

Markovian Elections

John Duggan Jean Guillaume Forand

October 9, 2013

Abstract

We establish existence and continuity properties of equilibria in a model of dynamic elections with a discrete (countable) state space and general policies and preferences. We provide conditions under which there is a representative voter in each state, and we give characterization results in terms of the equilibria of an associated “representative voting game.” When the conditions for these results are not met, we provide examples that uncover new classes of dynamic political failures.

Contents

1	Introduction	1
2	General Framework	5
3	Simple Markov Electoral Equilibria	10
4	General Properties of Equilibria	14
5	Representative Voters	21
6	Dynamic Core Convergence	29
A	Existence and Continuity Proof	43
B	Proof of Representative Voter Theorem	53
C	Proofs for Examples	55

1 Introduction

The development of dynamic models of elections is critical for our understanding of the interplay between politics and dynamic processes such as economic growth and cycles, the evolution of income inequality, and transitions to democracy. A common thread in these examples is the existence of a state that evolves over time and is, in principle, influenced by policy choices. The purpose of the present paper is to contribute to the analysis of elections in the presence of endogenous economic and political state variables, e.g., capital stock, the distribution of income, or the institutional rules governing the political system. Our goal is a framework that is general (amenable to a range of structure on preferences, policies, and states), viable (equilibria should exist widely and allow non-constructive characterizations), and practically useful (so we can solve special cases to generate novel insights). Accordingly, we define an appropriate selection of Markov perfect equilibria and prove existence and continuity in a general model with a discrete (countable) state space and a compact metric policy space. We provide sufficient conditions for the existence of a representative voter, and we give a dynamic form of the median voter theorem in which policy choices of politicians solve the dynamic programming problem of the representative voter; more generally, we allow the representative voter to be state-dependent, and we characterize equilibria in terms an associated “representative voting game.” The model is surprisingly rich, and we find that when the assumptions of our characterization results are violated, even in simple cases with two or three states, the model can reconcile seemingly paradoxical political outcomes.

The analysis of endogenous state variables has received considerable attention in the political economy literature, but it has been limited by the lack of a tractable framework. Early work assumed (implicitly) that candidates can commit to infinite sequences of taxes in the first period and invoked the median voter theorem to determine taxes prior to running the economy,¹ but because this work assumes *ex ante* commitment to sequences of policies, the political interaction is static. Klein et al. (2008), Krusell et al. (1997), and Krusell and Rios-Rull (1999) analyze endogenous taxation in a model of economic growth, where voting takes place in each period and policy is chosen by a representative voter. Battaglini and Coate (2007, 2008) consider a dynamic non-cooperative model in which the state variable is, respectively, a durable public good or public debt level, with a focus on incentives in the dynamic legislative bargaining game. Yared (2010) considers the optimal equilibrium for a representative voter in a model of surplus extraction where the government and the consumer can accumulate debt; again, politicians are homogeneous, and elections are not the focus. Although the economic environment evolves endogenously in this work, the political environment is fixed over time. Camara (2012) includes an extension to growth economies that preserves the stationary structure of his equilibrium in a model of repeated elections with adverse selection. Forand (2012) considers elec-

¹See Bertola (1993), Alesina and Rodrik (1994), and Bassetto and Benhabib (2006)).

tions in which an office holder’s policy choice determines her policy platform in the next period, and challengers balance policy gains today against losses in the future. In both papers, characteristics of the politician in office can change over time. Azzimonti (2011) and Battaglini (2011) study local public good provision in which a component of voter preferences varies—Azzimonti (2011) considers partisan candidates in the presence of incumbency advantage, while Battaglini (2011) assumes stochastic shocks on voters’ preferences—in addition to capital stock and debt levels. Closer to our work are recent papers that incorporate an economic state and broader variation in the political environment, including the evolution of institutional rules. Bai and Lagunoff (2011) allow current government policy to directly determine future public decision-makers. Acemoglu et al. (2012) analyze dynamic institutional choice in a finite framework and characterize Markov equilibrium outcomes in terms of a cooperative concept of stability. Duggan and Kalandrakis (2012) analyze a model of legislative bargaining in which policy in one period determines the status quo in the next, and recognition probabilities and the voting rule can themselves depend on the status quo.²

We consider a model of elections such that in each period, a state is given and an incumbent office holder chooses policy; then a challenger is drawn and an election is held; and then a new state is realized, and so on. In each period, the state determines preferences and the electoral rule, and the transition probabilities on the challenger’s type and next period’s state can depend on the current state and policy choice. States are discrete, policies lie in a compact metric space, and stage utilities and transition probabilities are assumed only to be continuous (no convexity properties are imposed). We assume that information is symmetric between voters and the office holder, so that the stage utility (or type) of a politician is observed by voters once the politician takes office; but a challenger’s type may not be observed. Thus, elections pit a known incumbent against a relatively unknown challenger. Within the literature on electoral accountability, our model is related to Barro (1973), which we obtain when policy is a public good level, and there is a single, fixed state.³ Our framework does not generalize most of the papers cited above in connection to endogenous state variables, because capital stock or debt are commonly formulated as continuous variables, but we can approximate continuous state-space models with a countable, dense set of states. Moreover, models of discrete states are often quite natural, as in the finite model of Acemoglu et al. (2012).

We prove existence and continuity of “simple Markov electoral equilibria,” and we provide sufficient conditions for existence of a representative voter in a given state: the key conditions are that stage utility is quadratic, and that the core (with respect to stage utilities) is nonempty in that state. When there is a representative voter in each state, we show that equilibria of an associated repre-

²See Duggan and Kalandrakis (2012) for a review of the related literature on dynamic bargaining with an endogenous status quo.

³See also Aragonés et al. (2007), in which policy is a one-dimensional ideological variable.

representative voting game can be replicated in the electoral model when politicians are sufficiently office-motivated. In case the representative voter is fixed across states, the result implies existence of an equilibrium that solves the dynamic programming problem of the representative voter, providing theoretical justification for the application of the median voter theorem in models of dynamic policy choice (e.g., Krusell et al. (1997) and Krusell and Rios-Rull (1999)).

We show that if there is a fixed representative voter, and if the state transition is independent of the office holder’s type, then there is always at least one politician type that can satisfy a majority of voters. However, if the transition depends on the politician’s type, then it is possible that at a given state, the game may transition to a new state in which the politician’s preferences diverge from a majority of the electorate. Because politicians cannot commit to policy in the new state, there may be equilibria such no politician has available a policy that can lead to re-election—even if the benefit of holding office is arbitrarily large—demonstrating a type of “dynamic political failure.” When the state transition is independent of policy choice as well as the office holder’s type, and politicians are sufficiently office-motivated, we show that all types of office holder implement the optimal policy rule of the voter in all equilibria of the electoral model. But if the state transition depends on policy choice, then there may be equilibria characterized by “dynamic policy extremism,” in which all politicians (including those whose type is the same as the representative voter) choose policies that are sub-optimal for the voter. Thus, our model presents the possibility of a new source of inefficiency that arises from the dynamic incentives of elections.⁴

A benefit of the discrete state space model is its comparative tractability: under otherwise quite general assumptions on policies and payoffs, we are able to prove existence of simple Markov electoral equilibria in our framework. In contrast, Azzimonti (2011), Bai and Lagunoff (2011), Krusell et al. (1997), Krusell and Rios-Rull (1999), and Klein et al. (2008), use a first order approach to analyze necessary conditions of equilibria and solve for equilibria in parameterized examples. The issue of existence of stationary Markovian equilibria in dynamic games is non-trivial, as examples of relatively well-behaved stochastic games that do not admit such equilibria are available (see Levy (2012a,b)). In models of elections, these issues can be exacerbated because of discontinuities introduced by the conditioning of voters’ strategies on policy choices. Unlike other general solutions to the existence problem that add noise to voter preferences (e.g., Duggan and Kalandrakis (2012), Duggan (2012)) to address these discontinuities, we do not require preference shocks or an atomless environmental variable.⁵ The cost of this parsimony is that the existence proof is involved

⁴This complements other explanations of inefficiency due to commitment problems (Besley and Coate (1998)), “tying the hands” of one’s successor (Persson and Svensson (1989), Alesina and Tabellini (1990), Aghion and Bolton (1990)), signaling competence (Rogoff and Sibert (1988)), interest groups (Coate and Morris (1995)), and pandering (Canes-Wrone et al. (2001), Maskin and Tirole (2004)).

⁵Our proof approach is adapted from Duggan (2011), which backs out voting strategies in

(voting strategies are problematic for reasons of compactness, so they are not included in the domain of the fixed point argument and must be backed out after the fact), but a benefit is that simple special cases of the model can be solved relatively easily.

To generate interesting and plausible policy choices, we assume a form of “state-by-state commitment,” whereby an incumbent politician is committed to her current policy for successive realizations of the current state. Although we do not model the commitment mechanism explicitly, it may proxy for the possibility that an office holder accumulates state-specific experience for a given policy, or that policy is implemented by making state-specific judicial or administrative appointments, or that the politician writes state-specific contracts with interest groups or other politicians. A more traditional approach, consistent with the classical Downsian model, would be to assume that in each period, the incumbent and a challenger can make binding policy commitments prior to an election. It is well-known that even in the static model this approach leads to equilibrium existence problems when the policy space is multidimensional, and these issues would be exacerbated (and would contaminate the one-dimensional model) in the dynamic setting. We view our commitment assumption as more plausible than the Downsian assumption, for in our framework a politician must choose a policy in a state (putting her “money where her mouth is”) to make a commitment. Another approach to generating a dynamic linkage between current and future policy choices would be to assume incomplete information in the form of adverse selection: then a policy choice in one period would affect voter beliefs about the office holder’s type, and this would present office holders with a trade off between short-run considerations of current policy and long-run considerations of re-election and future policy choice. Because a policy choice in one state has reputational consequences in all future states, however, existence of equilibrium is problematic in such models, and they are impractical to solve analytically.⁶

Despite technical differences in the form of linkage, however, our framework has interesting connections to a dynamic model of elections with adverse selection considered by Duggan (2000).⁷ In this model, politicians are privately informed about their preferences, but there is no state variable, and in equilibrium, after information is revealed by an office holder’s initial policy choice and (assuming she is re-elected), the politician’s policy choice is expected to remain the same over time. Because of this, the equilibria of the model with adverse selection are replicated in our model by specifying a single state: the

a bargaining model with a countable set of states.

⁶These issues are more problematic when adverse selection is combined with moral hazard. See Banks and Sundaram (1993) for the analysis of non-Markovian equilibria, and see Banks and Sundaram (1998) for the model with a two-period term limit.

⁷See Bernhardt et al. (2004) for the model with a two-period term limit. Further applications include the analysis of competence (Meirowitz (2007)), parties (Bernhardt et al. (2009)), valence (Bernhardt et al. (2011)), and taxation (Camara (2012)). Duggan (2013) provides a folk theorem for the model when non-Markovian equilibria are permitted.

assumption of state-by-state commitment acts as a substitute for voter beliefs in the adverse selection model. The parallel between the frameworks extends to the model with multidimensional policy space analyzed by Banks and Duggan (2008). But fundamental differences arise when we move beyond the degenerate model to allow multiple states and the more complex incentives they entail: whereas reputational issues blow up and render the adverse selection model unworkable, our model remains viable, and it is possible (as we demonstrate) to work with interesting special cases of the model.

