Bargaining over an endogenous agenda

Vincent Anesi
School of Economics, University of Nottingham

Daniel J. Seidmann
School of Economics, University of Nottingham

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. Our main goal is to study the policies that can be implemented from any initial default in a pure-strategy stationary Markov perfect equilibrium for an interesting class of environments including multidimensional and infinite policy spaces. It is convenient to start by characterizing the set of immovable policies that are implemented, once reached as default. These policies form a weakly stable set and, conversely, any weakly stable set is supported by some equilibrium. Using these results, we show that minimum-winning coalitions may not form and that a player who does not propose may nevertheless earn all of the surplus from agreement. We then consider how equilibrium outcomes change as we vary the order in which players propose, the identity of proposers, and the set of winning coalitions. First, 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, but this result does not generalize to other cases. We also show, surprisingly, that a player may prefer not to be given the opportunity to propose and that the set of immovable policies may shrink as the quota increases. Finally, we derive conditions under which immovable policies in semi-Markovian equilibria form a consistent choice set.

Keywords. Bargaining, committee voting, evolving default, stable set.

JEL classification. C78, D71, D72.

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 endogenously drawing policies/proposals and then implement a proposal according to a stopping rule. Congress and the Federal Open Market Committee (FOMC) instate committees that stop deliberating as soon as some proposal wins a final vote, while the European Union's (EU's) Council of Ministers reaches a decision by final vote when the issue must be addressed urgently or some government wants to signal to

its domestic audience (see 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 does not use a final vote stopping rule. In particular, the Council of Minis-
ters 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 European Central Bank
(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 vot-
ers 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 imple-
mented when all of the proposers have failed to amend it, either because they have cho-
sen 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.) Payoffs in
the game only depend on the policy implemented. 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 that can be reached in a pure-strategy station-
ary Markov equilibrium from any initial default and for any protocol. We use an algo-
rithmic 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 default to the imple-
mented policy. Stationarity implies that any policy in the range of this mapping is "im-
movable": 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 that 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.[1] 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).[2]

---

[1]Every vNM stable set is, therefore, weakly stable.

[2]This result provides microfoundations of weakly stable sets as a by-product of our analysis. Propositions
1 and 2 imply that the policies that can be implemented in any equilibrium are the union of weakly stable
sets.

Equilibrium outcomes have some surprising properties. We demonstrate by example that a winning coalition may amend a default to a policy that 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 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 that could destabilize policies are no longer winning with a larger quota.[3] 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

---

[3]Arguments of this sort have been repeatedly used during the prolonged debates over qualified majority voting in the Council of Ministers.

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.

Any weakly stable set is contained in a consistent choice set. We provide weaker conditions on the stationarity of strategies under which any semi-Markovian equilibrium's immovable policies are a consistent choice set, and we show that every consistent choice set is supported by some semi-Markovian equilibrium.

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 that the policy set supported by any equilibrium must satisfy internal stability. However, there may be equilibria that support policy sets that fail external stability.

We characterize equilibria in Section 3. 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 microfoundations for the consistent choice set, study equilibria in games with repeated implementation, and analyze our open rule bargaining model. Section 6 reviews the related literature. We conclude in Section 7. We relegate longer proofs to an Appendix. Additional arguments are available in a supplementary file on the journal website, http://econtheory.org/supp/1318/supplememt.pdf.

2. The model

We consider a finite committee that consists of $m \geq 1$ *proposers*, $M \equiv \{1, \ldots, m\}$, and $n \geq 2$ *voters*, $N \equiv \{1, \ldots, n\}$. The set of committee members, or *players*, is thus $C \equiv M \cup N$. A player may be both a proposer and a voter, but we also allow for the possibility that $M \cap N = \varnothing$. Our model therefore encompasses the Council of Ministers, where the European Commission makes all proposals but cannot vote, and most Annual General Meetings (AGMs), where management proposes 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 a 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.

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 $x^0 \in X$. The default can only be amended by a winning coalition of voters. We write $\mathbf{W} \subseteq 2^N \setminus \{\varnothing\}$ for the set of winning coalitions of voters. We make the following assumption throughout.

ASSUMPTION A1. *The set of winning coalitions* $\mathbf{W}$ *is both*

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

  (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.

Decision making takes place as follows. Each of a (possibly) infinite number of discrete rounds, indexed by $t = 1, 2, \ldots$, starts in the shadow of an ongoing default policy x^{t-1}. For each possible default $x \in X$, there is a fixed protocol $\pi_x : \{1, \ldots, m_x\} \to M$, $m_x \in \mathbb{N}$, that determines the order in which the proposers (i.e., the players in M) are given the opportunity to propose policies to amend the ongoing default. That is, when $x \in X$ is the current default, protocol π_x gives proposer $\pi_x(k)$ the kth opportunity to amend x for each $k \in \{1, \ldots, m_x\}$. 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 $\pi \equiv \{\pi_x\}_{x \in X}$. Proposers do not have to offer a policy when given the opportunity to do so; instead they can "pass." When a proposer passes, we assume, for expositional convenience, that she offers the current default x.

Consider a history that ends in round t with default x^{t-1} after $k - 1$ proposals have been made in that round. The kth proposer (i.e., player $\pi_{x^{t-1}}(k) \in M$) then makes an offer of $y \in X$. If $y \neq x^{t-1}$ (the proposal is to amend the default). then y is put to a vote by all players in N, who vote sequentially (in an arbitrary order).[4] If those who vote in favor constitute a winning coalition, then the default is amended to y and a new round starts. In all other cases, the consequences depend on whether there are any more opportunities to propose in that round. If $k < m_{x^{t-1}}$ and either $y = x^{t-1}$ or y loses the vote, then the round continues with proposer $\pi_{x^{t-1}}(k + 1)$ given the $k + 1$th opportunity to amend x^{t-1}. By contrast, the game ends (with implementation of default x^{t-1}) after the last proposer in a round either passes ($y = x^{t-1}$) or proposes an amendment that secures less than a winning majority of votes. In sum, bargaining ends if and only if all proposers either pass or fail in their attempt to amend the default.

It is convenient to spell this out formally.

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

Step 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. 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.

---

[4]It is readily checked that subgame perfection makes the voting order irrelevant. Furthermore, we show in Section C of the supplementary file that the set of equilibrium outcomes would be unchanged if we modified the game to allow for abstention.

Step 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 $y_k^t = x^{t-1}$, then (i) the game moves to Step 1 with k increased by 1 if $k < m_{x^{t-1}}$ or (ii) x^{t-1} is implemented and the game ends if $k = m_{x^{t-1}}$.

We assume that all players possess perfect recall and that the game is of perfect information: At any stage, each player observes and recalls everything that has previously transpired in the game.

Players only care about the policy that is eventually implemented, rather than the route from the initial default to the implemented policy. When comparing two different paths, each player $i \in C$ therefore prefers the one that yields 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.[5] This assumption is consistent with the conventional claim that the Council is worst off when it fails to decide (e.g., Thomson 2011). Our assumption that players do not discount is consistent with evidence that the Council of Ministers typically reaches urgent decisions by a final vote procedure (cf. Heisenberg 2005). 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. To avoid repeatedly having to include the relevant qualification, we will refer to stationary Markov pure-strategy equilibria of our game as equilibria throughout the paper—with the sole exception of Section 5.2, where we study history-dependent strategies.

