those that are immune to these farsighted coalitional deviations — should satisfy a consistency condition, which in the context of our paper is defined as follows.

Say that a set of policies $Z \subseteq X$ is *consistent* if and only if the following is true for all $z \in Z$: for any $x \in X$ and $S \in \mathcal{W}$, there exists $z' \in Z$, where $z' = x$ or $z'Rx$, such that $z \succeq_i z'$ for some $i \in S$.

In words, any element z of a consistent set Z is 'stable' in the sense that each winning coalition S anticipates that a deviation from z will eventually lead to another policy z' in Z which makes at least one member of S worse off. Interestingly, Chwe (1994) shows that there exists a *largest consistent set*, $\overline{Z}$: Z consistent implies $Z \subseteq \overline{Z}$ and $\overline{Z}$ is itself consistent. Thus, $\overline{Z}$ comprises all the policies that are immune to farsighted coalitional deviations.

Our next goal is to study the relationship of our bargaining model to the largest consistent set. We have analyzed the model in previous sections by characterizing its Markov stationary equilibria. We have shown that every Markov stationary equilibrium implements a weakly stable set. Although it is readily checked that a weakly stable must be consistent and, therefore, a subset of the largest consistent set, the converse is not true: a consistent set may not be weakly stable, so that in general $\overline{Z} \notin \mathcal{V}$. In this subsection, we weaken stationarity, and show that the ensuing set of equilibria implement the largest consistent set.

To do so, we first need some definitions. In general, a history at some stage of the game describes all that has transpired in the previous rounds and stages (the sequence of defaults and proposers, their respective proposals and the associated pattern of votes). We call a 'partial round-t history' any list $\left(x^0, S^1, x^1, \ldots, S^{t-1}, x^{t-1}\right)$ where $S^s \in \mathcal{W}$ stands for the winning coalition which amended x^{s-1} to x^s. Let H^t be the set of round-t partial histories — $H^1 \equiv \left\{x^0\right\}$ being the null history — and let $H \equiv \bigcup_{t=1}^{\infty} H^t$ be the set of partial histories. Call a strategy 'semi-Markovian' if the proposals and voting behavior it prescribes at any (complete) round-t history only depend on the partial history, i.e. on the sequence of defaults and coalitions that changed defaults: $\left(x^0, S^1, x^1, \ldots, S^{t-1}, x^{t-1}\right)$.[20]

As in the case of stationary Markov strategies, we can now associate outcome functions with semi-Markovian strategies. Any semi-Markovian strategy σ generates an outcome function ϕ^σ, which assigns to every history $h \in H$ and every $k \in \{1, \ldots, m_{x^{t-1}}\}$ the unique final outcome $\phi^\sigma(h, k)$ eventually implemented (given σ) when h is the current partial history and the kth proposer is about to move. We are particularly interested in $\phi^\sigma\left(x^0, 1\right)$, which describes the policy implemented in $\Gamma\left(\pi, x^0\right)$ if players act according to σ. We will sometimes abuse notation and write $\phi^\sigma\left(x^0\right)$ instead of $\phi^\sigma\left(x^0, 1\right)$.

We now turn to the characterization of semi-Markovian equilibria — i.e., subgame perfect equilibria of $\Gamma\left(\pi, x^0\right)$ in which all players use semi-Markovian strategies. It turns out that the tree construction introduced in Section 3 can also be applied to consistent sets to obtain semi-Markovian equilibria. More specifically, if Z be a consistent set then each length m_{x^0} path of tree $\mathfrak{T}\left(Z, x^0\right)$ ending with a policy in Z describes behavior in round 1

[20]More formally, let $\mathbb{H}^t(h)$ be the set of (complete) round-t histories of $\Gamma\left(\pi, x^0\right)$ that share the same partial round-t history h. Thus, $\mathbb{H}^t \equiv \left\{\mathbb{H}^t(h)\right\}_{h \in H^t}$ is a partition of the set of round-t histories and $\mathbb{H} \equiv \bigcup_{t=1}^{\infty} \mathbb{H}^t$ a partition of the set of histories of $\Gamma\left(\pi, x^0\right)$. A strategy σ is semi-Markovian if it is measurable with respect to $\mathbb{H}$ — i.e., σ prescribes the same behavior at any two histories that belong to the same partition element $\mathbb{H}^t(h)$.

in some semi-Markovian equilibrium. Hence, there exists a semi-Markovian equilibrium σ in which a policy in Z is 'agreed on' immediately: if the initial default x^0 belongs to Z, it is implemented at the end of round 1; otherwise, it is amended to some policy in Z that is implemented at the end of round 2. Our next result mirrors Proposition 1.

Proposition 10. *Suppose that Z is the closure of a consistent set, and let $f \in Z^X$ be any selection of $F^\pi(Z,\cdot)$: $f(x) \in F^\pi(Z,x)$ for all $x \in X$. There exists a collection $\{\sigma_x\}_{x\in X}$ such that, for all $x \in X$, σ_x is a semi-Markovian equilibrium of $\Gamma(\pi,x)$ and $\phi^{\sigma_x}(x) = f(x)$. Hence, $\bigcup_{x\in X} \phi^{\sigma_x}(x) = Z$.*

The last part of the statement in the proposition says that, for any consistent set Z and initial default x^0, we can construct an equilibrium of $\Gamma\left(\pi,x^0\right)$, σ in which the final policy outcome reached from x^0 must belong to Z. Inspection of the proof (in the Appendix) reveals that more is true: the final policy outcome reached from any partial history $h \in H$ must belong to Z; so that $\phi^\sigma(H) \equiv \bigcup_{h\in H} \phi^\sigma(h,1) = Z$. The next result establishes that the converse is also true.

Proposition 11. *If σ is a semi-Markovian equilibrium then $\phi^\sigma(H) \equiv \bigcup_{h\in H} \phi^\sigma(h,1)$ must be a consistent set.*

Thus, for any semi-Makovian equilibrium, the set of policy outcomes that can be reached from all possible partial histories is a consistent set and, therefore, a subset of the largest consistent set $\overline{Z}$: $\phi^\sigma(H) \subseteq \overline{Z}$ for all semi-Markovian equilibria σ. Furthermore, we know from Proposition 10 that any policy $z \in \overline{Z}$ is the outcome of a semi-Markovian equilibrium of $\Gamma(\pi,z)$. Consequently, we have

Corollary 2. *Let $\Sigma^{NM}\left(\pi,x^0\right)$ be the set of semi-Markovian equilibria of $\Gamma\left(\pi,x^0\right)$. The set of all semi-Markovian equilibrium policy outcomes that can be obtained from any initial default in X coincides with the largest consistent set:*

$$\bigcup_{x^0\in X} \bigcup_{\sigma\in\Sigma^{NM}(\pi,x^0)} \phi^\sigma\left(x^0\right) = \overline{Z}\ .$$

Thus, the predictions of our noncooperative bargaining framework coincide with those of Chwe's (1994) largest consistent set when we use semi-Markovian strategies. This result provides noncooperative foundations for the largest consistent set, extending Proposition 8 in Acemoglu et al. (forthcoming) to non-acyclic preferences.

6 Conclusion

We have presented a model of bargaining in which the committee takes a single policy seriously at any time, and implements this policy if none of the proposers is willing or able to amend it. We have characterized the policies which can reached from any initial default, and shown that every equilibrium of the model supports a weakly stable set. We have provided conditions for a chair to manipulate the protocol, showing that she cannot improve on proposing last if the policy space is well ordered. We have also shown, inter alia, that an increase in the quota can contract the set of stable policies. In the remainder of this section, we will discuss some directions in which our model and our analysis could be extended:

6.1 Changing the model

Round the table bargaining

According to our model, there is a fixed protocol at every default, specifying the order in which proposers move. This assumption and the Markovian solution concept preclude a natural stopping rule: proposers sit round a table, and the first proposer in any new round sits next to the player who amended the previous default. This is inconsistent with our approach because we identify 'states' with 'defaults' when defining Markovian strategies. We could obtain analogous results for bargaining round the table with another definition of a state.

Multi-issue bargaining

We have supposed that a proposal must be a single policy. In some negotiations, it seems natural to suppose that players can provisionally agree to subsets of the policy space: e.g. when each dimension of the policy space represents an issue. Problems of this sort have been analyzed in the literature (cf. Winter (1997)) on the additional supposition that issues which have been agreed upon are no longer on the table. The history of the Oslo Process suggests that this supposition is problematic: no partial agreement is finalized until all issues have been addressed. An extension of our model could address this feature: proposals are subsets of the policy space, but the game can only end when a proposal which specifies a single point is agreed (and not amended).

Random proposers

We have assumed that the protocol at any default specifies the exact order in which players propose in a given round. This simplifies our stopping rule: a default is implemented if it is not amended by any proposer. If the identity of the next proposer were determined randomly (as in Baron and Ferejohn (1989)) then the stopping rule could require that each proposer have an opportunity to amend before the default is implemented. This variant might have different properties from our model when there are more than two proposers because players do not know the remaining protocol when the first proposer is selected, precluding the backward induction arguments which we have used so extensively.

6.2 Changing the solution concept

Mixed strategy equilibria

We have argued that every (pure strategy) equilibrium supports a weakly stable set. In cases like the Condorcet Paradox, there is no weakly stable set, and therefore no equilibrium. However, there may be equilibria if we allow for mixed strategy equilibria. Consider, for example, a symmetric version of the Condorcet Paradox with three policies and three proposers/voters, each of whom earn 0, 1 or 2 from any policy. We can show that there is a mixed strategy Markov perfect equilibrium in which each player proposes her top ranked policy, and a single voter mixes between accepting and rejecting each proposal. According to this equilibrium, each policy is equally likely to be implemented at any default. Play on the equilibrium path almost surely ends with implementation of some policy. (We provide further details in the last section of the Appendix.)

