Existence of Markov Electoral Equilibria

John Duggan[*] Jean Guillaume Forand[†]

June 13, 2018

Abstract

We establish existence and continuity properties of equilibria in an infinite-horizon model of dynamic elections with a discrete (countable) state space and general policies and preferences.

1 Introduction

In this note, we introduce a dynamic model of elections and prove existence and continuity properties of its Markov electoral equilibria, which adapt the notion of stationary Markov perfect equilibrium to our political environment. The model features a sequence of elections such that in each period: a state is publicly observed; an incumbent politician chooses policy; voters (after observing this choice) decide between the incumbent and a challenger; and a new state is drawn at the beginning of the next period, and the process repeats. We assume that the incumbent's type is revealed to voters, but we allow for uncertainty regarding the challenger's type, so elections typically match a known incumbent against a relatively unknown challenger. We focus on a form of stationary equilibrium, in which voters compare the expected payoffs from the incumbent and challenger, and consistent with weak dominance, all voters cast their ballots for their preferred candidate; accordingly, an office holder weighs policy consequences against electoral incentives. To add depth to the dynamic incentives of politicians, and to capture stickiness of voter expectations about incumbents, we allow for the possibility of "ex post" commitment, i.e., if an incumbent chooses a policy in a certain state, and the state persists into the next period, then the policy remains in place. This form of commitment differs from the standard Downsian assumption, in which both candidates can costlessly make binding campaign promises. Here, the commitment is only on the incumbent's side and is after the politician has made a policy choice, reflecting fixed costs, personnel appointments, or institutional arrangements specific to that policy.

The model belongs to a broad literature on electoral accountability; see Duggan and Martinelli (2017) for a review. It includes as a special case the model of Duggan and Forand (2018), which assumes that a single representative voter is fixed, independently of the state, and that the state transition probability is independent of the incumbent's type. In that paper, we consider the possibility that equilibrium incentives lead to policy responsiveness, in the sense that policy choices solve (or approximately solve) the dynamic programming problem of the representative voter. Here, we allow for multiple voter types, and we allow the set of decisive coalitions to vary with the state, reflecting changes in the voting rule, the

[*]University of Rochester, dugg@ur.rochester.edu.

[†]University of Waterloo, jgforand@uwaterloo.ca. This author acknowledges support from a SSHRC IDG.

franchise, voter demographics, or differential turnout across types. Moreover, we allow the state transition to depend on the incumbent's type, so that types may capture competence or other politician-specific constraints on policy.

The existence of Markov electoral equilibria in this setting does not follow from existence results in the game-theoretic or bargaining literature. The closest point of contact is Duggan (2017), who proves existence of stationary bargaining equilibria in a dynamic bargaining model with a countable set of states. There is a correspondence between the bargaining and electoral models of these papers: in the latter, a politician chooses policy, utility from that policy is realized, and then voters approve or disapprove of the decision by casting ballots in an election; in the former, a proposer offers an alternative, other players approve or disapprove of the proposal in a vote, and then utility is accrued from either the proposal or a status quo, depending on the outcome of the vote. The key difference is in the timing of the realization of utility. In the electoral model, it occurs immediately, as the incumbent has the power to implement any feasible policy, but in the bargaining model, utility is accrued after the vote. In case the proposal is rejected, utility is derived from a status quo alternative (that may depend on the state) that does not have an obvious analogue in the electoral model. The availability of the status quo alternative in the bargaining model confers a technical benefit in that the proposer can always push through at least one alternative—if the status quo is proposed, then it prevails whether it passes or not—so the proposer's optimization problem is well-defined. In the electoral model, we address this by giving the incumbent the option of stepping down from office after choosing policy, so that, in effect, the continuation value of a challenger plays the role of the status quo. In contrast to the status quo, however, the continuation value of a challenger is endogenous, and thus additional arguments are needed to extend the approach to dynamic elections.

The remainder of this note is organized as follows: Section 2 presents the model of elections, Section 3 defines our equilibrium concept precisely, and Section 4 contains our results on existence and continuity of Markov electoral equilibrium.

2 A Dynamic Model of Elections

Political environment The model takes as given a set N of voters and a countably infinite set M of politicians, and we assume these sets are disjoint. The set $M \cup N$ of political actors is partitioned into a finite set T of types, typically denoted τ (for a voter) or t (for a politician). Politician types are initially private information and are given by the measurable type profile $\omega\colon M \to T$, and we assume that the voters' common prior beliefs about ω are such that politician types are independently (but not necessarily identically) distributed. Elections take place in discrete time over an infinite horizon. Each period begins with a state and an office holder, and the state and the office holder's type are observed by voters and politicians. The office holder chooses a policy; a challenger is selected; an election is held; a new state is realized and the winner's type is observed; and the process repeats. A type t office holder in state s chooses a policy from the feasible set $Y_t(s)$. We assume that states belong to a countable set S, that policies lie in a compact metric space Y, and that each feasible set $Y_t(s)$ is a nonempty, closed (and therefore compact) subset of Y.

In addition to choosing policy, the office holder also chooses whether to run for reelection: rather than model this decision using a separate variable, it is convenient to use Y to represent choices of policy and the decision to run for reelection, and to use a copy of Y, denoted Z, to represent policy choices and the decision not to run. We maintain the convention that $Y \cap Z = \varnothing$, we assume a mapping $\xi\colon Y \cup Z \to Z$ so that for all $y \in Y$, $\xi(y) = z$ is the element of Z corresponding to y and for all $z \in Z$, $\xi(z) = z$,[1] and we let $Z_t(s) = \xi(Y_t(s))$ be the feasible policy choices for a type t candidate who chooses not to seek reelection in state s. Let $X = Y \cup Z$ represent the space of simultaneous policy choices and campaign decisions, and let $x \in X$ denote a generic choice of policy and campaign decision.

Challengers After the office holder chooses policy, a challenger is drawn at large from the pool of politicians that have never held office. The challenger's type is not observed by voters. We maximize generality by allowing challenger selection to depend on the incumbent's type, the previous state and policy choice, and the newly realized state. Rather than explicitly deriving the challenger distribution by identifying challengers by name and using the voters' common prior over ω, we take a reduced form approach: let $q_t(t'|s,x)$ denote the probability that challenger is type t' given that a type t incumbent chooses policy x in state s. We assume that the challenger distribution $q_t\colon T \times S \times X \to [0,1]$ is continuous,[2] and that it is independent of the incumbent's campaign decision, i.e., $q_t(t'|s,y) = q_t(t'|s,\xi(y))$ for all $y \in Y$.

Elections We model elections in a parsimonious way, relying implicitly on the restriction to type-symmetric voting strategies. If the incumbent seeks reelection, then voters simultaneously cast ballots for the incumbent or the challenger. An *electoral outcome* for a type t incumbent in state s is $e \in \{0,1\}$, where $e = 1$ indicates that the incumbent seeks reelection ($x \in Y$) and is reelected, while $e = 0$ indicates that either the incumbent does not seek reelection ($x \in Z$) or that she is defeated by the challenger. This is determined by a set $\mathcal{D}_t(s) \subseteq 2^T \backslash \{\varnothing\}$ of *decisive* coalitions of types: if the coalition of voter types τ who vote for the incumbent belongs to $\mathcal{D}_t(s)$, then the incumbent retains office ($e = 1$) in the following period.[3] We assume that $\mathcal{D}_t(s)$ is monotonic, i.e., $C \in \mathcal{D}_t(s)$ and $C \subseteq C'$ imply $C' \in \mathcal{D}_t(s)$. Otherwise, if $x \in Z$ or if the set of voter types voting for the incumbent is not decisive, then the challenger assumes office in the following period ($e = 0$). Our formulation of the electoral rule is quite general and admits a wide variety of special cases as applications: weighed majority rule, more stringent quota rules, complex electoral systems such as the US Electoral College or domestic politics within an autocratic regime.

