



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS



The University of
Nottingham

Discussion Paper No. 2012-03

2012

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 MacCormick
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

March 22, 2012

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 a player who does not propose may nevertheless earn all of the surplus from agreement. The set of immovable 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 set of proposers may have surprising effects on the set of immovable policies. 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 immovable policies in semi-Markovian equilibria form the largest consistent set.

JEL classification: C78, D71, D72.

Keywords: bargaining, committee voting, 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 en-

*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.

dogenously drawing policies/proposals, and implement a proposal according to a stopping rule. Congress and the FOMC instance committees which stop deliberating as soon as some proposal wins a final vote, while the EU's Council of Ministers (henceforth 'the Council') reaches a decision by final vote when the issue must be addressed urgently or some government wants to signal to its domestic audience (cf. Heisenberg (2005)). Conventional bargaining models, like Rubinstein (1982) and Baron and Ferejohn (1989), focus on such committees: the game ends when some proposal wins against the status quo. However, an important class of committees do not use a final vote stopping rule. In particular, the Council reaches more than 80% of its decisions without a final vote, including on issues as controversial as trade and budgetary policy (cf. Hayes-Renshaw et al. (2006) and Mattila (2009)). Analogously, insiders report that the Bundesbank and the ECB decide on the interest rate without a final vote. Anecdotal evidence suggests that proceedings end (and a policy is implemented) in such committees when discussion grinds to a halt.

We present and analyze a model of decision making without final voting, in which the committee entertains a single policy at a time. The game is played by proposers and voters: some players may both propose and vote. The game starts with an initial default. Players have the opportunity to propose amendments to the default in a fixed sequence (the protocol), which can depend on the ongoing default. If a winning coalition of voters accepts the proposal then the default is amended; if you like, 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. (This is what we mean by discussion grinding to a halt.) In light of evidence that the Council takes urgent decisions by final vote, we suppose that payoffs in the game only depend on the policy implemented.¹ Two aspects of this noncooperative game are worth highlighting. First, preferences and the sets of winning coalitions and of proposers determine an underlying cooperative game (independently of the protocol and the initial default). Second, 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.

Our main results describe the policies which can be reached in a pure strategy stationary Markov equilibrium from any initial default and for any protocol. We use an algorithmic

¹This 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.

technique, which exploits a relationship between equilibria of the game and a solution concept for the underlying cooperative game:

Any equilibrium determines a function, which maps from any initial default to the implemented policy. Stationarity implies that any policy in the range of this mapping is ‘immovable’: the equilibrium prescribes that this policy is implemented whenever it is the initial default. We show that the immovable policies form a weakly stable set in a related simple game (Proposition 2). 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 (henceforth ‘vNM’) stable set, but external stability is weakened to allow for weak social preference.² We also obtain a converse result. Specifically, for any closed weakly stable set, we construct equilibria whose immovable policies are exactly that weakly stable set (Proposition 1).³

Equilibrium outcomes have some surprising properties. We demonstrate by example that a winning coalition may amend a default to a policy which is implemented, leaving all coalition members worse off than at the initial default; all players may earn a surplus in a majority-rule divide the dollar game; and a player who does not propose may earn all of the surplus from agreement.

Our algorithmic technique reveals that equilibrium outcomes depend on the order in which players can propose (the protocol) and on the set of winning coalitions, which are determined both by the set of proposers and by the voters who can amend a default. Accordingly, we exploit the characterization results (Propositions 1 and 2) to study how variations in the protocol and the winning coalition affect the policy implemented from any initial default.

We start by considering 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 we show, strikingly, that the chair can only affect the policy implemented by changing the order in which players propose *at the initial default*. Furthermore, 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 immovable policies are a

²Every vNM stable set is therefore weakly stable.

³This result provides micro-foundations of weakly stable sets as a byproduct of our analysis. Propositions 1 and 2 imply that the policies which can be implemented in any equilibrium is the union of weakly stable sets.

given weakly stable set. If the policy space is well ordered (no player is indifferent between any two policies) and there is a unique weakly stable set then a chair who proposes cannot improve on a protocol in which she proposes last. This result does not generalize to games in which a player may be indifferent between policies or there are several weakly stable sets: the chair may then be best off when another player proposes last.

We then study the implications of varying the set of winning coalitions via changes in the quota and in the set of proposers. According to a natural conjecture, increasing the quota expands the set of immovable policies because coalitions which could destabilize policies are no longer winning with a larger quota.⁴ We provide conditions for this conjecture to be true; but we also show that increasing the quota may *contract* the set of immovable policies. The basis for this surprising result is that 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. Variations in the set of proposers can have analogously surprising results. We show by example that a player may be worse off if she is given the opportunity to propose. The intuition again turns on the implications of changing winning coalitions for the structure of weakly stable sets.

We end the paper by extending our analysis in three 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 that accrue 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 any equilibrium's immovable policies are a consistent set, and the union of immovable policies coincides with the largest consistent set.

⁴Arguments of this sort have been repeatedly used during the prolonged debates over qualified majority voting in the Council.

Finally, we extend our results to an open rule bargaining game in the spirit of Baron and Ferejohn (1989), where a policy is implemented if and only if a proposer successfully moves the previous question. We show that any weakly stable set can be supported in an open rule bargaining game, and the policy set supported by any equilibrium must satisfy internal stability. However, there may be equilibria that support policy sets which fail external stability.

After reviewing the related literature in the next subsection, we present the model in Section 2. We characterize equilibria in Section 3, and use our results to explain policies chosen by the Council. In Section 4, we explore how the policy implemented varies with the protocol and with the set of winning coalitions. In Section 5, we provide micro-foundations for the largest consistent set, construct contemporaneous perfect ε -equilibria in games with repeated implementation, and analyze our open rule bargaining model. We conclude in Section 6, and briefly discuss variants on our model with bargaining round the table; 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 and a unique weakly stable set also exhibit the power of the last word. 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. 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.

Harsanyi (1974) provides micro-foundations for vNM 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 vNM 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 allows any policy in the vNM stable set which socially dominates the initial default to be implemented in equilibrium: for players and the chair respectively only care about the implemented policy and the number of amendments. 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 Appendix A.2.⁵

Harsanyi argues that vNM 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 randomly selected proposers can amend the existing default. In contrast to Bernheim et al (2006), this game can last indefinitely; but, unlike our model, the game only ends when a player proposes moving the previous question (viz. the current default). 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. As noted above, we study a version of open rule bargaining which extends our model, but is not restricted to the distributive problems which Baron and Ferejohn consider.⁶ We show that an immovable policy in our model is also immovable in the open rule bargaining game, but that the converse does not hold. In further contrast, Baron and Ferejohn show that a supra-minimal majority may earn a share of a fixed pie, provided that players are impatient enough: the proposer fears that an excluded player will be selected to propose. In our model, patient players may all earn a share of the pie.

Following Baron (1996), a recent literature has studied equilibria of games with repeated

⁵In contrast to Harsanyi (1974) and this paper, Hortala-Vallve (forthcoming) studies play in a related model without a weakly stable set.

⁶In this sense, our model also generalizes Volden and Wiseman's (2007) variant on Baron and Ferejohn (1989).

implementation (as described in the last subsection). Each round ends with a final vote, but the game continues with a new status quo/default. 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.⁷ Acemoglu et al prove existence when social preferences are acyclic and the policy space is finite, and demonstrate (in an online Appendix) that the limiting policies constitute the unique vNM stable set and the largest consistent set.⁸ 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 vNM stable; so existence of an equilibrium in our model is equivalent to existence of a vNM stable set in simple games. It is well known that vNM stable sets may not exist, even if the policy space is finite (e.g. when there are Condorcet cycles or in Bernheim et al's (2006) pork barrel model), but vNM stable sets have been characterized for some games with infinite policy spaces, such as three player divide the dollar games.⁹ 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 our approach, Anesi (2010) is related most closely to this paper. Anesi demonstrates that any vNM 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.¹⁰ 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 weakly stable sets and the class of absorbing sets of Markov perfect equilibria when the policy space is finite and well ordered.¹¹

⁷Acemoglu et al. allow for some exogenous proposals (and exploit this possibility in their proofs).

⁸Von Neumann and Morgenstern (1944) 65.7 prove existence and uniqueness of vNM stable sets under Acemoglu et al's conditions. See Duggan and Kalandrakis (2011, forthcoming) and Kalandrakis (2004) and (2010) for existence results in related models.

⁹Ordeshook (1986) discusses these issues; Hirai (2009) provides some recent results. More generally, Lucas (1992) surveys the literature on vNM stable sets.

¹⁰Acemoglu et al (forthcoming) show, in an online Appendix, that equilibria in their model support the unique stable set, which coincides with the largest consistent set.

¹¹Anesi only proves that the former is a subset of the latter, demonstrating by example that the legislature may choose policies outside stable sets when proposers are chosen randomly. Anesi (2006) obtains the equivalence between vNM 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

Diermeier and Fong (2011) study a game with repeated implementation in which a single player can propose (so the induced social preference relations are acyclic), and the policy space is finite, but not necessarily well ordered. Their game form is therefore the same as a special case of our model. In contrast to our approach, Diermeier and Fong focus on equilibria in which an indifferent voter always votes in favor of the proposal. This has important implications for predicted solutions, as instanced by their benchmark case, in which three players bargain over division of a pie. If any division were feasible then this game has a pair of weakly stable sets, in each of which the proposer shares the pie exclusively with one of the other players. Any equilibrium outcome in our model must lie in one of these sets. Diermeier and Fong’s tie-breaking rule eliminates both of these weakly stable sets, which immediately implies that any steady state policy must give each player a share of the pie. Diermeier and Fong also provide an existence proof for their game; von Neumann and Morgenstern’s (1944) argument (cf. footnote 8 above) implies that our version of Diermeier and Fong’s game also has a unique equilibrium.¹²

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.¹³ 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.¹⁴ 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.

policy space. Our model provides noncooperative bargaining foundations for vNM stable sets under these conditions.