The organization of the paper is as follows. In Sections 2 and 3, we specify the model and formulate our equilibrium concept. In Section 4, we establish existence and continuity of simple Markov electoral equilibria, and we consider the effect of removing commitment and the need for mixed voting strategies. In Section 5, we provide conditions under which some voter type is representative at a given state, and we illustrate the possibility of dynamic political failure. In Section 6, we isolate conditions under which simple Markov electoral equilibria simulate an associated representative voting game or, when one voter type is representative at all states, solve the voter’s dynamic programming problem. In addition, we show the possibility of dynamic policy extremism. Detailed analysis is presented in Appendices A–C.

2 General Framework

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 $\omega: M \rightarrow T$, 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 revealed to 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 revealed; and the process repeats. A type t office holder in state s chooses a policy y 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 . The dependence of the feasible set on the office holder’s type allows us to incorporate differences in competence (i.e., valence), and dependence on the state allows us to interpret s as a state of the economy, which can effect the range of available policies.

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 endow Z with the metric topology and maintain the convention that $Y \cap Z = \emptyset$; we assume a mapping $\xi: Y \cup Z \rightarrow Z$ so that for all $y \in Y$, $\xi(y) = z$ is the element of Z corresponding to y and for all $z \in Z$, $\xi(z) = z$; ⁸ 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. Conditional on the incumbent politician choosing to run, i.e., $x \in Y$, the electoral outcome is $e = 1$ if she wins the election; otherwise, if the politician does not seek election, i.e., $x \in Z$, or loses the election, then the electoral outcome is $e = 0$.

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 $q_t: T \times S \times X \rightarrow [0, 1]$ is continuous,⁹ 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. 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. When the incumbent seeks reelection, the electoral outcome is determined by a set $\mathcal{D}_t(s) \subseteq 2^T \setminus \{\emptyset\}$ 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.¹⁰ We assume that $\mathcal{D}_t(s)$ is nonempty and 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

⁸Technically, ξ restricted to Y is a homeomorphism.

⁹We give S and T the discrete topology, so our continuity assumption means that $q_t(t'|s, \cdot, s')$ is continuous 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, $\gamma_t: S \times X \rightarrow \Delta(M)$, that generates q_t .

¹⁰Assuming the electorate a measurable structure, (N, \mathcal{N}, ν) , 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.

decisive, then the challenger assumes office in the following period ($e = 0$).

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 $p_t(s'|s, x, e)$; thus, 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

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

of states, office holder names, types of office holders, policy choices, and electoral outcomes. Of course, an infinite public history is any infinite list $h \in (S \times M \times T \times X \times \{0, 1\})^\infty$. Note that there is some redundancy in the definition of a public history, as electoral outcomes can be inferred from names of office holders, and the type of a reelected politician is fixed over time.

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 . More precisely, if an office holder chooses x in state s , if $s' = s$ is realized, and if she is reelected, then the politician is *bound* to x . This commitment is binding until the state shifts to a different state $s' \neq s$, at which point the politician is *free*. The politician is also free upon any initial recurrence of the state s . Formally, given any complete finite history h^m with $e_m = 1$ and triple (s, j, t) such that $j = j_m$ and $t = t_m$, the action set available to office holder j in state s 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).

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 $w_t(s, x)$. With this formulation, we capture the standard special cases in the literature: for all s , all t , and all x , and for all t' with $q_t(t'|s, x) > 0$,

- *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*: $w_{t'}(s, x) = u_{t'}(s, x) + b$,

where $b > 0$ represents the benefits of holding office. We assume that running for office is costless, i.e., for all $y \in Y$, $u_t(s, y) = u_t(s, \xi(y))$ and $w_t(s, y) =$

$w_t(s, \xi(y))$, 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 payoff of a type τ voter or type t politician is

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

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

Timing and information To summarize the timing of moves and flow of information, suppose a politician of type t holds office at the beginning of a period in state s . Then political interaction proceeds as follows:

- the state s and the office holder's type t are revealed to all political actors,
- the office holder selects a policy $x \in X_t(s)$,
- a challenger's type is drawn from $q_t(\cdot | s, x)$ and is not observed by voters,
- if $x \in Y$, then an election is held and electoral outcome $e \in \{0, 1\}$ determined; otherwise, if $x \in Z$, then $e = 0$,
- if $e = 1$, then the incumbent is reelected; and if $e = 0$, then the challenger takes office,
- a new state s' is drawn from $p_t(\cdot | s, x, e)$; if $e = 1$ and $s' = s$, then the office holder (the incumbent) is bound to x , and otherwise, if $s' \neq s$, then the office holder is free,
- the new state s' and current office holder's type t' are revealed, and the process repeats.

Remarks

1. The model is isomorphic to that of Banks and Duggan (2001, 2008) when: (i) the state space is trivial, i.e., $|S| = 1$, (ii) Y_t is a compact, convex subset of Euclidean space independent of t , (iii) voter utilities are concave in policy, (iv) politicians have mixed motives, (v) the distribution of challenger types is independent of the incumbent's type, her policy choice, and her campaign decision. The models of Duggan (2000) and later papers in the literature are technically outside the framework because they assume a continuum of types, but they can be approximated to an arbitrary degree.

2. Our formulation of the electoral rule, $\{\mathcal{D}_t(s)\}_{s,t}$, is quite general. We obtain weighed majority rule, as well as more stringent quota rules, as special cases; and more generally we can incorporate complex electoral systems such as

the US Electoral College, in which a candidate is elected if she obtains a majority of voters in a majority of states. The electoral rule can also represent complex democratic election procedures, but it can reflect less formal, non-democratic politics as well. Under this interpretation, the type t office holder could be an authoritarian ruler, and in state s the ruler requires the support of at least one decisive coalition in $\mathcal{D}_t(s)$. It may be, for example, that the support of a fixed oligarchy, $C = \bigcap \mathcal{D}_t(s)$, is necessary and sufficient to remain in power.

3. We have shown that several standard formulations of payoffs are obtained as special cases of the model. In addition to the political motivations considered above, we can capture common models of rent-seeking and political agency in which, respectively, the office holder claims residual surplus or determines a level of public good through effort choice. In general, we can assume $Y \subseteq \mathbb{R}^d$, that $u_\tau(s, x)$ is independent of s and weakly increasing in the coordinates of $x = (x_1, \dots, x_d)$, and that $w_t(s, t)$ is decreasing in the coordinates of x . To capture rent-seeking models, we can let $d = |T|$ and interpret x as a vector of private goods allocated to voters of different types, so that $u_\tau(x) = x_\tau$ for all τ . There is an endowment $c_t(s) \geq 0$ of private good that may depend on the state and office holder's type, and we specify $w_t(s, x) = c_t(s) - \sum_\tau x_\tau$. In this case, the total allocation must satisfy the budget constraint, so that $Y_t(s) = \{x \in \mathbb{R}_+^{|T|} \mid \sum_\tau x_\tau = c_t(s)\}$. To capture the political agency model, assume $Y \subseteq \mathbb{R}^2$ and interpret $x = (e, g)$ as an effort choice and public good level, so that $u_\tau(x) = g$ and $w_t(s, x) = -e$. Given state s and effort choice e by a type t office holder, the output of public good is $f_t(s, e)$, where f_t is continuous and increasing in e . Finally, the choice of public good must respect production technology, so that $Y_t(s) = \{(e, g) \in Y \mid g = f_t(s, e)\}$. In both examples, the ability of an office holder (either the budget available or public good production technology) can depend on the politician's type.

4. An apparently more general model would allow the stage utility of voters to depend on the office holder's type; we would want to allow for such dependence if, for example, types include information about some characteristic of politicians that voters care about beyond policy outcomes. That added generality is redundant, as we can obtain it by suitable specification of the model. Given the model in which the utility of a type τ voter is $u_\tau(s, t, x)$ when a type t office holder chooses x in state s , we can map this into our framework by imbedding office holder types into policies, as in $\tilde{Y} = Y \times T$, setting feasible policies as $\tilde{Y}_t(s) = Y_t(s) \times \{t\}$, and specifying $\tilde{u}_\tau(s, (x, t)) = u_\tau(s, t, x)$. In short, feasible policies can be used to identify the office holder's type and to incorporate preferences over types. In fact, this observation underlies the previous remark, where we incorporate the ability of a type t office holder into the set of feasible policies.

5. The model is general enough to encompass electoral competition between infinitely lived parties. Let $\{A, B\} \subseteq T$, and let M consist of a countably infinite number of type A politicians and a countably infinite number of type B

politicians, so there are essentially two parties, A and B . We specify that, for all states s and policies x , $q_A(B|s, x) = q_B(A|s, x) = 1$, so that the challenger is always the party out of power. When $w_A(s, x) = u_A(s, x)$, we have the case of policy-motivated parties, whereas $w_A(s, x) = 1$ and $u_A(s, x) = 0$ corresponds to the case of purely office-motivated parties. By the above remark, which shows that the stage utility of an out of office politician can depend on the current office holder's type, we capture long-lived parties with mixed motivation as well: letting $\tilde{Y} = Y \times \{A, B\}$ and $\tilde{Y}_t(s) = Y_t(s) \times \{t\}$, we specify that for $\tau \in \{A, B\}$, $\tilde{u}_\tau(s, (x, t)) = u_\tau(s, t, x) + b$ if $\tau = t$ and $\tilde{u}_\tau(s, (x, t)) = u_\tau(s, t, x)$ otherwise. Hence, while we assume that a replaced incumbent politician is not returned to the candidate pool, our assumptions allow parties to internalize the office benefits they collect in their future terms in office.

6. We can obtain a more flexible form of commitment in which the politician is committed with some probability $\gamma_t(s, x)$ (again continuous), by suitable specification of the model. Given the model with states S and commitment probability γ_t , we can map this into our framework by doubling the states, i.e., $\tilde{S} = (S \times \{1\}) \cup (S \times \{-1\})$, and defining a transition probability $\tilde{p}_t((s', k)|(s, k), x) = p_t(s'|s)\gamma_t(s, x)$ and $\tilde{p}_t((s', -k)|(s, k), x) = p_t(s'|s)(1 - \gamma_t(s, x))$. Here, the first component of $(s, k) \in \tilde{S}$ records the payoff relevant information in the state, and the second merely indicates whether the office holder is bound to x following a transition to the same payoff relevant state.

3 Simple Markov Electoral Equilibria

Strategies A mixed behavioral strategy for politician j maps public histories h_m , states s , and types 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 Markovian 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 $\pi_j(\cdot|s, t)$. 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 *simple Markov policy strategy*, and $\pi = (\pi_t)_t$ denotes a profile of such strategies.