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.

Any stationary Markov strategy combination $\sigma = (\sigma_i)_{i \in C}$ generates an outcome function f^σ, which assigns to every $x \in X$ and every $k \in \{1, \ldots, m_x\}$ the unique final outcome $f^\sigma(x, k)$ eventually implemented (given σ) when x is the ongoing default and the kth proposer is about to move (in any round t). 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

---

[5]This 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.

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, **W**, so as to obtain a collection of coalitions that better accounts for the distribution of power among committee members. Let $\mathcal{W} \equiv \{S \subseteq C : (S \cap N) \in \mathbf{W} \ \& \ (S \cap M) \neq \varnothing\}$. 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 **W**.

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

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

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

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

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

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

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

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 worth discussing some of their properties. First of all, a vNM stable set is a weakly stable set that 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 1. Let $M = N = \{1, 2, 3\}$, let $X = \{x, y, z\}$, and let every pair of players be 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 $y \, P \, z$, and that $\{x, z\}$ and $\{y\}$ are weakly but not vNM stable. By contrast, only $\{x, y\}$ is vNM stable. $\diamond$

The predictive power of weak 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 games; see Example 2 below); there may be several weakly stable sets (e.g., in three-player divide the pie bargaining: cf. Ordeshook 1986, Chapter 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 $Z \subseteq X$, we use the notation $Q(x) \equiv \{y \in X : y \, Q \, x\}$, $Q_Z(x) \equiv \{y \in Z : y \, Q \, x\}$, and $\mathbf{M}(Q, Z) \equiv \{y \in Z : \forall y' \in Z \setminus \{y\}, y' \, Q \, y \text{ implies } y \, Q \, y'\}$.

3.2 *The algorithm*

We now turn to our main purpose in this section, which is to describe an algorithmic procedure that is 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 $V \in \mathcal{V}$ and an initial default $x \in X$. It then constructs a tree $\mathfrak{T}^\pi(V, x)$ inductively in (at most) m_x steps, starting with step m_x and then continuing in decreasing order. The initial node of $\mathfrak{T}^\pi(V, x)$ is always x. Step m_x consists of the construction—to be detailed below—of the set of immediate successors of x, which is denoted by $\mathfrak{s}^\pi_{m_x}(V, x)$. If the latter is empty, then the algorithm ends: the tree has a single node. Otherwise, we proceed inductively, moving backward from the $k + 1$th step to the kth step the same way. For each element y_{k+1} of $\mathfrak{s}^\pi_{k+1}(V, y_{k+2})$, we construct the set of immediate successors of y_{k+1}, denoted by $\mathfrak{s}^\pi_k(V, y_{k+1})$. As previously, the algorithm ends if this set is empty: y_{k+1} is then the last node of a path of length $m_x - k + 1$ (i.e., with $m_x - k + 1$ edges). The algorithm necessarily ends when we reach Step 1, so that no path in the tree has more than m_x edges.

We now explain how the successor nodes of x in $\mathfrak{T}^\pi(V, x)$—the $\mathfrak{s}^\pi_k(V, y_{k+1})$'s—are constructed.

Case 1. If $x \in V$, then the m_x steps of the procedure are trivial: the tree has a single path of length m_x in which all nodes are equal to x; that is,

$$s^\pi_k(V, x) = \{x\} \quad \text{for all } k = 1, \ldots, m_x.$$

Case 2. If $x \notin V$, then the successors of x in the tree are obtained in up to m_x steps $k = m_x, m_x - 1, \ldots, 1$:

- $k = m_x$: Recall that $P_V(x)$ is the set of policies in V that P-dominate x and that $R_V(x)$ is the set of policies in V that R-dominate x. For each subset $Y \subseteq R_V(x)$ (including $\varnothing$), we can determine the (possibly empty) set of policies that are ideal for the last proposer in $P_V(x) \cup \{x\} \cup Y$: $\mathbf{M}(\succeq_{\pi_x(m_x)}, P_V(x) \cup \{x\} \cup Y)$. The set of immediate successors of x is then obtained by taking the union of all such sets:

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

If $\mathfrak{s}_{m_x}^{\pi}(V, x) = \varnothing$, then the procedure stops.[6]

- $1 \le k \le m_x - 1$: If the procedure has not stopped before step $k + 1$ (i.e., there is a sequence $y_{k+1}, y_{k+2}, \ldots, y_{m_x+1}$ such that $y_{m_x+1} = x$ and $y_l \in \mathfrak{s}_l^{\pi}(V, y_{l+1})$ for each $l = k, k+1, \ldots, m_x$), then the set of immediate successors of y_{k+1} is

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

If $\mathfrak{s}_k^{\pi}(V, y_{k+1}) = \varnothing$, then the procedure stops.

Now suppose the tree $\mathfrak{T}^{\pi}(V, x)$ has paths of lengths m_x (i.e., the procedure has reached $k = 1$). The set of outcomes of our algorithm, denoted by $F^{\pi}(V, x)$, is defined as the set of terminal nodes of its paths of length m_x that belong to V. Put formally, $y \in F^{\pi}(V, x)$ if and only if there exists a sequence $(y_1, \ldots, y_{m_x+1})$ such that $y_1 = y \in V$, $y_{m_x+1} = x$, and $y_k \in \mathfrak{s}_k^{\pi}(V, y_{k+1})$ for each $k = 1, \ldots, m_x$. Observe that, by construction, if $x \in V$, then we must have $F^{\pi}(V, x) = \{x\}$. As $F^{\pi}(V, x)$ is by definition a subset of V, this implies that $\bigcup_{x \in X} F^{\pi}(V, x) = V$. If $\mathfrak{T}^{\pi}(V, x)$ has no path of length m_x (i.e., the procedure stopped before $k = 1$), then our algorithm has no outcome.

Observe that the tree construction greatly simplifies if the policy space is well ordered. Indeed, $R_V(x) = P_V(x) \cup \{x\}$ for every policy $x \in X$. As V is assumed to be P-internally stable—so that $P_V(x) = \varnothing$ for all $x \in V$—this implies that the set of immediate successors of any node $y_{k+1} \in V$ is y_{k+1} itself:

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

Applying the same logic recursively thus shows that all successors of node $y_{k+1} \in V$ are equal to y_{k+1}: $y_l = y_{k+1}$ for all $l = 1, \ldots, k+1$. We will return to the special case of well ordered X throughout the paper.

Before we proceed any further, it may be helpful to illustrate this construction by exploiting Example 1 above.

EXAMPLE 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, $P_{\{x,z\}}(y) = \varnothing$ 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 $P_{\{x,z\}}(x) = \varnothing$ 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,

---

[6]The tree $\mathfrak{s}_{m_x}^{\pi}(V, x)$ may indeed be empty when V is not closed.