[1] Technically, ξ restricted to Y is an isometric embedding. It suffices to set $Z = Y \times \{1\}$ and to specify that $\xi(y) = (y,1)$ for all $y \in Y$.

[2] We give S and T the discrete topology, so our continuity assumption means that $q_t(t'|s,x,s')$ is continuous in x for all $s,s' \in S$ and all $t' \in T$. Given any function q_t, we can specify the voters' prior and a randomized challenger selection rule, $\gamma_t\colon S \times X \to \Delta(M)$, that generates q_t.

[3] Assuming the electorate has a measurable structure, $(N, \mathcal{N}, \nu)$, with ν nonatomic, and assuming the voting rule is insensitive to measure zero sets of voters (see Banks et al. (2006)), our type-symmetric formulation of the voting rule is sufficient. This is also true if the electorate is finite and types are uniquely assigned to voters. In case the electorate is finite and two or more voters have the same type, however, we should define the voting rule to account for deviations that are not type-symmetric.

State transitions States are used to describe the political and/or economic environment in the current period. Given a type t office holder that chooses a policy x in state s and given a subsequent electoral outcome e, a new state s' is drawn with probability $p_t(s'|s, x, e)$: that is, states evolve according to a controlled Markov process. The new state s' is not initially observed. We assume that the transition probability $p_t: S \times S \times X \times \{0, 1\} \rightarrow [0, 1]$ is continuous and independent of the incumbent's campaign decision, i.e., $p_t(s'|s, y, 0) = p_t(s'|s, \xi(y), 0)$ for all $y \in Y$.

Histories A *complete finite public history* of length m is therefore a sequence

$$h^m \quad = \quad \{(s_\ell, j_\ell, t_\ell, x_\ell, e_\ell)\}_{\ell=1}^m \in (S \times M \times T \times X \times \{0, 1\})^m$$

of states, office holder names, types of office holders, policy choices, and electoral outcomes. A *partial finite public history* of length $m + 1$ is a complete finite public history of length m concatenated with a triple $(s_{m+1}, j_{m+1}, t_{m+1})$ representing the state and the office holder's name and type in period $m + 1$, prior to choice of a policy. An *infinite public history* is an infinite list $h \in (S \times M \times T \times X \times \{0, 1\})^\infty$.

State by state commitment We assume that if an office holder chooses a policy x in a state s, and if she is subsequently reelected, then she is committed to her policy choice if the state remains s in the following period. By implication, she remains committed in successive periods in which she is reelected and the state remains s. Formally, given any partial finite public history (h^m, s, j, t) such that $e_m = 1$, $j = j_m$, and $t = t_m$, the action set available to the office holder is $\{x_m\}$ if $s = s_m$ (she is bound to her previous choice) and is $X_t(s)$ if $s \neq s_m$ (she is free). An important remark is that as a special case we obtain the model in which commitment is probabilistic (including with probability 0, that is, a model without commitment), as in Duggan and Forand (2018).

Payoffs The stage utility of a type τ voter or out-of-office politician from policy x in state s is $u_\tau(s, x)$, while the stage utility of a type t office holder is $w_t(s, x)$. This formulation captures the standard special cases in the literature: for all s, all t, and all x,

- *office motivation:* $w_t(s, x) = 1$ and $u_t(s, x) = 0$
- *policy motivation:* $w_t(s, x) = u_t(s, x)$
- *mixed motivation:* $w_t(s, x) = u_t(s, x) + b_t$,

where $b_t > 0$ represents the benefits of holding office for type t. We assume for simplicity that running for office is costless, i.e., for all $x \in Y$, $u_t(s, x) = u_t(s, \xi(x))$ and $w_t(s, x) = w_t(s, \xi(x))$, and that $u_t: S \times X \rightarrow \Re$ and $w_t: S \times X \rightarrow \Re$ are bounded and continuous. Each voter and politician of type t discounts flows of payoffs by the factor $\delta_t < 1$. Thus, given the infinite public history $h = \{(s_\ell, j_\ell, t_\ell, x_\ell, e_\ell)\}_{\ell=1}^\infty$, the discounted payoffs of the type τ voter and politician j of type t are

$$\sum_{\ell=1}^{\infty} \delta_\tau^{\ell-1} u_\tau(s_\ell, x_\ell) \quad \text{and} \quad \sum_{\ell=1}^{\infty} \delta_t^{\ell-1} (I_j(j_\ell) w_t(s_\ell, x_\ell) + (1 - I_j(j_\ell)) u_t(s_\ell, x_\ell)),$$

respectively, where I_j is an indicator function taking value one if $j_\ell = j$ and zero otherwise.

Remark We can capture a number of interesting features by suitable specialization of the model. For example, in addition to the several standard formulations of payoffs described above, we can capture models of rent-seeking and political agency, and we can allow voters to have preferences over office holders' types. Using the dependence of the voting rule $\mathcal{D}_t(s)$ on the state, we can accommodate term limits for incumbent politicians. Using the dependence of state transitions on electoral outcomes, we can model incumbency effects, such as learning-by-doing or tenure-dependent opportunities for corruption. Finally, while it may seem that the model restricts attention policy choices generated by sequences of individual office holders, the model also admits competition between two long-lived parties by associating each party with a politician type, exploiting dependence of the challenger transition $q_t(t'|s,x)$ on the incumbent's type, and letting the parties' preferences depend on the current office holder's type.

3 Markov Electoral Strategies and Equilibria

Strategies A mixed behavioral strategy for politician j maps partial public histories (h_m, s, j, t) into probability distributions $\pi_j(\cdot|h_m, s, j, t)$ on policies that are feasible and respect binding commitments: (i) $\pi_j(\cdot|h_m, s, j, t)$ puts probability one on $X_t(s)$, and (ii) if $j = j_m$, $t = t_m$, $s = s_m$, and $e_m = 1$, then $\pi_j(\cdot|h_m, s, j, t)$ puts probability one on x_m. Note that the politician mixes only when transitioning from one state s to another $s' \neq s$; once the state has transitioned to s', the politician chooses the same policy for successive draws of s'. We restrict attention to stationary Markov strategies, in the sense that $\pi_j(\cdot|h_m, s, j, t)$ depends on past policies and states only through the commitment assumption (ii), and therefore we need only model the politician's mixing over policies at the initial transition to a state s. Thus, we can write simply $\pi_j(\cdot|s,t)$ for this mixture. We further restrict politicians to strategies that are type-symmetric, so henceforth we adopt the notational convention $\pi_t(\cdot|s)$ for the behavioral strategy of a type t politician, and we refer to π_t as a *Markov policy strategy,* and $\pi = (\pi_t)_t$ denotes a profile of such strategies.

We adopt a parsimonious view of voting strategies, letting $\rho(h_m, s, j, t, x)$ be the probability that politician j is reelected after after public history h_m, the realization of state s, being type t, and choosing policy $x \in Y$.[4] As with policy strategies, we need only consider mixed voting upon the initial transition to a state s and policy choice x: if $s = s_m$, $j = j_m$, $t = t_m$, and $x = x_m$, then $\rho(h_m, s, j, t, x) = e^m$. Also consistent with our formulation of policy strategies, we restrict attention to strategies that are stationary with respect to the state and policy choice of the preceding period and the incumbent's type: thus, we write simply $\rho(s,t,x)$ for the probability that a type t office holder is reelected following policy choice x in state s.[5] In contrast to policy strategies, however, we do *not* assume that the electorate is bound to previous reelection decisions. Although we focus attention on strategies for which

[4] If the politician chooses $x \in Z$, then the challenger automatically assumes office, and it is convenient to set $\rho(h_m, s, j, t, x) = 0$ for all $x \in Z$.