We adopt a parsimonious view of voting strategies, letting $\rho(h_m, s, j, t, y)$ 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 $y \in Y$.¹¹ As with policy strategies, we need only consider mixed voting upon the initial transition to a state s and policy choice y : if $s = s_m$, $j = j_m$, $t = t_m$, and $y = x_m$, then $\rho(h_m, s, j, t, y) = e^m$. 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 $\rho(s, t, y)$ for the probability that a type t office holder is reelected following policy choice y in state s .¹² 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 an incumbent reelected after choosing y in state s is again reelected after choosing y 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 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 are agnostic as to interpretation. We refer to ρ as a *simple Markov voting strategy*, and to $\sigma = (\pi, \rho)$ as a *simple Markov strategy profile*.

Continuation values Given a simple Markov 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 $x \in Y$,

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

or equivalently,

$$V_\tau^B(s, t, x) = \frac{p_t(s|s, x, 1)u_\tau(s, x) + \sum_{s' \neq s} p_t(s'|s, x)V_\tau^F(s', t)}{1 - \delta_\tau p_t(s|s, x, 1)}, \quad (1)$$

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. We adopt the convention that for all $x \in Z$, $V_\tau^B(s, t, x) = V_\tau^C(s, t, x)$, 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

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

¹¹If the politician chooses $z \in Z$, then the challenger automatically assumes office, and it is convenient to set $\rho(h_m, s, j, t, z) = 0$ for all $z \in Z$.

¹²We impose the standard restriction that $\rho: S \times T \times X \rightarrow [0, 1]$ is measurable.

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

$$V_\tau^F(s, t) = \int_x \left[u_\tau(s, x) + \delta_\tau [\rho(s, t, x) V_\tau^B(s, t, x) + (1 - \rho(s, t, x)) V_\tau^C(s, t, x)] \right] \pi_t(dx|s).$$

Intuitively, the expression for $V_\tau^B(s, t, x)$ 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. Obviously, the expression for $V_\tau^C(s, t, x)$ reflects that a newly elected challenger is free regardless of the state realized.

For future reference, we note that the expected discounted utility from electing an office holder who is free in s can be written as

$$V_\tau^F(s, t) = \int_x \left[\rho(s, t, x) [u_\tau(s, x) + \delta_\tau V_\tau^B(s, t, x)] + (1 - \rho(s, t, x)) [u_\tau(s, x) + \delta_\tau V_\tau^B(s, t, \xi(x))] \right] \pi_t(dx|s), \quad (3)$$

which follows from the convention that $V_\tau^B(s, t, \cdot) = V_\tau^C(s, t, \cdot)$ on Z .

Finally, 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 $x \in Y$,

$$W_t^B(s, x) = w_t(s, x) + \delta_t \left[p_t(s|s, x, 1) W_t^B(s, x) + \sum_{s' \neq s} p_t(s'|s, x, 1) \int_{x'} [\rho(s', t, x') W_t^B(s', x') + (1 - \rho(s', t, x')) W_t^C(s', x')] \pi_t(dx'|s') \right], \quad (4)$$

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 $x \in Y$,

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

By convention, for all $x \in Z$, let $W_t^B(s, x) = W_t^C(s, x)$. 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 $W_t^B(s, x)$; 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 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

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

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.

In fact, because we use the reduced form representation ρ of voter behavior, it is not immediately obvious how to formulate the expected discounted utility of a voter appropriately. We rely on intuition from the finite N case to motivate the above approach. We want to capture the idea that voters do not use weakly dominated strategies, and so the relevant calculation is that of a voter, say τ , conditional on her vote being pivotal given mixed voting strategies of the other voters in some state s . Then we hypothesize that after mixing, the coalition C comprises the other voters who vote for the incumbent, and that $C \cup \{\tau\}$ is decisive but C is not. Consistent with our focus on voting strategies for which mixing occurs only at the initial realization of a state at which the incumbent is bound, we further hypothesize that the voters in C continue to vote for the incumbent in s . By stationarity of voter τ 's decision problem, if it is optimal for her to vote for the incumbent, then it is always optimal to do so; likewise if it is optimal for her to vote for the challenger. Thus, it suffices to compare the

challenger payoff $V_\tau^C(s, t, x)$ with the expected discounted utility $V_\tau^B(s, t, x)$ of continuing to reelect the incumbent for successive realizations of s .

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

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

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

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

where $\rho(s, t, y)$ is unrestricted if $y \in R(s, t) \setminus P(s, t)$. In this case, in every 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 $\rho(s, t, z) = 0$ for all $z \in Z$.

Although mixed voting is required for general existence of equilibria, we say that σ is *deferential* if (i) for all s and all t , $R(s, t) \neq \emptyset$, and (ii) for all s , all t , and all $y \in R(s, t)$, we have $\rho(s, t, y) = 1$. In particular, a deferential equilibrium implies there is no mixing over electoral outcomes.

4 General Properties of Equilibria

Existence and continuity The starting point of our analysis is the next theorem, which provides a foundation for the model by establishing existence of equilibrium under general conditions of the framework.

Theorem 1. *There is a simple 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 $w = (w_{s,t})_{s,t} \in \mathfrak{R}^{S \times T}$ 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 $V_\tau^F(s, x, t)$. Then $v = (v_{s,t,\tau})_{s,t,\tau} \in \mathfrak{R}^{S \times T \times T}$ is the vector of expected voter payoffs. We endow $\mathfrak{R}^{S \times T} \times \mathfrak{R}^{S \times T \times T}$ with the

product topology. Define the correspondence $\mathcal{E}: \Gamma \rightrightarrows \mathfrak{R}^{S \times T} \times \mathfrak{R}^{S \times T \times T}$ so that $\mathcal{E}(\gamma)$ consists of vectors (w, v) such that in the model parameterized by γ , there exists a simple Markov electoral equilibrium $\sigma^* = (\pi^*, \rho^*)$ such that for all s and all t , we have

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

and for all s , all t , and all τ , we have $v_{s,t,\tau} = V_\tau^F(s, t; \sigma^*)$, where we now parameterize continuation values by the strategy profile generating them.

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

The following example illustrates the strategic incentives in a simple two-state model with state-contingent preferences. It embeds a version of the single-state model of Banks and Duggan (2008) but illustrates the critical role of state transitions for the characteristics of simple Markov electoral equilibria: equilibrium play in any one state depends on anticipated equilibrium behavior in other states that may be reached in future periods. In doing so, we introduce a type of equilibrium diagram used to explain the remaining examples in the paper.

Example 1. Let the state space be $S = \{\hat{s}, \check{s}\}$ and the type space be $T = \{\ell, \kappa, r\}$. The set of feasible policies is independent of states and politicians' types and is given by $Y = \{\hat{x}_\ell, \hat{x}_\kappa, \hat{x}_r\}$. Transition probabilities are independent of policies, incumbents' types and electoral outcomes and are such that $p(\hat{s}|\hat{s}) = \hat{p}$ and $p(\check{s}|\hat{s}) = \hat{p}$. Challenger selection probabilities are independent of states, policies and incumbents' types and are such that, for all types t , $q(t) = \frac{1}{3}$. Decisive coalitions are independent of states and incumbents' types and consist of all those sets $A \subseteq T$ such that $|A| \geq 2$, so that a politician is elected if and only if it obtains the support of at least two types of voters. Voters have state-independent ideal policies, with a voter of type τ having ideal policy \hat{x}_τ . In state \hat{s} , voters' stage utilities are single-peaked, with ideal policies ordered such that $\hat{x}_\ell < \hat{x}_\kappa < \hat{x}_r$, so that policy \hat{x}_κ is a Condorcet winner in the stage game. In state \check{s} , voters' preferences induce a Condorcet cycle, with

$$\begin{aligned} u_\ell(\check{s}, \hat{x}_\ell) &> u_\ell(\check{s}, \hat{x}_\kappa) > u_\ell(\check{s}, \hat{x}_r) \\ u_\kappa(\check{s}, \hat{x}_\kappa) &> u_\kappa(\check{s}, \hat{x}_r) > u_\kappa(\check{s}, \hat{x}_\ell) \\ u_r(\check{s}, \hat{x}_r) &> u_r(\check{s}, \hat{s}_\ell) > u_r(\check{s}, \hat{x}_\kappa). \end{aligned}$$

For all voter types and all states, let \hat{u} denote a voter's payoff from its ideal policy, u denote its payoff from its middle-ranked policy, and \check{u} denote its payoff from its third-ranked policy. We further assume that

$$\frac{u - \check{u}}{2} < \hat{u} - u < u - \check{u},$$

so that voters are “risk averse” but not to too great a degree. Elections are decided by majority rule, so that a politician is elected if and only if she obtains the support of at least two types of voters. Politicians have mixed motivations with office benefit $b \geq 0$, and all types have common discount factor δ .

We assume that the single-peaked state \hat{s} is absorbing, i.e., that $\hat{p} = 1$.¹³ This implies that simple Markov electoral equilibria in that state replicate the equilibria of the single state model of Banks and Duggan (2008) and in any equilibrium, politicians of type κ implement policy \hat{x}_κ . Since $u_\kappa(\hat{s}, \hat{x}_\kappa) = \hat{u} > u = u_\kappa(\hat{s}, \hat{x}_r) = u_\kappa(\hat{s}, \hat{x}_\ell)$, voters of type κ vote against any incumbent having implemented a policy other than \hat{x}_κ . Similarly, voters of type ℓ vote against any incumbent having implemented policy \hat{x}_r , and voters of type r vote against any incumbent having implemented policy \hat{x}_ℓ . Hence, \hat{x}_κ is the only policy that can lead to reelection for any politician. We focus on two types of pure strategy equilibria in state \hat{s} . If b is sufficiently high, then equilibrium displays compromise and all politician types implement policy \hat{x}_κ and are reelected. If b is sufficiently low, then equilibrium displays shirking and all politician types implement their ideal policies, with only politicians of type κ being reelected. The assumption that $u - \tilde{u} > \hat{u} - u$ ensures that voters of type $\tau \in \{\ell, r\}$ support politicians implementing their second-ranked policy in this equilibrium, while the assumption that $\hat{u} - u > \frac{1}{2}(u - \tilde{u})$ ensures that politicians of type $t \in \{\ell, r\}$ prefer to implement their ideal policy for a single term and be replaced by a challenger to implement their second-ranked policy \hat{x}_κ and retain office. See Figure 1 for a diagram of the equilibrium. Here, arrows emanating from the realization of a state indicate policy choices of different types, where dark arrows represent type ℓ , medium represent type κ , and red is type r ; and arrows emanating from a policy choice indicate the electoral outcome as a function of the office holder’s type and policy choice, so, e.g., a blue arrow pointing toward C from \hat{x}_ℓ indicates that when a type κ office holder chooses \hat{x}_ℓ , voters replace the politician with a challenger.

