CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

Online Appendix to Discussion Paper No. 2012-03

Vincent Anesi and Daniel J Seidmann

December 2012

Online Appendix to Accompany

Bargaining over an Endogenous Agenda

CeDEx Discussion Paper Series

ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

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

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

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

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

The full list of CeDEx Discussion Papers is available at

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

Online Appendix to Accompany
Bargaining over an Endogenous Agenda

Vincent Anesi
University of Nottingham

Daniel J. Seidmann
University of Nottingham

December 5, 2012

In this online appendix, we provide some details of proofs which are not included in the main text of the paper.

A The dynamic game with endogenous protocol

Let Π be the set of all protocols as defined in Section 2. In Section 4, we analyzed a game in which the chair first selects a protocol $\pi \in \Pi$, and $\Gamma\left(\pi, x^0\right)$ is then played. Call this game: $\Gamma^c\left(\pi, x^0\right)$: the superscript stands for commitment. In this section, we establish the claim that Corollary 3 also applies in a different 'dynamic' game, $\Gamma^d\left(\Pi, x^0\right)$, where the chair selects the next proposer immediately after each vote which does not end the game.

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

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

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

B History-dependent strategies

B.1 Proof of Proposition 7

Consider game $\Gamma(\pi, x)$. Suppose that Z is a consistent choice set, and let $g \in Z^X$ be any selection of $F^\pi(Z, \cdot)$ — where $F^\pi(Z, \cdot)$ is obtained from the tree construction described in Section 3. This implies that $g(x) = x$ for all $x \in Z$ (recall that, when $x \in Z$, tree $\mathfrak{T}^\pi(Z, x)$ has a single path whose nodes are all equal to x). Furthermore, as Z is a consistent choice set, $R_Z(x) \neq \emptyset$ for all $x \notin Z$. This implies that at least one proposer (weakly) prefers to amend x to some policy in Z. This in turn implies that $\{x\} \neq \mathbf{M}\left(\succeq_{\pi_x(k)}, R_Z(x)\right) \subseteq \mathfrak{s}_k^\pi(Z, x)$ for al least one proposer k. Furthermore, it is readily checked that a consistent choice set must be closed; so that $\mathbf{M}\left(\succeq_{\pi_x(k)}, R_Z(x)\right) \neq \emptyset$. Hence, tree $\mathfrak{T}^\pi(Z,x)$ has at least one final node in Z, so that $F^\pi(Z,x) \neq \emptyset$ for all $x \notin Z$. This proves that $g(x)$ is well defined.

Our next step is to describe the semi-Markovian strategy profile σ^x which, to lighten the notation, will henceforth be referred to as σ. To describe σ, we first construct a partition $\{H_z\}_{z\in Z}$ of H, where each element H_z of the partition will be interpreted as the set of histories at which z should be implemented according to σ.

The partition is constructed recursively, starting with the null history $h = x^0 = x$. The null history $h = x$ belongs to $H_{g(x)}$. From the construction of $F^\pi(Z,x)$, this implies that there exists a vector $(z_1(h), \ldots, z_{m_x+1}(h))$ such that:

- If $x \in Z$ then $g(x) = x = z_1(h) = \ldots = z_{m_x+1}(h)$;
- If $x \notin Z$ then $g(x) = z_1(h) \in Z$, $x = z_{m_x+1}(h)$, and $z_k(h) \in s_k^\pi\left(Z, z_{k+1}(h)\right)$ for each $k = 1, \ldots, m_x$. The latter condition implies that $z_k(h)$ is one of the kth proposer's ideal policies in a set $A_k\left(Z, z_{k+1}(h)\right) \equiv P_Z\left(z_{k+1}(h)\right) \cup \{z_{k+1}(h)\} \cup Y_k(h)$, where $Y_k(h) \subseteq R_Z\left(z_{k+1}(h)\right.$.

Next, for every $(z,x) \in Z \times X$, let $r(z,x)$ be an arbitrary element of $F^\pi\left(Z \setminus P(z), x\right)$ (which is well defined since Z is a consistent choice set). Note that, by definition of $F^\pi\left(Z \setminus P(z), \cdot\right)$, $r(z,x) = x$ whenever $x \in Z \setminus P(z)$. For all $z \in Z$, $h \in H_z$ and $x \in X$, history (h,x) belongs to H_x if $x = z_k(h)$ for some proposer k at h, and belongs to $H_{r(z,x)}$ otherwise. This implies that there is a vector $(z_1(h,x), \ldots, z_{m_x+1}(h,x))$ such that:

- If $x = z_k(h)$ for some proposer k at h then $x = z_1(h,x) = \ldots = z_{m_x+1}(h,x)$;
- If $x \neq z_k(h)$ for any proposer k at h then $r(z,x) = z_1(h,x)$, $x = z_{m_x+1}(h,x)$, and $z_k(h,x) \in s_k^\pi\left(Z \setminus P(z), z_{k+1}(h,x)\right)$ for each $k = 1, \ldots, m_x$. The latter condition implies that $z_k(h,x)$ is one of the kth proposer's ideal policies in a set $A_k\left(Z \setminus P(z), z_{k+1}(h,x)\right) \equiv P_{Z\setminus P(z)}\left(z_{k+1}(h,x)\right) \cup \{z_{k+1}(h,x)\} \cup Y_k(h)$, where $Y_k(h,x) \subseteq R_{Z\setminus P(z)}\left(z_{k+1}(h,x)\right)$.

We are now in a position to define strategies. Take any round-t history $h \in H_z$. If $x^{t-1} = z$ then the ongoing default should be implemented at h: σ_i prescribes player

$i = \pi_x(k)$ to pass. (For expositional convenience, we will sometimes say that i proposes $z_k(h) = x$.) If $x^{t-1} \neq z$ then σ_i prescribes player $i = \pi_x(k)$ to propose $z_k(h)$ if $z_k(h) \neq z_{k+1}(h)$, and to pass if $z_k(h) = z_{k+1}(h)$. Since $\{H_z\}_{z\in Z}$ is a partition of H, the description of proposal strategies is complete.