[5] We impose the standard restriction that $\rho\colon S \times T \times X \to [0,1]$ is measurable.

an incumbent reelected after choosing x in state s is again reelected after choosing x in state s, this is not a constraint imposed on voters: rather, by stationarity of the decision problem of the electorate, it will be consistent with the incentives of voters in equilibrium. When N is finite, the probability of reelection may be decomposed into mixed voting strategies of individual voters. When the electorate N is infinite, individual uncertainty generated by mixed voting strategies washes out due to the law of large numbers, in which case we may interpret reelection probabilities as the result of conditioning on a public randomization device. We refer to ρ as a *Markov voting strategy,* and to $\sigma = (\pi, \rho)$ as a *Markov electoral strategy profile.*

Continuation values Given a Markov electoral strategy profile σ, we can define continuation values for politicians and voters. The discounted expected utility of a type τ voter from electing a type t incumbent who chooses policy x in state s (and continuing to do so for successive realizations of s) satisfies: for all $x \in Y$,

$$\begin{aligned} V_\tau^I(s,t,x) &= p_t(s|s,x,1)[u_\tau(s,x) + \delta_\tau V_\tau^I(s,t,x)] \\ &\quad + \sum_{s' \neq s} p_t(s'|s,x,1) V_\tau^F(s',t), \end{aligned} \tag{1}$$

where $V_\tau^F(s,t)$ is the expected discounted utility to a type τ voter from a type t office holder who is free in state s, calculated before a policy is chosen. When an office holder chooses $x \in Z$ and thus not to stand for reelection, we have $V_\tau^I(s,t,x) = V_\tau^C(s,t,x)$, where $V_\tau^C(s,t,x)$ is the expected discounted utility of electing a challenger following the choice of x in state s by a type t incumbent and is defined by

$$V_\tau^C(s,t,x) = \sum_{t'} q_t(t'|s,x) \sum_{s'} p_t(s'|s,x,0) V_\tau^F(s',t'). \tag{2}$$

Finally, $V_\tau^F(s,t)$ is given by

$$\begin{aligned} V_\tau^F(s,t) &= \int_x \Big[u_\tau(s,x) + \delta_\tau[\rho(s,t,x) V_\tau^I(s,t,x) \\ &\quad + (1-\rho(s,t,x)) V_\tau^C(s,t,x)] \Big] \pi_t(dx|s). \end{aligned} \tag{3}$$

Intuitively, the expression for $V_\tau^I(s,t,x)$ reflects that if an incumbent is bound to policy x in state s and is reelected, then either s is realized again, in which case the politician is bound to x and is reelected, or a different state $s' \neq s$ is realized, in which case the politician is free in s'. The expression for $V_\tau^F(s,t)$ reflects that the office holder chooses a policy x according to the policy strategy $\pi_t(\cdot|s)$, and is either reelected or replaced by a challenger.

A type t office holder's expected discounted utility from choosing policy x in state s (and being bound to x if s is realized again), conditional on being re-elected (and continuing to be for successive realizations of s), is such that for all $x \in Y$,

$$\begin{aligned} W_t^I(s,x) &= w_t(s,x) \\ &+ \delta_t \Big[p_t(s|s,x,1) W_t^I(s,x) + \sum_{s' \neq s} p_t(s'|s,x,1) \int_{x'} [\rho(s',t,x') W_t^I(s',x') \\ &+ (1-\rho(s',t,x')) W_t^C(s',x')] \pi_t(dx'|s') \Big], \end{aligned} \tag{4}$$

where $W_t^C(s,x)$ is a type t office holder's expected discounted utility from choosing policy $x \in X$ in state s, conditional on being replaced by a challenger, and is such that for all $x \in X$,

$$W_t^C(s,x) \quad = \quad w_t(s,x) + \delta_t V_t^C(s,t,x).$$

By convention, for all $x \in Z$, let $W_t^I(s,x) = W_t^C(s,x)$. In words, the politician receives utility $w_t(s,x)$ from policy x in state s while holding office. If the office holder does not seek reelection, then a challenger takes office in the next period, and she receives the expected discounted utility of a challenger, $V_t^C(s,t,x)$. Otherwise, if the office holder is re-elected, then a new state s' is drawn, which may be equal to s or not. In the case $s' = s$, then the politician is bound to x, reelected, and receives her expected discounted utility $W_t^I(s,x)$, and in case $s' \neq s$, the politician is free and mixes over policies according to $\pi_t(\cdot|s')$, which may or may not lead to reelection in these states.

Reelection sets Given a Markov electoral strategy profile $\sigma = (\pi, \rho)$ and policy choice x in state s by a type t incumbent, the type τ voter must consider the expected discounted utility of retaining the incumbent and must decide between her and a challenger. We therefore define for all states s, all incumbent types t, and all voter types τ, the sets

$$\begin{aligned} P_\tau(s,t) &= \{x \in Y_t(s) : V_\tau^I(s,t,y) > V_\tau^C(s,t,y)\} \\ R_\tau(s,t) &= \{x \in Y_t(s) : V_\tau^I(s,t,y) \geqslant V_\tau^C(s,t,y)\} \end{aligned}$$

of policies that yield type τ voters an expected discounted utility strictly and weakly greater, respectively, than the expected discounted utility of a challenger. For all coalitions $C \subseteq T$, define

$$P_C(s,t) = \bigcap \{P_\tau(s,t) : \tau \in C\} \quad \text{and} \quad R_C(s,t) = \bigcap \{R_\tau(s,t) : \tau \in C\},$$

and let the *strict* and *weak re-election sets,* denoted

$$\begin{aligned} P(s,t) &= \bigcup \{P_C(s,t) : C \in \mathcal{D}_t(s)\} \\ R(s,t) &= \bigcup \{R_C(s,t) : C \in \mathcal{D}_t(s)\}, \end{aligned}$$

be the policies that yield the members of at least one decisive coalition of types an expected discounted utility strictly and weakly greater, respectively, than the continuation of an unknown challenger. Note that these definitions isolate subsets of Y, for we are only concerned here with the case in which the office holder seeks reelection.

Equilibrium concept A Markov electoral strategy profile σ is a *Markov electoral equilibrium* if policy strategies are optimal for all types of office holders and voting is consistent with incentives of voters in all states; formally, we require that (i) for all s and all t, $\pi_t(\cdot|s)$ puts probability one on solutions to

$$\max_{x \in X_t(s)} \rho(s,x,t) W_t^I(s,x) + (1 - \rho(s,x,t)) W_t^C(s,x),$$

and (ii) for all s, all t, and all x,

$$\rho(s,t,x) \quad = \quad \begin{cases} 1 & \text{if } x \in P(s,t) \\ 0 & \text{if } x \notin R(s,t), \end{cases}$$

where $\rho(s,t,x)$ is unrestricted if $x \in R(s,t)\backslash P(s,t)$. In this case, there is no decisive coalition in which all voter types strictly prefer the incumbent, and in some decisive coalition, all voter types weakly prefer the incumbent but there is some type that weakly prefers to elect a challenger and so is indifferent: then any distribution of electoral outcomes is consistent with voting incentives. Note that $R(s,t) \subseteq Y$ by construction, so in equilibrium we require that $\rho(s,t,x) = 0$ for all $x \in Z$.

4 Existence and Continuity of Equilibria

The next theorem provides a foundation for the model by establishing existence of equilibrium under the general conditions of our framework.