Equilibrium play in state \hat{s} depends on whether non- κ politicians compromise or shirk in state \hat{s} . When b is high and all politicians compromise in state \hat{s} , then there exists a simple Markov electoral equilibrium in which all politicians implement their ideal policies in state \hat{s} and are reelected. In this equilibrium, there is no disagreement expected in state \hat{s} , and all politicians deliver the same payoffs to all voters once a transition to that state occurs. Hence, voters’ decisions to reelect an incumbent in state \hat{s} depend only on the policy she implements while the state remains \hat{s} , and all politician types can garner the support of some majority of voters in that state. In state \hat{s} , only voters not of type κ vote in favor of their second-ranked policy. However, without a Condorcet winner in state \hat{s} , all voter types support incumbents having implemented their second-ranked policy (types ℓ support \hat{x}_κ , types κ support \hat{x}_r and types r support \hat{x}_ℓ).

¹³In Example 2, we let $\hat{p} > 1$ but have state \hat{s} be absorbing, i.e., $\hat{p} = 1$.




Figure 1: Equilibrium with state-dependent preferences

When b is low and all politicians shirk in state \hat{s} , then there exists \bar{p} such that when $\hat{p} \leq \bar{p}$, there exists an equilibrium in which all politicians implement their ideal policy in state \check{s} and only politicians of type κ are reelected in that state. Indeed, if \hat{p} is low, then in state \check{s} voters support candidates who, when freed from their policy commitments by a transition to state \hat{s} , implement policies they find acceptable. Hence, in state \check{s} , voters of type r are no longer willing to support incumbents of type ℓ and voters of type κ are no longer willing to support incumbents of type r . Meanwhile, incumbents of type κ retain the support of voters of type ℓ when they implement policy \hat{x}_κ . Hence, when disagreement is expected in state \hat{s} , only politicians of type κ , the median type in state \hat{s} , gain majority support in state \check{s} , so the remaining types simply choose their ideal policies before being removed from office. \square

In the remainder of this section, we present results that elaborate on some general properties of simple Markov electoral equilibria. First, we consider the role of state-by-state commitment in generating dynamic incentives in the equilibria of the model. We then turn to the importance of mixed voting and policy strategies and provide conditions sufficient for equilibria with pure voting strategies and limited mixing (in a sense defined below) by politicians.

Commitment We examine the role of state-by-state commitment by analyzing differential equilibria in the model where states never persist. In particular, if transitions p_t and q_t are independent of policies, and if the current state s transitions to a different state $s' \neq s$ with probability one if the incumbent is reelected, i.e., $p_t(s|s, 1) = 0$, then the dynamic linkage in the model is broken.¹⁴

¹⁴The restriction to differential equilibria precludes possibly interesting equilibria in which

Indeed, the payoffs to a voter of type τ from the incumbent or the challenger,

$$\begin{aligned} V_\tau^B(s, t, x) &= \sum_{s'} p_t(s'|s, 1) V_\tau^F(s', t) \\ V_\tau^C(s, t, x) &= \sum_{s'} p_t(s'|s, 0) \sum_{t'} q_t(t'|s, s') V_\tau^F(s', t'), \end{aligned}$$

are independent of x . Then in a deferential equilibrium, either $\rho(s, t, \cdot)$ is constant at zero or constant at one, so $W_t^B(s, x)$ and $W_t^C(s, x)$ are equivalent (up to constant terms) to the stage utility $w_t(s, x)$. Therefore,

$$\rho(s, t, x) W_t^B(s, x) + (1 - \rho(s, t, x)) W_t^C(s, x) = w_t(s, x) + \text{constant},$$

and the office holders trivially maximizes her stage utility.

Proposition 1. *Let σ be a deferential simple Markov electoral equilibrium. For given s and t , assume that $p_t(s'|s, x, 1)$ and $q_t(t'|s, x)$ are independent of x , and that $p_t(s|s, 1) = 0$. Then $\pi_t(\cdot|s)$ puts probability one on solutions to*

$$\max_{x \in X_t(s)} w_t(s, x).$$

Mixed voting Randomization over electoral outcomes, i.e., $\rho(s, t, x) \in (0, 1)$, is needed in some environments, such as when X is finite. It is less important under the following set of assumptions: for given s and t ,

- (A1) $Y_t(s)$ is convex,
- (A2) $u_t(s, \cdot)$ and $w_t(s, \cdot)$ are strictly quasi-concave,
- (A3) transitions $p_t(s'|s, x, e)$ and $q_t(t'|s, x)$ are independent of x for all s' , all t' , and all e ,
- (A4) $p_t(s|s, 1) > 0$.

These conditions impose typical convexity structure, along with the assumption that transitions are policy-independent and that the probability that a state persists is always positive. Then

$$V_\tau^B(s, t, x) = \frac{p_t(s|s, 1) u_\tau(s, x) + \sum_{s' \neq s} p_t(s'|s, 1) V_\tau^F(s', t)}{1 - \delta_\tau p_t(s|s, 1)} \quad (5)$$

is strictly quasi-concave in x , and

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

is independent of x . Thus, whenever $P(s, t) \neq \emptyset$, weak reelection sets $R(s, t)$ are the closure of strict $P(s, t)$, i.e., no policies in $R(s, t)$ are isolated from $P(s, t)$.

The proof of the following lemma is straightforward and is omitted.

voters impose a cutoff for re-electing the incumbent, although they are indifferent between all politicians; see, e.g., Barro (1973).

Lemma 1. *Let σ be a simple Markov electoral equilibrium. Assume (A1)–(A4) hold at s and t . Let $C \in \mathcal{D}_t(s)$ be such that $P_C(s, t) \neq \emptyset$. Then $P_C(s, t)$ is convex, $R_C(s, t)$ is convex, and $R_C(s, t) = \text{cl}P_C(s, t)$.*

We now provide conditions under which a simple Markov electoral equilibrium with randomized voting strategies can be replaced by an outcome-equivalent deferential simple Markov electoral equilibrium.

Proposition 2. *Let $\sigma = (\pi, \rho)$ be a simple Markov electoral equilibrium. Assume that for all s and all t , (A1)–(A4) hold; furthermore, assume that for all $C \in \mathcal{D}_t(s)$, we have $P_C(s, t) \neq \emptyset$. Then there is a deferential simple Markov electoral equilibrium $\sigma' = (\pi', \rho')$ that attains the same policy outcomes, and hence payoffs, as σ , i.e., for all s , all t , all τ , and all x , $V_\tau^F(s, t|\sigma') = V_\tau^F(s, t)$ and $W_t^B(s, x|\sigma') = W_t^B(s, x)$.*

Proof. Consider a simple Markov electoral equilibrium σ , state s , and type t office holder. Let $W_t(s)$ denote the equilibrium payoff of the office holder in state s under σ , i.e.,

$$W_t(s) \equiv \max \left\{ \max_{x \in R(s, t)} \left\{ \rho(s, t, x) W_t^B(s, x) + (1 - \rho(s, t, x)) W_t^C(s, x) \right\}, \right. \\ \left. \max_{x \in Z_t(s)} W_t^C(s, x) \right\}.$$

We claim that for all $x \in R(s, t)$,

$$W_t(s) \geq W_t^B(s, x). \quad (6)$$

Otherwise, there exists $C \in \mathcal{D}_t(s)$ such that $x \in R_C(s, t)$. Lemma 1 implies that $R_C(s, t) = \text{cl}P_C(s, t)$, and since $W_t^B(s, \cdot)$ is continuous and $P_C(s, t) \neq \emptyset$, given any $\epsilon > 0$, there exists $x^\epsilon \in P_C(s, t)$ such that $W_t^B(s, x^\epsilon) + \epsilon > W_t^B(s, x) > W_t(s)$. Since we can choose $\epsilon > 0$ small enough that $W_t^B(s, x^\epsilon) > W_t(s)$ and since $\rho(s, t, x^\epsilon) = 1$, this contradicts the optimality of $\pi_t(\cdot|s)$. Therefore, $W_t(s) \geq W_t^B(s, x)$, as claimed.

We define a voting strategy ρ' that is identical to ρ except that $\rho'(s, t, x) = 1$ for all $x \in R(s, t)$, and define proposal strategies $\pi'_t(\cdot|s)$ such that for all measurable $E \subseteq X_t(s)$,

$$\pi'_t(E|s) = \pi_t(E \cap Z_t(s)|s) + \int_{E \cap Y_t(s)} \rho(s, t, x) \pi_t(dx|s) \\ + \int_{\xi|_Y^{-1}(E \cap Z_t(s))} (1 - \rho(s, t, x)) \pi_t(dx|s),$$

where $\xi|_Y$ is the restriction of ξ to Y . Voters never randomize under profile (π', ρ') , and instead we rely on randomization by politicians to mimic the electoral outcomes of the equilibrium (π, ρ) . That is, whenever the office-holder

implements policy $x \in R(s, t)$ under π , under π' she implements policy x with probability $\rho(s, t, x)$ and policy $\xi(x)$ with probability $1 - \rho(s, t, x)$. First, note that under (π', ρ') all continuation values are the same as under (π, ρ) , and payoffs to office holders are still $W_t(s)$. Second, note that this implies that for all states s and types t , $R'(s, t) = R(s, t)$, and hence that ρ' is consistent with voter optimality. Third, to verify that π' is consistent with optimality of policy choices, consider $x \in R(s, t)$ with $\rho(s, t, x) < 1$, and note that if $x \notin \text{supp}(\pi_t(\cdot|s))$, then, by (6), specifying that $\rho'(s, t, x) = 1$ does not provide any incentives for type t politicians to deviate to policy x . If instead $x \in R(s, t) \cap \text{supp}(\pi_t(\cdot|s))$, then it must be that

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

Indeed, if $x \in \text{supp}(\pi_t(\cdot|s))$, then we have that $W_t^B(s, x) \geq w_t(s, x) + \delta_t V_t^C(s, t)$, since the politician can obtain the payoff $w_t(s, x) + \delta_t V_t^C(s, t)$ by implementing policy $\xi(x)$. Also, if $W_t^B(s, x) > w_t(s, x) + \delta_t V_t^C(s, t)$, then (6) and $\rho(s, t, x) < 1$ imply that $W_t(s) > \rho(s, t, x)W_t^B(s, x) + (1 - \rho(s, t, x))W_t^C(s, x)$, contradicting $x \in \text{supp}(\pi_t(\cdot|s))$. Given that the office-holder is indifferent between being reelected or not following policy x , the additional randomization introduced by π' relative to π is consistent with optimality. \square

A condition of Proposition 2 is that strict reelection sets be nonempty in all states. This ensures that voters can mix following policies in $R(s, t) \setminus P(s, t)$ only when incumbents are indifferent between being reelected or not, so that we rely on randomized campaign decisions by indifferent politicians in the deferential equilibrium to mimic randomized voting strategies in the original equilibrium. This assumption is not innocuous. Indeed, in Example 3 below, we provide a simple setting (with policy-dependent transitions) in which, in some states, the weak reelection sets of all politicians are empty.