We now turn to voting strategies. Consider first the null history $h = x^0$. Following a proposal $y \neq x^0$ by the kth proposer, σ_i prescribes voter i to act as follows:

(A0) If $h = x^0 \in Z$ then i votes 'yes' iff $z_1(h, y) \succ_i x^0$;

(B0) if $h = x^0 \notin Z$ and $z_1(h, y) \in A_k(Z, z_{k+1}(h))$ then i votes 'yes' iff $z_1(h, y) \succeq_i z_{k+1}(h)$;

(C0) if $h = x^0 \notin Z$ and $z_1(h, y) \notin A_k(Z, z_{k+1}(h))$ then i votes 'yes' iff $z_1(h, y) \succ_i z_{k+1}(h)$.

Take any round-t history h of the form $h = (h', x^{t-1})$ where $h' \in H_z$ for some $z \in Z$. Following a proposal $y \neq x^{t-1}$ by the kth proposer at h, σ_i prescribes voter i to act as follows:

(A1) If $h \in H_{x^{t-1}}$ then i votes 'yes' iff $z_1(h, y) \succ_i x^{t-1}$;

(B1) if $h \notin H_{x^{t-1}}$ and $z_1(h, y) \in A_k(Z \setminus P(z), z_{k+1}(h))$ then i votes 'yes' iff $z_1(h, y) \succeq_i z_{k+1}(h)$;

(C1) if $h \notin H_{x^{t-1}}$ and $z_1(h, y) \notin A_k(Z \setminus P(z), z_{k+1}(h))$ then i votes 'yes' iff $z_1(h, y) \succ_i z_{k+1}(h)$.

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

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

Proof: Let h be of the form $h = (h', x)$ where $h' \in H_z$ for some $z \in Z$. If $x = x^{t-1} = z_k(h')$ for some proposer k at h' then the claim is trivial: $h \in H_x$ and $z_{k+1}(h) = \ldots = z_{m+1}(h) = x^{t-1} = x$ (all the remaining proposers pass). Accordingly, suppose that $x = x^{t-1} \neq z_k(h')$ for any proposer k at h'. Since the kth proposer at h has not amended x^{t-1}, the $(k+1)$th proposer is given the opportunity to make a proposal. By definition of proposal strategies, she proposes $z_{k+1}(h)$ if $z_{k+1}(h) \neq z_{k+2}(h)$, and passes otherwise. Suppose first that $z_{k+1}(h) \neq z_{k+2}(h)$. If $z_{k+1}(h)$ were accepted then the history at the start of the next round would belong to $H_{z_{k+1}(h)}$, so that all proposers would pass and $z_{k+1}(h)$ would be implemented at the end of that round. Hence, $z_1(h, z_{k+1}(h)) = z_{k+1}(h) \in A_{k+1}(Z \setminus P(z), z_{k+2}(h)) \subseteq R(z_{k+2}(h))$. Condition (B1) in the definition of voting strategies then ensures that proposal $z_{k+1}(h)$ is accepted and then implemented in the next round.

Suppose now that $z_{k+1}(h) = z_{k+2}(h)$, so that the $(k+1)$th proposer passes. This implies that the $(k+2)$th proposer is given the opportunity to make a proposal. We can apply the same argument as above to show that either $z_{k+1}(h) = z_{k+2}(h)$ $(\neq z_{k+3}(h))$ is implemented in the next round or $z_{k+1}(h) = z_{k+2}(h) = z_{k+3}(h)$. Going on until the m_xth proposer, we obtain the claim.

A similar argument applies to the null history $h = x^0$.

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

Proof: If $z_k(h) \neq z_{k+1}(h)$ then, as demonstrated in the proof of the previous claim (end of the first paragraph), the kth proposer offers $z_k(h)$, which is accepted and implemented at the end of the next round.

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

Claim 3: $\phi^{\sigma}\left(x^{0}\right)=z_{1}\left(x^{0}\right)=g\left(x^{0}\right)$ for all $x^{0} \in X$.

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

Now suppose that x^{0} is not a member of Z. Since $z_{1}\left(x^{0}\right)=g\left(x^{0}\right) \in F^{\pi}\left(Z, x^{0}\right) \subseteq Z$, at least one proposer tries to amend x^{0}. The first proposer who does so, say $\pi_{x^{0}}(k)$, offers $z_{k}\left(x^{0}\right) R z_{k+1}\left(x^{0}\right)$ which, by condition (B0) in the definition of voting strategies, is accepted. This implies that $h=\left(x^{0}, z_{k}\left(x^{0}\right)\right) \in H_{z_{k}\left(x^{0}\right)}$, which in turn implies that $z_{k}\left(x^{0}\right)$ is never amended and is therefore implemented at the end of round 2. By definition of proposal strategies, $z_{l}\left(x^{0}\right)=z_{k}\left(x^{0}\right)$ for all proposers $l<k$ who do not try to amend x^{0}. Hence, $\phi^{\sigma}\left(x^{0}\right)=z_{k}\left(x^{0}\right)=z_{1}\left(x^{0}\right)=g\left(x^{0}\right)$.

As this is true for any $x^{0} \in X$, this proves that $\phi^{\sigma}(X) \equiv\left\{\phi^{\sigma}\left(x^{0}\right): x^{0} \in X\right\}=\left\{z_{1}\left(x^{0}\right): x^{0} \in X\right\}=Z$.

Claim 4: Let $h \in H$ be a round-t history. If the kth proposer has made proposal $y \neq x^{t-1}$ then σ_{i} prescribes i to vote 'yes' only if $\phi^{\sigma}(h, y ; 1) \succeq_{i} \phi^{\sigma}(h ; k+1)$, and to vote 'no' only if $\phi^{\sigma}(h ; k+1) \succeq_{i} \phi^{\sigma}(h, y ; 1)$.

Proof: Suppose h is of the form $h=\left(h^{\prime}, x^{t-1}\right)$ where $h^{\prime} \in H_{z}$ for some $z \in Z$. Claim 2 immediately implies that $\phi^{\sigma}(h, y ; 1)=z_{1}(h, y)$ for all $y \neq x^{t-1}$, and that $\phi^{\sigma}(h ; k+1)=z_{k+1}(h)$.