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 $\{x, z\}$).   ◇

3.3 *The equilibrium correspondence*

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

PROPOSITION 1. *Suppose that V is the closure of a weakly stable set, and let $g \in V^X$ be a selection of $F^\pi(V, \cdot)$: $g(x) \in F^\pi(V, x)$ for all $x \in X$. There exists an equilibrium σ such that $f^\sigma(x) = g(x)$ 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 $g(\cdot)$ of $F^\pi(V, \cdot)$ is an equilibrium outcome of $\Gamma(\pi, x^0)$.[7] 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 $\Gamma(\pi, x^0)$. Thus, Proposition 1 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.

---

[7]Inspection of the proof of reveals that this is actually true for all $V \in \mathcal{V}$ that satisfy $F^\pi(V, x) \neq \varnothing$ for every $x \in X$ (even if V is not closed).

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 $(y_1, \ldots, y_{m_x+1})$, 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_xth 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. In any subgame equilibrium, 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 $Y = \varnothing$, 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 $P_V(x) \cup Y \cup \{x\}$: $y_{m_x} \in \mathbf{M}(\succeq_{\pi_x(m_x)}, P_V(x) \cup Y \cup \{x\}) \subseteq \mathfrak{s}^\pi_{m_x}(V, x)$, 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_xth proposer, except that x must be replaced by y_{m_x}: players anticipate that if the $(m_x - 1)$th proposer's proposal is rejected, then y_{m_x} will be the final policy outcome. Hence, 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 $y_{m_x-1} \in \mathbf{M}(\succeq_{\pi_x(m_x-1)}, P_V(y_{m_x}) \cup Y \cup \{y_{m_x}\}) \subseteq \mathfrak{s}^\pi_{m_x-1}(V, y_{m_x})$. Moving backward, we can repeatedly apply the same reasoning to all proposers until the first, $\pi_x(1)$, whose choice 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 $S \in \mathcal{W}$ contains 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 $(x, y_{m_x}(x), \ldots, y_1(x))$ of tree $\mathfrak{T}^\pi(V, x)$ with $y_1(x) \in V$, there exists an equilibrium σ such that $f^\sigma(x, k) = y_k(x)$ for each $k \in \{1, \ldots, m_x\}$. Second, for every equilibrium σ of $\Gamma(\pi, x^0)$ and every $x \in X$, there exists a weakly stable set V and a length-m_x path $(x, y_{m_x}(x), \ldots, y_1(x))$ of $\mathfrak{T}^\pi(V, x)$, with $y_1(x) \in V$, such that $y_k(x) = f^\sigma(x, k)$ for each $k \in \{1, \ldots, m_x\}$. Thus, the construction of trees associated with weakly stable sets also provides a complete characterization of equilibrium behavior both on and off equilibrium paths.

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 subsequent propositions) in the Appendix.

Proposition 2 implies that the game only has an equilibrium if there is a weakly stable set, and, as noted above, even games with a finite, well ordered policy space may not have a weakly stable set. Nonexistence might be addressed by weakening our solution concept. We will consider the implications of history-dependent strategies in Section 5.2.

Propositions 1 and 2 also have a number of technical implications that 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. This does not mean, however, that determining the equilibrium policies of a given game $\Gamma(\pi, x^0)$ boils down to determining the class of weakly stable sets: the set of equilibrium policies of $\Gamma(\pi, x^0)$, $\bigcup_{V \in \mathcal{V}} F^\pi(V, x^0)$, may be a strict subset of the union of weakly stable sets. In most cases, $F^\pi(V, x^0)$ only selects a strict subset of the elements of V, this selection depending on the initial default x^0 and on the protocol π.

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.

On the other hand, Propositions 1 and 2 allow us to provide sufficient conditions for a unique equilibrium. In what follows, a policy is referred to as a *Condorcet winner* if and only if it *P*-dominates every other policy in X.

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

(b) The Condorcet winner (when it exists) 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 2. Consider a committee whose two members ($i = 1, 2$) can divide a unit-sized pie. Each player has one opportunity to propose in each round, with player 1 proposing first. The committee operates with a unanimity quota: a proposal is only successful if both players vote in favor. In other words, this is a standard two-player bargaining problem, modified by our rules for ending the game.

The policy space consists of the initial default (x^0), at which each player earns 0, and any division of the whole pie; that is,

$$X = \{(x_1, x_2) \in [0, 1]^2 : x_1 + x_2 = 1\} \cup \{x^0\}.$$

This policy space is not well ordered because each player is indifferent between x^0 and a division that yields her none of the pie. Consequently, the premise of Corollary 2(a) does not hold.

It is easy to confirm that there is a unique weakly stable set (V), consisting of every division of the pie. Corollary 1 therefore implies that every policy in V is immovable: any policy $(x_1, x_2) \in V$ would be implemented in an equilibrium if it were the initial default. However, $\Gamma(\pi, x^0)$ has a unique equilibrium outcome, in which player 2 earns the entire pie. To see this, consider first the proposal by player 2 (if she is given the opportunity to propose). Player 2's acceptance set includes every policy $(x_1, x_2) \in V$ such that $x_1 > 0$ and $x_2 > 0$ because x^0 would be implemented unless successfully amended, and the acceptance set only has a maximal element if player 1 accepts policy $(0, 1)$. Consequently, player 2 would successfully propose $(0, 1)$—her ideal policy—in any subgame equilibrium. The unanimity quota and no discounting therefore imply that $(0, 1)$ is the only policy that player 1 can successfully propose: player 2 would reject any other offer. Hence, the second proposer (i.e., player 2) would earn the entire pie in any equilibrium of $\Gamma(\pi, x^0)$. In particular, $\Gamma(\pi, x^0)$ has a unique equilibrium outcome, even though the policy space is not well ordered. ◇

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 that can be implemented across initial defaults:

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

However, variations in the order of proposers may affect the policies that can be implemented from a given initial default: in Example 2, for instance, the player who proposes second gets the entire pie, so switching the order of proposers changes the equilibrium outcome.

Interestingly, the next result states that the set of policies that can be implemented from a given default, $\bigcup_{V \in \mathcal{V}} F^\pi(V, x^0)$, 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 depends 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 that 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 kth 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 kth proposer's ideal policy in A_k, then she could profitably deviate by offering this policy, which would be accepted, directly in round 1.

3.4 *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. 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 two-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 50 cents: $\bar{x}_1 \leq \frac{1}{2}$. 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, Chapter 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, the following situation occurs.

- 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 $y_2 \le 1 - \bar{x}_1$.

The first proposal is successful in this equilibrium, and it secures the votes of exactly two players. A minimal winning coalition forms at the voting stage, yet the three players each earn a share of the dollar. ◇

EXAMPLE 1 CONTINUED. In the last subsection, we constructed tree $\mathfrak{T}^{\pi}(\{x, z\}, y)$ for Example 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 that is P-dominated by the initial default. Inspection of Example 1 reveals that this property relies on the supposition that some players are indifferent between policies, which allows policies in the weakly stable set to be amended. Indeed, Proposition 3 below will imply that the implemented policy must R-dominate the initial default if the policy space is well ordered. ◇

We record the arguments in this subsection as follows.

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.