Next, we give conditions under which an office holder will at most mix over the decisive coalitions that can be targeted to gain reelection through compromise and shirking. In such equilibria, policy strategies will not involve mixing over policies acceptable to a given decisive coalition. Specifically, under conditions (A1)–(A4), the expected discounted utility of a type t office holder from choosing policy x in state s and being reelected in state s (and continuing to be reelected in s) is for all $x \in Y$,

$$W_t^B(s, x) = \frac{1}{1 - \delta_t p_t(s|s, 1)} \left[w_t(s, x) + \delta_t \sum_{s' \neq s} p_t(s'|s, 1) \int_{x'} [\rho(s', x', t) W_t^B(s', x') + (1 - \rho(s', x', t)) W_t^C(s, x')] \pi_t(dx'|s') \right],$$

and for all $x \in Z$,

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

which are strictly quasi-concave in x . An implication, with Lemma 1, is that for each $C \in \mathcal{D}_t(s)$ such that $R_C(s, t) \neq \emptyset$, the function $W_t^B(s, x)$ has a unique maximizer on the set $R_C(s, t)$.

Proposition 3. *Let σ be a simple Markov electoral equilibrium. Assume (A1)–(A4) hold at s and t . Then for all $C \in \mathcal{D}_t(s)$, we have*

$$|\{x \in R(s, t) : \rho(s, t, x) > 0\} \cap \text{supp } \pi_t(\cdot|s)| \leq 1.$$

Proof. Given state s and a type t office-holder, consider any coalition $C \in \mathcal{D}_t(s)$ and let $x', x'' \in \{x \in R(s, t) : \rho(s, t, x) > 0\} \cap \text{supp } \pi_t(\cdot|s)$. If $x' \neq x''$, then strict quasi-concavity of $u_\tau(s, \cdot)$ for all voter types $\tau \in C$ implies that for all $\lambda \in (0, 1)$, $\lambda x' + (1 - \lambda)x'' \in P_C(s, t)$, and therefore $\rho(s, t, \lambda x' + (1 - \lambda)x'') = 1$. Assume without loss of generality that $w_t(s, x'') \geq w_t(s, x')$. Since $\rho(s, t, x') > 0$, it must be that $W_t^B(s, x') \geq W_t^C(s, x')$, otherwise a politician of type t could obtain a strictly higher payoff by deviating to $\xi(x')$ instead of x' in state s . Hence, the strict quasiconcavity of $w_t(s, \cdot)$ implies that

$$\begin{aligned} W_t^B(s, \lambda x' + (1 - \lambda)x'') &> W_t^B(s, x') \\ &\geq \rho(s, t, x')W_t^B(s, x') + (1 - \rho(s, t, x'))W_t^C(s, x'), \end{aligned}$$

contradicting the optimality of $\pi_t(\cdot|s)$. \square

Thus, if a single voter is representative in state s , in the sense considered in the next section, then the preceding result establishes that the office-holder will either compromise by choosing a policy acceptable to this voter, choose her ideal policy, or randomize over these two options.

5 Representative Voters

Given a simple Markov electoral equilibrium σ , a type τ voter is *representative* in state s for a type t office holder if $P(s, t) = P_\tau(s, t)$ and $R(s, t) = R_\tau(s, t)$.

Example 2. In all equilibria from Example 1, since the state \hat{s} with a Condorcet winner was absorbing, voter type κ was representative in that state. Note also that voter type κ was representative in state \hat{s} when the equilibrium in state \hat{s} called for shirking, whereas no voter type was representative in state \hat{s} when the equilibrium in state \hat{s} called for compromise. That voter type κ need not be representative in state \hat{s} in all equilibria is natural since policy \hat{x}_κ is not a Condorcet winner in that state. However, in a variant of the model from that example, we show that in some equilibria voter type κ need not even be representative in state \hat{s} .

To this end, suppose that the state with the Condorcet cycle is absorbing ($\hat{p} = 1$). If the probability \hat{p} of remaining in the state with a Condorcet winner is sufficiently low, then there exists a simple Markov electoral equilibrium in

which, in all states, politicians of all types implement their ideal policies and are reelected. In state \hat{s} , voters support politicians who will implement their second-ranked policy in state \hat{s} , even if these politicians implement their third-ranked policy in state \hat{s} . In particular, voters of type ℓ and κ support policy \hat{x}_κ when implemented by a politician of type κ , voters of type r and κ support policy \hat{x}_r when implemented by a politician of type r and voters of type r and ℓ support policy \hat{x}_ℓ when implemented by a politician of type ℓ . When transitions away from state \hat{s} are likely, \hat{x}_κ is a Condorcet winner in state \hat{s} , voters support incumbents based on the policies they will implement in state \hat{s} . The absence of a Condorcet winner in state \hat{s} is carried into state \hat{s} through dynamic incentives. \square

Given s , say $\mathcal{D}_t(s)$ is a *weighted majority rule* if there exist weights $n_\tau(s) \geq 0$ with $\sum_{\tau \in T} n_\tau(s) = 1$ such that $\mathcal{D}_t(s) = \{C : \sum_{\tau \in C} n_\tau(s) > \frac{1}{2}\}$. Assume policies lie in Euclidean space, so $Y \subseteq \mathbb{R}^d$, and to each type τ , associate a vector $\hat{x}_\tau \in \mathbb{R}^d$. The stage utility $u_\tau(s, \cdot)$ is *quadratic* if for all $x \in Y$, $u_\tau(s, x) = -\|x - \hat{x}_\tau\|^2$. A vector $x \in \mathbb{R}^d$ is a *total median* in state s for a type t office holder if for every non-zero $p \in \mathbb{R}^d$, we have

$$\sum \{n_\tau(s) : p \cdot (\hat{x}_\tau - x) > 0\} \leq \frac{1}{2}.$$

If $d = 1$, then there is at least one total median, although there may be no total median in higher dimensions. If $\mathcal{D}_t(s)$ is *strong*, in the sense that there is no coalition C with $\sum_{\tau \in C} n_\tau(s) = \frac{1}{2}$, then there is at most one total median, and if x is a total median, then there exists a type $\kappa \in T$ with $n_\kappa(s) > 0$ such that $x = \hat{x}_\kappa$, and we refer to κ as the *total median type* at s for t . A policy $x \in Y$ is a *core policy* in state s for a type t office holder if there is no $y \in Y$ such that $\{\tau \in T : u_\tau(s, y) > u_\tau(s, x)\} \in \mathcal{D}_t(s)$. Assuming Euclidean stage utilities, every core policy that belongs to the interior of Y (if any) is a total median, but the converse need not hold; in particular, a total median need not be a feasible policy.

We make use of the following conditions, where the last two restrictions hold for a given state s and type t :

- (B1) $Y \subseteq \mathfrak{R}^d$,
- (B2) $u_\tau(s, x)$ is quadratic for all τ ,
- (B3) $\delta_\tau = \delta$ for all τ ,
- (B4) $\mathcal{D}_t(s)$ is a strong weighted majority rule,
- (B5) there is a total median policy at s for t .

The assumption of quadratic utility, in (B2), may seem more stringent than it actually is. Although we write utility as a function of policy alone, we can exploit

state-dependence of the feasible policies to capture settings in which ideal points are subject to a common shock. To be more precise, consider a model in which utility functions have the form $u_\tau(x, s) = -\|x - \hat{x}_\tau - s\|^2$, where the state $s \in \mathbb{R}^d$ shifts the ideal points of all voters. This functional form is precluded by (B2), but we can reformulate the model so that the feasible set is $\tilde{Y}(s) = Y(s) - s$, the challenger distribution and state transition are $\tilde{q}_t(t'|s, \tilde{x}) = q_t(t'|s, \tilde{x} + s)$ and $\tilde{p}_t(s'|s, \tilde{x}, e) = p_t(s'|s, \tilde{x} + s, e)$, and utilities are $\tilde{u}_\tau(\tilde{x}) = -\|\tilde{x} - \hat{x}_\tau\|^2$; this satisfies (B2) and preserves the strategic incentives of the original model.

The next result establishes that the above conditions indeed imply that some voter type is representative in s for t .

Theorem 3. *Let σ be a simple Markov electoral equilibrium. Assume (B1)–(B3), and assume (B4) and (B5) hold for s and t . Let κ be the total median type at s for t . Then voter type κ is representative in s for t :*

$$P(s, t) = P_\kappa(s, t) \quad \text{and} \quad R(s, t) = R_\kappa(s, t).$$

Banks and Duggan (2001) establish the representative voter theorem for quadratic utilities in the adverse selection model, which corresponds to the single-state model in the present context. The approach uses a lemma of Banks and Duggan (2006) to the effect that the total median voter is decisive between pairs of lotteries over policies: there is a decisive coalition preferring one lottery to another if and only if this is the preference of the total median voter type. The application of the lemma relies on the observation that a voter’s expected discounted payoffs from electing the incumbent and challenger (normalized by $1 - \delta$) can be written as integrals with respect to appropriately specified probability measures. Given a deferential equilibrium in the single-state model with mixed motivations, we have $p_t(s|s, x, 1) = 1$, so $(1 - \delta)V_\tau^B(s, t, x) = u_\tau(s, x)$, which is trivially the integral with respect to the unit mass on x . For the continuation value of a challenger, we can solve explicitly for $V_\tau^C(s, t, x)$ to obtain

$$\begin{aligned} & (1 - \delta)V_\tau^C(s, t, x) \\ &= \frac{\sum_{t'} [q_t(t'|s, x)\pi_{t'}(X \setminus R_\tau(s, t'))(1 - \delta)u_\tau(\hat{x}_{t'}) + \int_{R_\tau(s, t')} u_\tau(x)\pi_{t'}(dx|s)]}{1 - \delta \sum_{t'} q_t(t'|s, x)\pi_{t'}(X \setminus R_\tau(s, t'))} \end{aligned}$$

By inspection, this is the integral with respect to the “continuation distribution” ν^* defined by

$$\nu^*(A) = \frac{\sum_{t'} [q_t(t'|s, x)\pi_{t'}(X \setminus R_\tau(s, t'))(1 - \delta)\Delta(\hat{x}_{t'}) + \pi_{t'}(A \cap R(s, t'))]}{1 - \delta \sum_{t'} q_t(t'|s, x)\pi_{t'}(X \setminus R_\tau(s, t'))}$$

for every Borel set $A \subseteq X$, where $\Delta(\hat{x}_{t'})$ is the unit mass on the ideal point of the type t' voter. Then decisiveness of the total median voter implies that $R(s, t)$ consists of policies that yield utility to the total median at least equal to the expected utility from the lottery ν^* , so the total median is representative.

In the general model with multiple states and transition probabilities that depend on policy, the application of the lemma of Banks and Duggan (2006) is not straightforward: the relevant lotteries will depend on the state s , the incumbent type t , and the incumbent's policy choice x , and they must be obtained by solving a (possibly infinite) system of equations. The proof of Theorem 3 applies the contraction mapping theorem in a space of measure-valued functions to establish the existence of families $\{\mu_X^*(\cdot|s, t, x) \mid s \in S, t \in T, x \in X\}$ and $\{\nu_X^*(\cdot|s, t, x) \mid s \in S, t \in T, x \in X\}$ of probability measures on X such that

$$\begin{aligned} V_\tau^B(s, t, x) &= \frac{1}{1 - \delta} \int_{x'} u_\tau(x') \mu_X^*(dx' | s, t, x) \\ V_\tau^C(s, t, x) &= \frac{1}{1 - \delta} \int_{x'} u_\tau(x') \nu_X^*(dx' | s, t, x) \end{aligned}$$

for all s , all t , and all x . Mathematically speaking, then, an election presents voters with the choice between two lotteries, the *incumbent lottery* $\mu_X^*(\cdot|s, t, x)$ and the *challenger lottery* $\nu_X^*(\cdot|s, t, x)$. Then the representative voter theorem follows from the social choice result on the decisiveness of the total median.