Suppose first that $h \in H_{x^{t-1}}$. If player i votes 'yes' then, by condition (A1), $z_{1}(h, y) \succ_{i} x^{t-1}$. Claim 2 implies that $x^{t-1}=z_{k+1}(h)=\phi^{\sigma}(h ; k+1)$ (proposal strategies prescribe all proposers to pass at all $h \in H_{x^{t-1}}$). Hence, $z_{1}(h, y) \succ_{i} x^{t-1}$ implies $\phi^{\sigma}(h, y ; 1) \succ_{i} \phi^{\sigma}(h ; k+1)$ and, therefore, that $\phi^{\sigma}(h, y ; 1) \succeq_{i} \phi^{\sigma}(h ; k+1)$. If player i votes 'no' then, by condition (A), $x^{t-1} \succ_{i} z_{1}(h, y)$. This in turn implies that $\phi^{\sigma}(h ; k+1) \succeq_{i} \phi^{\sigma}(h, y ; 1)$.

Now suppose that $h \notin H_{x^{t-1}}$ and that $z_{1}(h, y) \in A_{k}\left(Z \setminus P(z), z_{k+1}(h)\right)$. If player i votes 'yes' then, by condition (B1), $\phi^{\sigma}(h, y ; 1)=z_{1}(h, y) \succeq_{i} z_{k+1}(h)=\phi^{\sigma}(h ; k+1)$. If player i votes 'no' then, by condition (B1), $z_{k+1}(h) \succ_{i} z_{1}(h, y)$. This in turn implies that $\phi^{\sigma}(h ; k+1) \succ_{i} \phi^{\sigma}(h, y ; 1)$ and, therefore, $\phi^{\sigma}(h ; k+1) \succeq_{i} \phi^{\sigma}(h, y ; 1)$.

Finally, suppose that $h \in H_{x^{t-1}}$ and that $y \notin A_{k}\left(Z \setminus P(z), z_{k+1}(h)\right)$. If player i

votes 'yes' then, by condition (C1), $z_1(h, y) \succ_i z_{k+1}(h)$. This implies that $\phi^\sigma(h, y; 1) \succ_i \phi^\sigma(h; k+1)$ and, therefore, that $\phi^\sigma(h, y; 1) \succeq_i \phi^\sigma(h; k+1)$. Similarly, if i votes 'no' then (C1) implies that $z_{k+1}(h) \succeq_i z_1(h, y)$ and then $\phi^\sigma(h; k+1) \succeq_i \phi^\sigma(h, y; 1)$.

A similar argument applies to the null history $h = x^0$.

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

Let $i = \pi_x(k)$, and let h be of the form $h = (h', x)$ where $h' \in H_z$ for some $z \in Z$.

Suppose first that $h \in H_x$: viz. σ dictates all proposers to pass at h. Consequently, if i conforms to σ_i then the final policy outcome will be $x^{t-1} = x$. Hence, i can only profitably deviate by amending x with some policy $y \neq x$. By construction, however, history (h, y) belongs to $H_{r(x,y)}$ where $r(x, y) = z_1(h, y) \in Z \setminus P(x)$; so that $z_1(h, y) \notin P(x)$. From Condition (A1), this implies that i cannot amend x and, therefore, cannot profitably deviate.

Now suppose that $h \in H_w$ for some $w \neq x^{t-1}$. Any proposal y such that $z_1(h, y) \notin A_k(Z \setminus P(z), z_{k+1}(h))$ must be unsuccessful. Indeed, condition (C1) in the definition of voting histories implies that voters only vote 'yes' if they strictly prefer $z_1(h, y)$ to $z_{k+1}(h)$. As $P_{Z\setminus P(z)}(z_{k+1}(h)) \subseteq A_k(Z \setminus P(z), z_{k+1}(h)) \not\ni z_1(h, y)$, $z_1(h, y) \notin P_{Z\setminus P(z)}(z_{k+1}(h))$ and y must be voted down. Thus, as $z_k(h)$ is $\succeq_i$-maximal in $A_k(Z \setminus P(z), z_{k+1}(h)) \supseteq \{z_{k+1}(h)\}$, player i cannot improve on proposing $z_k(h)$ when $z_k(h) \neq z_{k+1}(h)$, and passing otherwise.

A similar argument applies to the null history $h = x^0$. This completes the proof of the Proposition.

B.2 Proof of Proposition 8

Let σ be a semi-Markovian equilibrium. Suppose that, contrary to the statement of the result, $\phi^\sigma(H)$ is not a consistent choice set. This implies that there exist $x \in \phi^\sigma(H)$, $y \in X$ such that, for all $z \in \phi^\sigma(H)$, one of the following conditions is true:

(a) $z \notin R(y)$;

(b) $z \in R(y) \cap P(x)$.

Now consider a history $h \in H$ at which, instead of following σ and implementing x at the end of the round, some players have deviated as follows: a proposer $\pi_x(k)$ has proposed to amend x with y and all members of some $S \in \mathcal{W}$ have voted 'yes'. This deviation yields a new outcome $z \in \phi^\sigma(H)$, which satisfies one of the conditions (a)-(b) above. Under assumptions (i) or/and (ii) in the statement of the result,[1] some winning coalition in $\mathcal{W}$ must find it (weakly) profitable to induce z from y in equilibrium and, therefore, z must satisfy (b). Hence, there exists $S \in \mathcal{W}$ such that $z \succ_i x$ for all $i \in S$.

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

B.3 Quasi-Markovian equilibria and quasi-consistent sets

We observed at the end of Subsection 5.2 that, by allowing strategies to depend not only on the sequence of previous defaults but also on the coalitions which amended previous defaults, we can obtain analogs of Propositions 7-8 in which "consistent choice set" is replaced by "quasi-consistent set." To prove that statement, we first need some definitions. In this subsection, we will indulge in a slight abuse of terminology and will call a 'round-t history' any list $\left(x^0, S^1, x^1, \ldots, S^{t-1}, x^{t-1}\right)$ where $S^s \in \mathcal{W}$ stands for the winning coalition which amended x^{s-1} to x^s. Let $\overline{H}^t$ be the set of round-t histories — $\overline{H}^1 \equiv \left\{x^0\right\}$ being the null history — and let $\overline{H} \equiv \bigcup_{t=1}^{\infty} \overline{H}^t$ be the set of histories. We define a 'quasi-Markovian' strategy as an analog of a stationary Markov strategy where histories play the

[1] Those conditions ensure that it is always the last amender (if any) who changes the current default x to another policy y. This in turn implies that, following the last amender's proposal, voters compare x with the final policy outcome induced by the move from x to y, say z. For that move to happen in equilibrium, therefore, it must be that z R-dominates y.

role of the ongoing default. More specifically: in proposal stages, strategies only depend on the history and the identity of the remaining proposers in the current round; in voting stages, strategies only depend on the history, the proposal just made, the votes already cast thereon, and the remaining proposers in the current round.