¹²Other related papers include Kalandrakis (2004, 2009), Battaglini and Coate (2007, 2008), Duggan and Kalandrakis (2011, forthcoming), Dziuda and Loeper (2010), Battaglini and Palfrey (2011).

¹³Dutta et al (2004) consider endogenous agenda formation in a less structured model, which is not based on a specific game form or protocol.

¹⁴Our model therefore integrates features of successive elimination and amendment agendas.

2 The model

We consider a finite committee consisting of $m \geq 1$ *proposers*, $M \equiv \{1, \dots, m\}$, and $n \geq 2$ *voters*, $N \equiv \{1, \dots, n\}$. 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 the Council of Ministers, where the European Commission makes all proposals but cannot vote, and most AGMs, where management propose corporate policies to shareholders who then vote (cf. Matsusaka and Ozbas (2012)).

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 $t = 1, 2, \dots$, starts in the shadow of an ongoing default policy x^{t-1} . For each possible default $x \in X$, there is a fixed protocol $\pi_x : \{1, \dots, m_x\} \rightarrow M$, $m_x \in \mathbb{N}$, that determines the order in which the proposers (i.e. the players in M) are given the opportunity to propose policies to amend the ongoing default. That is, when $x \in X$ is the current default, protocol π_x gives proposer $\pi_x(k)$ the k th opportunity to amend x for each $k \in \{1, \dots, m_x\}$. Each proposer $i \in M$ has at least one opportunity to amend the default in every round: $|\pi_x^{-1}(i)| \geq 1$ 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:¹⁵

¹⁵A similar bargaining process is used in Acemoglu et al. (forthcoming).

1. If the k th 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 $k < m_{x^{t-1}}$, then the game moves to Step 1 with k increased by 1; if $k = m_{x^{t-1}}$ then 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 . Bargaining indefinitely makes all players worse off than if any policy is implemented after a finite number of rounds.¹⁶ This assumption is consistent with the conventional claim that the Council is worst off when it fails to decide (e.g. Thomson (2011)). Let $\Gamma(\pi, x^0)$ 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

¹⁶This assumption precludes indefinite bargaining and ensures that the one-shot deviation principle applies even though the game is not continuous at infinity. The principle would hold in games with a finite policy space if payoffs were discounted by the number of rounds, and players were patient enough. Our main results would still be true in such games if, in addition, we assumed that the policy space is well ordered.

‘equilibria.’¹⁷

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 $k \in \{1, \dots, m_x\}$ the unique final outcome $f^\sigma(x, k)$ eventually implemented (given σ) when x is the ongoing default and the k th proposer is about to move (in any round t). We are particularly interested in $f^\sigma(x^0, 1)$, which describes 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 $f^\sigma(x^0)$ instead of $f^\sigma(x^0, 1)$. The characterization of this function for all possible equilibria of $\Gamma(\pi, x^0)$ is the subject matter of the next section.

3 Equilibrium characterization

3.1 Preliminaries

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 $\mathcal{W} \equiv \{S \subseteq C : (S \cap N) \in \mathbf{W} \text{ \& } (S \cap M) \neq \emptyset\}$. 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 P -dominance and R -dominance respectively, as follows: for all $x, y \in X$,

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

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

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

¹⁷In Section 5.2, we will consider strategies that are measurable with respect to other elements in the history of the game.

(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 vNM stable set is a weakly stable set which is P -externally stable: that is, it satisfies a variant of (ES_R) in which R is replaced by P . Conversely, if the policy space X is well ordered (i.e., if all the \succeq_i 's are linear orders) then \mathcal{V} corresponds to the collection of vNM stable sets. This is not true when X is not well ordered: there may be policy sets that are weakly stable but not vNM 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 $z \succ_1 x \succ_1 y$, $y \succ_2 x \sim_2 z$, and $x \sim_3 y \succ_3 z$. It is easy to confirm that yPz , and that $\{x, z\}$ and $\{y\}$ are weakly but not vNM stable. By contrast, $\{x, y\}$ is vNM stable.

□

The predictive power of weakly stability, like vNM stability, depends on other parameters of the model: there may be a unique and small weakly stable set (e.g. any Condorcet winner); 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, it is readily checked that, under our assumptions: a weakly stable set may contain weakly Pareto dominated policies; and the closure of a weakly stable set is itself weakly stable.

We end this subsection with some additional 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 $\mathbf{M}(Q, Y) \equiv \{y \in Y : \forall y' \in Y \setminus \{y\}, y'Qy \text{ implies } yQy'\}$. We will refer to the elements of the latter set as the Q -maximal policies in Y .

3.2 Computation

We now turn 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 $x^0 \in X$.

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 $k = m_x, m_x - 1, \dots, 1$:

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

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

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

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

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

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, \dots, y_{m_x+1}) such that $y_1 = y \in V$, $y_{m_x+1} = x$, and $y_k \in \mathfrak{s}_k^\pi(V, y_{k+1})$ for each $k = 1, \dots, m_x$.

Before we proceed any further, it may be helpful to illustrate this construction by exploiting Example 3.1 above.

Example 3.1 (continued). Suppose that the initial default is $x^0 = y$, and that the protocol is defined as $\pi_w(i) = i$ for all $w \in \{x, y, z\}$ and all $i \in M$ — in words: players propose in the order 1, 2, 3 at every default. The tree $\mathfrak{T}^\pi(\{x, z\}, y)$ is depicted in Figure 1 — recall that $\{x, z\}$ is a weakly stable set. The set of immediate successors of the initial node, y , is $\mathfrak{s}_3^\pi(\{x, z\}, y) = \{x, y\}$. Indeed, $P_{\{x, z\}}(y) = \emptyset$ and $R_{\{x, z\}}(y) = \{x\}$. As player 3 is the last proposer, the definition of $\mathfrak{s}_3^\pi(\{x, z\}, y)$ implies that $\mathfrak{s}_3^\pi(\{x, z\}, y) = \mathbf{M}(\succeq_3, \{y\}) \cup \mathbf{M}(\succeq_3, \{x, y\}) = \{y\} \cup \{x, y\} = \{x, y\}$ (recall that x and y are player 3's ideal policies). Following the dashed path in Figure 1, consider now the set of immediate successors of node $x \in \mathfrak{s}_3^\pi(\{x, z\}, y)$. Note first that $P_{\{x, z\}}(x) = \emptyset$ and $R_{\{x, z\}}(x) = \{x, z\}$. Hence, given that player 2 is the second proposer, $\mathfrak{s}_2^\pi(\{x, z\}, x) = \mathbf{M}(\succeq_2, \{x\}) \cup \mathbf{M}(\succeq_2, \{x, z\}) = \{x\} \cup \{x, z\} = \{x, z\}$ (proposer 2 is indifferent between the two policies in $\{x, z\}$). Finally,




Figure 1: Tree $\mathfrak{T}^\pi(\{x, z\}, y)$

the set of immediate successors of node $x \in \mathfrak{s}_2^\pi(\{x, z\}, x)$ is $\mathfrak{s}_1^\pi(\{x, z\}, x) = \mathbf{M}(\succeq_1, \{x\}) \cup \mathbf{M}(\succeq_1, \{x, z\}) = \{x\} \cup \{z\} = \{x, z\}$: x and z are final nodes of tree $\mathfrak{T}^\pi(\{x, z\}, y)$. This completes the description of the dotted paths in Figure 1. One could apply the same procedure and intuition to the other paths of $\mathfrak{T}^\pi(\{x, z\}, y)$, so as to obtain $F^\pi(\{x, z\}, y) = \{x, z\}$ (recall that $F^\pi(\{x, z\}, y)$ only selects the terminal nodes that belong to $\{x, z\}$).

□

Our first result states that the construction of tree $\mathfrak{T}^\pi(V, x^0)$, and therefore of $F^\pi(V, x)$, generates equilibrium policies of game $\Gamma(\pi, x^0)$.

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 $F^\pi(V, \cdot)$ is an equilibrium outcome of $\Gamma(\pi, x^0)$. Put differently, Proposition 1 says that the 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 $\Gamma(\pi, x^0)$.

Proposition 1 thus implies that the closure of any weakly stable set V is ‘immovable’ in the sense that there is an equilibrium σ of $\Gamma(\pi, x^0)$ such that the union of $f^\sigma(x)$ over $x \in X$ is V : that is, exactly the initial defaults in V are not amended in that equilibrium. We will say that the equilibrium *supports* V in such a case. 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 $\mathbf{M}(\succeq_{\pi_x(k)}, P_V(y_{k+1}) \cup \{y_{k+1}\} \cup Y)$ may be empty.

Inspection of the proof of Proposition 1 reveals that these equilibria have a no-delay property: a policy in V is implemented in no more than two rounds. The intuition behind the construction of these equilibria is as follows. Let $x \notin V$ be the ongoing default in a given round t , and let (y_1, \dots, y_{m_x+1}) , with $y_1 = y \in V$, be a path of tree $\mathfrak{T}^\pi(V, x)$. Suppose that all players believe that policies in V and only these policies are immovable: if a policy outside V is voted up and becomes the new default then it will be amended to a policy in V , which will never be amended. Hence, when considering possible amendments to the current default $x \notin V$, proposers only consider policies in V . Suppose that the m_x th proposer, $\pi_x(m_x)$, is given the opportunity to make a proposal in this 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 — her ‘acceptance set’ to use the language of the previous literature. 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. By subgame perfection, voters who strictly prefer y to x must therefore vote ‘yes’. Voters who are indifferent between x and y may vote either ‘yes’ or ‘no’. Thus, the last proposer’s acceptance set is of the form $P_V(x) \cup Y$, where the set $Y \subseteq R_V(x)$ describes 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 $Y = R_V(x)$. As V satisfies (ES_R) , there must exist a nonempty $Y \subseteq R_V(x)$ and therefore a nonempty acceptance set for the last proposer. In our equilibrium construction, Y is such that y_{m_x} is the proposer’s ideal policy in $P_V(x) \cup Y \cup \{x\}$: $y_{m_x} \in \mathbf{M}(\succeq_{\pi_x(m_x)}, P_V(x) \cup Y \cup \{x\}) \subseteq \mathfrak{s}_{m_x}^\pi(V, x)$, where $y_{m_x} = x$ stands either for ‘pass’ or for a proposal outside the acceptance set (which is then voted down). It is consequently optimal for her to choose y_{m_x} . Now consider the $(m_x - 1)$ th proposer’s choice. She faces the same problem as the m_x th 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, her acceptance set is of the form $P_V(y_{m_x}) \cup Y$, where $Y \subseteq R_V(y_{m_x})$ is such that she optimally chooses $y_{m_x-1} \in \mathbf{M}(\succeq_{\pi_x(m_x-1)}, P_V(y_{m_x}) \cup Y \cup \{y_{m_x}\}) \subseteq \mathfrak{s}_{m_x-1}^\pi(V, y_{m_x})$. Moving backward, we

can repeatedly apply the same reasoning to all proposers until the first, $\pi_x(1)$, whose choice y_1 is in $F^\pi(V, x) \subseteq V$.