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 ("the power of the penultimate word")*

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 that is constant across X.[8] Hence, we can without loss of generality restrict attention to constant protocols.

The chair's choice depends 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 $k \equiv \max\{l \in \{1, \ldots, m_x\} : \mathbf{M}(\succ_{\pi(l)}, R(x) \cap f^\sigma(X)) \succ_{\pi(l)} x\}$.

In Proposition 3, the kth proposer is the last proposer among those who have an incentive to amend the ongoing default x in equilibrium σ: namely those who strictly prefer some policy in their acceptance set 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 $\sigma \in \Sigma^*(\pi, x^0)$ that 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 is 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 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.

---

[8]Corollary 3 also applies in a different "dynamic" game where the chair selects the next proposer immediately after each vote that does not end the game (see Section A in the supplementary file for a formal proof).

- 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 that is also optimal.

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. Let $M = N = \{1, 2, 3, 4, 5, 6\}$ and $X = \{x, v_1, v_2, v_3\}$. Preferences over X are given by $v_3 \succ_1 v_2 \sim_1 v_1 \succ_1 x$, $v_2 \succ_2 v_1 \succ_2 x \succ_2 v_3$, $x \succ_3 v_2 \succ_3 v_3 \succ_3 v_1$, $x \succ_4 v_3 \succ_4 v_2 \succ_4 v_1$, $v_3 \sim_5 v_1 \succ_5 x \succ_5 v_2$, and $v_1 \succ_6 x \succ_6 v_2 \succ_6 v_3$. 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.[9] 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, two options exist.

(i) The chair may optimally choose to make the last two proposals, thereby strictly improving on a protocol in which she only proposes last.[10]

---

[9]If 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.

[10]This 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.

(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 "power of the last word" in Bernheim et al. (2006), which we elaborate on in Section 6.1.

We summarize the discussion above as follows.

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.*

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. Black (1958, p. 99) made the following proposal:

> 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.

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 $s \geq (n+1)/2$. 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, Black's 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 so 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 that 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 5. 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$_{P4}$), would fail: $w\, P^4\, x$, $z\, P^4\, w$, $x\, P^4\, y$, and $y\, P^4\, z$.

Now suppose that the quota becomes 5. As $w\, P^5\, x$ and $y\, P^5\, z$, internal stability (IS$_{P5}$) 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 $z\, P^4\, w$ and $x\, P^4\, y$, $\neg(z\, P^5\, w)$ and $\neg(x\, P^5\, y)$. 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 the following proposition.

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 or if X is a weakly stable set for the higher quota. However, an increase in the quota may contract the set of equilibrium decisions in other cases.*

Black's conjecture underlies 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 that 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 so as to compare our results with Diermeier and Fong (2011), who 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 a model with repeated implementation. 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 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 that 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 3.4, we used Example 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, 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 6. 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 $\mathcal{W}' = \{S \in \mathbf{W} : \{1, 2\} \cap S \neq \varnothing\}$, 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 6, 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 that 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 that 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 follows.

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.[11] 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. 2012 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 this subsection, 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 that 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 round T such that $x^t = x^T$ for every subsequent round: $t > T$. We abuse terminology in this subsection by identifying equilibria in our model with equilibria in the related model with continued (but payoff-irrelevant) bargaining.

---

[11]The proposal that 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.

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. Our next result confirms this intuition 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:

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

The assumption that X is finite and well ordered thus guarantees that there exists a sufficiently large $\delta < 1$ ($\bar{\delta}$) such that players evaluate sequences of defaults similarly in 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 CONTINUED. Consider again the divide-the-dollar game in Example 3. Recall that the set of policies $(x_1, x_2, x_3) \in \Delta$ at which player 1 earns 50 cents, $V(\frac{1}{2})$, is a weakly stable set. From Proposition 1, therefore, there are equilibria of $\Gamma(\pi, x^0)$ that support $V(\frac{1}{2})$. In particular, the equilibrium σ constructed in the proof of Proposition 1 prescribes the committee to implement a policy in $V(\frac{1}{2})$ without delay: in any subgame starting with an ongoing default $x \notin V(\frac{1}{2})$, x is amended to some $v_x \in V(\frac{1}{2})$ 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 $v_0 \equiv (\frac{1}{2}, x_2, \frac{1}{2} - x_2) \in V(\frac{1}{2})$. Suppose that player 1, if given the opportunity to amend v_0, proposes $y = (\frac{1}{2} + \epsilon/2, x_2 - \epsilon, \frac{1}{2} - x_2 + \epsilon/2)$ for any small and positive ϵ. If player 1's proposal were accepted, then, by construction of σ, a policy $v_y \in R(y) \cap V(\frac{1}{2})$ would be implemented in all future periods. As v_y must R-dominate y, it must be of the form $v_y = (\frac{1}{2}, x_2 - \epsilon + \gamma_2, \frac{1}{2} - x_2 + \epsilon/2 + \gamma_3)$, 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:

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

and

$$(1 - \delta)u_3(y) + \delta u_3(v_y) = \tfrac{1}{2}(1 + \epsilon) - x_2 + \delta\gamma_3 > \tfrac{1}{2} - x_2 = u_3(v_0).$$

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

5.2 *History-dependent strategies*

Our goal in this subsection is to study how the set of immovable policies is affected when we extend the strategy set to strategies that are more history-dependent than stationary Markov strategies. In particular, we observed in Section 3 that—as the archetypal example of the Condorcet paradox illustrates—a weakly stable set, and, therefore, a stationary Markov equilibrium, may not exist.

　　Our first step is to characterize equilibrium policy outcomes when strategies depend on the entire sequence of previous defaults. To do so, we first need some definitions. We call a *partial round-t history* any list of defaults $(x^0, x^1, \ldots, x^{t-1})$ such that $x^{t-1} \neq x^{t-2}$ (i.e., the game did not end in period $t - 1$).[12] 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. Our aim is to characterize semi-Markovian equilibria, i.e., subgame perfect equilibria of $\Gamma(\pi, x^0)$ in which all players use semi-Markovian strategies.

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

　　In Vartiainen's (2012) framework, which we elaborate on in the next section, coalitional moves in each period t are dictated by "policy programs," which may depend on the entire sequence of previous defaults $(x^0, \ldots, x^{t-1})$. He argues that in such a context, the set of implementable policies must be a "consistent choice set": A nonempty set of policies $Z \subseteq X$ is a *consistent choice set* if, for any $z \in Z$ and for any $x \in X$, there is $z' \in Z$ such that $z' \in R(x) \setminus P(z)$. It turns out that the tree construction introduced in Section 3 can also be applied to consistent choice sets to obtain semi-Markovian equilibria. More

---

[12]A (complete) "history" at some stage of a given round t would describe all that has transpired in the previous rounds and stages (the sequence of proposers, their respective proposals, and the associated pattern of votes).

specifically, if Z is a consistent choice 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 6. *Suppose that Z is a consistent choice set and let $g \in Z^X$ be any selection of $F^\pi(Z, \cdot)$: $g(x) \in F^\pi(Z, x)$ 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 $\phi^{\sigma^x}(x) = g(x)$. Hence, $\bigcup_{x \in X} \phi^{\sigma^x}(x) = Z$.*[13]

The last part of the statement in the proposition says that, for any consistent choice 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 reveals that more is true: the final policy outcome reached from any partial history $h \in H$ must belong to Z; so that $\phi^\sigma(H) \equiv \bigcup_{h \in H} \phi^\sigma(h, 1) = Z$.