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

The proofs of following results parallel those of Propositions 7 and 8.

Result 1 *Suppose that Z is the closure of a quasi-consistent 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 quasi-Markovian equilibrium of $\Gamma(\pi, x)$ and $\bar{\phi}^{\sigma^x}(x) = g(x)$. Hence, $\bigcup_{x \in X} \bar{\phi}^{\sigma^x}(x) = Z$.*

Proof: Consider game $\Gamma(\pi, x)$. Suppose that Z is the closure of a quasi-consistent set, and let $g \in Z^X$ be any selection of $F^{\pi}(Z, \cdot)$ — where $F^{\pi}(Z, \cdot)$ is obtained from the tree construction described in Section 3. It is readily checked that that the closure of a quasi-consistent is itself a quasi-consistent set; so that Z is quasi-consistent. Note that that $g(x) = x$ for all $x \in Z$ (recall that, when $x \in Z$, tree $\mathfrak{T}^{\pi}(Z, x)$ has a single path whose nodes are all equal to x). Furthermore, as Z is quasi-consistent, $R_Z(x) \neq \emptyset$ for all $x \notin Z$. This implies that at least one proposer (weakly) prefers to amend x to some policy in Z. This in turn implies that $\{x\} \neq \mathbf{M}\left(\succeq_{\pi_x(k)}, R_Z(x)\right) \subseteq \mathfrak{s}_k^{\pi}(Z, x)$ for al least one proposer k. Hence, tree $\mathfrak{T}^{\pi}(Z, x)$ has at least one final node in Z, so that $F^{\pi}(Z, x) \neq \emptyset$ for all $x \notin Z$. This proves that $g(x)$ is well defined.

Our next step is to describe the quasi-Markovian strategy profile σ^x which, to lighten the notation, will henceforth be referred to as σ. To describe σ, we first construct a partition $\{\overline{H}_z\}_{z\in Z}$ of $\overline{H}$, where each element $\overline{H}_z$ of the partition will be interpreted as the set of histories at which z should be implemented according to σ. In what follows, we will denote by $\left(h^t, S, y\right)$ the concatenation of the round-t history h^t followed by a round in which coalition $S \in \mathcal{W}$ amends x^{t-1} to y.

The partition is constructed recursively, starting with the null history $h = x^0 = x$. The null history $h = x$ belongs to $\overline{H}_{g(x)}$. From the construction of $F^\pi(Z, x)$, this implies that there exists a vector $(z_1(h), \ldots, z_{m_x+1}(h))$ such that:

- If $x \in Z$ then $g(x) = x = z_1(h) = \ldots = z_{m_x+1}(h)$;
- If $x \notin Z$ then $g(x) = z_1(h) \in Z$, $x = z_{m_x+1}(h)$, and $z_k(h) \in s_k^\pi\left(Z, z_{k+1}(h)\right)$ for each $k = 1, \ldots, m_x$. The latter condition implies that $z_k(h)$ is one of the kth proposer's ideal policies in a set $A_k\left(Z, z_{k+1}(h)\right) \equiv P_Z\left(z_{k+1}(h)\right) \cup \{z_{k+1}(h)\} \cup Y_k(h)$, where $Y_k(h) \subseteq R_Z\left(z_{k+1}(h)\right.$.

Next, for every $(z, S, x) \in Z \times \mathcal{W} \times X$, let

$$T(S, z) \equiv \left\{z' \in Z : z \succeq_i z \text{ for some } i \in S\right\} ,$$

and let $r(z, S, x)$ be an arbitrary element of $F^\pi\left(T(S, z), x\right)$ (which is well defined since Z is quasi-consistent). Note that, by definition of $F^\pi\left(T(S, z), \cdot\right)$, $r(z, S, x) = x$ whenever $x \in T(S, z)$. For all $z \in Z$, $h \in \overline{H}_z$ and $x \in X$, history (h, x) belongs to $\overline{H}_x$ if $x = z_k(h)$ for some proposer k at h, and belongs to $\overline{H}_{r(z,S,x)}$ otherwise. This implies that there is a vector $(z_1(h, S, x), \ldots, z_{m_x+1}(h, S, x))$ such that:

- If $x = z_k(h)$ for some proposer k at h then $x = z_1(h, S, x) = \ldots = z_{m_x+1}(h, S, x)$;
- If $x \neq z_k(h)$ for all proposers k at h then $r(z, x) = z_1(h, S, x)$, $x = z_{m_x+1}(h, S, x)$, and $z_k(h, S, x) \in s_k^\pi\left(T(S, z), z_{k+1}(h, S, x)\right)$ for each $k = 1, \ldots, m_x$. The latter condition implies that $z_k(h, S, x)$ is one of the kth proposer's ideal policies in a set

  $$A_k\left(T(S, z), z_{k+1}(h, S, x)\right) \equiv P_{T(S,z)}\left(z_{k+1}(h, S, x)\right) \cup \{z_{k+1}(h, S, x)\} \cup Y_k(h, S, x) ,$$

  where $Y_k(h, S, x) \subseteq R_{T(S,z)}\left(z_{k+1}(h, S, x)\right)$.

We are now in a position to define strategies. Take any round-t history $h \in \overline{H}_z$. If $x^{t-1} = z$ then the ongoing default should be implemented at h: σ_i prescribes player $i = \pi_x(k)$ to pass. (For expositional convenience, we will sometimes say that i proposes $z_k(h) = x$.) If $x^{t-1} \neq z$ then σ_i prescribes player $i = \pi_x(k)$ to propose $z_k(h)$ if $z_k(h) \neq z_{k+1}(h)$, and to pass if $z_k(h) = z_{k+1}(h)$. Since $\{\overline{H}_z\}_{z \in Z}$ is a partition of $\overline{H}$, the description of proposal strategies is complete.