When some voter type κ is a representative in state s for type t , we may consider the *optimal retention problem* in s for t for the representative voter: the problem is to retain a type t office holder or replace her with a challenger as a function of her policy choice $x \in Y$ at the current state s (and all future realizations of s), given the strategies of all players. The Bellman equation for this problem is

$$\overline{V}_\kappa^B(s, t, x) = \max \left\{ p_t(s|s, x, 1)[u_\kappa(s, x) + \delta_\kappa \overline{V}_\kappa^B(s, t, x)] + \sum_{s' \neq s} p_t(s'|s, x, 1) \overline{V}_\kappa^F(s', t), \overline{V}_\kappa^C(s, t, x) \right\}$$

for all $x \in Y$, where \overline{V}_κ^C and \overline{V}_κ^F are defined as in (2) and (3), with \overline{V}_κ^B substituted for V_κ^B in state s given politician type t . The outcome of an election in state s between a type t incumbent and a challenger will be determined, by definition, according to whether the incumbent or challenger offers the representative voter a higher expected discounted payoff. Although the representative voter's preferences indicate which candidate will win in state s , however, she plays a passive role in this definition: she cannot unilaterally overturn electoral outcomes in the current and future periods. Nevertheless, the next lemma establishes that electoral outcomes for a type t office holder in state s do solve the representative voter's optimal retention problem. Thus, in equilibrium, it is as if the representative voter in s for t can actively choose a strategy for reelecting a type t incumbent in all realizations of s given any policy choice.

Lemma 2. *Let σ be a simple Markov electoral equilibrium. Assume that the type κ voter is representative in s for t . Then for all $x \in Y$,*

$$\overline{V}_\kappa^B(s, t, x) = \rho(s, t, x) V_\kappa^B(s, t, x) + (1 - \rho(s, t, x)) V_\kappa^C(s, t, x).$$

Proof. Let the type κ voter be representative in s for t , and define

$$\tilde{V}_\kappa^B(s, t, x) = \rho(s, t, x)V_\kappa^B(s, t, x) + (1 - \rho(s, t, x))V_\kappa^C(s, t, x)$$

for all $x \in Y$. It suffices to show that

$$\tilde{V}_\kappa^B(s, t, x) = \max \left\{ \begin{array}{l} p_t(s|s, x, 1)[u_\kappa(s, x) + \delta_\kappa \tilde{V}_\kappa^B(s, t, x)] \\ + \sum_{s' \neq s} p_t(s'|s, x, 1)\tilde{V}_\kappa^F(s', t) \end{array} \right\}, \tilde{V}_\kappa^C(s, t, x) \left\},$$

where \tilde{V}_κ^C and \tilde{V}_κ^F are defined as in (2) and (3), with $\tilde{V}_\kappa^B(s, t, x)$ substituted for V_κ^B in state s given politician type t for each policy x . For future use, we refer to the modified equations as (2') and (3'). First, we claim that for all $x \in Y$,

$$\begin{aligned} \rho(s, t, x)\tilde{V}_\kappa^B(s, t, x) + (1 - \rho(s, t, x))V_\kappa^C(s, t, x) \\ = \rho(s, t, x)V_\kappa^B(s, t, x) + (1 - \rho(s, t, x))V_\kappa^C(s, t, x). \end{aligned}$$

Indeed, the claim follows from two observations: if $V_\kappa^B(s, t, x) \geq V_\kappa^C(s, t, x)$, then because κ is representative in s for t , $\rho(s, t, x) < 1$ implies $V_\kappa^B(s, t, x) = V_\kappa^C(s, t, x)$; and if $V_\kappa^C(s, t, x) > V_\kappa^B(s, t, x)$, then $\rho(s, t, x) = 0$. Second, the claim implies that equations (2') and (3') are in fact identical to (2) and (3), and we conclude that for all x , $\tilde{V}_\kappa^C(s, t, x) = V_\kappa^C(s, t, x)$, and that for all s' , $\tilde{V}_\kappa^F(s', t) = V_\kappa^F(s', t)$. Third, to show that $\tilde{V}_\kappa^B(s, t, \cdot)$ solves the desired functional equation, consider any $x \in Y$. In case $V_\kappa^B(s, t, x) \geq V_\kappa^C(s, t, x)$, the second step above implies that $\tilde{V}_\kappa^B(s, t, x) = V_\kappa^B(s, t, x)$. With the third step, we then have

$$\begin{aligned} V_\kappa^B(s, t, x) &= p_t(s|s, x, 1)[u_\kappa(s, x) + \delta_\kappa \tilde{V}_\kappa^B(s, t, x)] \\ &\quad + \sum_{s' \neq s} p_t(s'|s, x, 1)\tilde{V}_\kappa^F(s', t), \end{aligned}$$

and since $V_\kappa^B(s, t, x) \geq \max\{V_\kappa^B(s, t, x), V_\kappa^C(s, t, x)\}$, we are done. In the complementary case $V_\kappa^B(s, t, x) < V_\kappa^C(s, t, x)$, the second and third steps imply $\tilde{V}_\kappa^B(s, t, x) = V_\kappa^C(s, t, x) = \tilde{V}_\kappa^C(s, t, x)$. Suppose in order to deduce a contradiction that

$$\tilde{V}_\kappa^B(s, t, x) < \max \left\{ \begin{array}{l} p_t(s|s, x, 1)[u_\kappa(s, x) + \delta_\kappa \tilde{V}_\kappa^B(s, t, x)] \\ + \sum_{s' \neq s} p_t(s'|s, x, 1)\tilde{V}_\kappa^F(s', t) \end{array} \right\}, \tilde{V}_\kappa^C(s, t, x) \left\}.$$

Therefore, after substitution we have

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

or equivalently,

$$\begin{aligned} V_\kappa^C(s, t, x) &< \frac{p_t(s|s, x, 1)u_\kappa(s, x) + \sum_{s' \neq s} p_t(s'|s, x, 1)V_\kappa^F(s', t)}{1 - \delta_\kappa p_t(s|s, x, 1)} \\ &= V_\kappa^B(s, t, x), \end{aligned}$$

a contradiction. \square

The next result establishes conditions under which the weak reelection set is nonempty for at least one office holder type at any given state. The logic is fairly simple: assuming some voter type is representative in s for all t , if the best possible politician type chooses the best possible policy for the representative voter in state s , then no challenger type can offer a higher expected payoff to the voter. This argument relies on politicians feasibly mimicking the policy choices and associated continuation payoffs of other types, so that it requires state and challenger transitions to be independent of types. The analysis rests on the following assumptions, given a state s :

- (C1) $Y_t(s)$ is independent of t ,
- (C2) $p_t(s'|s, x, e)$ is independent of t and e for all s' ,
- (C3) $q_t(t'|s, x)$ is independent of t for all t' .

Thus, it is at least conceivable that some politician type can be reelected in equilibrium.

Proposition 4. *Let σ be a simple Markov electoral equilibrium. Assume (C1)–(C3), and assume that the type κ voter is representative in s for all t . Then there exists a type t such that $R(s, t) \neq \emptyset$.*

Proof. Fix state s , and suppose, in order to deduce a contradiction that for all types t , $R(s, t) = \emptyset$. Thus, for all types t and all policies x , $\rho(s, t, x) = 0$ and

$$V_\kappa^F(s, t) = \int_x [u_\kappa(s, x) + \delta_\kappa V_\kappa^C(s, x)] \pi_t(dx|s).$$

Let type t and policy x satisfy

$$\begin{aligned} x &\in \operatorname{argmax} \left\{ u_\kappa(s, x') + \delta_\kappa V_\kappa^C(s, x') : x' \in \bigcup_{t' \in T} \operatorname{supp}(\pi_{t'}(\cdot|s)) \right\}, \\ t &\in \operatorname{argmax} \left\{ V_\kappa^B(s, t', x) : t' \in T \right\}. \end{aligned}$$

Assume that $x \in Y(s)$. This is without loss of generality since if $x \in Z(s)$, we can consider $\xi^{-1}(x)$. For every type t' and every policy $x' \in \operatorname{supp}(\pi_{t'}(\cdot|s))$, we have

$$\begin{aligned} V_\kappa^B(s, t, x) &\geq V_\kappa^B(s, t', x) \\ &= p(s|s, x) \left[u_\kappa(s, x) + \delta_\kappa V_\kappa^C(s, x) \right] + \sum_{s' \neq s} p(s'|s, x) V_\kappa^F(s', t') \\ &\geq p(s|s, x) \left[u_\kappa(s, x') + \delta_\kappa V_\kappa^C(s, x') \right] + \sum_{s' \neq s} p(s'|s, x) V_\kappa^F(s', t'), \end{aligned}$$

which implies that $V_\kappa^B(s, t, x) \geq \sum_{s'} p(s'|s, x) V_\kappa^F(s', t')$. Therefore, since t' was arbitrary, we have

$$V_\kappa^B(s, t, x) \geq \sum_{t'} q(t'|s, x) \sum_{s'} p(s'|s, x) V_\kappa^F(s', t') = V_\kappa^C(s, x),$$

which implies that $R(s, t) \neq \emptyset$, yielding the desired contradiction. \square

If in equilibrium voters support politicians implementing acceptable policies and politicians place sufficiently great weight on holding office, then any office holder whose weak reelection set is nonempty will choose a policy from this set and be reelected. Under the conditions of Proposition 4, therefore, at least one politician type will compromise to win reelection.

Corollary 1. *Let σ be a simple Markov electoral equilibrium and for all t , assume mixed motivation, i.e., $w_t(s, x) = u_t(s, x) + b_t$, with b_t large, and $\delta_t > 0$.*

1. *Assume (C1)–(C3) and assume that the type κ voter is representative in s for all t and that, for all types t and all $x \in R(s, t)$, we have $\rho(s, t, x) > 0$. Then there is a type t such that $\pi_t(R(s, t)|s) = 1$.*
2. *If σ is deferential, then for all types t , $\pi_t(R(s, t)|s) = 1$.*

Proof. Without loss of generality, normalize stage utilities so that $0 \leq u_t \leq \bar{u}$. If the simple Markov electoral equilibrium has the property listed in part 1, then the payoff to a type t office-holder from choosing policy $x \in R(s, t)$ is at least $b_t[1 + \delta_t \rho(s, t, x)]$, while her payoff to choosing policy $x \notin R(s, t)$ is at most $b_t + \frac{\bar{u}}{1 - \delta_t}$. Since $\delta_t > 0$ and $\rho(s, t, x) > 0$ for all t and all $x \in R(s, t)$, politicians prefer reelection for sufficiently high office benefit b_t . In particular, the type t politician established in Proposition 4, for whom $R(s, t) \neq \emptyset$, will chose to compromise, so that $\pi_t(R(s, t)|s) = 1$, as required. Part 2 follows immediately, since in a deferential equilibrium $R(s, t) \neq \emptyset$ for all s and t and $\rho(s, t, x) = 1$ for all $x \in R(s, t)$. \square

If the assumptions of Proposition 4 are relaxed, then the result no longer holds. In particular, if state transitions can depend on the office holder's type, then it is possible that in some states, the weak reelection sets of all politicians are empty: voters have a strict incentive to remove all incumbents, and so politicians have no incentives to compromise in order to improve their reelection chances. This new form of “political failure,” which is driven by dynamic incentives, is illustrated in the next example.