Interestingly, Vartiainen (2012) establishes existence of a (nonempty) consistent choice set. An immediate consequence of our results, therefore, is that *a semi-Markovian equilibrium exists* in our game.[14]

Though every consistent choice set can be supported by a semi-Markovian equilibrium, the converse is not true; that is, *it is generally not true that the set of immovable policies in a semi-Markovian equilibrium is a consistent choice set*; Section B.4 in the supplementary file provides a counterexample. Nevertheless, our next result states that immovable policies form a consistent choice set whenever the policy space is well ordered and/or there is a single proposer in every round (as is usually assumed in the literature on bargaining games with repeated implementation).

PROPOSITION 7. *Suppose that (at least) one of the following assumptions holds: (i) X is well ordered; (ii) $m_x = 1$ for all $x \in X$. If σ is a semi-Markovian equilibrium, then $\phi^\sigma(H) \equiv \bigcup_{h \in H} \phi^\sigma(h, 1)$ is a consistent choice set.*[15]

Thus, when X is well ordered or there is a single proposer, the set of policy outcomes that can be reached from all possible partial histories in equilibrium is a consistent choice set. Furthermore, we know from Proposition 6 that any policy $z \in Z$ is the outcome of a semi-Markovian equilibrium of $\Gamma(\pi, z)$. Consequently, we have a corollary.

COROLLARY 4. *Suppose that (at least) one of the following assumptions holds: (i) X is well ordered; (ii) $m_x = 1$ for all $x \in X$. The set of all semi-Markovian equilibrium policy*

---

[13]The proof of Proposition 6 is similar to that of Proposition 1, so we provide the proof in Section B.1 of the supplementary file.

[14]Vartiainen (2012) assumes that R is complete, but this assumption is not needed for the existence result. We are grateful to him for discussing this point.

[15]We prove Proposition 7 in Section B.2 of the supplementary file because of similarities to the proof of Proposition 2.

outcomes that can be obtained from any initial default in X coincides with the union of consistent choice sets.

One can add more history dependence by allowing strategies to depend not only on the sequence of previous defaults, but also on the coalitions that amended previous defaults. Thus, a relevant history is now of the form $(x^0, S^1, x^1, \ldots, S^{t-1}, x^{t-1})$, where $S^s \in \mathcal{W}$ stands for the winning coalition that amended x^{s-1} to x^s. We prove in the supplementary file (Section B.3) that, by doing so, we obtain analogs to Propositions 6 and 7 and Corollary 4 in which "quasi-consistent set" replaces "consistent choice set." A nonempty set of policies $Z \subseteq X$ is *quasi-consistent* if and only if it satisfies the following condition.[16]

(QC) We have $z \in Z$ only if, for all $x \in X$ and $S \in \mathcal{W}$, there exists $z' \in Z \cap R(x)$ such that $z \succeq_i z'$ for some $i \in S$.

Hence, if strategies are allowed to depend on both the sequence of previous defaults and the coalitions that amended previous defaults, then (i) the closure of any quasi-consistent set can be supported by an equilibrium and (ii) when X is well ordered or there is a single proposer, the set of policy outcomes that can be reached from all possible partial histories in equilibrium is a quasi-consistent set.

5.3 *Open rule bargaining*

Thus far, we have focused on games that end when no proposer amends a default. We have demonstrated that there is an equilibrium that supports the closure of any weakly stable set and that every equilibrium supports a weakly stable set (Propositions 1 and 2). 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 $\{1, \ldots, m_x\}$ 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 that 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 let 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

---

[16]If "only if" is replaced by "if and only if" in condition (QC), then Z is a *consistent set* (Chwe 1994).

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 that is then implemented.

The argument above implies that the constructed strategy combination supports V, by analogy to Proposition 1. 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 that are not externally stable. To see this, consider Example 2, where two players bargain 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 ($\frac{1}{2}$) then a player who gets less than $\frac{1}{2}$ proposes amending x to $\frac{1}{2}$, while any other player passes, and both players vote in favor when $\frac{1}{2}$ is moved. Any player who gets less than $\frac{1}{2}$ at y vetoes amending x to y and also votes against if y is moved.

The following observation is a summary.

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 that fail external stability.*

6. RELATED LITERATURE

The literature encompasses various approaches to modeling bargaining with an endogenous default. Some papers explore the equilibrium correspondence of particular noncooperative models of committee policy making. By contrast, papers in the axiomatic tradition develop properties of bargaining that are robust to variations in the procedure. Various papers provide noncooperative foundations for cooperative solution concepts. Our paper falls in the first of these categories, but our results also relate to the literature on noncooperative foundations.

6.1 *Noncooperative models of committee policy making*

The literature on noncooperative bargaining with an endogenous default contains papers in which implementation is delayed until negotiations end and papers in which the default is implemented each period. This paper, like Bernheim et al. (2006), focuses on problems in the first category. However, in contrast to our model, Bernheim et al. (2006) suppose that 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. As some winning coalitions exclude players, the final proposer can play off putative members of the winning coalition. Our argument, by contrast, relies on our stopping rule: a default that is not amended by any proposer is implemented, ending the game. Proposition 3, by contrast, allows for cases in which only the grand coalition is winning, e.g., in variants on Example 2, where the pie can only be split in a finite number of proportions. In further contrast to Bernheim et al., we allow for an infinite policy space without requiring that it be well ordered. In such cases, the chair may prefer not to propose last (cf. Observation 2).

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. We have also studied a version of open rule bargaining that extends our model, but is not restricted to the distributive problems that Baron and Ferejohn consider.[17] 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, patient players may all earn a share of the pie in our model. This property fails 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 dissenting voters earn nothing.

Following Baron (1996), a recent literature has studied equilibria of games with repeated implementation. 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. (2012), which essentially shares our game tree, but allows the set of winning coalitions to depend on the default.[18] 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.[19] We focus on characterization, rather than existence results, but explore a much larger class of policy spaces that includes nonacyclic social preferences. Most of this literature supposes that players discount future returns. Our main results concern games without discounting, but we have shown that they also apply to games with repeated implementation, a finite and well ordered policy space, and patient enough players.

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

---

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

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

[19]von Neumann and Morgenstern (1944) 65.7 prove existence and uniqueness of vNM stable sets under Acemoglu et al.'s conditions. See Anesi and Seidmann (2012), Duggan and Kalandrakis (2011, 2012), and Kalandrakis (2004, 2010) for existence results in related models.

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. We do not impose this tie-breaking rule, so the size principle fails in our model for different reasons. Diermeier and Fong also provide an existence proof for their game; von Neumann and Morgenstern's (1944) argument (cf. footnote 19) implies that our version of Diermeier and Fong's game also has a unique equilibrium.