We now turn to voting strategies. Consider first the null history $h = x^0$. Suppose the kth proposer has made proposal $y \neq x^0$. Let S_i^- be the set of players who have already voted 'yes' when it is i's turn to vote, and let S_i^+ be the set of voters j who will vote after i and are prescribed to vote 'yes' by σ_j. If $S \equiv S_i^- \cup \{i\} \cup S_i^+$ is a winning coalition then σ_i prescribes voter i to act as follows:

(A0) If $h = x^0 \in Z$ then i votes 'yes' iff $z_1(h, S, y) \succ_i x^{t-1}$ for any winning coalition $S \ni i$;

(B0) if $h = x^0 \notin Z$ and $y \in A_k(Z, z_{k+1}(h))$ then i votes 'yes' iff $y \succeq_i z_{k+1}(h)$;

(C0) if $h = x^0 \notin Z$ and $y \notin A_k(Z, z_{k+1}(h))$ then i votes 'yes' iff $z_1(h, S, y) \succ_i z_{k+1}(h)$ for any winning coalition $S \ni i$.

If S is not a winning coalition then the voting behavior prescribed by σ_i at h is arbitrary.

Now take any round-t history h of the form $h = (h', S^{t-1}, x^{t-1})$ where $h' \in \overline{H}_z$ for some $z \in Z$. Suppose the kth proposer has made proposal $y \neq x^{t-1}$. Again, let S_i^- be the set of players who have already voted 'yes' when it is i's turn to vote, and let S_i^+ be the set of voters j who will vote after i and are prescribed to vote 'yes' by σ_j. If $S \equiv S_i^- \cup \{i\} \cup S_i^+$ is a winning coalition then σ_i prescribes voter i to act as follows:

(A1) If $h \in \overline{H}_{x^{t-1}}$ then i votes 'yes' iff $z_1(h, S, y) \succ_i x^{t-1}$;

(B1) if $h \notin \overline{H}_{x^{t-1}}$ and $z_1(h, S, y) \in A_k\left(T\left(S^{t-1}, z\right), z_{k+1}(h)\right)$ then i votes 'yes' iff $z_1(h, S, y) \succeq_i z_{k+1}(h)$;

(C1) if $h \notin \overline{H}_{x^{t-1}}$ and $z_1(h, S, y) \notin A_k\left(T\left(S^{t-1}, z\right), z_{k+1}(h)\right)$ then i votes 'yes' iff $z_1(h, S, y) \succ_i z_{k+1}(h)$ for any winning coalition $S \ni i$.

If S is not a winning coalition then the voting behavior prescribed by σ_i at h is arbitrary.

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

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

Proof: Let h be of the form $h = (h', S, x)$ where $h' \in \overline{H}_z$ for some $z \in Z$. If $x = x^{t-1} = z_k(h')$ for some proposer k at h' then the claim is trivial: $h \in \overline{H}_x$ and $z_{k+1}(h) = \ldots = z_{m+1}(h) = x^{t-1} = x$ (all the remaining proposers pass). Accordingly, suppose that $x = x^{t-1} \neq z_k(h')$ for all proposers k at h'. Since the kth proposer at h has not amended x^{t-1}, the $(k+1)$th proposer is given the opportunity to make a proposal. By definition of proposal strategies, she proposes $z_{k+1}(h)$ if $z_{k+1}(h) \neq z_{k+2}(h)$, and passes otherwise. Suppose first that $z_{k+1}(h) \neq z_{k+2}(h)$. If $z_{k+1}(h)$ were accepted by some winning coalition S' then the history at the start of the next round would belong to $\overline{H}_{z_{k+1}(h)}$, so that all proposers would pass and $z_{k+1}(h)$ would be implemented at the end of that round. Hence, $z_1(h, S', z_{k+1}(h)) = z_{k+1}(h) \in A_{k+1}(T(S,z), z_{k+2}(h)) \subseteq R(z_{k+2}(h))$. Condition (B1) in the definition of voting strategies then ensures that proposal $z_{k+1}(h)$ is accepted and then implemented in the next round.

Suppose now that $z_{k+1}(h) = z_{k+2}(h)$, so that the kth proposer passes. This implies that the $(k+2)$th proposer is given the opportunity to make a proposal. We can apply the same argument as above to show that either $z_{k+1}(h) = z_{k+2}(h)$ $(\neq z_{k+3}(h))$ is implemented in the next round or $z_{k+1}(h) = z_{k+2}(h) = z_{k+3}(h)$. Going on until the m_xth proposer, we obtain the claim.

The same argument applies if h is the null history.

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

Proof: If $z_k(h) \neq z_{k+1}(h)$ then, as demonstrated in the proof of the previous claim (end of the first paragraph), the kth proposer offers $z_k(h)$, which is accepted and implemented at the end of the next round.

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

Claim 3: $\bar{\phi}^\sigma\left(x^0\right) = z_1\left(x^0\right) = g\left(x^0\right)$ for all $x^0 \in X$.

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

Now suppose that x^0 is not a member of Z. Since $z_1\left(x^0\right) = g\left(x^0\right) \in F^\pi\left(Z, x^0\right) \subseteq Z$, at least one proposer tries to amend x^0. The first proposer who does so, say $\pi_{x^0}(k)$, offers $z_k\left(x^0\right) R z_{k+1}\left(x^0\right)$ which, by condition (B0) in the definition of voting strategies, is accepted by some winning coalition S_0. This implies that $h = \left(x^0, S_0, z_k\left(x^0\right)\right) \in \overline{H}_{z_k(x^0)}$, which in turn implies that $z_k\left(x^0\right)$ is never amended and is therefore implemented at the end of round 2. By definition of proposal strategies, $z_l\left(x^0\right) = z_k\left(x^0\right)$ for all proposers $l < k$ who do not try to amend x^0. Hence, $\phi^\sigma\left(x^0\right) = z_k\left(x^0\right) = z_1\left(x^0\right) = g\left(x^0\right)$.