Theorem 1. *There is a Markov electoral equilibrium.*

Next, we establish upper hemicontinuity of equilibria. We parameterize the stage utility functions and state transition by the elements γ of a metric space Γ, as in $u_t(s,x,\gamma)$ and $p_t(s'|s,x,e,\gamma)$, and we assume u_t and p_t are jointly continuous in their arguments. In what follows, $w_{s,t}$ represents the expected discounted utility of a type t office holder evaluated at the first time s is realized during her term of office, where $w = (w_{s,t})_{s,t} \in \Re^{S\times T}$ is the vector of expected politician payoffs, and $v_{s,t,\tau}$ represents the expected discounted utility of a type τ voter from a type t office holder who is free in state s and before a policy is chosen, i.e., it corresponds to $V^F_\tau(s,x,t)$. Then $v = (v_{s,t,\tau})_{s,t,\tau} \in \Re^{S\times T\times T}$ is the vector of expected voter payoffs. We endow $\Re^{S\times T} \times \Re^{S\times T\times T}$ with the product topology. Define the correspondence $\mathcal{E}\colon \Gamma \rightrightarrows \Re^{S\times T} \times \Re^{S\times T\times T}$ so that $\mathcal{E}(\gamma)$ consists of vectors (w,v) such that in the model parameterized by γ, there exists a Markov electoral equilibrium $\sigma^* = (\pi^*,\rho^*)$ such that for all s and all t, we have

$$w_{s,t} \quad = \quad \int_x [\rho^*(s,t,x)W^I_t(s,x;\sigma^*) + (1-\rho^*(s,t,x))W^C_t(s,x;\sigma^*)]\pi^*_t(dx|s),$$

and for all s, all t, and all τ, we have $v_{s,t,\tau} = V^F_\tau(s,t;\sigma^*)$, where we now parameterize continuation values by the strategy profile generating them. The following result establishes upper hemi-continuity of the equilibrium payoff correspondence.

Theorem 2. *The correspondence $\mathcal{E}\colon \Gamma \rightrightarrows \Re^{S\times T} \times \Re^{S\times T\times T}$ has closed values and is upper hemicontinuous.*

To prove these results, we use a fixed point argument that takes place in a space consisting of the product of policy strategies, ex ante expected payoffs for voters, and ex ante expected payoffs for politicians.[6] Normalize utilities so that the images of $\frac{u_t}{1-\delta_t}$ and $\frac{w_t}{1-\delta_t}$ lie in $[0,\overline{u}]$ for all types. Recall that $v_{s,t,\tau}$ and $w_{s,t}$ denote expected discounted payoffs to voters and office holders, so we can assume $v_{s,t,\tau}, w_{s,t} \in [0,\overline{u}]$ for all s and all t. Let $\pi_{s,t}$ represents the mixture over policies played by an office holder who is free at s upon the

[6]This proof follows the lines of Duggan (2017), which establishes existence of equilibrium in a model of dynamic bargaining. See the latter paper for an informal discussion of the proof technique.

initial transition to that state, i.e., given an equilibrium σ, $\pi_{s,t}$ corresponds to $\pi_t(\cdot|s)$. Then $\pi = (\pi_{s,t}) \in \Delta(X)^{S\times T}$ is the vector mixing probabilities, where $\Delta(\cdot)$ denotes the space of Borel probability measures endowed with the weak* topology.

Define the nonempty, convex product space

$$\Theta \;\; = \;\; \left(\Delta(X)^{S\times T}\right) \times \left([0,\overline{u}]^{S\times T}\right) \times \left([0,\overline{u}]^{S\times T\times T}\right),$$

with elements $\theta = (\pi, w, v)$. As usual, we imbed $\Delta(X)$ in the vector space $\mathcal{M}$ of signed Borel measures with the weak* topology (as the topological dual of the space of bounded, continuous, real-valued functions on X), which is Hausdorff and locally convex. As is well-known, $\Delta(X)$ is compact in the weak* topology. Of course, we imbed $[0,\overline{u}]$ in the real line with the Euclidean topology. Then the product topology on $(\mathcal{M}^{S\times T}) \times (\Re^{S\times T}) \times (\Re^{S\times T\times T})$ makes it a locally convex, Hausdorff topological space, and Θ is a non-empty, compact, convex subset of this space. Finally, let $\Theta^+ = \Theta \times \Gamma$ be Θ augmented by the parameters of the model. Denote a generic element of Θ^+ by $\theta^+ = (\pi, w, v, \gamma)$.

We will define a correspondence $\mathcal{F}\colon \Theta^+ \rightrightarrows \Theta$ such that for all $\gamma \in \Gamma$, $\mathcal{F}(\cdot, \gamma)$ has a fixed point $\theta^* = (\pi^*, w^*, v^*) \in F(\theta^*, \gamma)$; each fixed point θ^* corresponds to a Markov electoral equilibrium in the model parameterized by γ; and conversely, each Markov electoral equilibrium corresponds to a fixed point; and the correspondence of fixed points has closed graph. Write $\mathcal{F}$ as a product correspondence $\mathcal{F} = \mathcal{P} \times \mathcal{W} \times \mathcal{V}$.

For the construction of the component correspondences, we must consider the induced expected discounted utilities of voters and politicians that will parallel the continuation values defined in the setup of the model. The induced expected discounted utility of a type τ voter from electing a type t incumbent who is bound to policy x in state s (and continuing to reelect the politician after choosing x in s), calculated before the next state s' is realized, satisfies: for all $x \in Y$,

$$\begin{aligned} V^I_\tau(s,t,x,\theta^+) \;\; &= \;\; p_t(s|s,x,1,\gamma)[u_\tau(s,x,\gamma) + \delta_\tau V^I_\tau(s,t,x,\theta^+)] \\ &\qquad + \sum_{s'\neq s} p_t(s'|s,x,1,\gamma) v_{s',t,\tau}, \end{aligned}$$

or equivalently,

$$V^I_\tau(s,t,x,\theta^+) \;\; = \;\; \frac{p_t(s|s,x,1,\gamma)u_\tau(s,x,\gamma) + \sum_{s'\neq s} p_t(s'|s,x,1,\gamma) v_{s',t,\tau}}{1 - p_t(s|s,x,1,\gamma)\delta_\tau},$$

and for all $x \in Z$, we adopt the convention that $V^I_\tau(s,t,x,\theta^+) = V^C_\tau(s,t,x,\theta^+)$. As before, the induced expected discounted utility of a type τ voter from electing a challenger given policy choice x by office holder type t is:

$$V^C_\tau(s,t,x,\theta^+) \;\; = \;\; \sum_{t'} q_t(t'|s,x,\gamma) \sum_{s'} p_t(s'|s,x,0,\gamma) v_{s',t',\tau}.$$

The induced expected discounted utility of a type t office holder from choosing x in state s and being reelected (and continuing to choose x in s and being reelected if $x \in Y$) satisfies:

for all $x \in Y$,

$$W_t^I(s,x,\theta^+) \;=\; w_t(s,x,\gamma) + \delta_t p_t(s|s,x,1,\gamma) W_t^I(s,x,\theta^+) + \delta_t \sum_{s' \neq s} p_t(s'|s,x,1,\gamma) w_{s',t}, \tag{5}$$

or equivalently,

$$W_t^I(s,x,\theta^+) \;=\; \frac{w_t(s,x,\gamma) + \delta_t \sum_{s' \neq s} p_t(s'|s,x,1,\gamma) w_{s',t}}{1 - \delta_t p_t(s|s,x,1,\gamma)},$$

and for all $x \in Z$,

$$W_t^I(s,x,\theta^+) \;=\; w_t(s,x,\gamma) + \delta_t V_t^C(s,t,x,\theta^+).$$

Note that all of the above induced payoffs are continuous in (x, θ^+).

To define $\mathcal{P}$, for all states s, all office holder types t, and voter types τ, let

$$\begin{aligned} R_\tau(s,t,\theta^+) &= \{y \in Y_t(s) \mid V_\tau^I(s,t,y,\theta^+) \geqslant V_\tau^C(s,t,y,\theta^+)\} \\ P_\tau(s,t,\theta^+) &= \{y \in Y_t(s) \mid V_\tau^I(s,t,y,\theta^+) > V_\tau^C(s,t,y,\theta^+)\} \end{aligned}$$

and for each coalition C, define the correspondences

$$\begin{aligned} P_C(s,t,\theta^+) &\equiv \bigcap \{P_\tau(s,t,\theta^+) : \tau \in C\} \\ R_C(s,t,\theta^+) &\equiv \bigcap \{R_\tau(s,t,\theta^+) : \tau \in C\}, \end{aligned}$$

and as well define the correspondences

$$\begin{aligned} R_t(s,\theta^+) &\equiv \bigcup \{R_C(s,t,\theta^+) : C \in \mathcal{D}_t(s)\} \\ P_t(s,\theta^+) &\equiv \bigcup \{P_C(s,t,\theta^+) : C \in \mathcal{D}_t(s)\}. \end{aligned}$$

Continuity of V_τ^I and V_τ^C implies that the correspondence R_t has closed graph (and, by compactness of $Y_t(s)$, is therefore upper hemicontinuous) and that for each s and t, $P_t(s, \cdot)$ has open graph in $\Theta^+ \times Y_t(s)$ with the relative topology on $Y_t(s)$ induced by Y.