Refinements

Although weakly stable sets may not exist, simple games often have multiple weakly stable sets, implying equilibrium multiplicity in our noncooperative game. As policies in weakly stable sets are (weakly) Pareto efficient, commonly used refinements which are based on Pareto perfection and renegotiation-proofness have no bite in our bargaining game without discounting. Acemoglu et al (2009) have recently developed an equilibrium refinement concept for voting and agenda-setting games like ours: Markov Trembling Hand Perfect Equilibrium (MTHPE). A (stationary Markov) equilibrium σ is Markov trembling-hand perfect if and only if there is some sequence of totally mixed stationary Markov strategies $\left\{\sigma^k\right\}$ such that $\sigma^k \to \sigma$ and σ dictates each 'agent' — MTHPE is defined in the agent-strategic form — a best response to her opponents' perturbed strategies in σ^k for all $k = 1, \ldots, \infty$.

We conclude this paper with an observation, which shows that restricting attention to MTHPEs will typically not reduce the set of equilibrium outcomes in our game.

Observation 2. *If X is finite and well ordered then the set of (pure strategy) MTHPE policies coincides with the set of equilibrium policies (and is therefore the union of stable sets).*

The proof of this observation (which is provided in the last section of the Appendix) shows that something even stronger is true: *for every stable set* $V \in \mathcal{V}$*, there is a (pure strategy) MTHPE* σ *that supports* V: $f^\sigma = V$. This reinforces the noncooperative foundations our bargaining model provides for stable sets.

Appendix

Proof of Proposition 1

Let $V \in \mathcal{V}$, and let $f \in V^X$ be a selection of $F^\pi(V, \cdot)$. By construction of $F^\pi(V, x)$, for every $x \in X$, there exists a vector $(y_1(x), \ldots, y_{m+1}(x))$ such that:

- if $x \in V$, then $f(x) = y_1(x) = \ldots = y_{m+1}(x) = x$;
- if $x \notin V$, then $f(x) = y_1(x) \in V$, $x = y_{m+1}(x)$, and $y_k(x) \in \mathfrak{s}_k^\pi(V, y_{k+1}(x))$ for each $k = 1, \ldots, m$. The latter condition implies that $y_k(x)$ is one the kth proposer's ideal policies in a set $A_k(V, y_{k+1}(x)) \equiv P_V(y_{k+1}(x)) \cup \{y_{k+1}(x)\} \cup Y$, where $Y \subseteq R_V(y_{k+1}(x))$.

We now define the strategy profile $\sigma = (\sigma_i)_{i \in N}$. If the ongoing default is $x \in X$ then player $i = \pi_x(k)$ proposes $y_k(x)$ (if given the opportunity) with $y_k(x) = x$ being interpreted as 'pass'. Therefore, all proposers pass when the current default belongs to V.

When the ongoing default is x and the kth proposer has just proposed to change x to $y \neq x$, σ_i prescribes legislator i to vote 'yes' if and only if one of the following conditions hold:

(A) $x \in V$ and $y_1(y) \succ_i x$;

(B) $x \notin V$, $y_1(y) \in A_k(V, y_{k+1}(x))$, and $y_1(y) \succeq_i y_{k+1}(x)$;

(C) $x \notin V$, $y_1(y) \notin A_k(V, y_{k+1}(x))$, and $y_1(y) \succ_i y_{k+1}(x)$.

To prove the proposition, we proceed in three steps. The first step shows that $f^\sigma(x, 1) = f(x)$ for all $x \in X$. Step 2 shows that there is no voting stage in which a voter, say i, has a profitable one-shot deviation from σ_i. Step 3 demonstrates that there is no proposal stage in which a proposer, say j, has a profitable one-shot deviation from σ_j. Combined, Steps 2 and 3 show that no player has a profitable one-shot deviation from σ. This proves that no player can profitably deviate from σ in a finite number of stages. Finally, as infinite bargaining sequences constitute the worst outcomes for all players, this proves that σ is an equilibrium.

Step 1: $f^\sigma(x) \equiv f^\sigma(x, 1) = y_1(x)$ *for all* $x \in X$ *and, in particular, and* $f^\sigma(x) = x$ *for all* $x \in V$.

Consider an arbitrary round t starting with default $x^{t-1} = x$. If $x \in V$, then the result is trivial: all proposers pass and x is implemented at the end of the round. Suppose then that $x \notin V$. Let $l = \max\{k \in \{1, \dots, m_x\} : y_k(x) \neq y_{k+1}(x)\}$ (external stability with respect to R^π ensures that this set is nonempty), and suppose that the lth proposer is given the opportunity to make a proposal. By construction of $(y_k(x))$, this implies that $y_l(x) \in V$ and therefore $y_l(x) = y_1(y_l(x)) \in A_l(V, y_{l+1}(x))$. The definition of voting strategies (condition (B)) then implies that all members of $\{i \in N : y_l(x) \succeq_i y_{l+1}(x)\}$ vote 'yes', so that $y_l(x) = x^t$. As $x^t = y_l(x) \in V$, all proposers pass in round $t+1$ and $y_l(x)$ is implemented.

Now consider the $(l-1)$th proposer. Suppose that she is given the opportunity to make a proposal. By definition of proposal strategies, she must propose $y_{l-1}(x) \in A_{l-1}(V, y_l(x))$. By definition of $A_{l-1}(V, y_l(x))$, this implies that $\{i \in N : y_{l-1}(x) \succeq_i y_l(x)\} \in \mathbf{W}$. Therefore, condition (B) implies that $y_{l-1}(x)$ is accepted and implemented at the end of the next round.

Repeating this argument recursively for every $l = 1, \dots, l-2 \dots$, we obtain that $f^\sigma(x, 1) = y_1(x)$. This proves that $f^\sigma(x, 1) = f(x)$ for all $x \in X$.

Step 2: Consider a proposal y by the kth proposer when the ongoing default is $x \neq y$. σ_i prescribes $i \in N$ to vote 'yes' whenever $f^\sigma(y, 1) \succ_i f^\sigma(x, k+1)$, and to vote 'no' whenever $f^\sigma(x, k+1) \succ_i f^\sigma(y, 1)$.

From Step 1, we know that $f^\sigma(y, 1) = y_1(y) \in V$.

If $x \in V$, then it will be implemented at the end of round t if the kth proposer fails to amend it: by definition of the proposal strategies, all the remaining proposers will pass. Hence, $f^\sigma(x, k+1) = x$. As a consequence, $f^\sigma(y, 1) \succ_i f^\sigma(x, k+1)$ implies $y_1(y) \succ_i x$, which in turn implies that i must vote 'yes' (condition (A) in the definition of voting strategies). Similarly, $f^\sigma(x, k+1) \succ_i f^\sigma(y, 1)$ implies $x \succ_i y_1(y)$. Hence, i must vote 'no'.

If $x \notin V$, then $f^\sigma(x, k+1) = y_{k+1}(x)$. To see this, suppose first that no proposer $l > k$ amends x. We then have $y_{k+1}(x) = \dots = y_m(x) = x = f^\sigma(x, k+1)$. Now suppose that the lth proposer is the next proposer (after the kth) who makes a successful proposal, $y_l(x) \neq x$. By construction, this implies that $y_{k+1}(x) = \dots = y_l(x) \in V$. As a consequence, $f^\sigma(x, k+1) = f^\sigma(y_l(x), 1) = y_l(x) = y_{k+1}(x)$.

Thus, $f^\sigma(y, 1) \succ_i f^\sigma(x, k+1)$ implies that $y_1(y) \succ_i y_{k+1}(x)$. Conditions (B) and (C) in the definition of voting strategies then imply that legislator i votes "yes." Similarly, $f^\sigma(x, k+1) \succ_i f^\sigma(y, 1)$ implies she votes 'no'.

Step 3: In any proposal stage with ongoing default x, the kth proposer cannot gain by offering some $y \neq y_k(x)$ and conforming to $\sigma_{\pi_x(k)}$ thereafter.

Suppose first that $x \in V$. In such a case, $\sigma_{\pi_x(k)}$ prescribes the kth proposer to pass (i.e., $y_k(x) = x$). If she has a profitable deviation at this stage, then she must be able to successfully make some proposal $y \neq x$ such that $f^\sigma(y,1) = y_1(y) \succ_{\pi_x(k)} x$. Indeed, if she does not deviate then all the remaining proposers will pass ($y_l(x) = x$ for all l) and x will then be the final outcome. Nevertheless, if proposal y is successful, then condition (A) in the definition of voting strategies implies that there is a winning coalition whose members all strictly prefer $y_1(y) \in V$ to $x \in V$; a contradiction with V satisfying (IS$_P$).

Now suppose that $x \notin V$. In such a case, $\sigma_{\pi_x(k)}$ prescribes the kth proposer to propose $y_k(x) \in A_k^V(y_{k+1}(x))$ (where $y_k(x) = x$ means that she should pass). Suppose that, instead, she chooses some $y \neq y_k(x)$. The resulting outcome will be $f^\sigma(y,1) = y_1(y)$ if y is a successful proposal (i.e.: $y_1(y) \in R_V(y_{k+1}(x))$), or $f^\sigma(y,1) = y_{k+1}(x)$ if she passes or y is an unsuccessful proposal. Such a deviation, however, cannot be profitable. Indeed, $y_k(x)$ is by definition $\succeq_{\pi_x(k)}$-maximal in $[R_V(y_{k+1}(x)) \cup \{y_{k+1}(x)\}]$.

Proof of Proposition 2

The proof of Proposition 2 hinges on the following lemma.

Lemma 1. *If σ is an equilibrium of $\Gamma(\pi, x^0)$ then $f^\sigma(X) \equiv \bigcup_{x\in X} f^\sigma(x)$ is a weakly stable set.*

Proof: Let σ be an equilibrium of $\Gamma(\pi, x^0)$. To prove the lemma, we must show that $f^\sigma(X)$ satisfies (IS$_P$) and (ES$_R$).