Now suppose that the ongoing default x belongs to V . Voters anticipate that amending x to any other policy y will eventually lead to the implementation of a policy x_y in V . In our equilibrium, proposers only make and voters only accept a proposal y if they strictly prefer x_y to x . As V satisfies (IS_P) , every coalition $S \in \mathcal{W}$ comprises a player who weakly prefers x to x_y . This makes any proposal y unsuccessful, so that each proposer k optimally passes (i.e., she chooses $y_k(x) = x$). This confirms players' beliefs (assumed at the start of the previous paragraph) that policies in V are immovable, thus completing the description of the equilibrium. As the same construction applies in any round t , and in particular in round 1, this equilibrium is no-delay.

Proposition 1 prompts the following question: Can there be equilibria of $\Gamma(\pi, x^0)$ (including equilibria with delay) 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 all $x \in X$. 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 $(x, y^m(x), \dots, y_1(x))$ 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 $k \in \{1, \dots, m_x\}$. 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 $(x, y^m(x), \dots, y_1(x))$ of $\mathfrak{T}^\pi(V, x)$, with $y_1(x) \in V$, such that $y_k(x) = f^\sigma(x, k)$ for each $k \in \{1, \dots, m_x\}$. Thus, the construction of trees associated with weakly stable sets also provides a complete characterization of equilibrium behavior both on and off equilibrium paths.

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

These two propositions provide, as a by-product of our analysis, a new bargaining interpretation for (weakly) stable sets in voting games. In contrast to the existing literature (e.g., Harsanyi (1974) and Anesi (2010)), these microfoundations extend to situations with both non-transferable utility and an infinite policy space.

Furthermore, they have a number of implications which will prove useful below. We end this subsection by detailing these properties:

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

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

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.

Our analysis above reveals 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 $\mathfrak{T}^\pi(V, x^0)$ is the policy outcome of some equilibrium of $\Gamma(\pi, x^0)$. Hence, each weakly stable set may contain several equilibrium policies. The issue of equilibrium refinement is discussed more extensively in Section 6.2. In particular, Observation 5 shows that Markov trembling hand perfection (Acemoglu et al, 2009) leaves the set of equilibrium policies unchanged.

On the other hand, Propositions 1 and 2 allow us to provide sufficient conditions for a unique equilibrium:

Corollary 2. *a) $\Gamma(\pi, x^0)$ has a unique equilibrium outcome if there is a unique weakly stable set and X is well ordered;*

b) Any Condorcet winner is implemented in every equilibrium of $\Gamma(\pi, x^0)$, for every $x^0 \in X$.

Both parts follow from our algorithm. If the premise of part a) holds then each node in the tree has a unique successor; and any Condorcet winner must constitute the unique weakly stable set. The premise of part a) is sufficient, rather than necessary, as our next example demonstrates:

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. $\Gamma(\pi, x^0)$ then has a unique equilibrium outcome in which player 2 takes the whole pie.

This is, of course, an equilibrium outcome of the conventional final-vote bargaining game with patient players. More interestingly, it is also the only equilibrium outcome of such a game when payoffs are discounted after each round (rather than each proposal): for any common discount factor.

□

The collection of weakly stable sets in a game only depends 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 Section 4, variations in the protocol π may affect the policies which can be implemented from a given initial default — i.e., from Corollary 1: $\bigcup_{V \in \mathcal{V}} F^\pi(V, x^0)$. Interestingly, the next result states that this set only depends on the protocol at the initial default: π_{x^0} . To see this, observe that, for any weakly stable set V and any initial default x^0 , the tree $\mathfrak{T}^\pi(V, x^0)$, and therefore the selection of terminal nodes $F^\pi(V, x^0)$, only depend on π_{x^0} . Indeed, the construction of the tree reveals that all equilibrium policies can be reached in one bargaining round and, in that round, each proposer $i \in M$ chooses her ideal policies from a set which is independent of the protocol.

Corollary 3. *Let $\pi^1 \equiv \{\pi_x^1\}$ and $\pi^2 \equiv \{\pi_x^2\}$ be two protocols. If $\pi_{x^0}^1 = \pi_{x^0}^2$ then*

$$F^{\pi^1}(V, x^0) = F^{\pi^2}(V, x^0) .$$

Corollaries 1 and 3 thus jointly imply that *the set of equilibrium policies only depends on the protocol at the initial default*. In all equilibria supporting $V \in \mathcal{V}$, proposers and voters anticipate in round 1 that policies in V , and only those policies, are immovable: once reached, they must be implemented. In particular, each proposer k faces an ‘acceptance set’ of the form $A_k = P_V(y_{k+1}) \cup \{y_{k+1}\} \cup Y$, where y_{k+1} is the policy that will be implemented if she fails to amend x^0 and Y is some subset of $R_V(y_{k+1})$: any proposal in A_k is accepted. If the k th proposer amends x^0 in round 1 then the equilibrium path must lead to the implementation of her ideal policy in A_k (which is independent of the protocol). If protocols in future rounds (i.e. $\{\pi_x\}_{x \neq x^0}$) induced equilibrium paths not leading to the implementation of the k th proposer’s ideal policy in A_k , then she could profitably deviate by offering this policy — which would be accepted — directly in round 1.

3.3 Some quirky properties of the equilibrium correspondence

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

Example 3.3. Suppose that $M = N = \{1, 2, 3\}$, and that any two players can agree to any division of a dollar: $\mathbf{W} = \mathcal{W} = \{S \subseteq N : |S| \geq 2\}$. Take any point $\bar{x} = (\bar{x}_1, \bar{x}_2, \bar{x}_3)$ in the 2-dimensional simplex $\Delta \equiv \{(x_1, x_2, x_3) \in \mathbb{R}_+^3 : x_1 + x_2 + x_3 = 1\}$ at which some player (say, 1) earns less than 50c: $\bar{x}_1 \leq 1/2$. It is well known that

$$V(\bar{x}_1) = \{x \in \Delta : x_2 + x_3 = 1 - \bar{x}_1\}$$

is a vNM (and therefore weakly) stable set: cf. Ordeshook (1986) Ch 9.2; so the union of weakly stable sets for this game is the entire simplex. This would remain true if we changed M to $\{1, 2\}$: \mathcal{W} — and therefore \mathcal{V} — remains unchanged as long as at least two players can propose. Combined with Proposition 1, 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 ($x^0 = (0, 0, 0)$). By contrast, a player who cannot propose earns 0 in Baron and Ferejohn's (1989) closed rule model, and in their open rule game with patient enough players. Furthermore, any policy in the interior of the triangle may be implemented in an equilibrium. Specifically, suppose (without loss of generality) that \bar{x} belongs to the interior of the simplex and that players propose in the order 1, 2, 3, at every default. Using tree $\mathfrak{T}^\pi(x^0, V(\bar{x}_1))$, it is readily checked that there is an equilibrium in which, at the initial default,

- player 3 (the last proposer) would propose $(\bar{x}_1, 0, 1 - \bar{x}_1)$, player 2 would propose $(\bar{x}_1, 1 - \bar{x}_1, 0)$, and player 1 (the first proposer) proposes \bar{x} ;
- the first proposal is accepted: players 1 and 3 vote in favor of x , while player 2 votes against any proposal y such that $y_2 \leq 1 - \bar{x}_1$.

The first proposal is successful in this equilibrium, and it secures the votes of exactly two players. The voting pattern on this proposal satisfies the size principle, yet the three players each earn a share of the dollar. These properties also hold in Baron and Ferejohn's (1989) open rule game with impatient players, and in Diermeier and Fong's (2011) three-player bargaining game with a single proposer (because of their tie-breaking rule: cf. Section 1.2 above). By contrast, the two properties fail in Baron and Ferejohn's (1989) closed rule game and the open rule game with *patient* players, where the final vote satisfies the size principle, but the dissenting voter earns nothing.

□

Example 3.1 (continued). In the last subsection, we constructed tree $\mathfrak{T}^\pi(\{x, z\}, y)$ for Example 3.1 when players propose in the order 1, 2, 3. Two of the paths in Figure

1 have z as a terminal node. Proposition 1 then implies that $\Gamma(\pi, y)$ has an equilibrium in which z is implemented. Preferences in this example imply that y P -dominates z . We therefore conclude that a committee can implement a policy which is P -dominated by the initial default. This property must hold in Bernheim et al's (2006) benchmark pork barrel model as the last proposer's ideal policy is implemented; but it does not hold in Diermeier and Fong (2011), and 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) All players in a majority-rule divide the dollar game may earn a positive surplus in equilibrium;

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

3.4 The model at work

We end this section by illustrating how our model can serve as a workhorse to generate insights for interesting real-world problems. Specifically, we show that our results can explain the power of a proposer who cannot vote in a model based on the Council of Ministers. According to this model, the European Commission (EC) cannot vote, but is the only proposer; all other members of the Council ($N = \{1, \dots, n\}$) can vote. The default is amended if a proposal secures a qualified majority of votes. We take the quota to be a bare majority to simplify exposition: $\mathbf{W} = \{S \subseteq N : |S| > n/2\}$, with n odd.¹⁸

According to Thomson (2011) Ch 7, controversial issues addressed by the Council can be modeled as scalars — so that $X \subseteq \mathbb{R}$ — and, for typical issues, the initial default and the EC 's ideal policy span all of the voters' ideal policies.¹⁹ We adopt these assumptions, and additionally suppose that each player's preference ordering is single-peaked on X . We write player i 's ideal policy as x_i , with countries ordered such that $i < j \Rightarrow x_i \leq x_j$ for

¹⁸In fact, contrary to our model, voting members can propose a policy, which amends the default if it secures all votes.