Similarly, W_t^I is continuous, and the correspondence $P_t(s, \cdot)$ is lower hemi-continuous, since it has open graph. Then Aliprantis and Border's (2006) Lemma 17.29 implies that the extended real-valued function

$$\overline{W}_t(s,\theta^+) \;\equiv\; \sup\{W_t^I(s,y,\theta^+) \mid y \in P_t(s,\theta^+)\}$$

is lower semi-continuous. Note also that the maximized value of $W_t^I(s,z,\theta^+)$ over $z \in Z_t(s)$, denoted

$$\overline{Z}_t(s,\theta^+) \;=\; \max\{W_t^I(s,z,\theta^+) : z \in Z_t(s)\},$$

is well-defined by nonemptiness and compactness of $Z_t(s)$ and continuity of $W_\tau^I(s, \cdot, \theta^+)$; and that by the theorem of the maximum, this maximized value is continuous. Then, as the pointwise maximum of two lower semi-continuous functions, it follows that

$$f_t(s,\theta^+) \;\equiv\; \max\{\overline{W}_t(s,\theta^+), \overline{Z}_t(s,\theta^+)\}.$$

is lower semi-continuous. Now define

$$\hat{\mathcal{P}}_t(s,\theta^+) \quad = \quad \{x \in R_t(s,\theta^+) \cup Z_t(s) \mid W_t^I(s,x,\theta^+) \geqslant f_t(s,\theta^+)\}$$

to consist of any policy x such that her expected payoff meets or exceeds $f_t(s,\theta^+)$ if the office holder is reelected after choosing x in s, if the office holder steps down after choosing x in s. This set is non-empty (see Duggan (2011)). Furthermore, by continuity of $W_t^I(s,\cdot)$ and lower semi-continuity of f_t, $\hat{\mathcal{P}}_\tau(s,\cdot)$ has closed graph in $\Theta^+ \times X$. Define $\mathcal{P}\colon \Theta^+ \rightrightarrows \Delta(X)^{S\times T}$ by

$$\mathcal{P}(\theta^+) \quad = \quad \prod_{s,t} \Delta(\hat{\mathcal{P}}_t(s,\theta^+)).$$

By Aliprantis and Border's (2006) Theorem 17.13, this correspondence has non-empty, convex values and has closed graph.

To define $\mathcal{W}$, let $\mathrm{supp}(\pi_{s,t})$ denote the support of $\pi_{s,t}$, and note that the correspondence $\mathrm{supp}\colon \Delta(X) \rightrightarrows X$ is lower hemi-continuous (see Aliprantis and Border's (2006) Theorem 17.14). By Aliprantis and Border's (2006) Lemma 17.29, the mapping

$$g_t(s,\theta^+) \quad \equiv \quad \min\{W_t^I(s,x,\theta^+) \mid x \in \mathrm{supp}(\pi_{t,s})\}$$

is upper semi-continuous. Define the (possibly empty) set

$$\hat{\mathcal{W}}_t(s,\theta^+) \quad = \quad [f_t(s,\theta^+), g_t(s,\theta^+)].$$

For each state s, since $f_t(s,\cdot)$ is lower semi-continuous and $g(s,\cdot)$ is upper semi-continuous in θ^+, the correspondence $\hat{\mathcal{W}}_t(s,\cdot)$ has closed, in fact, compact graph in $\Theta^+ \times [0,\overline{u}]$. Since the projection mapping from $\mathrm{graph}(\hat{\mathcal{W}}_t(s,\cdot))$ to Θ^+ is continuous, the set

$$\hat{\Theta}_t(s) \quad = \quad \{\theta^+ \in \Theta^+ \mid f_t(s,\theta^+) \leqslant g_t(s,\theta^+)\}$$

is compact. To see that $\hat{\Theta}_t(s) \neq \varnothing$, choose any $\theta^+ = (\pi,w,v,\gamma)$ such that $\pi_{s,t}$ puts probability one on an outcome that maximizes $W_t^I(s,x,\theta^+)$ over $x \in X_t(s)$ for a type t office holder in model γ. By Lemma A1 of Duggan (2011), we can extend $\hat{\mathcal{W}}_t(s,\cdot)$ from $\hat{\Theta}_t(s)$ to a correspondence (still denoted $\hat{\mathcal{W}}_t(s,\cdot)$) on Θ^+ that has non-empty, convex values and has closed graph. Then define the correspondence $\mathcal{W}\colon \Theta^+ \rightrightarrows [0,\overline{u}]^{S\times T}$ by

$$\mathcal{W}(\theta^+) \quad = \quad \prod_{s,t} \hat{\mathcal{W}}_t(s,\theta^+),$$

which has non-empty, convex values and has closed graph.

To define $\mathcal{V}$, note that given state s and office holder type t, a type τ voter's expected discounted utility depends on the probability that the incumbent is reelected in future states, and these probabilities are not explicitly given in the argument θ^+. To back out these probabilities, we use the expected discounted utility of the office holder. We are concerned with the case in which the type t office holder chooses $y \in R_t(s,\theta^+)\backslash P_t(s,\theta^+)$, for then the equilibrium conditions on voting strategies impose no restrictions on the probability of reelection. Specifically, we use the observation that if $y \in \mathrm{supp}(\pi_{s,t})$, then the proposal should generate the payoff $w_{s,t}$ for the office holder, providing a restriction on voting strategies. Indeed, the probability, say $\hat{r}$, that the office holder is reelected must be such that for all $y \in \mathrm{supp}(\pi_{s,t})$,

$$w_{s,t} = \hat{r} W_t^I(s,y,\theta^+) + (1-\hat{r}) W_t^I(s,\xi(y),\theta^+),$$

so, assuming $W_t^I(s,y,\theta^+) > W_\tau^I(s,\xi(y),\theta^+)$, we must have

$$\hat{r} = \frac{w_{s,t} - W_t^I(s,\xi(y),\theta^+)}{W_t^I(s,y,\theta^+) - W_t^I(s,\xi(y),\theta^+)}.$$

More generally, for all y such that $W_t^I(s,y,\theta^+) \neq W_t^I(s,\xi(y),\theta^+)$, define

$$\hat{\rho}_t(s,y,\theta^+) = \max\left\{0, \min\left\{1, \frac{w_{s,t} - W_t^I(s,\xi(y),\theta^+)}{W_t^I(s,y,\theta^+) - W_t^I(s,\xi(y),\theta^+)}\right\}\right\},$$

which is continuous in (s,y,θ^+). This is not defined when $W_t^I(s,y,\theta^+) = W_t^I(s,\xi(y),\theta^+)$, in which case $\hat{r}$ is not pinned down uniquely.