(IS$_P$). If $|f^\sigma(X)| = 1$ then P-internal stability is trivial; so suppose that $|f^\sigma(X)| \geq 2$. Imagine that $f^\sigma(X)$ does not satisfy (IS$_P$). This implies that there are two policies in $f^\sigma(X)$, say x and y, such that xPy. By definition of $f^\sigma(X)$, x and y are fixed points of $f^\sigma(\cdot,1)$. An immediate consequence of xPy is therefore that there is a winning coalition $S \in \mathcal{W}$ such that $f^\sigma(x,1) \succ_i f^\sigma(y,1)$ for every $i \in S$. But this implies that any proposer in S could successfully propose to amend y to x; a contradiction with σ being an equilibrium of $\Gamma(\pi, x^0)$.

(ES$_R$). Suppose that $f^\sigma(X)$ does not satisfy (ES$_R$). This implies that there exists a policy $x \notin f^\sigma(X)$ such that, for all $y \in f^\sigma(X)$, $\neg(yRx)$. In particular, $\neg[f^\sigma(y,1)Rx]$ for all $y \in f^\sigma(X)$. Consequently, in any $S \in \mathcal{W}$ and any $y \in f^\sigma(X)$, there is at least one player who strictly prefers x to $f^\sigma(y,1)$.

Now consider the continuation game that starts with x as the ongoing default policy. Suppose the last potential proposer, $\pi_x(m_x)$, is given the opportunity to amend x with some policy $y \neq x$. Players anticipate that $f^\sigma(y,1) \in f^\sigma(X)$ will eventually be implemented if x is amended, and that x will be implemented otherwise. As no winning coalition including proposer $\pi_x(m_x)$ would support the amendment, x should be implemented. As a consequence, another proposer must amend x in equilibrium.

Now consider $\pi_x(m_x - 1)$. We can repeat the same reasoning as with $\pi_x(m_x)$. If $\pi_x(m_x - 1)$ offers to change x to some policy $y \neq x$, all committee members will anticipate that this will lead to $f^\sigma(y, 1)$ being the final outcome if the amendment is voted up, and to x being implemented otherwise. Again, no winning coalition would support the amendment and x would be implemented. Repeating this argument recursively until committee proposer $\pi_x(1)$, we obtain the desired contradiction.

□

We now return to the main proposition. Let $\sigma = (\sigma_i)_{i \in N}$ be an equilibrium of $\Gamma(\pi, x^0)$. From Lemma 1, we know that there exists $V \in \mathcal{V}$ such that $f^\sigma(X) = V$. Evidently, for all $x \in V$, we have $\{f^\sigma(x, 1)\} = \{x\} = F^\pi(V, x)$.

Now consider an arbitrary $x \notin V$, and an arbitrary round starting with x as the ongoing default. Suppose that (possibly off the equilibrium path) the m_xth proposer is given the opportunity to amend x. When she offers a policy $y \neq x$, voters compare $f^\sigma(y, 1) \in V$ with x. Therefore, voter i must vote 'yes' if $f^\sigma(y, 1) \succ_i x$, may vote either 'yes' or 'no' if $f^\sigma(y, 1) \sim_i x$, and must vote 'no' otherwise. The acceptance set faced by the m_xth proposer — i.e., the set of policies $y \neq x$ that would be accepted by a winning coalition to amend x — must then be the set of policies y such that $f^\sigma(y, 1)$ belongs to $[P_V(x) \cup Y] \subseteq V$, where Y is some (possibly empty) subset of $R_V(x)$. As a consequence, if $\sigma_{\pi_x(m_x)}$ prescribes the m_xth proposer to amend x with $\bar{y}_{m_x} \neq x$, then $f^\sigma(x, m_x) = f^\sigma(\bar{y}_{m_x}, 1)$ must be $\succeq_{\pi_x(m_x)}$-maximal in $[P_V(x) \cup Y \cup \{x\}]$ (x is always feasible to the m_xth proposer, for she can always pass). If $\sigma_{\pi_x(m_x)}$ prescribes the m_xth proposer not to amend x — i.e. to pass or to make an unsuccessful proposal — then $f^\sigma(x, m_x) = x$ must be $\succeq_{\pi_x(m_x)}$-maximal in $[P_V(x) \cup Y \cup \{x\}]$. This proves that $y_{m_x} \equiv f^\sigma(x, m_x) \in \mathfrak{s}^\pi_m(V, x)$.

Proceeding recursively, one can use the same argument to show that, for each $k = 1, \ldots, m_x - 1$, $y_k \equiv f^\sigma(x, k) \in \mathfrak{s}^\pi_k(V, x)$: just substitute y_{k+1} for x in the argument above. Since $x \notin V = f^\sigma(X)$, there must be some proposer k who amends x, so that $f^\sigma(x, k) \neq x$. This proves that the finite sequence $(y_1, \ldots, y_{m_x}, x) \equiv (f^\sigma(x, 1), \ldots, f^\sigma(x, m_x), x)$ constitutes a path of tree $\mathfrak{T}^\pi(V, x)$ whose terminal node belongs to V. Hence, $f^\sigma(x) \in F^\pi(V, x)$.

Proof of Proposition 3

Let σ be an equilibrium. Proposition 2 implies that there must be some weakly stable set V such that $f^\sigma(X) = V$. Consider the kth proposer as defined in the statement of the proposition. If she failed to amend the ongoing default x then nobody else would, and x would be implemented at the end of the round. As she strictly prefers her ideal policy in the set of equilibrium policy outcomes that dominate x — i.e.: $\mathbf{M}\left(\succ_{\pi(l)}, R(x) \cap f^{\sigma}(X)\right)$ — to x, she must successfully propose that policy in equilibrium.

We therefore need to show that no proposer who is given the opportunity to amend x before the kth proposer can successfully do so. Suppose first that the $(k-1)$th proposer successfully offers some policy y. This implies there is a winning coalition in $\mathbf{W}$ whose members all strictly prefer $f^{\sigma}(y,1) \in V$ to $\mathbf{M}\left(\succ_{\pi(l)}, R(x) \cap f^{\sigma}(X)\right) \in V$: a contradiction with V satisfying (IS_P). Applying this argument recursively from the $(k-2)$th proposer until the first, we obtain the result.

Proof of Proposition 4

(Necessity) Suppose that there exists some $V \in \mathcal{V}$, a policy $x \notin V$, and a pair of proposers (i,j) in $M(V,x)$ such that (1) does not hold. Let π_i and π_j be two elements of Π such that $\pi_i(m_x) = i$ and $\pi_j(m_x) = j$. As $V \in \mathcal{V}$ for both orders, Proposition 1 implies that there exist equilibria σ_i and σ_j of $\Gamma(\pi_i, x)$ and $\Gamma(\pi_j, x)$ respectively, such that $V = f^{\sigma_i}(X) = f^{\sigma_j}(X)$. (Since $X \supseteq V$ is finite, the closure of V is V itself.) As i and j are both amenders of x in V, Proposition 3 implies that

$$F^{\pi_i}(V,x) = f^{\bar{\sigma}_i}(x) = M\left(P_i, R(x) \cap V\right) \neq M\left(P_j, R(x) \cap V\right) = f^{\bar{\sigma}_j}(x) = F^{\pi_j}(V,x) \ ,$$

for all $\bar{\sigma}_i \in \Sigma_V^*(\pi_i, x)$ and $\bar{\sigma}_j \in \Sigma_V^*(\pi_j, x)$ such that $V = f^{\bar{\sigma}_i}(X) = f^{\bar{\sigma}_j}(X)$ — the inequality comes from (1) not being satisfied. This proves that order independence is not satisfied.

(Sufficiency) Let $V \in \mathcal{V}$ and $x \notin V$. We know from Proposition 3 that, for all $\pi' \in \Pi$ and all $\sigma \in \Sigma_V^*(\pi', x)$, the policy implemented when the default is x, $f^{\sigma}(x)$, must be the ideal policy of the last amender of x within $R(x) \cap V$. But condition (1) implies that all amenders have the same ideal policy in $R(x) \cap V$. As a consequence, whoever the last amender is, and therefore whatever the order of proposers $\pi' \in \Pi$, the outcome must always be this common ideal policy. Hence

$$\sigma_1 \in \Sigma_V^*(\pi_1, x) \ \& \ \sigma_2 \in \Sigma_V^*(\pi_2) \Longrightarrow f^{\sigma_1}(x) = F^{\pi_1}(V,x) = F^{\pi_2}(V,x) = g^{\sigma_2}(x)$$

for all $\pi_1, \pi_2 \in \Pi$. This condition is also true when $x \in V$ since, in such a case, $f^{\sigma}(x) = F^{\pi'}(V,x) = x$ for all $\pi' \in \Pi$ and all $\sigma \in \Sigma_V^*(\pi', x)$.

Proof of Proposition 5

By definition of Π, the same players propose at every default $x \in X$ for each protocol $\pi \in \Pi$; so the sets of winning coalitions, $\mathcal{W}$, and weakly stable sets, $\mathcal{V}$, are each constant across Π. Inspection of the proof reveals that Proposition 2 holds irrespective of the protocols π in Π. In other words, for any $\pi \in \Pi$ and any equilibrium σ of $\Gamma\left(\pi, x^0\right)$, $f^\sigma(X)$ is a weakly stable set.