¹⁹Strictly speaking, this ordering refers to observable policy positions; but Thomson argues (Ch 6) that these positions can be explained by government's characteristics, so ideal policies are ordered in the same way as policy positions. For example, positions on integration issues are explained by domestic opinion on EU membership.

all voters i and j , and write μ for the median voter. Finally, we suppose for concreteness that $x^0 < x_1 < x_n < x_{EC}$; and, more substantively, that players' preferences satisfy the single-crossing property.

We characterize the set of equilibrium policies in this model by following the procedure described above. First, we modify the collection of winning coalitions to account for the EC 's proposal power, thus obtaining $\mathcal{W} = \{S \subseteq C : (S \cap N) \in \mathbf{W} \text{ \& } S \ni EC\}$. Using this collection of winning coalitions, it is easy to confirm that $V = [x_\mu, x_{EC}]$ is the unique weakly stable set. Corollary 1 then implies that we can find the set of equilibrium policies by using tree $\mathfrak{T}^\pi(V, x^0)$.

The immediate successor node of x^0 in the tree must be EC 's ideal policy in $R_V(x^0) \cup \{x^0\}$: in the last proposal of a round, EC (who is the sole and therefore last proposer) would amend the initial default to its ideal policy among those in V that dominate the default. By construction, this policy is uniquely defined as $x^* = \min\{x_{EC}, x'_\mu\}$, where $x'_\mu > x_\mu$ is the policy that makes the median voter indifferent to the initial default (i.e. $x'_\mu \sim_\mu x^0$.) If x^* has a successor node in $\mathfrak{T}^\pi(V, x^0)$ (i.e., if EC can propose more than once) then this node must be unique and equal to x^* itself because $x^* \succ_{EC} x$ for all $x \in R_V(x^*) \setminus \{x^*\}$. Consequently, x^* is the unique equilibrium policy. This example illustrates the possible power of a player who can propose, but not vote. (Example 3.3 above illustrates the power of a player who can vote, but not propose.) Though the EC cannot vote, the equilibrium policy is closer to its ideal policy than the median voter's ideal policy x_μ (indeed: $x_\mu < x^* \leq x_{EC}$). In particular, if $x'_\mu \geq x_{EC}$ then the equilibrium policy coincides with the EC 's ideal policy and, therefore, exceeds all voters' ideal policies. These results are consistent with Schmidt (2000), who argues that EC proposal power can explain why implemented policy has been more integrationist than the ideal policies of the Council's voting members.

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 π after observing the initial default x^0 ; and the game $\Gamma(\pi, x^0)$ is then played. As the chair's choice of protocols only affects her payoff via π_{x^0} (Corollary 3), the chair cannot improve on selecting a protocol which is constant across X .²⁰ Hence, we can without loss of generality restrict attention to constant protocols.

The chair's choice will depend on her expectation of behavior in $\Gamma(\pi, x^0)$ and, therefore, on her predictions of equilibrium policies in that game for every protocol π . Though the potential multiplicity of equilibria makes the comparative statics analysis of equilibrium policies delicate, the following result allows us to draw general conclusions about the chair's preferences over protocols in the well ordered case.

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

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

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

In Proposition 3, the k th 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. We will refer to any such proposer as an 'amender' of x . Proposition 3 says that the ideal policy in $R(x) \cap V$ of the last amender according to π is implemented in every equilibrium $\sigma \in \Sigma^*(\pi, x^0)$ which supports V .

Take any fixed protocol in which the chair is not the last proposer, and consider some equilibrium σ of that game. Now consider a protocol in which the chair proposes last. Proposition 3 implies that this game has an equilibrium in which the chair is at least as well off, and better off if she is an amender of x . The relevant equilibrium supports the same weakly stable set as σ so existence of an equilibrium in the first game implies existence of an equilibrium when the chair proposes last.

Proposition 3 does not imply that the chair cannot lose if she changes to a protocol in which she proposes last because there may be several weakly stable sets, and therefore a multiplicity of equilibrium outcomes. However, we can strengthen results in the last

²⁰Corollary 3 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 Appendix A.2 for a formal proof).

paragraph if there is a unique weakly stable set (say, V); so Corollary 2 implies that there is a unique equilibrium policy for every protocol. In such cases, there is room for protocol manipulation, except in the rather unlikely case where all amenders of x^0 share the same ideal policy among those in $R(x^0) \cap V$. Furthermore, proposing last is an advantageous position in the following sense:

- If the chair is an amender of x^0 then she can never improve on any protocol in which she proposes last: she is at least as well off when she proposes last as when she proposes earlier;
- If the chair is not an amender of x^0 then she does not lose anything by selecting herself as the last proposer: if another protocol is optimal for her then there is a protocol in which she proposes last which is also optimal.

These features are reminiscent of the ‘power of the last word’ in Bernheim et al’s (2006) pork barrel model: they show that the last proposer’s ideal policy 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 R -dominates the initial default rather than her best policy. 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.

Do the conclusions above carry over to cases in which the policy space is not well ordered? As the next example reveals, the answer is no, even if there is a unique weakly stable set: *a chair who is an amender can improve on a protocol in which she is the last proposer.*

Example 4.1. Let $M = N = \{1, 2, 3, 4, 5, 6\}$ and $X = \{x, v_1, v_2, v_3\}$. Preferences over X are given by: $v_3 \succ_1 v_2 \sim_1 v_1 \succ_1 x$; $v_2 \succ_2 v_1 \succ_2 x \succ_2 v_3$; $x \succ_3 v_2 \succ_3 v_3 \succ_3 v_1$; $x \succ_4 v_3 \succ_4 v_2 \succ_4 v_1$; $v_3 \sim_5 v_1 \succ_5 x \succ_5 v_2$; $v_1 \succ_6 x \succ_6 v_2 \succ_6 v_3$. Assume that preferences are aggregated by majority rule: \mathcal{W} is the collection of majority coalitions. It is easily checked that the unique weakly stable set here is $V = \{v_1, v_2, v_3\}$.

Suppose the initial default is $x^0 = x$. Consider first a protocol π_1 in which players 1 and 2 are the penultimate and last proposers, respectively. If player 2 were given the opportunity to make the last proposal in the first round then she would amend the default x with v_1 : v_1 is the only policy in V that P -dominates (and is not P -dominated by) x , and $v_1 \succ_2 x$. If player 1 were given the opportunity to make the penultimate proposal then voters would only vote ‘yes’ if they preferred player 1’s proposal to v_1 . Assuming that indifferent voters vote ‘yes’, player 1 would successfully propose her ideal policy in $R(v_1) \cap V = V$, which is v_3 . Since no policy in $V \setminus \{v_3\}$ R -dominates v_3 (except v_3 itself), proposers who appear before player 2 in protocol π_1 cannot prevent v_3 from being implemented (using the language of the tree: all successor nodes of v_3 are equal to v_3). Thus, *the worst policy of the last amender of x is implemented in equilibrium.*

Now consider another protocol, say π_2 , in which players 2 and 1 are the penultimate and last proposers, respectively. Using the same argument as in the previous paragraph, one can show that v_2 is the unique equilibrium policy when indifferent voters vote ‘yes’. Thus, if player 2 is the chair (and she anticipates that indifferent voters accept proposals) then she strictly prefers protocol π_2 to protocol π_1 .²¹ By the same logic, she strictly prefers any protocol in which she makes the last two proposals to π_1 : the chair’s ideal policy is also implemented by such a protocol.

□

This example shows that when the policy space is not well ordered:

- (i) The chair may optimally choose to make the last two proposals, thereby strictly improving on a protocol in which she only proposes last.²²
- (ii) The chair may optimally choose to (only) make the penultimate proposal, thereby strictly improving on a protocol in which she only proposes last. This contrasts with the equilibria in Bernheim et al (2006) and with our results for well ordered X and a unique weakly stable set, where each amender is at least as well off proposing last.

We summarize the discussion above in

Observation 2. *Suppose that there is a unique weakly stable set. If the policy space is well ordered then the chair is at least as well off proposing last as in any other protocol. However, 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.*

²¹If she anticipates that indifferent voters reject proposals then she is indifferent between π_1 and π_2 , for v_1 is the unique equilibrium policy in both cases.

²²This result should not be confused with Diermeier and Fong’s (2011) demonstration that a single proposer is at least as well off making a take it or leave it offer as playing their game.

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. In Black’s (1958) words:

“The larger the size of majority needed to arrive at a new decision on a topic, the smaller will be the likelihood of the committee selecting a decision that alters the existing state of affairs” p. 99.

We will discuss this conjecture in the context of our model with an arbitrary policy space.

We say that $\Gamma(\pi, x^0)$ 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. In light of our characterization results, the conjecture above can be reformulated as: $q > r$ implies that $\bigcup \mathcal{V}^r \subseteq \bigcup \mathcal{V}^q$ (where $\bigcup \mathcal{V}^s \equiv \{v \in X : v \in V \text{ for some } V \in \mathcal{V}^s\}$). In other words, an increase in the quota (weakly) expands the union of immovable policies. We will refer to this property as *conventional wisdom*.

It is easy to show that conventional wisdom holds if there is enough conflict of interest that X is a weakly stable set or if there is enough common interest that there is a Condorcet winner (which must be the only weakly stable set) with the higher quota. On the other hand, conventional wisdom would fail if some policy is in a weakly stable set if and only if the quota is lower. We provide an example below with a finite, well ordered policy space which satisfies the stronger property that $\bigcup \mathcal{V}^q \subset \bigcup \mathcal{V}^r$: an increase in the quota *contracts* the union of immovable policies.