Next, define the correspondence $\mathcal{R}_t \colon S \times X \times \Theta^+ \rightrightarrows [0,1]$ by

$$\mathcal{R}_t(s,x,\theta^+) = \begin{cases} \{\hat{\rho}_\tau(s,x,\theta^+)\} & \text{if } W_t^I(s,x,\theta^+) \neq W_t^I(s,\xi(x),\theta^+) \\ [0,1] & \text{else} \end{cases}$$

for $x \in Y$, and by $\mathcal{R}_t(s,x,\theta^+) = \{0\}$ for $x \in Z$. Note that $\mathcal{R}_t$ has non-empty, convex values. In particular, the office holder's reelection probability is pinned down if she chooses a policy in Z and decides not to run or she chooses a policy in $x \in Y$ such that the induced expected discounted utility from winning with x is different from that of losing, e.g., $W_t^I(s,x,\theta^+) \neq W_t^I(s,\xi(x),\theta^+)$. It is unrestricted if she chooses a policy $x \in Y$ such that she is indifferent between winning or losing following x, e.g., $W_t^I(s,x,\theta^+) = W_t^I(s,\xi(x),\theta^+)$. Moreover, $\mathcal{R}_t$ has closed graph because ρ_t and W_t^I are continuous (using the convention that $Y \cap Z = \varnothing$). Given s and θ^+, the correspondence $\mathcal{R}_t(s,\cdot,\theta^+)$ gives the reelection probabilities, as a function of the policy choice in s, that are consistent with the office holder's payoff $w_{s,t}$ in θ^+, but note that these reelection probabilities will not generally satisfy the conditions required in a Markov electoral equilibrium: it may be that $\hat{\rho}_\tau(s,x,\theta^+) < 1$ for some $x \in P_t(s,\theta^+)$, and it may be that $\hat{\rho}_\tau(s,x,\theta^+) > 0$ for some $x \in Y \backslash R_t(s,\theta^+)$. This discrepancy will be resolved after the fixed point argument. In any case, a voter's or politician's induced expected discounted utilities will be determined by the precise way that reelection probabilities depend on policies, i.e., by a selection from $\mathcal{R}_t(s,\cdot,\theta^+)$.

Define $\hat{\mathcal{V}}_t(s,\theta^+)$ to be the set of possible vectors of expected discounted voter utilities in state s from a free politician of type t induced by measurable selections from $\mathcal{R}_t(s,\cdot,\theta^+)$ as follows: given each measurable section $\hat{\rho}$ from $\mathcal{R}_t(s,\cdot,\theta^+)$, we specify that the vector $v' = (v'_{s,t,\tau})_\tau \in [0,\overline{u}]^T$ of induced expected discounted utilities defined by

$$\begin{aligned} v'_{s,t,\tau} = \int_x \Big[& \hat{\rho}(x)[u_\tau(s,x,\gamma) + \delta_\tau V_\tau^I(s,t,x,\theta^+)] \\ & +(1-\hat{\rho}(x))[u_\tau(s,x,\gamma) + \delta_\tau V_\tau^C(s,t,x,\theta^+)]\Big] \pi_{s,t}(dx), \end{aligned}$$

for $\tau \in T$, belongs to $\hat{\mathcal{V}}_t(s,\theta^+)$. Note that $\hat{\mathcal{V}}_t(s,\theta^+)$ is non-empty. Furthermore, since $\mathcal{R}_t(s,\cdot,\theta^+)$ is convex-valued, convexity of $\hat{\mathcal{V}}_t(s,\theta^+)$ follows. That $\hat{\mathcal{V}}_t(s,\cdot)$ has closed graph in $\Theta^+ \times [0,\overline{u}]^T$ follows from a version of Fatou's lemma in Lemma A2 of Duggan (2011). Indeed, to apply that result, let X (in the lemma) be the policy space X, let Y (in the lemma) be $([0,\overline{u}]^{S\times T}) \times ([0,\overline{u}]^{S\times T\times T}) \times \Gamma$, let $k = 1$, and let $\Phi = \mathcal{R}_t(s,\cdot)$. Note that the countable product of metric spaces is metrizable in the product topology (see Theorem 3.36 of Aliprantis and Border (2006)), so Y is metric. Let $f = (f_\tau)_\tau$ (in the lemma) be defined by

$$\begin{aligned} f_\tau(x,r,y) &= r[u_\tau(s,x,\gamma) + \delta_\tau V_\tau^I(s,t,x,\theta^+)] \\ &\quad +(1-r)[u_\tau(s,x,\gamma) + \delta_\tau V_\tau^C(s,t,x,\theta^+)] \end{aligned}$$

for all $x \in X$, all $y = (w,v,\gamma) \in Y$, and all $r \in [0,1]$.[7] Let the correspondence F consist of integrals of f with respect to $\mu = \pi_{s,t}$, i.e.,

$$F(y,\mu) = \left\{ \int f_\tau(x,\hat{\rho}(x),y)\pi_{s,t}(dx) \mid \begin{array}{c} \hat{\rho} \text{ is a Borel mble selection} \\ \text{from } \mathcal{R}_t(s,\cdot,\theta^+) \end{array} \right\},$$

so that $\hat{\mathcal{V}}_t(s,\theta^+) = F(y,\mu)$. Then closed graph of $\hat{\mathcal{V}}_t(s,\cdot)$ follows from Lemma A2 of the above-mentioned paper. Finally, define $\mathcal{V}\colon \Theta^+ \rightrightarrows [0,\overline{u}]^{S\times T\times T}$ by

$$\mathcal{V}(\theta^+) = \prod_{s,t} \hat{\mathcal{V}}_t(s,\theta^+),$$

which, following the above argument, has non-empty, convex values and has closed graph.

These components together define $\mathcal{F} = \mathcal{P} \times \mathcal{W} \times \mathcal{V}$, a correspondence with non-empty, convex values and closed graph. By Glicksberg's theorem, for each $\gamma \in \Gamma$, $\mathcal{F}(\cdot,\gamma)$ has a fixed point θ^*. Furthermore, the correspondence from parameters γ to the set of fixed points of $\mathcal{F}(\cdot,\gamma)$ has closed (in fact, compact) graph. The next lemma establishes a close relationship between the fixed points of $\mathcal{F}(\cdot,\gamma)$ and the Markov electoral equilibria of the model parameterized by γ: in fact, $\mathcal{E}(\gamma)$ is just the projection of the fixed points of $\mathcal{F}(\cdot,\gamma)$ onto $[0,\overline{u}]^{S\times T} \times [0,\overline{u}]^{S\times T\times T}$. This immediately delivers existence of equilibria and non-empty values of the correspondence $\mathcal{E}$. Closed graph follows as well, because the projection of a compact set is compact. And since $\mathcal{E}$ has compact range, closed graph implies upper hemicontinuity, as required.

Lemma 1. *For all (w,v,γ), there exists π such that (π,w,v) is a fixed point of $\mathcal{F}(\cdot,\gamma)$ if and only if there is a Markov electoral equilibrium $\sigma^* = (\pi^*,\rho^*)$ of the game parameterized by γ such that for all s and all t,*

$$w_{s,t} = \int_x [\rho^*(s,t,x)W_t^I(s,x;\sigma^*) + (1-\rho^*(s,t,x))W_t^C(s,x;\sigma^*)]\pi_t^*(dx|s),$$

and for all s, all t, and all τ, $v_{s,t,\tau} = V_\tau^F(s,t;\sigma^*)$.

[7]Note that the definition of $f_\tau(x,r,y)$ makes use of the induced expected utility $\hat{V}_\tau(s,x,\theta^+)$, which formally depends on θ^+, but this dependence is through $y = (w,v,\gamma)$ only; it does not depend on policy strategies π.