Now fix some weakly stable set V and some initial default $x^0 \notin V$, and consider any pair π^1 and π^2 in Π such that $\pi^1_{x^0} = \pi^2_{x^0} \equiv \pi_{x^0}$. Write x^0_k for any $\succeq_k$-maximal policy in $\left[V \cap R\left(x^0\right)\right] \cup \left\{x^0\right\}$, and define l as the last player in π_{x^0} such that $x^0_k = x^0$ for every $k > l$. Any player $k > l$ would pass at x^0 because she weakly prefers the default over any policy in V which weakly dominates x^0. This property does not depend on the protocol at other defaults; so if player k passes at x^0 for protocol π^1 then she also passes at x^0 for protocol π^2. Player l would amend x^0, with the result that x^0_l is implemented, because she rightly believes that players $k > l$ would otherwise pass, x^0_l is $\succeq_l$-maximal policy in $\left[V \cap R\left(x^0\right)\right] \cup \left\{x^0\right\}$, and members of a winning coalition weakly prefer x^0_l over x^0. These properties do not depend on the protocol at other defaults; so if player l amends x^0 such that x^0_l is eventually implemented for protocol π^1 then x^0_l is also eventually implemented when player l proposes at x^0 for protocol π^2. This observation implies that the same policies are implemented in subgames which start with players $k < l$ proposing at x^0 with protocols π^1 and π^2. The result then follows by taking the union of equilibrium policies, $f^\sigma(X)$, over every equilibrium of $\Gamma\left(\pi^i, x^0\right)$ for $i = 1$ and 2 (and applying Corollary 1).

Proof of Proposition 8

Let $V \in \mathcal{V}$. For every $x \notin V$, let K_x be the last proposer according to π who would be willing to amend x if given the opportunity; that is:

$$K_x \equiv \max\left\{k \in \{1, \ldots, m_x\} : v \succ_{\pi(k)} x \text{ for some } v \in R_V(x)\right\} \ .$$

Since V is R-externally stable, K_x is well-defined for all $x \notin V$. Next, let the function $\upsilon \in V^X$ be defined by: (i) for all $x \in V$, $\upsilon(x) = x$, and (ii) for all $x \notin V$, $\upsilon(x) = \bar{\upsilon}$, where $\bar{\upsilon}$ is $\succ_{\pi(K_x)}$-maximal in $R_V(x)$ for all $x \notin V$. That is, $\upsilon(x)$ is equal to x whenever x belongs to V, and to some of the K_xth proposer ideal policies in $R_V(x)$ otherwise.

We now define σ as follows. If the ongoing default is $x \in X$, player $i = \pi(k)$ proposes $\upsilon(x)$ (with $\upsilon(x) = x$ interpreted as 'pass') if $k \geq K_x$, and passes otherwise. Therefore, all proposers pass when the current default belongs to V.

When the ongoing default is x and the kth proposer has just proposed to change x to $y \neq x$, σ_i prescribes voter i to vote 'yes' if and only if one of the following conditions holds:

(A) $x \in V$ and $u_i\left(\upsilon(y)\right) > u_i(x)$;

(B) $x \notin V$, $k \geq K_x$, and $u_i\left(\upsilon(y)\right) \geq u_i\left(x\right)$;

(C) $x \notin V$, $k < K_x$, and $u_i\left(\upsilon(y)\right) > u_i\left(\upsilon(x)\right)$.

Applying the same argument as in the proof of Proposition 1 above, it is readily checked that $\sigma \in \Sigma_V^*(\pi)$.[21]

Now fix $\varepsilon > 0$ and suppose that

$$\delta \geq \delta_\varepsilon \equiv \frac{\Delta - \varepsilon}{\Delta} \in (0, 1) \ ,$$

where $\Delta \equiv \max_{i \in N} \max_{(x,y) \in X^2} \left[u_i(x) - u_i(y)\right]$. To prove the result, we proceed in two steps.

Step 1: In any voting stage, each voter i's continuation payoffs after any profitable one-shot deviation from σ_i must be within ε of her continuation payoff from conforming to σ_i.

Note first that, given σ, the policy implemented in any round starting with ongoing default x is $\upsilon(x)$ for all $x \in X$: If $x \in V$, then all proposers pass and, therefore, $\upsilon(x) = x$ is implemented at the end of the round; if $x \notin X$, then the proposals by the first $(K_x - 1)$ proposers are voted down (by condition (C) in the definition of voting strategies and internal stability), the K_xth proposer offers $\upsilon(x)$ which is voted up (condition (B)), and all the remaining proposers pass. Once $\upsilon(x) \in V$ has been implemented, it is never amended and is therefore implemented in all future periods. This implies that continuation payoff to player i under the strategy profile σ at the start of any round with default x is $u_i\left(\upsilon(x)\right)$.

Suppose first that σ_i prescribes voter i to vote 'yes' when the kth proposer offers $y \neq x$. Voter i profitably deviates by playing 'no' (and incurring an extra cost of ε) if and only if

$$(1 - \delta)u_i(y) + \delta u_i\left(\upsilon(y)\right) \geq (1 - \delta)u_i(x) + \delta u_i\left(\upsilon(x)\right) - \varepsilon \ .$$

The definition of voting strategies implies that $u_i\left(\upsilon(y)\right) \geq u_i\left(\upsilon(x)\right)$ whenever σ_i prescribes i to vote 'yes'. Consequently, the inequality above is satisfied whenever

$$(1 - \delta)\left[u_i(x) - u_i(y)\right] \leq \varepsilon \ ,$$

which must be true since $\delta > \delta_\varepsilon$.

Now suppose that σ_i prescribes player i to vote 'no' when the kth proposer offers $y \neq x$. Voter i cannot profitably deviate by voting 'yes' (and incurring an extra cost of ε) if and only if

$$(1 - \delta)u_i(x) + \delta u_i\left(\upsilon(x)\right) \geq (1 - \delta)u_i(y) + \delta u_i\left(\upsilon(y)\right) - \varepsilon \ .$$

[21] Just use the following sequence $(y_1(x), \ldots, y_{m+1}(x))$ for each $x \in X$: (i) if $x \in V$, then $y_k(x) = x$ for each k, and (ii) if $x \notin V$, then $y_k(x) = x$ for each $k \leq K_x$, and $y_k(x) = x$ for each $k > K_x$.

The definition of voting strategies implies that $u_i(\upsilon(y)) \leq u_i(\upsilon(x))$ whenever σ_i prescribes i to vote 'no'. Consequently, the inequality above is satisfied whenever

$$(1-\delta)[u_i(y) - u_i(x)] \leq \varepsilon \ ,$$

which must be true since $\delta > \delta_\varepsilon$.

Step 2: No proposer has a profitable deviation.

Let $k \in \pi^{-1}(i)$; that is, i is the kth proposer. Suppose, first, that $x \in V$. In such a case, σ_i prescribes i to pass. If she has a profitable deviation at this stage, then she must be able to amend x with some proposal $y \neq x$. Indeed if she does not deviate then all the remaining proposers will pass and x will then be the final outcome. Nevertheless, if proposal y is successful, then condition (A) in the definition of voting strategies implies that there is a winning coalition whose members all strictly prefer $\upsilon(y) \in V$ to $x \in V$; a contradiction with V satisfying (IS_P).

Now suppose that $x \notin V$. If $k > K_x$ then i has evidently no profitable deviation. To profitably deviate, she would have to amend x with some proposal $y \neq x$. But such a proposal can only be successful if $y \in R_V(x)$; and $K_x < k \in \pi^{-1}(i)$ implies that i strictly prefers x to any point in $R_V(x)$.

If $k = K_x$ then i can deviate either by proposing $y \neq \upsilon(x)$ or by passing. Suppose, first, that she proposes $y \neq \upsilon(x)$. Such a deviation cannot be profitable since $\upsilon(x)$ is by definition i's ideal policy in $R_V(x)$ and $u_i(\upsilon(x)) \geq u_i(x)$: if proposal y is successful, then it must belong to $R_V(x) \subseteq V$ (so that $\upsilon(y) = y$) and therefore cannot improve upon $\upsilon(x)$; if proposal y is unsuccessful, then the payoff from the deviation would be $(1-\delta)u_i(x) + \delta u_i(\upsilon(x))$ which must be lower than $u_i[\upsilon(x)] \geq u_i(x)$.

Finally, if $k < K_x$, any proposal y by i which she strictly prefers to $\upsilon(x)$, is voted down (condition (C) in the definition of voting strategies and P-internal stability of V.

Proof of Proposition 9

We first construct $\bar{\delta}$. For each $i \in N$ and every pair $(x, y) \in X^2$ such that $u_i(x) > u_i(y)$, let

$$\Delta_i(x, y, \delta) \equiv \min_{T_x, T_y \in \{1,\dots,|X|\}} \delta^{T_x} u_i(x) + \left(1 - \delta^{T_x}\right) \underline{u}_i - \delta^{T_y} u_i(y) - \left(1 - \delta^{T_y}\right) \bar{u}_i \ ,$$

where $\bar{u}_i \equiv \max_{x \in X} u_i(x)$ and $\underline{u}_i \equiv \min_{x \in X} u_i(x)$. Since $\Delta_i(x, y, \delta) \to u_i(x) - u_i(y) > 0$ as $\delta \to 1$, $\delta_i(x, y) \equiv \min\{d \geq 0 : \Delta_i(x, y, \delta) \geq 0\}$ is well-defined. From now on, we assume that

$$\delta > \bar{\delta} \equiv \max_{i \in N} \max_{x, y \in X : x \succ_i y} \delta_i(x, y) \in (0, 1) \ .$$

Suppose, first, that σ is an equilibrium of $\Gamma\left(\pi, x^0\right)$. This implies that, at any stage of this game, no player i has a profitable one-shot deviation from σ_i (given σ_{-i}). Consider an arbitrary stage of $\Gamma\left(\pi, x^0\right)$, and let x be the final policy outcome if i does not deviate from σ_i in that stage. Hence, any other policy outcome $y \neq x$ she could induce by a one-shot deviation satisfies: $u_i(y) < u_i(x)$. Suppose that, contrary to the statement of the result, i has a profitable one-shot deviation at the same stage in $\Gamma^\delta\left(\pi, x^0\right)$. This implies that there are two finite sequences $\{x_t\}_{t=1,\ldots,T_x}$ and $\{y_t\}_{t=1,\ldots,T_y}$, and a policy $y \in X$ such that

$$(1-\delta)\sum_{t=1}^{T_y}\delta^{t-1}u_i\left(y_t\right) + \delta^{T_y}u_i(y) > (1-\delta)\sum_{t=1}^{T_x}\delta^{t-1}u_i\left(x_t\right) + \delta^{T_x}u_i(x)$$

and $u_i(y) < u_i(x)$ (recall that a one-stage deviation from an equilibrium strategy in $\Gamma\left(\pi, x^0\right)$ must converge in a finite number of rounds). This is impossible when $\delta > \bar{\delta}$. By the one-shot deviation principle, σ is then an absorbing stationary Markov equilibrium of $\Gamma^\delta\left(\pi, x^0\right)$.