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

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

Proof: Suppose h is of the form $h = \left(h', x^{t-1}\right)$ where $h' \in \overline{H}_z$ for some $z \in Z$. Claim 2 immediately implies that $\bar{\phi}^\sigma(h, S, y; 1) = z_1(h, S, y)$ for all $y \neq x^{t-1}$, and $\bar{\phi}^\sigma(h; k+1) = z_{k+1}(h)$.

Suppose first that $h \in \overline{H}_{x^{t-1}}$. If player i votes 'yes' then, by condition (A1), $z_1(h, S, y) \succ_i x^{t-1}$. Claim 2 implies that $x^{t-1} = z_{k+1}(h) = \bar{\phi}^\sigma(h; k+1)$ (proposal strategies prescribe all proposers to pass at all $h \in H_{x^{t-1}}$). Hence, $z_1(h, S, y) \succ_i x^{t-1}$ implies $\bar{\phi}^\sigma(h, S, y; 1) \succ_i \bar{\phi}^\sigma(h; k+1)$ and, therefore, that $\bar{\phi}^\sigma(h, S, y; 1) \succeq_i \bar{\phi}^\sigma(h; k+1)$. If player i votes 'no' then, by condition (A1), $x^{t-1} \succ_i z_1(h, S, y)$. This in turn implies that $\bar{\phi}^\sigma(h; k+1) \succeq_i \bar{\phi}^\sigma(h, S, y; 1)$.

Now suppose that $h \notin \overline{H}_{x^{t-1}}$ and that $z_1(h, S, y) \in A_k(T(S, z), z_{k+1}(h))$. If player i votes 'yes' then, by condition (B1) and Claim 2, $\bar{\phi}^\sigma(h, S, y; 1) = z_1(h, S, y) \succeq_i z_{k+1}(h) = \bar{\phi}^\sigma(h; k+1)$. If player i votes 'no' then, by condition (B1), $z_{k+1}(h) \succ_i z_1(h, S, y)$. This in turn implies that $\bar{\phi}^\sigma(h; k+1) \succ_i \bar{\phi}^\sigma(h, S, y; 1)$ and, therefore, $\bar{\phi}^\sigma(h; k+1) \succeq_i \bar{\phi}^\sigma(h, y; 1)$.

Finally, suppose that $h \in \overline{H}_{x^{t-1}}$ and that $z_1(h, S, y) \notin A_k(T(S, z), z_{k+1}(h))$. If player i votes 'yes' then, by condition (C1), $z_1(h, S, y) \succ_i z_{k+1}(h)$. This implies that $\bar{\phi}^\sigma(h, S, y; 1) \succ_i \bar{\phi}^\sigma(h; k+1)$ and, therefore, that $\bar{\phi}^\sigma(h, S, y; 1) \succeq_i \bar{\phi}^\sigma(h; k+1)$. Similarly, if i votes 'no' then (C1) implies that $z_{k+1}(h) \succeq_i z_1(h, S, y)$ and then $\bar{\phi}^\sigma(h; k+1) \succeq_i \bar{\phi}^\sigma(h, S, y; 1)$.

A similar argument applies if h is the null history.

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

Let $i = \pi_x(k)$, and let h be of the form $h = (h', x)$ where $h' \in \overline{H}_z$ for some $z \in Z$.

Suppose first that $h \in H_x$: viz. σ dictates all proposers to pass at h. Consequently, if i conforms to σ_i then the final policy outcome will be $x^{t-1} = x$. Hence, i can only profitably deviate by amending x with some policy $y \neq x$. By construction, however, history (h, S, y) belongs to $\overline{H}_{r(x,S,y)}$ for all $S \in \mathcal{W}$, where $r(x, S, y) = z_1(h, S, y) \in T(S, x)$; so at least one member of S weakly prefers x to $z_1(h, S, y)$. From Condition (A1), this implies that i cannot amend x and, therefore, cannot profitably deviate.

Now suppose that $h \in \overline{H}_w$ for some $w \neq x^{t-1}$. Any proposal y such that $z_1(h, S, y) \notin A_k\left(T(S,z), z_{k+1}(h)\right)$ for all $S \in \mathcal{W}$ must be unsuccessful. Indeed, condition (C1) in the definition of voting histories implies that voters in $S \in \mathcal{W}$ only vote 'yes' if they strictly prefer $z_1(h, S, y)$ to $z_{k+1}(h)$. As $P_{T(S,z)}\left(z_{k+1}(h)\right) \subseteq A_k\left(T(S,z), z_{k+1}(h)\right) \not\ni z_1(h,y)$, $z_1(h,S,y) \notin P_{T(S,z)}\left(z_{k+1}(h)\right)$ and y must be voted down (i.e. at least one member of S votes no). Thus, as $z_k(h)$ is $\succeq_i$-maximal in $A_k\left(T(S,z), z_{k+1}(h)\right)$, player i cannot improve on proposing $z_k(h)$ when $z_k(h) \neq z_{k+1}(h)$, and passing otherwise.

A similar argument applies if h is the null history. This ends the proof of Result 1.

$\square$

Result 2 *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 quasi-Markovian equilibrium then $\bar{\phi}^{\sigma}\left(\overline{H}\right) \equiv \bigcup_{h \in \overline{H}} \bar{\phi}^{\sigma}(h,1)$ is a quasi-consistent set.*

Proof: Let σ be a quasi-Markovian equilibrium. Suppose that, contrary to the statement of Result 2, $\bar{\phi}^{\sigma}(\overline{H})$ is not a consistent set. This implies that there exist $o \in \bar{\phi}^{\sigma}(\overline{H})$, $x \in X$, and $S \in \mathcal{W}$ such that, for all $o' \in \bar{\phi}^{\sigma}(\overline{H})$, one of the following conditions is true:

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

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

(c) $\neg\left(o' R x\right)$.

Now consider a history $h \in \overline{H}$ at which, instead of following σ and implementing o at the end of the round, some players have deviated as follows: a proposer $\pi_o(k)$ in S has proposed to amend o with x and all members of S have voted 'yes'. This deviation yields a new outcome $o' \in \phi^{\sigma}(\overline{H})$, which satisfies one of the conditions (a)-(c) above. Under assumptions (i) or/and (ii) in the statement of Result 2,[2] some winning coalition in $\mathcal{W}$ must find it (weakly) profitable to induce o' from x in equilibrium and, therefore, o' cannot satisfy (c). As a consequence, o' must satisfy either (a) or (b).