The intuition for this result is that an increase in the quota makes it easier for a given set of policies to satisfy internal stability, but more difficult for that set to satisfy external stability. However, this intuition does not fully explain the result because the union of weakly stable sets is not necessarily weakly stable:

Example 4.2. Let $M = N = \{1, 2, 3, 4, 5, 6\}$, $X = \{w, x, y, z\}$, and $w \succ_i x \succ_i y \succ_i z$ for $i = 1, 2$, $y \succ_i z \succ_i w \succ_i x$ for $i = 3, 4$, $z \succ_5 w \succ_5 x \succ_5 y$, and $x \succ_6 y \succ_6 z \succ_6 w$. Suppose first that the quota is 4. Applying the definition of weakly stable sets, it is readily checked that there are two weakly stable sets in this case: $\{w, y\}$ and $\{x, z\}$; so $\bigcup \mathcal{V}^4 = X$. Note that x and z cannot form a weakly stable set with w and y because the internal stability condition, (IS_{P_4}) , would fail: wP_4x , zP_4w , xP_4y , and yP_4z .

Now suppose that the quota becomes 5. As wP_5x and yP_5z , internal stability (IS_{P_5}) does not allow for more weakly stable sets than before the increase in the quota. Furthermore, the increase in the quota implies that $\{x, z\}$ no longer satisfies external stability: while zP_4w and xP_4y , $\neg(zP_5w)$ and $\neg(xP_5y)$. As a result, $\{w, y\}$ is now the only weakly stable set (as w and y P_5 -dominate x and z , respectively). Thus, conventional wisdom fails in this example in the strong sense that an increase in the quota contracts the union of immovable policies: $\{w, y\} = \bigcup \mathcal{V}^5 \subset \bigcup \mathcal{V}^4 = X$.

□

We summarize the arguments above in

Proposition 4. *Suppose that $\Gamma(\pi, x^0)$ is a quota game. An increase in the quota weakly expands the set of equilibrium decisions if there is a Condorcet winner for the higher quota or if X is a weakly stable set. However, an increase in the quota may contract the set of equilibrium decisions in other intermediate cases.*

Black's conjecture underlay a series of reforms in the EU, which introduced qualified majority voting (QMV) in the Council of Ministers to expedite legislation. Our analysis above suggests an alternative perspective: if Council proceedings corresponded to our model then introduction of QMV might, perversely, have prevented the Council from amending the initial default. This may help to explain why the introduction of QMV has not significantly reduced the 80% of legislation that the Council has passed without a final vote: cf. Heisenberg (2005).

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. Indeed, Example 4.1 above demonstrates that a player may prefer to make the last two proposals than just the last proposal in each round. 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 various reasons:

- In Section 3.3, we used Example 3.3 to demonstrate that a player who does not propose may earn all of the surplus in an equilibrium. Adding another proposer does not change the weakly stable sets; and there are then equilibria in which the new proposer earns less than all of the surplus.
- In Section 4.1.2, we demonstrated 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, and there is a unique weakly stable set for each set of proposers (so that each associated game has a unique equilibrium policy):.

Example 4.3. Let $N = \{1, 2, 3, 4, 5\}$ and $X = \{w, x, y, z\}$. Preferences over X are given by: $z \succ_1 y \succ_1 w \succ_1 x$; $x \succ_2 y \succ_2 z \succ_2 w$; $z \succ_3 w \succ_3 x \succ_3 y$; $w \succ_4 x \succ_4 y \succ_4 z$;

and $x \succ_5 w \succ_5 y \succ_5 z$. Assume that the initial default is $x^0 = x$ and that defaults are amended by majority rule: \mathbf{W} is the collection of majority coalitions.

Suppose first that player 1 is the only proposer (i.e. $M = \{1\}$). This implies that \mathcal{W} is the collection of majority coalitions that include 1: $\mathcal{W} = \{S \in \mathbf{W} : 1 \in S\}$. It is readily checked that the unique weakly stable set here is $V = \{x, y, z\}$. Proposition 2 then implies that x is the unique equilibrium policy. Intuitively, since no policy in V is preferred to x by a majority coalition, it is impossible for 1 to amend it. Consequently, the proposer's worst policy must be implemented.

Now suppose that player 2 is given the opportunity to make a proposal after player 1, so that the set of proposers becomes $M' = \{1, 2\}$. The collection of winning coalitions $\mathcal{W}' = \{S \in \mathbf{W} : \{1, 2\} \cap S \neq \emptyset\}$, thus yielding a unique weakly stable set $V' = \{w, y\}$. The second (and last) proposer, player 2, would never amend the initial default x , which is her ideal policy, in equilibrium. Anticipating this, voters accept the first proposer's proposal if and only if they prefer the latter to x . Player 1 will therefore amend x to her ideal policy in $R(x) \cap V'$, which is w . This implies that w , player 2's worst policy in X , is now the unique equilibrium policy. In contrast, player 1 is now strictly better off, as $w \succ_1 x$.

□

Examples in which a player may be worse off when given the opportunity to propose are easy to concoct in final voting games with a finite horizon. In such games, the intuition is simple. The last amender cannot commit to pass. Her predecessor may therefore amend the default to the last amender's disadvantage, knowing that the latter would otherwise amend the default. In Example 4.3, however, the logic is different: for the initial default x is player 2's ideal policy and, therefore, player 1 does not expect her to amend it.

The problem for player 2 is that she cannot commit to pass *at all possible defaults*. Adding player 2 to the set of proposers changes the P -dominance relation and, consequently, the set of immovable policies which player 1 can successfully propose in equilibrium. In particular, player 1's ideal policy z , which was initially immovable, would now be amended to y by player 2 (off the equilibrium path). This in turn makes player 2's worst policy w immovable: changing w to z would lead to y being implemented, thus making a majority of voters (i.e. players 4, 5, and 6) worse off. Moreover, w is the only immovable policy which is majority preferred to the initial default x . In the first round, it is thus optimal for player 1 to propose w , which is accepted and never amended.

We record the conclusion from this example as

Observation 3. *A player may be strictly worse off if she is given the opportunity to*

propose.

Corporate governance reform often aims to extend shareholders' scope to propose policies. Opponents (e.g. the Business Roundtable) argue that it is disadvantageous to give such power to uninformed shareholders. Observation 3 shows that this property may also hold when shareholders are informed.²³ This is consistent with evidence that allowing shareholders to propose has a negative and/or insignificant effect on shareholder value. (See, for example, Akyol et al (2010) and Larcker et al (2011) on proxy access.)

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 is determined by 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 common discount factor $\delta \simeq 1$ is close to that in the variant on our model. This observation suggests that there is $\delta < 1$ such that an equilibrium strategy combination in our model (or, more precisely, in the related model) might be an equilibrium in the repeated implementation model. We explore such an intuition in this subsection.

More specifically, we consider a variant of $\Gamma(\pi, x^0)$ in which the bargaining process continues *ad infinitum*. At the end of each round $t \in \mathbb{N}$, default policy x^t is implemented and each player i receives an instantaneous payoff $(1 - \delta)u_i(x^t)$, 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 $\{x^t\}_{t=1}^\infty$ is $(1 - \delta) \sum_{t=1}^\infty \delta^{t-1} u_i(x^t)$. We will refer to such a game as $\Gamma^\delta(\pi, x^0)$, and will say that an equilibrium of $\Gamma^\delta(\pi, x^0)$ is 'absorbing' if there is a

²³The proposal which is put to the AGM for a final vote is typically negotiated by shareholders and management prior to the meeting. Our model refers to these negotiations, rather than to the final vote itself.

round T such that $x^t = x^T$ for every subsequent round: $t > T$.²⁴ We abuse terminology in this subsection by identifying equilibria in our model with equilibria in the related model with continued (but payoff-irrelevant) bargaining.

Our next result confirms the intuition above in the finite, well ordered case.

Proposition 5. *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 $\Gamma(\pi, x^0)$ if and only if it is an absorbing stationary Markov equilibrium of $\Gamma^\delta(\pi, x^0)$.*

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

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

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 absorbing equilibria of $\Gamma^\delta(\pi, x^0)$ and $\Gamma(\pi, x^0)$: only final (or limit) policies matter. Put differently, $x \succ_i y$ if and only if player i strictly prefers any sequence of defaults converging to x to any sequence converging to y in the repeated implementation model. This may not be true if X comprises a continuum: even though δ is close to 1 and $x \succ_i y$, $u_i(x) - u_i(y)$ may be so small that player i prefers the sequence of defaults leading to y over that leading to x .

Furthermore, as the following example illustrates, Proposition 5 does not hold when X is not well ordered.

Example 3.3 (continued). Consider again the divide-the-dollar game from Example 3.3. Recall that the set of policies $(x_1, x_2, x_3) \in \Delta$ at which player 1 earns 50c, $V(1/2)$, is a weakly stable set. From Proposition 1, therefore, there are equilibria of $\Gamma(\pi, x^0)$ that support $V(1/2)$. In particular, the equilibrium σ constructed in the proof of Proposition 1 prescribes the committee to implement a policy in $V(1/2)$ without delay: in any subgame starting with an ongoing default $x \notin V(1/2)$, x is amended to some $v_x \in V(1/2)$ in the first round.

Nevertheless, irrespective of the value of δ , σ is not an equilibrium of $\Gamma^\delta(\pi, x^0)$. To see this, consider a round in which the ongoing default is $v_0 \equiv (1/2, 1/2, 1/2) \in$

²⁴Existence 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).

$V(1/2)$. Suppose that player 1 (if given the opportunity to amend v_1) proposes $y = (1/2 + \epsilon/2, 1/2 - \epsilon, 1/2 + \epsilon/2)$ for any small and positive ϵ . If player 1's proposal were accepted then, by construction of σ , a policy $v_y \in R(y) \cup V(1/2)$ would be implemented in all future periods. As v_y must R -dominate y , it must be of the form $v_y = (1/2, 1/2 - \epsilon + \gamma_2, 1/2 + \epsilon/2 + \gamma_3)$ where $\gamma_i \geq 0$ and $\gamma_2 + \gamma_3 = \epsilon/2$. Hence, players 1 and 3 receive higher payoffs when they accept 1's offer to amend v_0 to y than when they reject it:

$$(1 - \delta)u_1(y) + \delta u_1(v_y) = \frac{1 + (1 - \delta)\epsilon}{2} > \frac{1}{2} = u_1(v_0) ,$$

and

$$(1 - \delta)u_3(y) + \delta u_3(v_y) = \frac{1 + \epsilon}{2} + \delta\gamma_3 > \frac{1}{2} = u_3(v_0) .$$

This proves that player 1 has a profitable deviation from σ_1 and, therefore, that σ is not an equilibrium of $\Gamma^\delta(\pi, x^0)$.