Within the same strand of literature, but more distantly related, Duggan and Kalandrakis (2012) study pure-strategy stationary Markov perfect equilibria of an infinite-horizon model of legislative policy making with an evolving default. They establish equilibrium existence and derive conditions under which proposers form minimal winning coalitions. In contrast to our framework with no uncertainty, however, they assume that players' preferences and the default are subject to random shocks. Other related papers in this literature include Baron (1996), Baron and Herron (2003), Kalandrakis (2004, 2010), Battaglini and Coate (2007, 2008), Dziuda and Loeper (2010), Duggan and Kalandrakis (2011), Zápal (2011), Battaglini and Palfrey (2012), Anesi (2012), and Battaglini et al. (2012).

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 then the committee votes over the agenda, so the agenda is endogenous in our sense.[20] 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.[21] 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.

6.2 *Cooperative models of dynamic collective choice*

Another strand of literature uses cooperative frameworks to study dynamic collective choice with an endogenous default. These models abstract away from procedural details: there is no protocol that determines the choices of proposers and voters. Consequently, they have little to say about proposal and agenda power. This stands in sharp

---

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

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

contrast to our model, in which the order and identity of proposers play a central role in determining policy outcomes. Nevertheless, it is worth explaining the paper's relationship to those contributions in terms of predictions about policy outcomes.

Konishi and Ray (2003) develop a model of coalition formation that transplants the blocking approach from traditional cooperative game theory to dynamic environments with repeated implementation of the defaults. In particular, they study a related notion of absorbing sets of Markovian equilibrium processes of coalition formation with far-sighted players and a finite policy space. They prove that those absorbing states must be contained in Chwe's (1994) largest consistent set and that a stochastic equilibrium process exists. Anesi (2006) shows that this solution concept is equivalent to a vNM stable set (and may therefore not exist) when one focuses on deterministic processes. However, he only studies games in which the policy space is finite and well ordered.

Vartiainen (2011) studies absorbing outcomes of a deterministic non-Markovian version of Konishi and Ray's coalition formation processes, assuming (like us) that players only care about final outcomes. Focusing on finite policy spaces, he establishes existence of equilibrium processes. In that paper, coalitional moves in each period may depend on the entire sequence of previous moves and on the identities of the coalitions that initiated those moves. In Section 5.2, we also characterized the set of immovable policies in our model when strategies are similarly history dependent. Our analysis revealed that the set of immovable polices is then a superset of Chwe's (1994) largest consistent set, which itself contains the set of absorbing equilibrium outcomes in Vartiainen (2011). Hence, every equilibrium policy outcome in Vartiainen (2011) can be implemented in a history-dependent equilibrium of our model, but the converse may not be true.

In a more recent paper that allows for a compact policy space, Vartiainen (2012) characterizes deterministic, absorbing policy sequences that can only depend on the history of defaults and that satisfy a one-deviation property. The set of absorbing history-dependent policies (the consistent choice set (CCS)) exists. Proposition 6 asserts that every CCS is the set of immovable policies for some semi-Markovian equilibrium of our model. In addition, the converse is true if the policy space is well ordered or there is a single proposer (Proposition 7), but may not be true otherwise. Vartiainen also shows that policies that are implementable by a Markovian policy sequence form a vNM stable set and that any vNM stable set is the outcome correspondence of such a Markov perfect policy sequence. These results recall Propositions 1 and 2 above. Despite these parallels, our results differ in one important respect. Vartiainen's one-deviation property (like Acemoglu et al.'s (2012) desirability axiom) implies that no winning coalition can prefer the initial default (x) over the policy implemented according to equilibrium σ ($\phi^\sigma(x)$). This property holds in our model if X is well ordered or if there is a single proposer (the premise of Proposition 7). However, as Observation 1(c) states, a winning coalition may otherwise prefer the initial default, the intuition being that a winning coalition must prefer $\phi^\sigma(x)$ to the policy that would be implemented if they did not amend x, rather than preferring $\phi^\sigma(x)$ to x. It is therefore interesting that Propositions 1, 2, and 6 apply even if the one-deviation property fails. The fundamental difference between our model

and Vartiainen (2011, 2012) is that we do not allow winning coalitions to stop the bargaining process. While winning coalitions can choose any alternative in $X \cup \{STOP\}$ in Vartiainen (2011, 2012), STOP is not an option for coalitions in our framework. Bargaining only ends when proposers are unable or unwilling to amend the existing default.

6.3 *Noncooperative foundations*

In common with a related literature, several of our results provide noncooperative foundations for cooperative solution concepts in voting games, though, in contrast to this literature, we develop this equivalence so as to explore how variations in our noncooperative model affect outcomes.

Propositions 1 and 2 show that the set of immovable policies in our model coincides with the closure of the union of weakly stable sets. In this sense, our results are related to Harsanyi (1974) and Anesi (2010, 2012).

Harsanyi (1974) provides microfoundations 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 that 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 that 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 Section A of the supplementary file.[22]

Our results on weakly stable sets are also related to Anesi (2010, 2012) and Diermeier and Fong (2012). Anesi (2010) demonstrates that any vNM stable set is the absorbing set of some pure-strategy Markov perfect equilibria in a legislative bargaining game with a finite, well ordered policy space, random proposers, and repeated implementation. 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 pure-strategy Markov perfect equilibria.[23] Anesi (2012) and Diermeier and Fong (2012) also obtain such an equivalence, but in different settings—the former in a dynamic game of electoral competition between two office-motivated parties; the latter in the bargaining game of Diermeier and Fong (2011) with a unique proposer—and only for finite policy spaces and well ordered preferences. While existence of a stable set (and therefore of a pure-strategy

---

[22]In contrast to Harsanyi (1974) and this paper, Hortala-Vallve (2011) studies play in a related model without a weakly stable set.

[23]Anesi 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.

Markov perfect equilibrium) is not guaranteed in Anesi (2010, 2012), Diermeier and Fong (2012) establish existence in their single-proposer framework.

Finally, one may also interpret the results in Section 5.2 as providing noncooperative foundations for consistent choice sets and for quasi-consistent sets (and therefore the largest consistent set). Vartiainen (2012) also presents a noncooperative model in which players vote each round over X and implementation of the default. If the latter secures the most votes, then the game ends; otherwise, the default is amended to a plurality winner (so, unlike our model, the game is not simple). Vartiainen shows that every strong subgame perfect equilibrium satisfies the one-deviation property. Observation 1(c) would also fail in our model if we allowed for coalitional deviations.

7. 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 no proposer is willing or able to amend it. We have characterized the policies that 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.

There are various natural extensions of the model. First, our assumption that a proposal is a single policy could be replaced by the supposition that players provisionally agree to subsets of the policy space, thereby allowing us to address multi-issue negotiations in which no partial agreement is finalized until all issues have been addressed. Second, we have focused on pure-strategy solutions, at the cost of nonexistence. It is easy to see that mixed strategy equilibria may exist in situations like the Condorcet paradox, where there is no stable set. On the other hand, there are multiple stable sets and, therefore, a multiplicity of (pure-strategy) equilibria in other situations. It would be interesting to develop refinements that reduce this multiplicity. However, we can show that Acemoglu et al.'s (2009) refinement has no power in our model if the policy space is finite and well ordered. (Proof of this claim is available from the authors on request.)