Now suppose that σ is an absorbing stationary Markov equilibrium of $\Gamma^\delta\left(\pi, x^0\right)$. This implies that no player i has a profitable one-shot deviation from σ_i (given σ_{-i}) at any stage of this game. Consider an arbitrary stage of $\Gamma^\delta\left(\pi, x^0\right)$, and let $\{x_t\}_{t=1,\ldots,T_x+1}$ be the finite sequence of policy outcomes (with $x = x_{T_x+1}$ being implemented indefinitely) if i does not deviate from σ_i at that stage. Hence, any other sequence $\{y_t\}_{t=1,\ldots,T_y+1}$ (with $y = y_{T_y+1}$ being implemented indefinitely) she could induce by a one-shot deviation satisfies:

$$(1-\delta)\sum_{t=1}^{T_y}\delta^{t-1}u_i\left(y_t\right) + \delta^{T_y}u_i\left(y\right) \leq (1-\delta)\sum_{t=1}^{T_x}\delta^{t-1}u_i\left(x_t\right) + \delta^{T_x}u_i(x)\ .$$

This inequality implies that $u_i(x) > u_i(y)$. To see this, suppose instead that $u_i(y) > u_i(x)$. $\delta > \bar{\delta}$ then implies that $\Delta_i(y, x, \delta) > 0$, so that

$$(1-\delta)\sum_{t=1}^{T_y}\delta^{t-1}u_i\left(y_t\right) + \delta^{T_y}u_i\left(y\right) - \left[(1-\delta)\sum_{t=1}^{T_x}\delta^{t-1}u_i\left(x_t\right) - \delta^{T_x}u_i(x)\right] \geq \Delta_i(y, x, \delta) > 0\ ;$$

a contradiction. At the equivalent stage in game $\Gamma\left(\pi, x^0\right)$, $u_i(x) > u_i(y)$ clearly implies that player i has no profitable one-shot deviation in this stage. This in turn implies that player i cannot profitably deviate from σ_i in a finite number of stages. Finally, as infinite bargaining sequences constitute the worst outcomes for all legislators in $\Gamma\left(\pi, x^0\right)$, this proves that σ is an equilibrium of $\Gamma\left(\pi, x^0\right)$.

Proof of Proposition 10

The first part of the proof puts in place some mathematical machinery that will be handy when we come to construct the equilibrium σ. In what follows, we will indulge in a slight abuse of terminology by referring to partial histories as 'histories'.

Let Z be the closure of a consistent set, and let $f \in Z^X$ be a selection of $F^\pi(Z, \cdot)$. We will use a sequence $(\tau_l) \in (\mathbb{N} \cup \{\varnothing\})^\infty$ to construct σ. For a given history,[22] each element of this sequence must be thought of as a round in which players (both proposers and voters) changed the default in accordance with σ. Given a round-t history $h \in H^t$, we define the sequence (τ_l) and proposal strategies as follows:

• $l = 1$: τ_1 is the first round in which an element of Z became the new default; if that has not happened so far, then we write $\tau_1 = \varnothing$ and say that $h \in H_1$. That is, H_1 is the set of histories in H at which no element of Z has ever been offered and accepted. (Note that this was the case at the start of round τ_1, so that $h \in H_1$ when $t = \tau_1$.)

We now define proposal strategies at any history $h \in H_1$. Let $x = x^{t-1}$ be the ongoing default at history h. By construction of H_1, therefore, $x \notin Z$. From the construction of $F^\pi(Z, x)$, there exists a vector $(z_1(h), \ldots, z_{m_x+1}(h))$ such that: $f(x) = z_1(h) \in Z$, $x = z_{m_x+1}(h)$, and $z_k(h) \in s_k^\pi(Z, z_{k+1}(h))$ for each $k = 1, \ldots, m_x$. The latter condition implies that $z_k(h)$ is one of the kth proposer's ideal policies in a set $A_k(Z, z_{k+1}(h)) \equiv P_Z(z_{k+1}(h)) \cup \{z_{k+1}(h)\} \cup Y_k(h)$, where $Y_k(h) \subseteq R_Z(z_{k+1}(h))$.

If $h \in H_1$ then σ_i prescribes player $i = \pi_x(k)$ to propose $z_k(h)$ if $z_k(h) \neq z_{k+1}(h)$, and to pass if $z_k(h) = z_{k+1}(h)$.

• $l \geq 2$: τ_l is the first round after τ_{l-1} at which an element of

$$Z_l \equiv \left\{z \in Z : x^{\tau_{l-1}} \succeq_i z \text{ for some } i \in S^{\tau_l+1}\right\}$$

became the new default; if that has not happened so far then we let $\tau_l = \varnothing$. In particular, if $\tau_l = \varnothing \neq \tau_{l-1}$ then we write $h \in H_l$. By definition of τ_{l-1}, $x^{\tau_{l-1}} \in Z$. Since Z is consistent, $Z_l \cap \{z\ inZ : zRx\}$ is nonempty for all $x \in X$. Using the tree $\mathfrak{T}^\pi\left(Z_l, x^{t-1}\right)$, we can then obtain a vector $(y_1(h), \ldots, y_{m_x+1}(h))$ such that: $x = z_{m_x+1}(h)$, and $z_k(h) \in s_k^\pi(Z_l, z_{k+1}(h))$ for each $k = 1, \ldots, m_x$. The latter condition implies that $z_k(h)$ is one of the kth proposer's ideal policies in a set $A_k(Z_l, z_{k+1}(h)) \equiv P_{Z_l}(z_{k+1}(h)) \cup \{z_{k+1}(h)\} \cup Y_k(h)$, where $Y_k(h) \subseteq R_{Z_l}(z_{k+1}(h))$.

If $h \in H_l$ then σ_i prescribes player $i = \pi_x(k)$ to propose $z_k(h)$ if $z_k(h) \neq z_{k+1}(h)$, and to pass if $z_k(h) = z_{k+1}(h)$. The idea behind this construction is that the kth proposer tries to "punish" at least one of the "deviators" in S^{τ_l+1} for not rejecting the k^{τ_l+1}th proposer's offer to amend x^{τ_l}.

So far, we have been silent about proposals at period-t histories such that $t = \tau_l + 1$ (so that $x^{t-1} = x^{\tau_l}$). We denote the set of such histories by H_0. At any history $h \in H_0$,

[22] To lighten the notation, we voluntarily omit the sequence's dependence on the history under consideration.

the ongoing default should be implemented: σ_i prescribes player $i = \pi_x(k)$ to pass. For expositional convenience, we will sometimes say that i proposes $z_k(h) = x$. Since $\{H_l\}_{l=0}^{\infty}$ is a partition of H, the description of proposal strategies is complete.

We now turn to voting strategies. At a round-t history $h \in H^t$, following a proposal $y \neq x^{t-1}$ by the kth proposer, σ_i prescribes voter i to act as follows:

(A) If $h \in H_0$ (i.e.: $t = \tau_l + 1$ for some $l \in \mathbb{N}$) then i votes 'yes' iff $z_1(h, k, S, y) \succ_i x^{t-1}$ for any winning coalition $S \ni i$;

(B) if $h \in H_l$ (i.e.: $\tau_{l-1} + 1 < t \leq \tau_l$), $l \neq 0$, and $y \in A_k(Z_l, z_{k+1}(h))$ then i votes 'yes' iff $y \succeq_i z_{k+1}(h)$;

(C) if $h \in H_l$ (i.e.: $\tau_{l-1} + 1 < t \leq \tau_l$), $l \neq 0$, and $y \notin A_k(Z_l, z_{k+1}(h))$ then i votes 'yes' iff $z_1(h, k, S, y) \succ_i z_{k+1}(h)$ for any winning coalition $S \ni i$;

where $Z_1 \equiv Z$.

We establish the statement of Proposition 10 via a series of claims. The first two claims provide useful characterization results about equilibrium policy outcomes. Claim 3 shows that $f^{\sigma}(x) = f(x)$ for all $x \in X$. Claim 4 shows that there is no voting stage in which a voter, say i, has a profitable one-shot deviation from σ_i. Claim 5 demonstrates that there is no proposal stage in which a proposer, say j, has a profitable one-shot deviation from σ_j. Claims 4 and 5 jointly show that no voter has a profitable one-shot deviation from σ. This proves that no player can profitably deviate from σ in a finite number of stages. Finally, as infinite bargaining sequences constitute the worst outcomes for all legislators, this proves that σ is an equilibrium.

Claim 1: Consider the round following a history $h \in H$, and suppose the kth proposer has just moved. If she has made no proposal or if her proposal is rejected, then the final outcome will be $z_{k+1}(h)$.