□

If we weaken the equilibrium concept by only requiring approximate best responses, however, we can obtain an analog of Proposition 5 in terms of equilibrium policies. Indeed, Proposition 6 below states that the set of immovable policies in any equilibrium of $\Gamma(\pi, x^0)$ is also the set of absorbing policies in some contemporaneous perfect ε -equilibrium (Mailath et al., 2005) of $\Gamma^\delta(\pi, x^0)$ when δ is close enough to 1.

Proposition 6. *For any $\varepsilon > 0$, there exists $\delta_\varepsilon < 1$ such that the following statement is true for all $\delta > \delta_\varepsilon$: If σ is an equilibrium of $\Gamma(\pi, x^0)$ then there is an absorbing contemporaneous perfect ε -equilibrium of $\Gamma^\delta(\pi, x^0)$ whose absorbing policy from any ongoing default $x \in X$ is $f^\sigma(x)$.*

Proposition 6 thus 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.

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 outcomes 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 $i \in S$. The closure of any consistent set is therefore consistent.

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*, \bar{Z} : Z consistent implies $Z \subseteq \bar{Z}$ and \bar{Z} is itself consistent. Thus, \bar{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 $\bar{Z} \notin \mathcal{V}$. In this subsection, we weaken stationarity, and show that the ensuing set of equilibria supports 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 $(x^0, S^1, x^1, \dots, S^{t-1}, x^{t-1})$ 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 — $H^1 \equiv \{x^0\}$ being the null history — and let $H \equiv \bigcup_{t=1}^{\infty} H^t$ be the set of partial histories. We define a ‘semi-Markovian’ strategy as an analog of a stationary Markov strategy where partial histories play the role of the ongoing default. More specifically: in proposal stages, strategies only depend on the partial history and the identity of the remaining proposers in the current round; in voting stages, strategies only depend on the partial history, the proposal just made, the votes already cast thereon, and the remaining proposers in the current round.

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 partial history $h \in H$ and every $k \in \{1, \dots, m_{x^{t-1}}\}$ the unique final outcome $\phi^\sigma(h, k)$ eventually implemented (given σ) when h is the current

partial history and the k th proposer is about to move. We are particularly interested in $\phi^\sigma(x^0, 1)$, which describes the policy implemented in $\Gamma(\pi, x^0)$ if players act according to σ . We will sometimes abuse notation by writing $\phi^\sigma(x^0)$ instead of $\phi^\sigma(x^0, 1)$.

We now turn to the characterization of semi-Markovian equilibria — i.e., subgame perfect equilibria of $\Gamma(\pi, x^0)$ 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 is a consistent set then each length m_{x^0} path of tree $\mathfrak{T}(Z, x^0)$ ending with a policy in Z describes behavior in round 1 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 7. *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 $\Gamma(\pi, x^0)$, σ , such that the final policy outcome reached from x^0 must belong to Z . Inspection of the proof (in Appendix A.1) 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 8. *If σ is a semi-Markovian equilibrium then $\phi^\sigma(H) \equiv \bigcup_{h \in H} \phi^\sigma(h, 1)$ is a consistent set.*

Thus, for any semi-Markovian 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 7 that any policy $z \in \overline{Z}$ is the outcome of a semi-Markovian equilibrium of $\Gamma(\pi, z)$. Consequently, we have

Corollary 4. *Let $\Sigma^{NM}(\pi, x^0)$ be the set of semi-Markovian equilibria of $\Gamma(\pi, x^0)$. 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:*

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

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.

5.3 Open rule bargaining

Thus far, we have focused on games which end when no proposer amends a default. We have demonstrated that there is an equilibrium which supports the closure of any weakly stable set, and that every equilibrium supports a weakly stable set (Propositions 1 and 2 above). Baron and Ferejohn's (1989) open rule model has a different stopping rule. Their game only ends when a proposer successfully 'moves the previous question': putting the existing default to an up-down vote. In further contrast to our model, a new round starts at default x if no proposer in $\{1, \dots, m_x\}$ has amended x or successfully moved x (the previous question). In this subsection, we argue that Proposition 1 holds, but that Proposition 2 fails in this variant on our model, which we dub open rule bargaining. (In contrast to Baron and Ferejohn's version, where proposers are selected at random, the protocol determines the fixed order in which players propose in any round.)

We can prove the analog of Proposition 1 by constructing an equilibrium strategy combination which supports any weakly stable set V :

If the default (x) is outside V then it is R -dominated by some policy $y^*(x) \in V$, so let any $k \in M$ who prefers $y^*(x)$ over x propose the former, and any other proposer pass; and if $x \in V$ then let every proposer move the previous question. To simplify subsequent exposition, write $y^*(x) = x$ whenever default $x \in V$. This means, in particular, that $y^*(x)$ is in V for any default x . We now turn to voting behavior. Suppose, first, that some $k \in M$ has proposed to amend x to y . If $y = y^*(x)$ then let i vote for y if and only if i weakly prefers y over x ; and if $y \neq y^*(x)$ then let i vote for y if and only if i strictly prefers $y^*(y)$ over $y^*(x)$. Finally, if some $k \in M$ has moved the previous question then i votes in favor if and only if i weakly prefers $y^*(x)$ over x . It is easy to confirm that this strategy combination forms an equilibrium, at which defaults in V are implemented, and defaults outside V are amended to a policy in V which is then implemented.

The argument above implies that the constructed strategy combination supports V , by analogy to Proposition 1 above. Furthermore, any policy set supported by an equilibrium must satisfy internal stability, else a proposer could profitably deviate to amending some policy in the set. However, equilibria may support sets of policies which are not externally stable. To see this, consider Example 3.2, where two players bargaining over division

of a pie. The only weakly stable set is the set of divisions, but the following strategy combination is an equilibrium. If the default (x) does not entail equal division of the pie ($1/2$) then a player who gets less than $1/2$ proposes amending x to $1/2$, while any other player passes; and both players vote in favor when $1/2$ is moved. Any player who gets less than $1/2$ at y vetoes amending x to y , and also votes against if y is moved.

In sum,

Observation 4. *Any weakly stable set can be supported in an open rule bargaining game, and the policy set supported by any equilibrium must satisfy internal stability. However, there may be equilibria that support policy sets which fail external stability.*

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 be 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 and there is a unique weakly stable set. We have also shown, *inter alia*, that an increase in the quota can contract the union of immovable 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 by appropriately redefining 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).

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, mixed strategy equilibria may exist. 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. 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 Appendix A.3.) It would be interesting to extend our analysis to mixed strategy equilibria.

Refinements

Although weakly stable sets may not exist, simple games often have multiple weakly stable sets, implying equilibrium multiplicity in our noncooperative game. One possible approach is to focus on equilibria which support cooperative refinements of V , such as Wilson's (1971) main simple stable sets. However, this approach seems unattractive in a noncooperative setting. 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 $\{\sigma^k\}$ such that $\sigma^k \rightarrow \sigma$ and σ prescribes each ‘agent’ — MTHPE is defined in the agent-strategic form — a best response to her opponents’ perturbed strategies in σ^k for all $k = 1, \dots, \infty$.

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

Observation 5. *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 weakly stable sets).*

The proof of this observation (which is provided in Appendix A.4) shows that something even stronger is true: for *every* weakly stable set $V \in \mathcal{V}$, there is a (pure strategy) MTHPE σ that supports V : $f^\sigma = V$. This reinforces the noncooperative foundations that our bargaining model provides for weakly stable sets. We conjecture that equilibria in our model would survive other plausible refinements.

A Appendix

A.1 Proofs of Propositions

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 $(y_1(x), \dots, y_{m_x+1}(x))$ such that:

- if $x \in V$, then $f(x) = y_1(x) = \dots = y_{m_x+1}(x) = x$;
- if $x \notin V$, then $f(x) = y_1(x) \in V$, $x = y_{m_x+1}(x)$, and $y_k(x) \in \mathfrak{s}_k^\pi(V, y_{k+1}(x))$ for each $k = 1, \dots, m_x$. The latter condition implies that $y_k(x)$ is one the k th 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 voting behavior in the putative equilibrium strategy profile σ . 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 k th proposer has just proposed to change x to $y \neq x$, σ prescribes player 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 has a profitable one-shot deviation from σ . Step 3 demonstrates that there is no proposal stage in which a proposer has a profitable one-shot deviation from σ . Steps 2 and 3 jointly imply 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, $f^\sigma(x) = x$ 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 ensures that this set is nonempty), and suppose that the l th proposer is given the opportunity to make a proposal. By construction of $(y_1(x), \dots, y_{m_x}(x))$, this implies that $y_l(x) \in V$ and therefore $y_l(x) = y_1(y_l(x)) \in A_l(V, y_{l+1}(x))$, where $y_{m_x}(x) \equiv x$. 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. If she passes (so $y_{l-1}(x) = y_l(x)$) then, by construction, $y_{l-1}(x)$ is implemented at the end of the next round. If she proposes an amendment then 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 $\{i \in N : y_{l-1}(x) \succeq_i y_l(x)\} \in \mathbf{W}$; so condition (B) implies that $y_{l-1}(x)$ is accepted and implemented at the end of the next round. In sum, $y_{l-1}(x)$ is accepted and implemented at the end of the next round.

Repeating this argument recursively for every $l = 1, \dots, l - 2$, 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 k th 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 $f^\sigma(y, 1) = y_1(y) \in V$.