Let (w, v, γ) be given. We first prove the "only if" direction, and to this end we consider π such that $(\pi, w, v) \in \mathcal{F}(\pi, w, v, \gamma)$. For all s and all t, we have $\pi_{s,t} \in \Delta(\hat{\mathcal{P}}_t(s, \theta^+))$, so that $\mathrm{supp}(\pi_{s,t}) \subseteq \hat{P}_t(s, \theta^+)$, and therefore $f_t(s, \theta^+) \leqslant g_t(s, \theta^+)$. It follows that $w_{s,t} \in \hat{\mathcal{W}}_t(s, \theta^+) = [f_t(s, \theta^+), g_t(s, \theta^+)]$. In particular, this implies that for all $x \in \mathrm{supp}(\pi_{s,t})$, we have $W_t^I(s, x, \theta^+) \geqslant w_{s,t} \geqslant f_t(s, \theta^+)$. Let $\hat{\rho}_t(s, \cdot, \theta^+)$ be the selection of reelection probabilities such that for all s, all t, and all τ,

$$v'_{s,t,\tau} = \int_x [\hat{\rho}_t(s, x, \theta^+)\hat{V}_\tau(s, t, x, \theta^+) + (1 - \hat{\rho}_t(s, x, \theta^+))\hat{V}_\tau(s, x, t, \theta^+)]\pi_{s,t}(dx).$$

We claim that every proposal x in the support of $\pi_{s,t}$ yields the induced expected payoff $w_{s,t}$ to a type t office holder in state s:

$$w_{s,t} = \hat{\rho}_t(s, x, \theta^+)W_t^I(s, x, \theta^+) + (1 - \hat{\rho}_t(s, x, \theta^+))W_t^I(s, \xi(x), \theta^+). \tag{6}$$

Indeed, consider $x \in \mathrm{supp}(\pi_{s,t})$. If $x \in Y$ and $W_t^I(s, x, \theta^+) \neq W_t^I(s, \xi(x), \theta^+)$, then the claim is true by construction of the correspondence $\mathcal{R}_t(s, \cdot, \theta^+)$ and the fact that $\hat{\rho}_t(s, \cdot, \theta^+)$ selects from it. If $x \in Y$ and $W_t^I(s, x, \theta^+) = W_t^I(s, \xi(x), \theta^+)$, then the claim holds regardless of the specification $\hat{\rho}_t(s, x, \theta^+)$ of the politician's reelection probability. And if $x \in Z$, then the right-hand side of (6) reduces to $W_t^I(s, x, \theta^+)$. We have noted that $W_t^I(s, x, \theta^+) \geqslant w_{s,t} \geqslant f_t(s, \theta^+)$, and furthermore, $f_t(s, \theta^+) \geqslant \overline{Z}_t(s, \theta^+) \geqslant W_t^I(s, x, \theta^+)$. Combining these two inequalities, we have $w_{s,t} = W_t^I(s, x, \theta^+)$, as claimed.

To construct a Markov electoral equilibrium, we first take state s and office holder type t as given, and we define the voting strategy $\rho^*(s, t, \cdot)$ as a function of policy by modifying the selections $\hat{\rho}_t(s, \cdot, \theta^+)$ in two ways: we require that an office holder is reelected with probability one after choosing $x \in P_t(s, \theta^+)$, and we require that the office holder is reelected with probability zero after choosing $x \in Y \backslash R_t(s, \theta^+)$. We then define policy strategies $\pi_t^*(\cdot|s)$ using $\pi_{s,t}$, with care to resolve possible inconsistencies created by the former modification of $\hat{\rho}_t(s, \cdot, \theta^+)$, completing the specification of the Markov strategy profile $\sigma^* = (\pi^*, \rho^*)$.

Case 1: Policy choice x belongs to $P_t(s, \theta^+)$. We specify that $\rho^*(s, t, x) = 1$. Note that it is possible that the selection $\hat{\rho}_t(s, \cdot, \theta^+)$ specifies that the office holder is reelected with probability less than one, i.e., $\hat{\rho}_t(s, x, \theta^+) < 1$. The modification could potentially create an inconsistency in the calculation of continuation values if $\pi_{s,t}$ puts positive probability on such policies, but the latter can occur only under special conditions. Since we consider $x \in P_t(s, \theta^+)$, we have $f_t(s, \theta^+) \geqslant \overline{W}_t(s, \theta^+) \geqslant W_t^I(s, x, \theta^+)$. But if $x \in \mathrm{supp}(\pi_{s,t})$, then we have noted that $W_t^I(s, x, \theta^+) \geqslant w_{s,t} \geqslant f_t(s, \theta^+)$. Combining these inequalities, we have $W_t^I(s, x, \theta^+) = w_{s,t}$. Thus, $W_t^I(s, x, \theta^+) > W_t^I(s, \xi(x), \theta^+)$ would imply $\hat{\rho}_t(s, x, \theta^+) = 1$ by definition of $\mathcal{R}_t(s, x, \theta^+)$. We conclude that $\hat{\rho}_t(s, t, x) < 1$ is only possible if $W_t^I(s, x, \theta^+) \leqslant W_t^I(s, \xi(x), \theta^+)$, and since $x \in \mathrm{supp}(\pi_{s,t})$, we also have

$$W_t^I(s, x, \theta^+) \geqslant f_t(s, \theta^+) \geqslant \overline{Z}_t(s, \theta^+) \geqslant W_t^I(s, \xi(x), \theta^+).$$

Combining these inequalities, we see that the problem described above can only arise if $W_t^I(s, x, \theta^+) = W_t^I(s, \xi(x), \theta^+)$, i.e., the office holder is indifferent between being reelected and stepping down from office after choosing x. When we define equilibrium policy choice strategies, below, we correct the inconsistency highlighted here by specifying that with probability $1 - \hat{\rho}_t(s, x, \theta^+)$, the office holder choose $\xi(x)$ instead of x.

Case 2: The policy choice belongs to $R_t(s,\theta^+)\backslash P_t(s,\theta^+)$. We specify that $\rho^*(s,t,x) = \hat{\rho}_t(s,x,\theta^+)$.

Case 3: The policy choice belongs to $X\backslash R_t(s,\theta^+)$. We specify $\rho^*(s,t,x) = 0$. It is possible that $\hat{\rho}_t(s,x,\theta^+) > 0$ for $x \in Y\backslash R_t(s,\theta^+)$, but since $\text{supp}(\pi_{s,t}) \subseteq \hat{\mathcal{P}}_t(s,\theta^+) \subseteq R_t(s,\theta^+) \cup Z_t(s)$, we have $\pi_{s,t}(Y\backslash R_t(s,\theta^+)) = 0$, so policies outside $R_t(s,\theta^+)$ are never chosen if the office holder seeks reelection. Thus, the modification here does not affect continuation values in this case and is immaterial.

To define policy choice strategies, consider any state s and office holder of type t. We specify that the politician mixes according to $\pi_{s,t}$, modified to correct the discrepancy in Case 1 above. For x in the support of $\pi_{s,t}$ with $x \in P_t(s,\theta^+)$, so that $W_t^I(s,x,\theta^+) = W_t^I(s,\xi(x),\theta^+)$, we require that the politician choose $\xi(x)$ with probability $1-\hat{\rho}_t(s,x,\theta^+)$, and otherwise the politician chooses according to $\pi_{s,t}$. Formally, define $\pi_t^*(\cdot|s)$ so that for all Borel measurable $A \subseteq X$,

$$\pi_t^*(A|s) \;=\; \pi_{s,t}(A\backslash P_t(s,\theta^+)) + \int_{A\cap P_t(s,\theta^+)} \hat{\rho}_t(s,x,\theta^+)\pi_{s,t}(dx)$$

and

$$\pi_t^*(\xi(A)|s) \;=\; \pi_{s,t}(\xi(A)) + \int_{A\cap P_t(s,\theta^+)} (1-\hat{\rho}_t(s,x,\theta^+))\pi_{s,t}(dx).$$

This maintains the continuation values generated from the fixed point, and in particular, we have

$$\begin{aligned} v_{s,t,\tau} \;&=\; \int_x [\rho^*(s,t,x)\hat{V}_\tau(s,t,x,\theta^*) \\ &\qquad + (1-\rho^*(s,t,x))\hat{V}_\tau(s,t,\xi(x),\theta^*)]\pi_t^*(dx|s) \end{aligned} \tag{7}$$

$$w_{s,t} \;=\; \rho^*(s,t,x)W_t^I(s,x,\theta^+) + (1-\rho^*(s,t,x))W_t^I(s,\xi(x),\theta^+) \tag{8}$$

for all s, all t, all τ, and all $x \in \text{supp}(\pi_t^*(\cdot|s))$.