Proof: If $h \in H_0$, then the claim is trivial: $z_{k+1}(h) = \ldots = z_{m+1}(h) = x^{t-1}$ (all the remaining proposers pass). Accordingly, suppose that $h \in H_l$ with $l \neq 0$. Since the kth proposer has not amended x, the $(k+1)$th proposer is given the opportunity to make a proposal. By definition of proposal strategies, she proposes $z_{k+1}(h)$ if $z_{k+1}(h) \neq z_{k+2}(h)$, and passes otherwise. If $z_{k+1}(h) \neq z_{k+2}(h)$ then $z_{k+1}(h) R z_{k+2}(h)$. Condition (B) in the definition of voting strategies then ensures that proposal $z_{k+1}(h) \in A_k(Z_l, z_{k+1}(h))$ is accepted. As a consequence, the history at the start of the next round belongs to H_0, so that all proposers pass and $z_{k+1}(h)$ is implemented at the end of that round.

If $z_{k+1}(h) = z_{k+2}(h)$ then the $(k+2)$th proposer is given the opportunity to make a proposal. We can apply the same argument as above to show that either $z_{k+1}(h) = z_{k+2}(h)$

($\neq z_{k+3}(h)$) is implemented in the next round or $z_{k+1}(h) = z_{k+2}(h) = z_{k+3}(h)$. Going on until the m_xth proposer, we obtain the claim.

Claim 2: Let $\phi^\sigma(h;k)$ be the unique final outcome eventually enacted (given σ) when, after history $h \in H$, the kth proposer is about to move. For all $h \in H$, $\phi^\sigma(h;k) = z_k(h)$. In particular, if $h \in H_0$ then $\phi^\sigma(h;k) = z_k(h) = x^{t-1}$.

Proof: If $z_k(h) \neq z_{k+1}(h)$ then $z_k(h) \in Z_l$. Condition (B) in the definition of voting strategies then ensures that the kth proposer's offer, $z_k(h) \in A_k(Z_l, z_{k+1}(h))$, is accepted. Therefore, the history at the start of the next round belongs to H_0, so that all proposers pass and $z_k(h)$ is implemented at the end of that round.

If $z_k(h) = z_{k+1}(h)$ then, by definition of proposal strategies, the kth proposer passes. From Claim 1, $z_k(h) = z_{k+1}(h)$ is then the final outcome.

Claim 3: $f^\sigma(x^0) = z_1(x^0) = f(x^0)$ for all $x^0 \in X = H^1$.

Proof: Suppose first that the initial default (x^0) is an element of Z: viz. $z_k(x^0) = x^0$ for any proposer k. No proposer then offers to amend x^0, which is implemented at the end of round 1: $f^\sigma(x^0) = x^0 = z_1(x^0) = f(x^0)$.

Now suppose that x^0 is not a member of Z, so that $x^0 \in H_1$. Since $z_1(x^0) = f(x^0) \in F^\pi(Z, x^0) \subseteq Z$, at least one proposer tries to amend x^0. The first proposer who does so, say $\pi_{x^0}(k)$, offers $z_k(x^0) R z_{k+1}(x^0)$ which, by condition (B) in the definition of voting strategies, is accepted. This implies that $\tau_1 = 1$, which in turn implies that $z_k(x^0)$ is never amended and is therefore implemented at the end of round 2. By definition of proposal strategies, $z_l(x^0) = z_k(x^0)$ for all proposers $l < k$ who do not try to amend x^0. Hence, $f^\sigma(x^0) = z_k(x^0) = z_1(x^0) = f(x^0)$.

As this is true for any $x^0 \in X$, this proves that $f^\sigma(X) \equiv \{f^\sigma(x^0) : x^0 \in X\} = \{z_1(x^0) : x^0 \in X\} = Z$.

Claim 4: Let $h \in H^t$. Suppose the kth proposer has made proposal $y \neq x^{t-1}$. Let S_i^- be the set of players who have already voted 'yes' when it is i's turn to vote, and let S_i^+ be the set of voters j who will vote after i and are prescribed to vote 'yes' by σ_j. If $S \equiv S_i^- \cup \{i\} \cup S_i^+$ is a coalition then σ_i prescribes i to vote 'yes' only if $\phi^\sigma(h,k,S,y;1) \succeq_i \phi^\sigma(h;k+1)$, and to vote 'no' only if $\phi^\sigma(h;k+1) \succeq_i \phi^\sigma(h,k,S,y;1)$.

Proof: Claim 2 immediately implies that $\phi^\sigma(h,k,S,y;1) = z_1(h,k,S,y)$ for all $y \neq x^{t-1}$, and $\phi^\sigma(h;k+1) = z_{k+1}(h)$.

Suppose first that $h \in H_0$. If player i votes 'yes' then, by condition (A), $z_1(h,k,S,y) \succ_i x^{t-1}$. Claim 2 implies that $x^{t-1} = z_k(h) = \phi^\sigma(h;k)$. Hence, $z_1(h,k,S,y) \succ_i x^{t-1}$ implies $\phi^\sigma(h,k,S,y;1) \succ_i \phi^\sigma(h;k+1)$ and, therefore, that $\phi^\sigma(h,k,S,y;1) \succeq_i \phi^\sigma(h;k+1)$. If player i votes 'no' then, by condition (A), $x^{t-1} \succ_i z_1(h,k,S,y)$. This in turn implies that $\phi^\sigma(h;k+1) \succeq_i \phi^\sigma(h,k,S,y;1)$.

Now suppose that $h \in H_l$ for some $l \in \mathbb{N}$ and that $y \in A_k(Z_l, z_{k+1}(h))$. If player i votes 'yes' then, by condition (B), $y \succeq_i z_{k+1}(h) = \phi^\sigma(h;k+1)$. Since $y \in A_k(Z_l, z_{k+1}(h)) \subseteq Z_l$, history $(h,k,S,y) \in H_0$, which in turn implies that $\phi^\sigma(h,k,S,y;1) = y$ (all proposers will pass at a history in H_0). Hence, $\phi^\sigma(h,k,S,y;1) \succeq_i \phi^\sigma(h;k+1)$. If player i votes 'no' then, by condition (B), $z_{k+1}(h) \succ_i y$. This in turn implies that $\phi^\sigma(h;k+1) \succ_i \phi^\sigma(h,k,S,y;1)$ and, therefore, that $\phi^\sigma(h;k+1) \succeq_i \phi^\sigma(h,k,S,y;1)$.

Finally, suppose that $h \in H_l$ for some $l \in \mathbb{N}$ and that $y \notin A_k(Z_l, z_{k+1}(h))$. If player i votes 'yes' then, by condition (C), $z_1(h,k,S,y) \succ_i z_{k+1}(h)$. This implies that $\phi^\sigma(h,k,S,y;1) \succ_i \phi^\sigma(h;k+1)$ and, therefore, that $\phi^\sigma(h,k,S,y;1) \succeq_i \phi^\sigma(h;k+1)$. Similarly, if i votes 'no' then (C) implies that $z_{k+1}(h) \succeq_i z_1(h,k,S,y)$ and then $\phi^\sigma(h;k+1) \succeq_i \phi^\sigma(h,k,S,y;1)$.

Claim 5: Let $h \in H^t$ be a history ending with default $x^{t-1} = x$. At this history, the kth proposer cannot gain by deviating from $\sigma_{\pi_x(k)}$ at that stage and conforming to $\sigma_{\pi_x(k)}$ thereafter.

Let $i = \pi_x(k)$.

Suppose first that $h \in H_0$ (or, equivalently, $t-1 = \tau_l$): viz. σ dictates all proposers to pass at h. Consequently, if i conforms to σ_i then the final policy outcome will be $x^{t-1} = x^{\tau_l} = z_{k+1}(h)$. Hence, i can only profitably deviate by amending x^{t-1} with some policy y such that $y \succ_i x^{t-1}$. However, for any $S \in \mathcal{W}$, history (h,k,S,y) belongs to H_{l+1}. Claim 2 then implies that

$$\phi^\sigma(h,k,S,y) = z_1(h,k,S,y) \in Z_{l+1} \equiv \{z \in Z : x^{\tau_l} = z_{k+1}(h) \succeq_j z \text{ for some } j \in S\} .$$

Therefore, for each coalition $S \in \mathcal{W}$, there is at least one member of S who weakly prefers $z_{k+1}(h)$ to $z_1(h,k,S,y)$. Condition (C) guarantees that any proposal $y \neq x^{t-1}$ would be rejected; so that i cannot profitably deviate from passing.

Now suppose that $h \in H_l$ for some $l \in \mathbb{N}$. Any proposal y such that $\phi^\sigma(h,k,S,y;1) = z_1(h,k,S,y) \notin A_k(Z_l, z_{k+1}(h))$ must be unsuccessful. Indeed, condition (C) in the definition of voting histories implies that voters only vote 'yes' if they strictly prefer $z_1(h,k,S,y) \in Z_l$ to $z_{k+1}(h)$. As $P_{Z_l}(z_{k+1}(h)) \subseteq A_k(Z_l, z_{k+1}(h))$, every winning coalition includes at least one player who votes 'no'. Thus, as $z_k(h)$ is $\succeq_i$-maximal in $A_k(Z_l, z_{k+1}(h)) \supseteq \{z_{k+1}(h)\}$, player i cannot improve upon proposing $z_k(h)$ when $z_k(h) \neq z_{k+1}(h)$ and passing otherwise.

Proof of Proposition 11

Let σ be a semi-Markovian equilibrium. Suppose that, contrary to the statement of Proposition 11, $\phi^{\sigma}(H)$ is not a consistent set. This implies that there exist $o \in \phi^{\sigma}(H)$, $x \in X$, and $S \in \mathcal{W}$ such that, for all $o' \in \phi^{\sigma}(H)$, one of the following conditions is true:

(a) $o' = x$ and $o' \succ_i o$ for all $i \in S$;

(b) $o'Rx$ and $o' \succ_i o$ for all $i \in S$;

(c) $\neg(o'Rx)$.