Example 3 (Dynamic political failure). Assume the state space is $S = \{s_1, s_{-1}, \underline{s}\}$, and the type space is $T = \{1, -1, \kappa\}$. Assume that type κ voters are representative in all states and that politicians can only be of type 1 or -1 , and let t range over $\{1, -1\}$. Sets of feasible policies are independent of politicians' types and are such that $Y(s_t) = \{x_1, x_{-1}\}$ and $Y(\underline{s}) = \{\underline{x}\}$. Transition




Figure 2: Dynamic political failure

probabilities are such that $p_t(s_{-t}|s_t, x_1) = 1$, $p_t(s_t|s_t, x_{-1}) = p_t(s_{-t}|s_{-t}, x_{-1}) = p_t(s_{-t}|s_{-t}, x_1) = p \in (0, 1)$, and $p_t(\underline{s}|s_t, x_{-1}) = p_t(\underline{s}|s_{-t}, x_{-1}) = p_t(\underline{s}|s_{-t}, x_1) = 1 - p$, where we assume that p is sufficiently small. State \underline{s} is absorbing. Challenger selection probabilities are independent of states, policies and incumbents' types and are such that $q(t) = q(-t) = \frac{1}{2}$. Note that, other than the part of (C2) that requires transition probabilities be independent of politicians' types, all conditions (C1)–(C3) are respected. The payoffs to type κ voters are independent of states and are such that $u_\kappa(x_1) > u_\kappa(x_{-1}) > u_\kappa(\underline{x})$. Politicians have mixed motivations with type-independent office benefit $b \geq 0$ and stage utilities such that $u_t(s_t, x_1) = u_t(s_{-t}, x_{-1}) > u_t(s_t, x_{-1}) = u_t(s_{-t}, x_1) > u_t(\underline{s}, \underline{x})$.

We claim that there exists a simple Markov electoral equilibrium in which all type t politicians choose policy x_1 in state s_t and policy x_{-1} in state s_{-t} and such that for all states $s \in \{s_1, s_{-1}\}$, we have $R(s, t) = \emptyset$. To see that voting strategies are optimal on the equilibrium path, note that in state s_t , a politician of type t who implements the optimal policy x_1 for the type κ voter induces a transition to state s_{-t} . If the voter κ reelects the incumbent, then this politician would choose her ideal policy x_{-1} in state s_{-t} and not be reelected. If instead the voter opts for the challenger, then this politician may be of type $-t$, in which case she would choose her ideal policy in state s_{-t} , which is the optimal policy x_1 for the type κ voter. Hence, voters of type κ have a strict incentive to opt for the challenger in order to target a politician that better fits the next period's state. See Figure 2 for the equilibrium diagram depicting transition probabilities and the policy choices of the type 1 politicians.

The surprising feature of this example is that politicians of type t cannot implement any policy in states s_t or s_{-t} that leads to reelection. Hence, it

must be that in state s_t , the voters of type κ have strict incentives to replace a type t office holder who chooses policy x_{-1} . If the state does not transition to the “bad” absorbing state \underline{s} , then it remains at s_t . A reelected incumbent is committed to x_{-1} , and so, in that state, she mimics the behavior of a type $-t$ office holder in that state, and they are always replaced. It must also be that in state s_{-t} , voters of type κ have strict incentives to replace an incumbent of type t who has implemented the optimal policy x_1 of the type κ voter. Again, the state can either remain in s_{-t} or transition to the bad state \underline{s} . If it remains in s_{-t} , then the incumbent is committed to the ideal policy of type κ voters. However, a challenger of type $-t$ would choose the same policy, but with the added benefit of steering future transitions away from the bad state \underline{s} . The risk for type κ voters of opting for the challenger is that if she is type t , then she chooses policy x_{-1} . If p , the probability of remaining in state s_{-t} , is sufficiently small, then supporting the challenger is strictly optimal for type κ voters. \square

6 Dynamic Core Convergence

In the basic Downsian model, two candidates simultaneously choose platforms in a one-dimensional policy space, an odd number of voters have single-peaked preferences, and candidates are motivated solely by winning the election. The classical median voter theorem is that the unique equilibrium of this game is that each candidate promises, and if elected implements, the ideal policy of the median voter. The result persists even if candidates are policy motivated or have mixed motivations. There are three aspects of the current model that undercut the possibility of such a result. First, candidates cannot make binding campaign promises; rather, an incumbent “competes” by choice of policy with the prospect of an unknown challenger. Second, even if a fixed voter type is representative across all states, there is not generally a single “ideal point” of the representative type; rather, the optimal policy choices of the voter will be state-dependent and would be obtained as the solution to a hypothetical dynamic programming problem in which this voter could choose policies directly. Third, it is possible that different voter types are representative in different states, so the hypothetical scenario is not as simple as solving a dynamic programming problem; rather, we hypothesize a game among representative voter types, and we must address the possibility of multiple equilibria in this “representative voting game.”

The first aspect is addressed by Banks and Duggan (2008), who establish in the single-state model that when players are sufficiently patient, or when office benefits are sufficiently high, the policies chosen by office holders of all types will converge to the ideal point of the median type. Here, we address the second and third aspects. We assume that in each state, some voter type is representative, and we do not initially assume that this type is the same in each state. We give relatively general conditions for “weak core convergence”

results, which state that equilibria in the representative voting game can be replicated in the electoral model; in case a fixed voter type is representative in each state, this means that there is a simple Markov electoral equilibrium that solves the representative voter’s dynamic programming problem. But these results leave the possibility that the electoral model can admit equilibria that do not correspond to equilibria of the representative voting game. We then take up the possibility of “strong core convergence” and provide more restrictive conditions under which full equivalence obtains.

We maintain the following assumptions in the electoral model, in order to formulate the hypothetical game among representative voters:

- (D1) $p_t(s'|s, x, e)$ is independent of t and e ,
- (D2) mixed motives, i.e., $w_t(s, x) = u_t(s, x) + b_t$,
- (D3) $\delta_t > 0$ for all t ,
- (D4) voter type $\kappa(s)$ is representative in s for all s ,
- (D4') the representative voter $\kappa = \kappa(s)$ is independent of s .

Given restrictions (D1)–(D4), we can define the associated *representative voting game* as follows:

- the set of players is $K = \{\kappa(s) : s \in S\}$,
- in state s , player $\kappa(s)$ chooses any $x \in Y_{\kappa(s)}(s)$,
- a new state s' is drawn from $p(\cdot|s, x)$,
- the per-period payoff to player $\tau \in K$ is $u_\tau(s, x)$,

and then $\kappa(s')$ chooses from $Y_{\kappa(s')}$, and so on. As is standard, payoffs from infinite histories are the sum of per-period payoffs discounted by δ_τ , for $\tau \in K$, and a strategy for player $\kappa(s)$ is a profile $\tilde{\pi}_{\kappa(s)} = (\tilde{\pi}_{s,\tau})_{\tau=\kappa(s)}$, where $\tilde{\pi}_{s,\kappa(s)}$ is a probability measure on $Y_{\kappa(s)}(s)$ that represents the mixture over policies by $\kappa(s)$ in state s . We consider stationary Markov perfect equilibria of the game. Under (D4'), when the representative voter type is fixed across states, the representative voting game reduces to a *representative dynamic programming problem* for agent κ . In this case, we drop the subscript on $\tilde{\pi}_\kappa$, and we let $V_\kappa^*(s)$ denote the optimal value function for κ in this problem.

Weak core convergence We first provide a result showing that under conditions (D1)–(D4), every equilibrium of the representative voting game can be replicated in the electoral game. Although we do not obtain full equivalence, the weak core convergence result can be viewed as providing a selection of simple Markov electoral equilibria that may be useful in applications. Specifically, we

show that if $\tilde{\pi}$ is an equilibrium of the representative voting game, then there is a simple Markov electoral equilibrium such that in each state s , office holders of all types t will use the mixed strategy $\tilde{\pi}_{s,\kappa(s)}$ of the representative voter in s . Because the type t politician may have very different preferences than the representative voter, this result may seem surprising. In particular, to induce this mixing, we must devise equilibrium strategies so that the type t politician is indifferent over all policies in the support of $\tilde{\pi}_{s,\kappa(s)}$. We use mixed voting strategies in combination with the assumption of large office benefit to achieve this. However, mixed voting strategies are critical only when the equilibrium $\tilde{\pi}$ is in mixed strategies: when $\tilde{\pi}$ is pure, our construction delivers an electoral equilibrium with pure voting strategies. Logically speaking, preconditions of the theorem are that politicians of all types be able to choose the optimal policies of the representative voters, for which a reasonable sufficient condition is simply that feasible policy sets are independent of politicians' types.

Theorem 4. *In addition to (D1)–(D4), assume:*

- b_t large for all t ,
- $Y_t(s)$ is independent of t for all s .

Let $\tilde{\pi} = (\tilde{\pi}_{\kappa(s)})$ be a stationary Markov perfect equilibrium in the representative voting game. Then there is a simple Markov electoral equilibrium $\sigma = (\pi, \rho)$ such that politicians implement the equilibrium from the voting game: for all s and all t , $\pi_t(\cdot|s) = \tilde{\pi}_{s,\kappa(s)}$. Furthermore, if $\tilde{\pi}$ is in pure strategies, then the corresponding simple Markov electoral equilibrium σ has pure voting strategies.

Proof. Let $\tilde{\pi} = (\tilde{\pi}_{s,\kappa(s)})$ be a stationary Markov perfect equilibrium of the representative voting game. Given state s and a choice of $x \in Y(s)$, the payoff to player $\kappa(s)$ is $u_{\kappa(s)}(s, x) + \delta_{\kappa(s)} \tilde{V}_{\kappa(s)}(s, x)$, where $\tilde{V}_{\kappa(s)}$ satisfies

$$\tilde{V}_{\kappa(s)}(s, x) = \sum_{s'} p(s'|s, x) \int_{x'} [u_{\kappa(s)}(s', x') + \delta_{\kappa(s)} \tilde{V}_{\kappa(s)}(s', x')] \tilde{\pi}_{s'}(dx'). \quad (7)$$

By assumption, $\tilde{\pi}_{s,\kappa(s)}$ places probability one on best response policies for player $\kappa(s)$, i.e., on solutions to

$$\max_{x \in Y(s)} u_{\kappa(s)}(s, x) + \delta_{\kappa(s)} \tilde{V}_{\kappa(s)}(s, x).$$

To define $\sigma = (\pi, \rho)$, we specify that for all s and all t , $\pi_t(\cdot|s) = \tilde{\pi}_{s,\kappa(s)}$. We define the voting strategy in a way that ensures each politician type is indifferent among all policies in the support of $\tilde{\pi}_{s,\kappa(s)}$ and has no profitable deviations outside the support set. To do so, we use the reelection probability to bring the office holder's payoff down to the minimum possible over $\text{supp}(\tilde{\pi}_{s,\kappa(s)})$: in this equilibrium an office holder is reelected with probability one after choosing her

least preferred policies in the support of $\tilde{\pi}_{s,\kappa(s)}$ and we assume that b_t is large to ensure that it is still optimal for this politician to compromise.