By construction, and using the expression in (7), the values $V_\tau^I(\cdot,\theta^+)$, $V_\tau^C(\cdot,\theta^+)$, and $\{v_{s,t,\tau}\}_{s,t}$ fulfill the recursive conditions (1)–(3), that uniquely define $V_\tau^I(\cdot;\sigma^*)$, $V_\tau^C(\cdot;\sigma^*)$, and $V_\tau^F(\cdot;\sigma^*)$ in the model parameterized by γ. Furthermore, substituting (8) into (5), the values $W_t^I(\cdot,\theta^+)$ fulfill the recursive condition (4) that uniquely defines $W_t^I(\cdot;\sigma^*)$. Therefore,

$$V_\tau^I(\cdot,\theta^+) = V_\tau^I(\cdot;\sigma^*),\; V_\tau^C(\cdot,\theta^+),\; v_{s,t,\tau} = V_\tau^F(s,t),\; W_t^I(\cdot,\theta^+) = W_t^I(\cdot;\sigma^*)$$

for all s, all t, and all τ. As required for the lemma, we then have for all s and all t,

$$w_{s,t} \;=\; \int_x [\rho^*(s,t,x)W_t^I(s,x;\sigma^*) + (1-\rho^*(s,t,x))W_t^C(s,t,x;\sigma^*)]\pi_t^*(dx|s),$$

and for all s, all t, and all τ, $v_{s,t,\tau} = V_\tau^F(s,t;\sigma^*)$.

Next, we argue that the Markov strategy profile $\sigma^* = (\pi^*, \rho^*)$ satisfies the conditions for equilibrium. Indeed, ρ^* clearly satisfies condition (ii) in the definition of Markov electoral equilibrium. To verify that $\pi_t^*(\cdot|s)$ fulfills condition (i), we must show that no proposal yields an expected discounted payoff greater than $w_{s,t}$: for all $x \in X$,

$$w_{s,t} \;\; \geqslant \;\; \rho(s,x,\theta^+)W_t^I(s,x;\sigma^*) + (1-\rho(s,x,\theta^+))W_t^I(s,\xi(x);\sigma^*).$$

Indeed, the latter inequality holds (in fact, with equality) for $x \in \mathrm{supp}(\pi_t^*(\cdot|s))$. For $x \in P_t(s,\theta^+)\backslash\mathrm{supp}(\pi_t^*(\cdot|s))$, we have $\rho^*(s,t,x) = 1$, and the inequality follows from

$$w_{s,t} \;\; \geqslant \;\; f_t(s,\theta^+) \;\; \geqslant \;\; \overline{W}_t(s,\theta^+) \;\; \geqslant \;\; W_t^I(s,x;\sigma^*).$$

For $x \in X\backslash[(P_t(s,\theta^+) \cup \mathrm{supp}(\pi_t^*(\cdot|s))]$, we have $\rho^*(s,t,x) = 0$, and the inequality follows from

$$w_{s,t} \;\; \geqslant \;\; f_t(s,\theta^+) \;\; \geqslant \;\; \overline{Z}_t(s,\theta^+) \;\; \geqslant \;\; W_t^I(s,\xi(x);\sigma^*),$$

as required.

For the "if" direction of hte lemma, consider a Markov electoral equilibrium σ^* of the game parameterized by γ satisfying conditions of Lemma 1, so that for all s and all t, we have

$$w_{s,t} \;\; = \;\; \int_x [\rho^*(s,t,x)W_t^I(s,x;\sigma^*) + (1-\rho^*(s,t,x))W_t^C(s,x;\sigma^*)]\pi_t^*(dx|s),$$

and for all s, all t, and all τ, we have $v_{s,t,\tau} = V_\tau^F(s,t)$. Note by optimality of policy choices, we have $w_{s,t} \geqslant W_t^I(s,z;\sigma^*)$ for all s, all t, and all $z \in Z_t(s)$. Define $\pi = (\pi_{s,t})_{s,t}$ by modifying π^* so that for all s and all t, an office holder of type t chooses $\xi(x) \in Z_t(s)$ whenever the original policy strategy dictates a choice of $x \in Y\backslash R(s,t;\sigma^*)$, i.e., we specify that

$$\begin{aligned} \pi_{s,t}(A) &= \pi_t^*(A \cap R(s,t;\sigma^*)|s) \\ \pi_{s,t}(\xi(A)) &= \pi_t^*(\xi(A)|s) + \pi_t^*(A\backslash R(s,t;\sigma^*)|s) \end{aligned}$$

for all Borel measurable $A \subseteq Y$.

To establish that $(\pi,w,v) \in \mathcal{F}(\pi,w,v,\gamma)$, define $\hat{\rho}_t \colon S \times X \to [0,1]$ as follows. Fix a state s. First, we specify that $\hat{\rho}_t(s,x) = \rho^*(s,t,x) = 0$ for all $x \in Z$. Second, for $x \in Y$ such that $W_t^I(s,x;\sigma^*) = W_t^I(s,\xi(x);\sigma)$, we specify that $\hat{\rho}_t(s,x) = \rho^*(s,t,x)$. Third, for $x \in Y$ such that $W_t^I(s,x;\sigma^*) > W_t^I(s,\xi(x);\sigma)$, we require: (i) if $w_{s,t} \geqslant W^I(s,x;\sigma^*)$, then $\hat{\rho}_t(s,x) = 1$, (ii) if $w_{s,t} = W_t^I(s,\xi(x);\sigma^*)$, then $\hat{\rho}_t(s,x) = 0$, and (iii) if $W_t^I(s,x;\sigma^*) > w_{s,t} > W_t^I(s,\xi(x);\sigma^*)$, then the politician's expected discounted utility is exactly $w_{s,t}$, i.e.,

$$w_{s,t} \;\; = \;\; \hat{\rho}_t(s,x)W_t^I(s,x;\sigma^*) + (1-\hat{\rho}_t(s,x))W_t^I(s,\xi(x);\sigma^*).$$

Fourth, for $x \in Y$ such that $W_t^I(s,x;\sigma^*) < W_t^I(s,\xi(x);\sigma)$, we specify that $\hat{\rho}_t(s,x) = 0$, completing the definition. Note that $\hat{\rho}_t(s,x) = \rho^*(s,t,x)$ for all $x \in \mathrm{supp}(\pi_t^*(\cdot|s))$, except perhaps on a set of $\pi_t^*(\cdot|s)$-measure zero. And with the above modification of π^*, the same equality holds for $\pi_{s,t}$-almost all x. Thus, letting $\theta^+ = (\pi,w,v,\gamma)$, we have

$$V_\tau^I(\cdot,\theta^+) = V_\tau^I(\cdot;\sigma^*),\; V_\tau^C(\cdot,\theta^+),\; v_{s,t,\tau} = V_\tau^F(s,t),\; W_t^I(\cdot,\theta^+) = W_t^I(\cdot;\sigma^*)$$

for all s, all t, and all τ. It follows that $\hat{\rho}_t$ is a selection from $\mathcal{R}_t(\cdot, \theta^*)$ for all t, which implies that $v \in \mathcal{V}(\theta^*)$. Furthermore, we have $\pi \in \mathcal{P}(\theta^+)$. And finally, we have $w \in \mathcal{W}(\theta^*)$. Therefore, (π, w, v) is a fixed point of $\mathcal{F}(\cdot, \gamma)$, completing the proof.

References

Aliprantis, C., Border, K., 2006. Infinite Dimensional Analysis: A Hitchhiker's Guide. Springer: New York, NY.

Banks, J., Duggan, J., Breton, M. L., 2006. Social choice and electoral competition in the general spatial model. Journal of Economic Theory 126, 194–234.

Duggan, J., 2011. Coalitional bargaining equilibria, unpublished paper.

Duggan, J., 2017. Existence of stationary bargaining equilibria. Games and Economic Behavior 102, 111–126.

Duggan, J., Forand, J. G., 2018. The power of commitment in dynamic elections, unpublished paper.

Duggan, J., Martinelli, C., 2017. The political economy of dynamic elections: Accountability, commitment, and responsiveness. Journal of Economic Literature 55 (3), 916–84.