Now consider a history $h \in H$ at which, instead of following σ and implementing o at the end of the round, some players have deviated as follows: a proposer $\pi_o(k)$ in S has proposed to amend o with x and all members of S have voted 'yes'. This deviation yields a new outcome $o' \in \phi^{\sigma}(H)$, which satisfies one of the conditions (a)-(c) above. As σ is an equilibrium, some winning coalition in $\mathcal{W}$ must find it (weakly) profitable to induce o' from x and, therefore, o' cannot satisfy (c). As a consequence, o' must satisfy either (a) or (b).

Denote the last player in $\pi_o(\{1, \ldots, m_o\}) \cap S$ by m_S, and suppose that this player has proposed amending o to x. Members of S anticipate that voting 'yes' will induce some $o' \in \phi^{\sigma}(H)$. As σ is semi-Markovian, it must still specify outcome o after an unsuccessful attempt to amend it. All players in S, including m_S, must then be strictly better off voting for x if o' satisfies either (a) or (b). Consequently, all voters in S would vote for x, and player m_S could profitably deviate from σ by proposing x, contrary to the supposition that σ is a semi-Markovian equilibrium.

The Dynamic Game with Endogenous Protocol (footnote 17)

In this section, we establish the claim that Proposition 5 also applies in a different 'dynamic' game, $\Gamma^d(\pi, x^0)$, where the chair selects the next proposer immediately after each vote which does not end the game.

$\Gamma^d(\Pi, x^0)$ starts with the chair selecting a proposer from M. This player either passes or proposes a policy in X, after which the players vote. A round necessarily ends if the default is amended. If the default has yet to be amended then the chair can either select a proposer from M or end the round, implementing the default. However, the chair can only end the game if the protocol in the final round is an element of Π. In particular, all M proposers have had an opportunity to propose. We construct payoffs as for $\Gamma^c(\Pi, x^0)$: players, including the chair, only care about the implemented decision. We again characterize play via the equilibria of $\Gamma^d(\Pi, x^0)$. Markov stationarity now requires that the chair's selection of proposer only depends on history via the default and the number of proposals by each player thus far in the current round.

The dynamic structure of $\Gamma^d\left(\Pi, x^0\right)$ is reminiscent of Harsanyi's (1974) model, where the chair solicits proposals at each default. By contrast, Harsanyi assumes that the chair's payoff is increasing in the number of amendments; so the equilibrium protocol in $\Gamma^d\left(\Pi, x^0\right)$ typically differs from that in Harsanyi (1974).

Proposition 5 implies that equilibrium proposals and voting in the dynamic game only depend on history via the default and the selected protocol in the current round. Consequently, the chair's selection in any equilibrium only depends on the default and on her previous selections that round. In equilibrium, the chair can anticipate whether and how any player, selected as proposer, would amend the default. Fix an equilibrium, and write the sequence of selections which the chair makes at x^0 when the default is not amended as $\pi^d\left(x^0, \Pi\right)$. Let $\pi^c\left(x^0, \Pi\right)$ be an equilibrium choice in $\Gamma^c\left(\pi, x^0\right)$. A chair who could commit to protocols could always do at least as well as the chair in $\Gamma^d\left(\pi, x^0\right)$ by choosing $\pi^c\left(x^0, \Pi\right) = \pi^d\left(x^0, \Pi\right)$. Conversely, the chair in $\Gamma^d\left(\Pi, x^0\right)$ could always do at least as well as the chair in $\Gamma^c\left(\Pi, x^0\right)$ by replicating $\pi^c\left(x^0, \Pi\right)$. We therefore conclude that the same set of policies can be implemented in an equilibrium of $\Gamma^c\left(\Pi, x^0\right)$ as in an equilibrium of $\Gamma^d\left(\Pi, x^0\right)$. In each case, an equilibrium protocol at x^0 is a best protocol in the class of games analyzed in Section 3.

Mixed Strategy Equilibria and MTHPEs

In this section, we provide proofs of our claims regarding mixed strategy equilibria and Markov trembling-hand perfect equilibria in the Conclusion of the paper.

Mixed strategy equilibria

In this subsection, we substantiate a claim in the Conclusion: that a mixed strategy Markov perfect equilibrium supports all three policies in a game which exhibits a Condorcet cycle: where 'supports' means that the process converges almost surely to implementing some policy. In light of the Condorcet cycle, there is no weakly stable set, and therefore no pure strategy Markov perfect equilibrium.

Suppose that three proposers = voters $i \in \{1, 2, 3\}$ have preferences over a policy space $\{x, y, z\}$ which are represented by utility functions u_i:

		Policies (w)		
		x	y	z
	1	2	1	0
Players (i)	2	0	2	1
	3	1	0	2

Utilities $u_i(w)$

and that the protocol is given by

$$(\pi_x(1), \pi_x(2), \pi_x(3)) = (2, 3, 1) \ ;$$
$$(\pi_y(1), \pi_y(2), \pi_y(3)) = (3, 1, 2) \ ;$$
$$(\pi_z(1), \pi_z(2), \pi_z(3)) = (1, 2, 3) \ .$$

Consider the following strategy combination. At any default, each player proposes her top-ranked policy; so, given the protocol, the default is implemented if it is not amended by either of the first two proposers. At any default and after any proposal, the player who top [resp. bottom] ranks the policy votes "yes" [resp. "no"], and the other player mixes.

In light of the symmetry across players, we write u for the initial default, U for the player who top-ranks u, and whose preferences satisfy $u \succ v \succ w$. The players who top rank v and w are respectively denoted by V and W. Thus, according to the protocol, the order of proposers is V, W, U.

Write p_u^v for the probability that v is eventually implemented at the beginning of a round with default u and Y_u^v for the probability that the decisive player votes "yes" to proposal v at default u.

If W proposes w [resp. v] then she is indifferent if and only if $2p_w^w + p_w^u = 1$ [resp. $2p_v^w + p_v^u = 1$]. It is easy to confirm that U and V would respectively vote "no" and "yes" if $p_t^s = 1/3$. W then proposes w if and only if

$$Y_u^w(2p_w^w + p_u^w - 1) \geq \max\{0, Y_u^v(2p_v^w + p_v^u - 1)\}$$

These arguments imply that, if V does not amend then u is amended to w with probability Y_u^w, and u is otherwise implemented. V then earns $Y_u^w \, (2p_w^v + p_w^w)$.

If V proposes v then W is indifferent as a voter if and only if

$$Y_u^w \, (2p_w^w + p_w^u - 1) = 2p_v^w + p_v^u - 1$$

V then earns

$$Y_u^v \, (2p_v^v + p_v^w) + (1 - Y_u^v) \, Y_u^w \, (2p_w^v + p_w^u)$$

Analogously, it is easy to confirm that W is decisive if V proposes w, and is indifferent if and only if

$$2p_w^w + p_w^u - 1 = Y_u^w \, (2p_w^w + p_w^u - 1)$$

V then earns

$$Y_u^w \, (2p_w^v + p_w^w) + (1 - Y_u^w) \, Y_u^w \, (2p_w^v + p_w^w)$$

if she proposes w. Hence, V cannot profitably deviate if and only if

$$Y_u^v\left(2p_v^v+p_v^w\right)\geq\max\left\{Y_u^w\left(2p_w^v+p_w^w\right)+\left(Y_u^v-Y_u^w\right)Y_u^w\left(2p_w^v+p_w^w\right),Y_u^vY_u^w\left(2p_w^v+p_w^w\right)\right\}\ .$$

All of these conditions are satisfied if $p_t^s=1/3$ for every $s,t\in X$. Accordingly, we will construct $\{Y_t^s\}$ such that every p_t^s satisfies this condition:

Given the strategy combination above, we have

$$\begin{aligned}p_u^u &= Y_u^vp_v^u+(1-Y_u^v)(Y_u^wp_w^u+1-Y_u^w)\\ p_u^v &= Y_u^vp_v^v+(1-Y_u^v)Y_u^wp_w^v\\ p_u^w &= Y_u^vp_v^w+(1-Y_u^v)Y_u^wp_w^w\end{aligned}$$

These equations hold when $p_u^s=1/3$ as long as $Y_u^v+Y_u^w=1+Y_u^vY_u^w$.

In sum, we have constructed a mixed strategy Markov perfect equilibrium for a game with no weakly stable set (and therefore no pure strategy equilibrium). This equilibrium supports the entire policy space.

Markov trembling-hand perfect equilibria

In this subsection, we provide a proof of Observation 2:

Observation 2. If X is finite and well ordered then the set of MTHP equilibrium policies coincide with the set of equilibrium policies (and is therefore the union of stable sets).

Proof: To prove this result, it suffices to show that, for every weakly stable set $V\in\mathcal{V}$, there is an MTHP equilibrium σ which supports V. To do so, we will use the construction described in the proof of Proposition 1. Consider the equilibrium described in that proof, say $\tilde{\sigma}$, which is obtained by setting $Y=\emptyset$. In this equilibrium, all proposers pass if the default x belongs to V. If $x\notin V$ then, for each $k\in\{1,\ldots,m_x\}$, the kth proposer offers $y_k(x)$ — i.e.: the $\succ_{\pi_x(k)}$-maximal element in $P_V\left(y_{k+1}(x)\right)\cup\{y_{k+1}(x)\}$ — and voter $i\in N$ accepts this proposal if and only if $y_1\left(y_k(x)\right)\succ_i y_{k+1}(x)$ — where, for all $x\notin V$, $y_1(x)$ is the ideal policy of the last amender of x in $P_V(x)\cup\{x\}$ and, for all $v\in V$, $y_1(v)=v$. Thus, if the current default x does not belong to V: all proposers who move before the last amender of x make an unsuccessful proposal (by internal stability of V); the last amender amends x to $y_1(x)$; and (off the equilibrium path) proposers k who move after the last amender choose $y_k(x)=x$ (i.e., they pass).