Any $x \in V$ is implemented at the end of round t if the k th proposer fails to amend it: for, by definition of the proposer strategies, all the remaining proposers will pass. Hence, $f^\sigma(x, k+1) = x$. Consequently, $f^\sigma(y, 1) \succ_i f^\sigma(x, k+1)$ is equivalent to $y_1(y) \succ_i x$, which in turn implies that player 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 that $x \succ_i y_1(y)$. 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 $y_{k+1}(x) = \dots = y_{m_x}(x) = x = f^\sigma(x, k+1)$. Now suppose that the l th proposer is the next proposer (after the k th) to make a successful proposal, $y_l(x) \neq x$. By construction, this implies that $y_{k+1}(x) = \dots = y_l(x) \in V$. Consequently, $f^\sigma(x, k+1) = f^\sigma(y_l(x), 1) = y_l(x) = y_{k+1}(x)$.

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 player i votes ‘yes’. Similarly, $f^\sigma(x, k+1) \succ_i f^\sigma(y, 1)$ implies that she votes ‘no’.

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

If $x \in V$ then σ prescribes the k th proposer to pass (i.e., $y_k(x) = x$). If she has a profitable deviation at this stage, then she must be able to amend x to some y 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. As proposal y is successful, 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 $x \in V$. This is impossible because V satisfies (IS_P) .

If $x \notin V$ then σ prescribes the k th 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 proposes 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. $y_1(y) \in R_V(y_{k+1}(x))$), and $f^\sigma(y, 1) = y_{k+1}(x)$ otherwise. Such a deviation cannot be profitable because $y_k(x)$ is $\succeq_{\pi_x(k)}$ -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 amend y to x , contrary to our supposition that σ is 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 for 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 which starts with x as the ongoing default policy. Suppose that 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 the first 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_x th 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 . Voter i must therefore 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_x th 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_x th 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_x th proposer, for she can always pass). If $\sigma_{\pi_x(m_x)}$ prescribes the m_x th 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 $y_{m_x} \equiv f^\sigma(x, m_x) \in \mathfrak{s}_m^\pi(V, x)$.

Proceeding recursively, one can use the same argument to show that, for each $k = 1, \dots, m_x - 1$, $y_k \equiv f^\sigma(x, k) \in \mathfrak{s}_k^\pi(V, x)$: 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 $(y_1, \dots, y_{m_x}, x) \equiv (f^\sigma(x, 1), \dots, f^\sigma(x, m_x), x)$ 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 k th 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.: $\mathbf{M}(\succ_{\pi(l)}, R(x) \cap f^\sigma(X))$ — 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 k th 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 $\mathbf{M}(\succ_{\pi(l)}, R(x) \cap f^\sigma(X)) \in V$: 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 5

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

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

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 $\Psi_i(x, y, \delta) \rightarrow u_i(x) - u_i(y) > 0$ as $\delta \rightarrow 1$, there exists $\delta_i(x, y) \in [0, 1)$ such that $\Psi_i(x, y, \delta) > 0$ for all $\delta > \delta_i(x, y)$. 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 $\Gamma(\pi, x^0)$. 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 $\Gamma(\pi, x^0)$, 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 $\Gamma^\delta(\pi, x^0)$. This implies that there are two finite sequences $\{x_t\}_{t=1, \dots, T_x}$ and $\{y_t\}_{t=1, \dots, T_y}$, and a policy $y \in X$ such that

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

and $u_i(y) < u_i(x)$ (recall that a one-stage deviation from an equilibrium strategy in $\Gamma(\pi, x^0)$ 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 $\Gamma^\delta(\pi, x^0)$.

Now suppose that σ is an absorbing stationary Markov equilibrium of $\Gamma^\delta(\pi, x^0)$. 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 $\Gamma^\delta(\pi, x^0)$, and let $\{x_t\}_{t=1, \dots, T_x+1}$ 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 $\{y_t\}_{t=1, \dots, T_y+1}$ (with $y = y_{T_y+1}$ being implemented indefinitely) she could induce by a one-shot deviation satisfies:

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

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 $\Psi_i(y, x, \delta) > 0$, so that

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

a contradiction. At the equivalent stage in game $\Gamma(\pi, x^0)$, $u_i(x) > u_i(y)$ 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 $\Gamma(\pi, x^0)$, this proves that σ is an equilibrium of $\Gamma(\pi, x^0)$.

Proof of Proposition 6

Let σ be an equilibrium of $\Gamma(\pi, x^0)$. By Proposition 2, the outcome function induced by σ , $f^\sigma(\cdot)$, is a selection of $F^\pi(V, \cdot)$ where V is the weakly stable set supported by σ . Using the construction in the proof of Proposition 1, we can construct an equilibrium $\bar{\sigma}$ that implements the same selection of $F^\pi(V, \cdot)$ without any delay; that is, $f^{\bar{\sigma}}(x) = f^\sigma(x)$ for all $x \in X$. In other words, while σ may prescribe a sequence of amendments that leads to $f^\sigma(x, k)$, $\bar{\sigma}$ prescribes the k th proposer to directly offer $f^\sigma(x, k)$ for any ongoing default $x \in X$. This offer is always accepted by a winning coalition of voters, so that it never takes more than a round to reach an agreement in any subgame. To prove Proposition 6, we will now show that $\bar{\sigma}$ is a contemporaneous perfect ε -equilibrium of $\Gamma^\delta(\pi, x^0)$ for all $\varepsilon > 0$.

Fix $\varepsilon > 0$. We first construct δ_ε . For each $i \in N$ and every pair $(x, y) \in X^2$ such that $u_i(x) \geq u_i(y)$, let

$$\Lambda_i(x, y, \delta) \equiv \delta u_i(x) + (1 - \delta) \underline{u}_i - \delta u_i(y) - (1 - \delta) \bar{u}_i .$$

Since $\Lambda_i(x, y, \delta) \rightarrow u_i(x) - u_i(y) \geq 0$ as $\delta \rightarrow 1$, $\delta_\varepsilon^i(x, y) \equiv \max \{d : \Lambda_i(x, y, d) + \varepsilon \leq 0\}$ is well defined and less than 1. From now on, we assume that

$$\delta > \delta_\varepsilon \equiv \max_{i \in N} \max_{x, y \in X: x \succeq_i y} \delta_i^i(x, y) < 1 .$$

(As X is compact and \succeq_i is continuous, $\{(x, y) \in X^2 : x \succeq_i y\}$ is compact.)

As $\bar{\sigma}$ is an equilibrium of $\Gamma(\pi, x^0)$, no player i has a profitable one-shot deviation from $\bar{\sigma}_i$ (given $\bar{\sigma}_{-i}$) at any stage of this game. Consider an arbitrary stage of $\Gamma(\pi, x^0)$ with current default x_0 , and let x be the final policy outcome if i does not deviate from $\bar{\sigma}_i$ in that stage. Hence, any other policy outcome $y \neq x$ she could induce by a one-shot deviation satisfies: $u_i(y) \leq u_i(x)$. Suppose that, contrary to the statement of the result, i has a profitable one-shot deviation, which amends the default to y_0 at the same stage in $\Gamma^\delta(\pi, x^0)$ (taking the ‘deviation cost’ ε into account). Let $\bar{\sigma}$ prescribe implementation of policy y in this subgame. This implies that

$$(1 - \delta)u_i(y_0) + \delta u_i(y) - \varepsilon > (1 - \delta)u_i(x_0) + \delta u_i(x)$$

and $u_i(y) \leq u_i(x)$. This is impossible when $\delta > \delta_\varepsilon$. Hence, $\bar{\sigma}$ is a contemporaneous perfect ε -equilibrium of $\Gamma^\delta(\pi, x^0)$.

Proof of Proposition 7

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’. Moreover, we denote by (h^t, S, y) the concatenation of the round- t history h^t followed by a round in which coalition $S \in \mathcal{W}$ amends x^{t-1} to y .

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 $(\tau_l) \in (\mathbb{N} \cup \{\emptyset\})^\infty$ to construct σ . For a given history,²⁵ 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 $\tau_1 = \emptyset$ 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 $(z_1(h), \dots, z_{m_x+1}(h))$ 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 $k = 1, \dots, m_x$. The latter condition implies that $z_k(h)$ is one of the k th 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

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

became the new default; if that has not happened so far then we let $\tau_l = \emptyset$. In particular, if h is a round- t history such that $\tau_l = \emptyset \neq \tau_{l-1}$ and $t \neq \tau_{l-1} + 1$ then we write $h \in H_l$. By definition of τ_{l-1} , $x^{\tau_{l-1}} \in Z$. Since Z is consistent, $Z_l \cap \{z \in Z : zRx\}$ is nonempty for all $x \in X$. Using the tree $\mathfrak{T}^\pi(Z_l, x^{\tau_{l-1}})$, we can then obtain a vector $(y_1(h), \dots, y_{m_x+1}(h))$ such that: $x = z_{m_x+1}(h)$, and $z_k(h) \in s_k^\pi(Z_l, z_{k+1}(h))$ for each $k = 1, \dots, m_x$. The latter condition implies that $z_k(h)$ is one of the k th 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 k th proposer tries to “punish” at least one of the “deviators” in $S^{\tau_{l-1}+1}$ for not rejecting the proposer’s offer to amend $x^{\tau_{l-1}}$ in round $\tau_{l-1} + 1$.

²⁵To simplify the notation, we omit the sequence’s dependence on the history under consideration.

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 $h \in H_0$, 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 k th 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 $z_1(h, S, y) \succ_i x^{t-1}$ 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 $z_1(h, S, y) \succ_i z_{k+1}(h)$ for any winning coalition $S \ni i$;

where $Z_1 \equiv Z$.

We establish the statement of Proposition 7 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 players, this proves that σ is an equilibrium.

Claim 1: Consider the round following a history $h \in H$, and suppose the k th 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: $z_{k+1}(h) = \dots = z_{m+1}(h) = x^{t-1}$ (all the remaining proposers pass). Accordingly, suppose that $h \in H_l$ with $l \neq 0$. Since the k th 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_x th 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 k th 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 k th 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 k th 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 k th 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 winning coalition then σ_i prescribes i to vote 'yes' only if $\phi^\sigma(h, S, y; 1) \succeq_i \phi^\sigma(h; k+1)$, and to vote 'no' only if $\phi^\sigma(h; k+1) \succeq_i \phi^\sigma(h, S, y; 1)$.