[2]Those conditions ensure that it is always the last amender (if any) who changes the current default x to another policy y. This in turn implies that, following the last amender's proposal, voters compare x with the final policy outcome induced by the move from x to y, say o'. For that move to happen in equilibrium, therefore, it must be that o' R-dominates x.

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

□

Combining Results 1 and 2, we obtain the following analog to Corollary 4:

Result 3 *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 quasi-Markovian equilibrium policy outcomes that can be obtained from any initial default in X coincides with the union of quasi-consistent sets.*

B.4 Semi-Markovian equilibrium policies and consistent choice sets: a counterexample

Suppose $C = N = \{1,2,3,4\}$, $M = \{2,3\}$, and $\mathbf{W} = \{S \subseteq N : |S| \geq 3\}$ (majority voting). The set of policies is $X = \{x, z_1, z_2, z_3\}$, and players' preferences over X are given by:

$$z_1 \succ_1 z_2 \sim_1 z_3 \succ_1 x \quad , \quad z_2 \sim_2 z_3 \succ_2 z_1 \succ_2 x \ ,$$

$$z_3 \succ_3 x \succ_3 z_1 \succ_3 z_2 \quad , \quad x \succ_4 z_2 \sim_4 z_3 \succ_4 z_1 \ .$$

Let π be the constant protocol in which, in every round, player 2 makes the first proposal and player 3 makes the second (and last) proposal — so that $\mathbf{W} = \mathcal{W}$. We want to show that the bargaining game $\Gamma(\pi, z_1)$ has a semi-Markovian equilibrium σ in which $\{z_1, z_2, z_3\}$ — though neither a consistent choice set nor a quasi-consistent set — is the set of immovable policies. To define σ, we first partition the set of (non-null) partial histories into the class $\{H(a,b)\}_{(a,b)\in X^2}$ where, for each ordered pair $(a,b) \in X^2$ and every partial history $h \in H$, $h \in H(a,b)$ if and only if a and b are the penultimate and last defaults in h, respectively.

The voting behavior prescribed by σ is described in the table below. It tells us, for all $(a,b) \in X^2$, which players vote 'yes' when proposer i proposes y_i at partial history $h \in H(a,b)$ ('z_1' stands for the null history).

	$y_3 = x$	$y_3 = z_1$	$y_3 = z_2$	$y_3 = z_3$	$y_2 = x$	$y_2 = z_1$	$y_2 = z_2$	$y_2 = z_3$	ψ
z_1	2,4	no vote	2,4	2,4	2,4	no vote	2,4	2,4	z_1
(z_1, z_1)	2,4	no vote	2,4	2,4	2,4	no vote	2,4	2,4	z_1
(x, z_1)	2,4	no vote	2,4	2 ,4	2,4	no vote	2,4	2,4	z_1
(z_2, z_1)	2, 4	no vote	2,4	2,4	2, 4	no vote	2,4	2,4	z_1
(z_3, z_1)	2, 4	no vote	2,4	2,4	2, 4	no vote	2,4	2,4	z_1
(z_1, x)	no vote	1,2	1,2	1,2,3	no vote	1	1,2,4		z_2
(z_1, z_2)	3	1,3	no vote	3	3	1,3	no vote		z_2
(z_2, z_2)	3	1,3	no vote	3	3	1,3	no vote		z_2
(z_1, z_3)				no vote			1,2,4	no vote	z_2
(z_3, z_2)	3	1,3	no vote	3	3	1,3	no vote		z_2
(x, z_2)	3	1,3	no vote	3	3	1,3	no vote		z_2
(x, x)	no vote	1 , 2	1,2	1,2,3	no vote	1			z_3
(z_3, z_3)	4	1	1,2	no vote	4	1	1,2	no vote	z_3
(z_2, x)	no vote	1 , 2	1,2	1,2,3	no vote	1			z_3
(z_3, x)	no vote	1 , 2	1,2	1,2,3	no vote	1			z_3
(z_2, z_3)	4	1	1,2	no vote	4	1	1,2	no vote	z_3
(x, z_3)	4	1	1,2	no vote	4	1	1,2	no vote	z_3

Moreover, at any $h \in H$, player 2 always proposes z_2 and player 3 always proposes z_3. It is readily checked that, under σ, $x^{t-2} \neq x^{t-1}$ implies that $x^t = \psi\left(x^{t-2}, x^{t-1}\right)$, where

$\psi : X^2 \cup \{z_1\} \to X$ is defined as: $\psi(z_1) = z_1$ and

$$\psi(a,b) \equiv \begin{cases} z_1 & \text{if } (a,b) = \{(z_1,z_1),(x,z_1),(z_2,z_1),(z_3,z_1)\} \ , \\ z_2 & \text{if } (a,b) \in \{(z_1,x),(z_1,z_2),(z_2,z_2),(z_1,z_3),(x,z_2),(z_3,z_2)\} \ , \\ z_3 & \text{if } (a,b) \in \{(x,x),(z_3,z_3),(z_2,x),(z_3,x),(z_2,z_3),(x,z_3)\} \ . \end{cases}$$

Thus, for all $(a,b) \in X^2$ such that $a \neq b$, and all $h \in H(a,b)$, we have

$$\phi^\sigma(h) = \lim\left(b, \psi(a,b), \psi(b,\psi(a,b)), \psi(\psi(a,b), \psi(b,\psi(a,b))), \ldots\right) \ ; \tag{1}$$

that is, $\phi^\sigma(h)$ is the limit of sequence of defaults generated by ψ from (a,b) — it is easy to verify that this limit always exists.

Inspection of the table above reveals that proposer i is confronted with a social acceptance set which is either empty — in which case, proposing z_i is trivially optimal — or equal to $\{z_i\}$ — in which case, it is optimal for player i to propose her ideal policy z_i. This shows that there is no profitable one-shot deviation in a proposal stage.