In equilibrium $\tilde{\sigma}$, as X is finite and well ordered, 'agents' (we are using the agent-strategic form) play strict best responses in all voting stages and in proposal stages where they are the last amenders. In proposal stages where they move before the last amender, they are indifferent between all proposals in X since (by internal stability of V) all proposals are voted down. In proposal stages where they move after the last amender, they are indifferent between all proposals in X that are rejected. Let σ be a stationary Markov strategy profile defined as follows:

- in stages where $\tilde{\sigma}$ prescribes strict best responses, σ coincides with $\tilde{\sigma}$;
- in proposal stages where the proposer moves before the last amender, σ prescribes that proposer to offer her ideal policy in V;
- in proposal stages where the proposer moves after the last amender, σ prescribes that proposer to offer her ideal policy in $V \cup \{x\}$, where x is the ongoing default.

By construction, σ must be an equilibrium of $\Gamma\left(\pi, x^0\right)$ and $f^\sigma\left(X\right) = V$. (Either σ dictates the same behavior as $\tilde{\sigma}$ or it dictates behavior that yield the same consequences as $\tilde{\sigma}$.) We will now prove that it is Markov trembling-hand perfect (MTHP).

To do so, we first construct a sequence of strategy profiles $\{\sigma^m\}$ as follows. At every voting history, σ^m is defined as

$$\sigma^m(h) = \frac{1}{m}\bar{\upsilon} + \left(1 - \frac{1}{m}\right)\sigma(h)$$

where $\bar{\upsilon}$ is a (completely mixed) voting profile such that the probability for every element of V to be accepted is the same (for all defaults and proposers). At all proposal histories h, σ^m is defined as

$$\sigma^m(h) = \frac{1}{m}\sigma'(h) + \left(1 - \frac{1}{m}\right)\sigma(h)$$

where σ' is an arbitrary stationary Markov (completely) mixed strategy. Evidently, $\sigma^m \to \sigma$ as $m \to \infty$.

To establish the result, we now have to show that for each player $i \in N$ and every history of the game h, the action prescribed by σ_i to the agent representing i at h, $i(h)$, is a best response to σ^m for all sufficiently large m. By construction of σ, this is obvious in all voting stages and in the proposal stages where the agent is the last amender. We can therefore concentrate on proposal stages in which proposers are indifferent between proposals (given σ).

Let h be such a proposal stage (or history) with ongoing default x, and consider the choice of the agent representing the kth proposer at this history, $i = \pi_x(k)$. Let $p_k^m(y)$ be the probability that proposal y by i is accepted, $V_i^m(y)$ the expected payoff of i when her proposal y is accepted, and v_i^m her expected payoff when her proposal is rejected, given that all players play according to σ^m. Denoting by y_i i's ideal policy in V, the action prescribed by σ_i to $i(h)$ is a best response to σ^m if and only if

$$p_k^m(y_i)\,V_i^m(y_i) + [1 - p_k^m(y_i)]\,v_i^m \geq p_k^m(y)\,V_i^m(y) + [1 - p_k^m(y)]\,v_i^m$$

or, equivalently,

$$p_k^m(y_i)\,[V_i^m(y_i) - v_i^m] \geq p_k^m(y)\,[V_i^m(y) - v_i^m]\ . \tag{2}$$

for all $y \in X$.

Suppose first that $x \in V$. In this case, the voting behavior dictated by $\tilde{\sigma}$, and therefore σ, makes any proposal in X unsuccessful. This implies that σ^m prescribes the same voting behavior as $\bar{v}$. As a consequence, $v_i^m \to u_i(x)$ and $p_k^m(y) = p_k^m(y')$ for all $y, y' \in V$. Moreover, by construction of σ, $V_i^m(y) \to y_1(y) \in V$ for all $y \in X$. As X is finite and well ordered, $i(h)$ cannot improve on proposing i's ideal policy in V when m is arbitrarily large: $V_i^m(y_i) \to u_i(y_i) > u_i(y) \leftarrow V_i^m(y)$ for all $y \in V \setminus \{y_i\}$.

Suppose now that $x \notin V$ and that i moves before the last amender (at h). Under strategy profile σ, every proposal by player i is rejected with a probability of 1. Therefore, all proposals in V are accepted with the same probability under σ^m (i.e., the same probability as under $\bar{v}$): $p_k^m(y) = p_k^m(y')$ for all $y, y' \in V$. We can then use the same argument as in the previous paragraph to show that (2) holds for sufficiently large m.

Finally, suppose that $x \notin V$ and that i moves after the last amender (at h). As explained above, we can concentrate on proposals in V. We distinguish between three different cases:

(1) $i(h)$ proposes $y \in P_V(x)$. In this case, the resulting expected payoff to player i when all agents play according to σ^m is given by $p_k^m(y)V_i^m(y) + [1 - p_k^m(y)]\,v_i^m$.

(2) $i(h)$ passes. The resulting expected payoff to player i when all agents play according to σ^m is then v_i^m.

(3) $i(h)$ proposes $v \notin P_V(x)$. In this case, the resulting expected payoff to player i when all agents play according to σ^m is given by $p_k^m(v)V_i^m(v) + [1 - p_k^m(v)]\,v_i^m$.

When m becomes arbitrarily large, σ^m becomes arbitrarily close to σ, so that $v_i^m \to u_i(x)$ and $V_i^m(y) \to u_i(y)$ for any proposal $y \in V$. Inspection of the three cases above (and the corresponding payoffs) reveals that $i(h)$ cannot improve on proposing player i's ideal policy in $V \cup \{x\}$ (which, i moving after the last amender, cannot be in $P_V(x)$) when is arbitrarily large, thus completing the proof of the result.

References

Acemoglu, D., G. Egorov and K. Sonin (2009), "Equilibrium Refinement in Dynamic Voting Games" mimeo.

Acemoglu, D., G. Egorov and K. Sonin (forthcoming), “Dynamics and Stability of Constitutions, Coalitions, and Clubs” *American Economic Review.*

Anesi, V. (2006), “Committees with Farsighted Voters: A New Interpretation of Stable Sets” *Social Choice and Welfare* 27, 595-610.

Anesi, V. (2010), “Nooncooperative Foundations of Stable Sets in Voting Games” *Games and Economic Behavior* 70, 488-493.

Austen-Smith, D. and J.S. Banks (2005), *Positive Political Theory II: Strategy and Structure.* University of Michigan Press, Ann Arbor.

Banks, J.S. and J. Duggan (2000), “A Bargaining Model of Collective Choice” *American Political Science Review* 94, 73-88.

Banks, J.S. and J. Duggan (2006), “A General Bargaining Model of Legislative Policy-Making” *Quarterly Journal of Political Science* 1, 49-85.

Baron, D.P. (1996), “A Dynamic Theory of Collective Goods Provision” *American Political Science Review* 90, 216-330.

Baron, D.P. and J.A. Ferejohn (1989), “Bargaining in Legislatures” *American Political Science Review* 83, 181-1206.

Bernheim, B.D., A. Rangel and L. Rayo (2006), “The Power of the Last Word in Legislative Policy Making” *Econometrica* 74, 1161-1190.

Chwe, M. (1994), “Farsighted Coalitional Stability" *Journal of Economic Theory* 63, 299-325.

Compte, O. and P. Jehiel (2010), “Bargaining and Majority Rules: A Collective Search Perspective” *Journal of Political Economy* 118, 189-221.

Diermeier, D. and P. Fong (2011), “Legislative Bargaining with Reconsideration” *Quarterly Journal of Economics* 126, 947-985.

Duggan, J. (2006), “Endogenous Voting Agendas” *Social Choice and Welfare* 27, 495-530.

Duggan, J. and T. Kalandrakis (2007), “Dynamic Legislative Policy Making” W. Allen Wallis Institute Working Paper No. 45.

Duggan, J. and T. Kalandrakis (2011), “A Newton Collocation Method for Solving Dynamic Bargaining Games” *Social Choice and Welfare* 36, 611-650.

Harsanyi, J.C. (1974), "An Equilibrium-Point Interpretation of Stable Sets and a Proposed Alternative Definition" *Management Science* 20, 1472-1495.

Hirai, T. (2009), "The Objective and the Discriminatory Solutions of Majority Voting Games with Single-Peaked, Ordinal Preferences", mimeo.

Hortala-Vallve, R. (forthcoming), "Generous Legislators? A Description of Vote Trading Agreements" *Quarterly Journal of Political Science.*

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

Kalandrakis, T. (2010), "Minimum Winning Coalitions and Endogenous Status Quo" *International Journal of Game Theory* 39, 617-643.

Lucas, W.F. (1992), "Von Neumann-Morgenstern Stable Sets" *Handbook of Game Theory with Economic Applications* 1, North Holland, Amsterdam.

Mailath, G.J., A. Postlewaite and L. Samuelson (2005), "Contemporaneous Perfect Epsilon-Equilibria" *Games and Economic Behavior* 53, 126-140.

McKelvey, R.D. (1986), "Covering, Dominance, and Institution-Free Properties of Social Choice" *American Journal of Political Science* 30, 283-314.

Moldovanu, B. and E. Winter (1995), "Order Independent Equilibria" *Games and Economic Behavior* 9, 21-34.

Roberts, K. (2007), "Condorcet Cycles? A Model of Intertemporal Voting" *Social Choice and Welfare* 29, 383-404.

Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model" *Econometrica* 50, 97-109.

Shepsle, K.A. and B. R. Weingast (1984), "Uncovered Sets and Sophisticated Voting Outcomes with Implications for Agenda Institutions" *American Journal of Political Science* 28, 49-74.

Wilson, R. (1971), "Stable Coalition Proposals in Majority-Rule Voting" *Journal of Economic Theory* 3, 254-271.

Winter, E. (1997), "Negotiations in Multi-Issue Committees" *Journal of Public Economics* 65, 323-342.