Appendix

PROOF OF PROPOSITION 1. Let $V \in \mathcal{V}$ and let $g \in V^X$ be a selection of $F^\pi(V, \cdot)$. By construction of $F^\pi(V, x)$, for every $x \in X$, there exists a vector $(y_1(x), \ldots, y_{m_x+1}(x))$ such that two scenarios exist:

- If $x \in V$, then $g(x) = y_1(x) = \cdots = y_{m_x+1}(x) = x$.

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

We now define 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 kth 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 holds:

(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 $f^\sigma(x, 1) = g(x)$ 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. We have $f^\sigma(x) \equiv f^\sigma(x, 1) = y_1(x)$ for all $x \in X$ and, in particular, $f^\sigma(x) = x$ for all $x \in V$; hence, $\bigcup_{x \in X} f^\sigma(x) = V$.

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 $l = \max\{k \in \{1, \ldots, m_x\} : y_k(x) \neq y_{k+1}(x)\}$ (external stability ensures that this set is nonempty) and suppose that the lth proposer is given the opportunity to make a proposal. By construction of $(y_1(x), \ldots, y_{m+1}(x))$, 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 $l = 1, \ldots, l - 2$, we obtain that $f^\sigma(x, 1) = y_1(x)$. This proves that $f^\sigma(x) \equiv f^\sigma(x, 1) = g(x)$ for all $x \in X$. By construction, $g(x) = x$ for all $x \in V$ and $g(x) \in F^\pi(V, x) \subseteq V$ for all $x \in X$. Hence, $\bigcup_{x \in X} f^\sigma(x) = \bigcup_{x \in X} g(x) = V$.

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

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

Any $x \in V$ is implemented at the end of round t if the kth 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 $y_{k+1}(x) = \cdots = y_{m_x}(x) = x = f^\sigma(x, k+1)$. Now suppose that the lth proposer is the next proposer (after the kth) to make a successful proposal, $y_l(x) \neq x$. By construction, this implies that $y_{k+1}(x) = \cdots = y_l(x) \in V$. 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 kth 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 kth proposer to pass (i.e., $y_k(x) = x$). If she has a profitable deviation at this stage then she must be able to 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 kth proposer to propose $y_k(x) \in A_k^V(y_{k+1}(x))$ (where $y_k(x) = x$ means that she should pass). Suppose that, instead, she 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 will be $f^\sigma(y, 1) = y_{k+1}(x)$ 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 $x P y$. By definition of $f^\sigma(X)$, x and y are fixed points of $f^\sigma(\cdot, 1)$. An immediate consequence of $x P y$ 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 $y \in f^\sigma(X)$, $\neg(y \, R \, x)$. In particular, $\neg[f^\sigma(y,1) \, R \, x]$ 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 that 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 $y \neq x$, then all committee members anticipate that this will lead to $f^\sigma(y,1)$ being the final outcome if the amendment is voted up and lead 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 $\sigma = (\sigma_i)_{i \in C}$ be an equilibrium of $\Gamma(\pi, x^0)$. From Lemma 1, we know that there exists $V \in \mathcal{V}$ such that $f^\sigma(X) = V$. Evidently, for all $x \in V$, we have $\{f^\sigma(x,1)\} = \{x\} = F^\pi(V,x)$.

Now consider an arbitrary $x \notin V$ and an arbitrary round starting with x as the ongoing default. Suppose that (possibly off the equilibrium path) the m_xth proposer is given the opportunity to amend x. When she offers a policy $y \neq x$, voters compare $f^\sigma(y,1) \in V$ with x. 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_xth proposer, i.e., the set of policies $y \neq x$ that would be accepted by a winning coalition to amend x, must then be the set of policies y such that $f^\sigma(y,1)$ belongs to $[P_V(x) \cup Y] \subseteq V$, where Y is some (possibly empty) subset of $R_V(x)$. As a consequence, if $\sigma_{\pi_x(m_x)}$ prescribes the m_xth proposer to amend x with $\bar{y}_{m_x} \neq x$, then $f^\sigma(x, m_x) = f^\sigma(\bar{y}_{m_x}, 1)$ must be $\succeq_{\pi_x(m_x)}$-maximal in $[P_V(x) \cup Y \cup \{x\}]$ (x is always feasible to the m_xth proposer, for she can always pass). If $\sigma_{\pi_x(m_x)}$ prescribes the m_xth proposer not to amend x, i.e., to pass or to make an unsuccessful proposal, then $f^\sigma(x, m_x) = x$ must be $\succeq_{\pi_x(m_x)}$-maximal in $[P_V(x) \cup Y \cup \{x\}]$. This proves that $y_{m_x} \equiv f^\sigma(x, m_x) \in \mathfrak{s}_m^\pi(V, x)$.

Proceeding recursively, one can use the same argument to show that, for each $k = 1, \ldots, m_x - 1$, $y_k \equiv f^\sigma(x, k) \in \mathfrak{s}_k^\pi(V, x)$: 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 $(y_1, \ldots, y_{m_x}, x) \equiv (f^\sigma(x, 1), \ldots, f^\sigma(x, m_x), x)$ constitutes a path of tree $\mathfrak{T}^\pi(V, x)$ whose terminal node belongs to V. Hence, $f^\sigma(x) \in F^\pi(V, x)$.   □

PROOF OF PROPOSITION 3. Let σ be an equilibrium. Proposition 2 implies that there must be some weakly stable set V such that $f^\sigma(X) = V$. Consider the kth proposer as defined in the statement of the proposition. If she failed to amend the ongoing default x, then nobody else would and x would be implemented at the end of the round. As

she strictly prefers her ideal policy in the set of equilibrium policy outcomes that dominate x (i.e., $\mathbf{M}(\succ_{\pi(l)}, R(x) \cap f^\sigma(X)))$ 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 kth proposer can successfully do so. Suppose first that the $(k-1)$th proposer successfully offers some policy y. This implies there is a winning coalition in $\mathbf{W}$ whose members all strictly prefer $f^\sigma(y, 1) \in V$ to $\mathbf{M}(\succ_{\pi(l)}, R(x) \cap f^\sigma(X)) \in V$: 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

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

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 $\Psi_i(x, y, \delta) \to u_i(x) - u_i(y) > 0$ as $\delta \to 1$, 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 $\{x_t\}_{t=1,\ldots,T_x}$ and $\{y_t\}_{t=1,\ldots,T_y}$, and a policy $y \in X$ such that

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

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 $\{x_t\}_{t=1,\ldots,T_x+1}$ be the finite sequence of policy outcomes (with $x = x_{T_x+1}$ being implemented indefinitely) if i does not deviate from σ_i at that stage. Hence, any other sequence $\{y_t\}_{t=1,\ldots,T_y+1}$ (with $y = y_{T_y+1}$ being implemented indefinitely) she could induce by a one-shot deviation satisfies