One can easily check that, following player 3's proposal y_3 at the voting stage of any history $h \in H(a,b)$, each player $i \in N$ votes 'yes' only if she weakly prefers the final policy induced from accepting y_3 over b — which, by equation (1), is equivalent to

$$\lim\left(b, \psi(a,b), \psi(b,\psi(a,b)), \psi(\psi(a,b), \psi(b,\psi(a,b))), \ldots\right) \succeq_i b$$

— and votes 'no' only if she weakly prefers b over the final policy induced from accepting y_3. Take for instance a partial history $h \in H(z_1, x)$:

- If player 3 passes (i.e. $y_3 = x$) then there is no vote.
- If player 3 proposes $y_3 = z_1$ then player i anticipates that amending x to z_1 would lead to the implementation of z_1 since

  $$\phi^\sigma(h) = \lim\left(z_1, \psi(x,z_1) = z_1, \psi(z_1,z_1) = z_1, \psi(z_1,z_1) = z_1, \ldots\right) = z_1 \ .$$

  It is therefore optimal for players 1 and 2 [resp. 3 and 4] to vote 'yes' [resp. 'no'] — as they both strictly prefer z_1 to x [resp. x to z_1].

- If player 3 proposes $y_3 = z_2$ then player i anticipates that amending x to z_2 would lead to the implementation of z_1 since

$$\phi^{\sigma}(h) = \lim\left(z_2, \psi\left(x, z_2\right) = z_2, \psi\left(z_2, z_2\right) = z_2, \psi\left(z_2, z_2\right) = z_2, \ldots\right) = z_2 \ .$$

  It is therefore optimal for players 1 and 2 [resp. 3 and 4] to vote 'yes' [resp. 'no'] — as they both strictly prefer z_2 to x [resp. x to z_2].

- If player 3 proposes $y_3 = z_3$ then player i anticipates that amending x to z_3 would lead to the implementation of z_1 since

$$\phi^{\sigma}(h) = \lim\left(z_3, \psi\left(x, z_3\right) = z_3, \psi\left(z_3, z_3\right) = z_3, \psi\left(z_3, z_3\right) = z_3, \ldots\right) = z_3 \ .$$

  It is therefore optimal for players 1, 2 and 2 [resp. player 4] to vote 'yes' [resp. 'no'] — as they both strictly prefer z_3 to x [resp. she strictly prefers x to z_3]. As player 3 (optimally) proposes policy z_3 at h, z_3 would be the final policy outcome if player 3 is given the opportunity to propose at h: $\phi^{\sigma}(h, 2) = z_3$.

The same argument applies to all other partial histories.

Furthermore, one can easily check that, following player 2's proposal y_2 at the voting stage of any history $h \in H(a, b)$, each player $i \in N$ votes 'yes' only if she weakly prefers the final policy induced from accepting y_2 over the final outcome induced from rejecting y_2. Take again the example of a partial history $h \in H\left(z_1, x\right)$:

- If player 2 passes (i.e. $y_2 = x$) then there is no vote.

- If player 2 proposes $y_2 = z_1$ then player i anticipates that, while rejecting z_1 would lead to the implementation of z_3 (see above), amending x to z_1 would lead to the implementation of z_1 (same argument as above). It is therefore optimal for player 1 to vote 'yes' — as $z_1 \succ_1 z_3$ — and for player $i \in \{2, 3, 4\}$ to vote 'no' — as $z_3 \succ_i z_1$.

- If player 2 proposes $y_2 = z_2$ then player i anticipates that, while rejecting z_2 would lead to the implementation of z_3 (see above), amending x to z_1 would lead to the implementation of z_1 (same argument as above). It is therefore optimal for players

1,2 and 4 — who are indifferent between z_2 and z_3 to vote 'yes', and for player 3 — who strictly prefers z_3 to z_2 — to vote 'no'.

- If player 2 proposes $y_2 = z_3$ then player i anticipates that, whether y_3 is accepted or rejected, the final policy outcome will be z_3. All voters are therefore indifferent between voting 'yes' or 'no' and optimally vote 'no' in equilibrium.

This proves that σ prescribe optimal voting behavior at any $h \in H(a, b)$ (i.e. the seventh row in the table above). A similar argument shows that it also prescribes optimal voting behavior at all the other partial histories.

C Mixed strategy equilibria

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

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

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

Utilities $u_i(w)$

and that the protocol is given by

$$(\pi_x(1), \pi_x(2), \pi_x(3)) = (2, 3, 1) \ ;$$

$$(\pi_y(1), \pi_y(2), \pi_y(3)) = (3, 1, 2) \ ;$$

$$(\pi_z(1), \pi_z(2), \pi_z(3)) = (1, 2, 3) \ .$$

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

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

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

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

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

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

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

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

V then earns

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

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

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

V then earns

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

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

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

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

Given the strategy combination above, we have

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

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

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

D Markov trembling-hand perfect equilibria

In this section, we provide a proof of Observation 5. To prove this result, it suffices to show that, for every weakly stable set $V\in\mathcal{V}$, there is an MTHPE equilibrium σ which supports V. To do so, we will use the construction described in the proof of Proposition 1. Consider the equilibrium described in that proof, say $\tilde{\sigma}$, which is obtained by setting $Y=\emptyset$. In this equilibrium, all proposers pass if the default x belongs to V. If $x\notin V$ then, for each $k\in\{1,\ldots,m_x\}$, the kth proposer offers $y_k(x)$ — i.e.: the $\succ_{\pi_x(k)}$-maximal element in $P_V\left(y_{k+1}(x)\right)\cup\{y_{k+1}(x)\}$ — and voter $i\in N$ accepts this proposal if and only if $y_1\left(y_k(x)\right)\succ_i y_{k+1}(x)$ — where, for all $x\notin V$, $y_1(x)$ is the ideal policy of the last amender of x in $P_V(x)\cup\{x\}$ and, for all $v\in V$, $y_1(v)=v$. Thus, if the current default x does not belong to V: all proposers who move before the last amender of x make an unsuccessful proposal (by internal stability of V); the last amender amends x to $y_1(x)$; and (off the equilibrium path) proposers k who move after the last amender choose $y_k(x)=x$ (i.e., they pass).

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

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

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

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

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

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

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

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

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

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

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

or, equivalently,

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

for all $y \in X$.

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

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

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

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

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

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

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