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

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 $(h, S, y) \in H_0$, which in turn implies that $\phi^\sigma(h, S, y; 1) = y$ (all proposers will pass at a history in H_0). Hence, $\phi^\sigma(h, S, y; 1) \succeq_i \phi^\sigma(h; k+1)$. If player i votes ‘no’ then, by condition (B), $z_{k+1}(h) \succ_i y$. This in turn implies that $\phi^\sigma(h; k+1) \succ_i \phi^\sigma(h, S, y; 1)$ and, therefore, that $\phi^\sigma(h; k+1) \succeq_i \phi^\sigma(h, S, y; 1)$.

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), $z_1(h, S, y) \succ_i z_{k+1}(h)$. This implies that $\phi^\sigma(h, S, y; 1) \succ_i \phi^\sigma(h; k+1)$ and, therefore, that $\phi^\sigma(h, S, y; 1) \succeq_i \phi^\sigma(h; k+1)$. Similarly, if i votes ‘no’ then (C) implies that $z_{k+1}(h) \succeq_i z_1(h, S, y)$ and then $\phi^\sigma(h; k+1) \succeq_i \phi^\sigma(h, S, y; 1)$.

Claim 5: Let $h \in H^t$ be a history ending with default $x^{t-1} = x$. At this history, the k th 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

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

Therefore, for each coalition $S \in \mathcal{W}$, there is at least one member of S who weakly prefers $z_{k+1}(h)$ to $z_1(h, S, y)$. 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 $\phi^\sigma(h, S, y; 1) = z_1(h, S, y) \notin A_k(Z_l, z_{k+1}(h))$ must be unsuccessful. Indeed, condition (C) in the definition of voting histories implies that voters only vote ‘yes’ if they strictly prefer $z_1(h, S, y) \in Z_l$

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 on proposing $z_k(h)$ when $z_k(h) \neq z_{k+1}(h)$, and passing otherwise.

Proof of Proposition 8

Let σ be a semi-Markovian equilibrium. Suppose that, contrary to the statement of Proposition 8, $\phi^\sigma(H)$ 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 $\pi_o(\{1, \dots, m_o\}) \cap S$ 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.

A.2 The Dynamic Game with Endogenous Protocol (footnote 20)

Let Π be the set of all protocols as defined in Section 2. In Section 4, we analyzed a game in which the chair first selects a protocol $\pi \in \Pi$, and $\Gamma(\pi, x^0)$ is then played. Call this game: $\Gamma^c(\pi, x^0)$: the superscript stands for commitment. In this section, we establish the claim that Corollary 3 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 $\Gamma^d(\Pi, x^0)$ 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 $\Gamma^d(\Pi, x^0)$ typically differs from that in Harsanyi (1974).

Corollary 3 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 $\pi^d(x^0, \Pi)$. Let $\pi^c(x^0, \Pi)$ be an equilibrium choice in $\Gamma^c(\pi, x^0)$. A chair who could commit to protocols could always do at least as well as the chair in $\Gamma^d(\pi, x^0)$ by choosing $\pi^c(x^0, \Pi) = \pi^d(x^0, \Pi)$. Conversely, the chair in $\Gamma^d(\Pi, x^0)$ could always do at least as well as the chair in $\Gamma^c(\Pi, x^0)$ by replicating $\pi^c(x^0, \Pi)$. We therefore conclude that the same set of policies can be implemented in an equilibrium of $\Gamma^c(\Pi, x^0)$ as in an equilibrium of $\Gamma^d(\Pi, x^0)$. In each case, an equilibrium protocol at x^0 is a best protocol in the class of games analyzed in Section 3.

A.3 Mixed strategy equilibria

In this section, 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
Players (i)	1	2	1	0
	2	0	2	1
	3	1	0	2
		Utilities $u_i(w)$		

and that the protocol is given by

$$\begin{aligned}
 (\pi_x(1), \pi_x(2), \pi_x(3)) &= (2, 3, 1) ; \\
 (\pi_y(1), \pi_y(2), \pi_y(3)) &= (3, 1, 2) ; \\
 (\pi_z(1), \pi_z(2), \pi_z(3)) &= (1, 2, 3) .
 \end{aligned}$$

Consider the following strategy combination. At any default, each player proposes her ideal 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 $Y_u^w(2p_w^v + p_w^u)$.

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

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

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

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

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

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

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

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.

A.4 Markov trembling-hand perfect equilibria

In this section, we provide a proof of Observation 5. To prove this result, it suffices to show that, for every weakly stable set $V \in \mathcal{V}$, there is an MTHPE 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 $k \in \{1, \dots, m_x\}$, the k th proposer offers $y_k(x)$ — i.e.: the $\succ_{\pi_x(k)}$ -maximal element in $P_V(y_{k+1}(x)) \cup \{y_{k+1}(x)\}$ — and voter $i \in N$ accepts this proposal if and only if $y_1(y_k(x)) \succ_i y_{k+1}(x)$ — 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 the proposer offer her ideal policy in $V \cup \{x\}$, where x is the ongoing default.

By construction, σ must be an equilibrium of $\Gamma(\pi, x^0)$ and $f^\sigma(X) = V$. (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.

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

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

where \bar{v} is a (completely mixed) voting profile such that the probability that each element of V is 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, $\sigma^m \rightarrow \sigma$ as $m \rightarrow \infty$.

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 k th proposer at this history, $i = \pi_x(k)$. Let $p_k^m(y)$ be the probability that i 's proposal y is accepted, $V_i^m(y)$ i 's expected payoff 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 player i 's ideal policy in V by y_i , 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,

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

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, $v_i^m \rightarrow u_i(x)$ and $p_k^m(y) = p_k^m(y')$ for all $y, y' \in V$. Moreover, by construction of σ , $V_i^m(y) \rightarrow y_1(y) \in V$ 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: $V_i^m(y_i) \rightarrow u_i(y_i) > u_i(y) \leftarrow V_i^m(y)$ 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 (1) 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 $v_i^m \rightarrow u_i(x)$ and $V_i^m(y) \rightarrow u_i(y)$ 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 m is arbitrarily large, thus completing the proof.

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* forthcoming.
- Akyol, A., W. Lim and P. Verwijmeren (2010), "Shareholders in the Boardroom", mimeo.
- Anesi, V. (2006), "Committees with Farsighted Voters: A New Interpretation of Stable Sets" *Social Choice and Welfare* 27, 595-610.
- Anesi, V. (2010), "Noncooperative 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.
- Battaglini, M. and S. Coate (2007), "Inefficiency in Legislative Policy-Making: A Dynamic Analysis" *American Economic Review* 97, 118-149.
- Battaglini, M. and S. Coate (2008), "A Dynamic Theory of Public Spending, Taxation, and Debt" *American Economic Review* 98, 201-236.
- Battaglini, M. and T. Palfrey (forthcoming), "The Dynamics of Distributive Politics" *Economic Theory* forthcoming.
- Bernheim, B.D., A. Rangel and L. Rayo (2006), "The Power of the Last Word in Legislative Policy Making" *Econometrica* 74, 1161-1190.
- Black, D. (1958), *The Theory of Committees and Elections* Cambridge University Press, Cambridge.
- 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 (2011), "A Newton Collocation Method for Solving Dynamic Bargaining Games" *Social Choice and Welfare* 36, 611-650.
- Duggan, J. and T. Kalandrakis (forthcoming), "Dynamic Legislative Policy Making" *Journal of Economic Theory*.
- Dutta, B., M.O. Jackson and M. Le Breton (2004), "Equilibrium Agenda Formation" *Social Choice and Welfare* 23, 21-57.

- Dziuda, W. and A. Loeper (2010), "Dynamic Collective Choice with Endogenous Status Quo" mimeo.
- Harsanyi, J.C. (1974), "An Equilibrium-Point Interpretation of Stable Sets and a Proposed Alternative Definition" *Management Science* 20, 1472-1495.
- Hayes-Renshaw, F., W. van Aken and H. Wallace (2006), "When and Why the EU Council of Ministers Votes Explicitly" *Journal of Common Market Studies* 44, 161-194.
- Heisenberg, D. (2005), "The Institution of 'Consensus' in the European Union" *European Journal of Political Research* 44, 65-90.
- 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. (2009), "Minimum Winning Coalitions and Endogenous Status Quo" *International Journal of Game Theory* 39, 617-643.
- Konishi, H., D. Ray (2003), "Coalition Formation as a Dynamic Process" *Journal of Economic Theory* 100, 1-41.
- Larcker, D., G. Ormazabal and D. Taylor (2011), "The Market Reaction to Corporate Governance Regulation" *Journal of Financial Economics* 101, 431-448.
- Lewis, J. (2008), "Strategic Bargaining, Norms, and Deliberation" in Naurin, D. and H. Wallace eds. *Unveiling the Council of the EU* New York, Palgrave Macmillan.
- Lucas, W.F. (1992), "Von Neumann-Morgenstern Stable Sets" in R. Aumann and S. Hart eds. *Handbook of Game Theory with Economic Applications Volume 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.

- Matsusaka, J. and O. Ozbas (2012), "Shareholder Empowerment: The Right to Approve and the Right to Propose", mimeo.
- Mattila, M. (2009), "Roll Call Analysis of Voting in the European Union Council of Ministers after the 2004 Enlargement" *European Journal of Political Research* 48, 840-857.
- Ordeshook, P.C. (1986) *Game Theory and Political Theory*. New York: Cambridge University Press.
- 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.
- Schmidt, S. (2000), "Only an Agenda Setter? The European Commission's Power over the Council of Ministers" *European Union Politics* 1, 37-61.
- 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.
- Thomson, R. (2011), *Resolving Controversy in the European Union* Cambridge University Press, Cambridge.
- Volden, C. and A. Wiseman (2007), "Bargaining in Legislatures over Particularistic and Collective Goods" *American Political Science Review* 101, 79-92.
- von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior* Princeton University Press, Princeton.
- 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.