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

This inequality implies that $u_i(x) > u_i(y)$. To see this, suppose instead that $u_i(y) > u_i(x)$. Then $\delta > \bar{\delta}$ 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 $\Gamma(\pi, x^0)$.   □

References

Acemoglu, Daron, Georgy Egorov, and Kostantin Sonin (2009), "Equilibrium refinement in dynamic voting games." Unpublished paper. [475]

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin (2012), "Dynamics and stability of constitutions, coalitions, and clubs." *American Economic Review*, 102, 1446–1476. [471, 473]

Akyol, Ali C., Wei F. Lim, and Patrick Verwijmeren (2012), "Shareholders in the boardroom: Wealth effects of the SEC's proposal to facilitate director nominations." *Journal of Financial and Quantitative Analysis*, 47, 1029–1057. [465]

Anesi, Vincent (2006), "Committees with farsighted voters: A new interpretation of stable sets." *Social Choice and Welfare*, 27, 595–610. [473]

Anesi, Vincent (2010), "Noncooperative foundations of stable sets in voting games." *Games and Economic Behavior*, 70, 488–493. [474, 475]

Anesi, Vincent (2012), "A new old solution for weak tournaments." *Social Choice and Welfare*, 39, 919–930. [472, 474, 475]

Anesi, Vincent and Daniel J. Seidmann (2012), "Bargaining in standing committees." Discussion Paper 2012-09, CeDEx. [471]

Austen-Smith, David and Jeffrey S. Banks (2005), *Positive Political Theory II: Strategy and Structure*. University of Michigan Press, Ann Arbor. [472]

Baron, David P. (1996), "A dynamic theory of collective goods programs." *American Political Science Review*, 90, 316–330. [471, 472]

Baron, David P. and John A. Ferejohn (1989), "Bargaining in legislatures." *American Political Science Review*, 83, 1181–1206. [446, 448, 459, 469, 471]

Baron, David P. and Michael C. Herron (2003), "A dynamic model of multidimensional collective choice." In *Computational Models of Political Economy* (Ken Kollman, John H. Miller, and Scott E. Page, eds.), MIT Press, Cambridge. [472]

Battaglini, Marco and Stephen Coate (2007), "Inefficiency in legislative policy-making: A dynamic analysis." *American Economic Review*, 97, 118–149. [472]

Battaglini, Marco and Stephen Coate (2008), "A dynamic theory of public spending, taxation, and debt." *American Economic Review*, 98, 201–236. [472]

Battaglini, Marco, Salvatore Nunnari, and Thomas R. Palfrey (2012), "Legislative bargaining and the dynamics of public investment." *American Political Science Review*, 106, 407–429. [472]

Battaglini, Marco and Thomas R. Palfrey (2012), "The dynamics of distributive politics." *Economic Theory*, 49, 739–777. [472]

Bernheim, B. Douglas, Antonio Rangel, and Luis Rayo (2006), "The power of the last word in legislative policy making." *Econometrica*, 74, 1161–1190. [462, 470, 471, 474]

Black, Duncan (1958), *The Theory of Committees and Elections*. Cambridge University Press, Cambridge. [462]

Chwe, Michael S.-Y. (1994), "Farsighted coalitional stability." *Journal of Economic Theory*, 63, 299–325. [469, 473]

Compte, Olinier and Philippe Jehiel (2010), "Bargaining and majority rules: A collective search perspective." *Journal of Political Economy*, 118, 189–221. [445]

Diermeier, Daniel and Pohan Fong (2011), "Legislative bargaining with reconsideration." *Quarterly Journal of Economics*, 126, 947–985. [461, 463, 471, 474]

Diermeier, Daniel and Pohan Fong (2012), "Characterization of the von Neumann–Morgenstern stable set in a non-cooperative model of dynamic policy-making with a persistent agenda setter." *Games and Economic Behavior*, 76, 349–353. [474, 475]

Duggan, John (2006), "Endogenous voting agendas." *Social Choice and Welfare*, 27, 495–530. [472]

Duggan, John and Tasos Kalandrakis (2011), "A Newton collocation method for solving dynamic bargaining games." *Social Choice and Welfare*, 36, 611–650. [471, 472]

Duggan, John and Tasos Kalandrakis (2012), "Dynamic legislative policy making." *Journal of Economic Theory*, 147, 1653–1688. [471, 472]

Dutta, Bhaskar, Matthew O. Jackson, and Michel Le Breton (2004), "Equilibrium agenda formation." *Social Choice and Welfare*, 23, 21–57. [472]

Dziuda, Wioletta and Antoine Loeper (2010), "Dynamic collective choice with endogenous status quo." Unpublished paper. [472]

Harsanyi, John C. (1974), "An equilibrium-point interpretation of stable sets and a proposed alternative definition." *Management Science*, 20, 1472–1495. [474]

Hayes-Renshaw, Fiona, Wim Van Aken, and Helen Wallace (2006), "When and why the EU Council of Ministers votes explicitly." *Journal of Common Market Studies*, 44, 161–194. [446]

Heisenberg, Dorothee (2005), "The institution of 'consensus' in the European Union." *European Journal of Political Research*, 44, 65–90. [446, 450, 463]

Hortala-Vallve, Rafael (2011), "Generous legislators? A description of vote trading agreements." *Quarterly Journal of Political Science*, 6, 179–196. [474]

Kalandrakis, Anastassios (2004), "A three-player dynamic majoritarian bargaining game." *Journal of Economic Theory*, 116, 294–322. [471, 472]

Kalandrakis, Tasos (2010), "Minimum winning coalitions and endogenous status quo." *International Journal of Game Theory*, 39, 617–643. [471, 472]

Konishi, Hideo and Debraj Ray (2003), "Coalition formation as a dynamic process." *Journal of Economic Theory*, 110, 1–41. [473]

Larcker, David F., Gaizka Ormazabal, and Daniel J. Taylor (2011), "The market reaction to corporate governance regulation." *Journal of Financial Economics*, 101, 431–448. [465]

Matsusaka, John G. and Oguzhan Ozbas (2012), "Shareholder empowerment: The right to approve and the right to propose." Unpublished paper. [448]

Mattila, Mikko (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. [446]

Ordeshook, Peter C. (1986), *Game Theory and Political Theory*. Cambridge University Press, New York. [452, 458]

Rubinstein, Ariel (1982), "Perfect equilibrium in a bargaining model." *Econometrica*, 50, 97–109. [446]

Thomson, Robert (2011), *Resolving Controversy in the European Union*. Cambridge University Press, Cambridge. [450]

Vartiainen, Hannu (2011), "Dynamic coalitional equilibrium." *Journal of Economic Theory*, 146, 672–698. [473, 474]

Vartiainen, Hannu (2012), "One-deviation principle and dynamic voting." Unpublished paper. [467, 468, 473, 474, 475]

Volden, Craig and Alan E. Wiseman (2007), "Bargaining in legislatures over particularistic and collective goods." *American Political Science Review*, 101, 79–92. [471]

von Neumann, John and Oskar Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton University Press, Princeton. [471, 472]

Zápal, Jan (2011), "Explicit and implicit status-quo determination in dynamic bargaining: Theory and application to FOMC directive." Unpublished paper. [472]