Normalize stage utilities so that $0 \leq u_t \leq \bar{u}$. Let $\underline{w}_{s,t}$ represent the minimum expected discounted utility in state s for a type t office holder from choosing a policy in the support of $\tilde{\pi}_{s,\kappa(s)}$ and being reelected (and doing so in all other states s'), so $\underline{w} = (\underline{w}_{s,t}) \in [\frac{b_t}{1-\delta_t}, \frac{\bar{u}+b_t}{1-\delta_t}]^{S \times T}$. Let $\tilde{w}_{s,t}$ represent the expected discounted utility to a type t politician from electing a challenger in state s , so that $\tilde{w} = (\tilde{w}_{s,t}) \in [0, \frac{\bar{u}}{1-\delta_t}]^S$. We identify the minimum payoff of the type t office holder in the support set by the recursion

$$\underline{w}_{s,t} = \min_{x \in \text{supp}(\tilde{\pi}_{s,\kappa(s)})} u_t(s, x) + b_t + \delta_t \sum_{s'} p(s'|s, x) \underline{w}_{s',t}.$$

Note that, as required, $\frac{b_t}{1-\delta_t} \leq \underline{w}_{s,t} \leq \frac{\bar{u}+b_t}{1-\delta_t}$. And we define the challenger payoff to the type t office holder following policy x by the recursion

$$\tilde{w}_{s,t} = \int_x [u_t(s, x) + \delta_t \sum_{s'} p(s'|s, x) \tilde{w}_{s',t}] \tilde{\pi}_{s,\kappa(s)}(dx).$$

Note that, as required, $0 \leq \tilde{w}_{s,t} \leq \frac{\bar{u}}{1-\delta_t}$. In the following, we choose b_t large enough that $\frac{b_t}{1-\delta_t} \geq \frac{\bar{u}}{1-\delta_t} + b_t$, where we use the assumption that $\delta_t > 0$.

To construct reelection probabilities for politicians in different states following different policy choices, we set the expected payoff of a type t politician in state s following policy choice x in the support of $\tilde{\pi}_{s,\kappa(s)}$ equal to $\underline{w}_{s,t}$. Specifically, given $x \in \text{supp}(\tilde{\pi}_{s,\kappa(s)})$, we set $\rho(s, t, x) = r$, where r solves

$$\underline{w}_{s,t} = u_t(s, x) + b_t + \delta_t \sum_{s'} p(s'|s, x) [r \underline{w}_{s',t} + (1-r) \tilde{w}_{s',t}].$$

Note that for all $x \in \text{supp}(\tilde{\pi}_{s,\kappa(s)})$ and all s' , we have

$$\underline{w}_{s',t} \geq \frac{b_t}{1-\delta_t} > \frac{\bar{u}}{1-\delta_t} + b_t \geq u_t(s, x) + b_t + \delta_t \tilde{w}_{s',t},$$

so that $r \in [0, 1]$ is uniquely defined by the above equation, with $r = 1$ if and only if x achieves payoff $\underline{w}_{s,t}$ for politician type t . That is, we define

$$\rho(s, t, x) = \begin{cases} \frac{\underline{w}_{s,t} - (u_t(s, x) + b_t + \delta_t \sum_{s'} p(s'|s, x) \tilde{w}_{s',t})}{\delta_t \sum_{s'} p(s'|s, x) (\underline{w}_{s',t} - \tilde{w}_{s',t})} & \text{if } x \in \text{supp}(\tilde{\pi}_{s,\kappa(s)}) \\ 0 & \text{else.} \end{cases}$$

Thus, the type t office holder is indifferent over policies in the support of $\tilde{\pi}_{s,\kappa(s)}$, and since $\underline{w}_{s,t} > \frac{\bar{u}}{1-\delta_t} + b_t$, there are no profitable deviations for the politician outside the support. Finally, because all politicians use the same policy strategy, voters are indifferent between political candidates, so the voting strategy ρ satisfies our equilibrium condition.

Note that if $\tilde{\pi}$ is in pure strategies, then for each state s and politician type t , $x \in \text{supp}(\tilde{\pi}_{s,\kappa(s)})$ attains $\underline{w}_{s,t}$ and the voting strategy specified by ρ is pure. \square

When the equilibrium of the representative voting game is pure, the voting strategy specified in the second part of Theorem 4 is pure but not deferential: there can be cases in which a representative voter is indifferent between the incumbent and challenger, yet rejects the incumbent. In case the representative voter is fixed across states, however, the representative voting game reduces to a dynamic programming problem, and we can restrict attention to deferential equilibria in which an office holder is reelected if and only if she chooses optimally for the voter. Moreover, we can use pure strategy equilibria in which a politician's policy strategy selects the best policy for the politician from the voter's optimal policies. With a single representative voter type, this ensures the politician's reelection.

Theorem 5. *In addition to (D1)–(D3) and (D4'), assume:*

- b_t large for all t ,
- $Y_t(s)$ is independent of t for all s .

Then there exists a pure optimal policy rule for the representative voter's dynamic programming problem, $\pi^ = (\pi_s^*)$, along with a deferential simple Markov electoral equilibrium $\sigma = (\pi, \rho)$ in pure strategies such that politicians implement this policy rule in the equilibrium: for all s and all t , $\pi_t(\cdot|s) = \pi_s^*$.*

Proof. Fix a state s , and recall that $V_\kappa^*(s)$ is the optimal value of the representative voter's dynamic programming problem in this state. Let

$$X^*(s) = \arg \max_{x \in Y(s)} u_\kappa(x, s) + \delta_\kappa \sum_{s'} p(s'|s, x) V_\kappa^*(s'),$$

be the optimal policies for voter κ in state s . For each politician type t , consider the dynamic program in which the politician chooses policy from $X^*(s)$ in each state and is always reelected. The Bellman equation for this program is

$$V_t^*(s) = \max_{x \in X^*(s)} u_t(s, x) + \delta_t \sum_{s'} p(s'|s, x) V_t^*(s').$$

For each state s , let $x_t^*(s)$ be a selection from the policies solving this program for the type t politician. Define $\sigma = (\pi, \rho)$ so that for all s and all t , $\pi_t(\{x_t^*(s)\}|s) = 1$, and $\rho(s, t, x) = 1$ if $x \in X^*(s)$ and $\rho(s, t, x) = 0$ otherwise. Obviously, σ is a simple Markov strategy profile in pure strategies. Note that $V_\kappa^F(s, t) = V_\kappa^*(s)$ and that, recalling value function \tilde{V}_κ from (7), $V_\kappa^B(s, x) \leq \tilde{V}_\kappa(s, x)$ for all x , with equality if and only if $x \in X^*(s)$. Then for all s , all t , and all x ,

$$\begin{aligned} V_\kappa^B(s, x, t) &= p(s|s, x)[u_\kappa(s, x) + \delta_\kappa V_\kappa^B(s, x, t)] + \sum_{s' \neq s} p(s'|s, x) V_\kappa^F(s', t) \\ &\leq p(s|s, x)[u_\kappa(s, x) + \delta_\kappa \tilde{V}_\kappa(s, x)] + \sum_{s' \neq s} p(s'|s, x) V_\kappa^F(s', t) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{s'} p(s'|s, x) V_{\kappa}^*(s'), \\
V_{\kappa}^C(s, x, t) &= \sum_{t'} q_t(t'|s, x) \sum_{s'} p(s'|s, x) V_{\kappa}^F(s', t) \\
&= \sum_{s'} p(s'|s, x) V_{\kappa}^*(s'),
\end{aligned}$$

so that $V_{\kappa}^B(s, x, t) \leq V_{\kappa}^C(s, x, t)$, with equality if and only if $x \in X^*(s)$. Thus, the voting rule ρ satisfies the conditions for equilibrium, and σ is in fact deferential. By construction, in each state s , no type t politician can deviate profitably from $\pi_t(\cdot|s)$ by choosing a policy in $X^*(s) = R(s, t)$. Given state s and politician type t , normalizing stage utilities so that $0 \leq u_t \leq \bar{u}$, assuming b_t sufficiently large, and using $\delta_t > 0$, we have

$$W_t^B(s, x_t^*(s)) \geq \frac{b_t}{1 + \delta_t} \geq b_t + \frac{\bar{u}}{1 - \delta_t} \geq W_t^C(s, x),$$

for all $x \in Y(s) \setminus X^*(s)$, fulfilling the optimality condition for politicians. \square

Strong core convergence Whereas Theorem 4 allows the representative voter to vary with the state and establish weak core convergence results — that every equilibrium of the representative voting game can be replicated in the electoral model — we now assume that the representative voter type is constant across states, and we establish the stronger conclusion that every equilibrium of the electoral model achieves (or almost achieves) this voter’s payoff in the representative voting game. Thus, all simple Markov electoral equilibria provide (close to) the same payoff for the representative voter, and all politicians (almost) solve the representative voter’s dynamic programming problem.

We first assume policy-motivated politicians and provide an asymptotic result: as voters and politicians become patient, the representative voter’s optimal payoff is approached in a strong sense. The logic of this is the following. By Lemma 2, equilibrium electoral outcomes solve the representative voter’s optimal retention problem. One voting strategy, which may not be optimal, is to simply retain any type κ office holder and reject all other types. Since type κ voters and politicians have perfectly aligned preferences, the latter’s policy choices are solutions to the representative dynamic programming problem. Intuitively, as the type κ voter and politician become arbitrarily patient, the loss from this strategy becomes negligible. The equilibrium voting strategy can do no worse than this simple rule, and therefore the representative voter’s payoffs approach the optimal level. Two assumptions are key for the representative voter to implement the previous voting strategy. First, any given state must recur under any profile of policy strategies by politicians. Second, the probability that a politician of type κ is selected as the challenger must be uniformly bounded away from zero.

We define a state s to be *strongly recurrent* if, starting from s , the probability of returning to s is equal to one, regardless of implemented policies. More formally, for all m , let

$$\Psi^m(s) = \{(s_0, \dots, s_m) \mid s_0 = s = s_j \text{ for some } j = 1, \dots, m\}$$

be the set of paths $\mathbf{s} = (s_0, \dots, s_m)$ of states of length $m + 1$ such that s recurs at least once. For all sequences $\mathbf{s} = (s_0, \dots, s_m)$ of states, let

$$\Phi^m(\mathbf{s}) = \left\{ (x_0, \dots, x_m) \mid \begin{array}{l} x_0 \in Y(s_0) \text{ and } x_j \in Y(s_j) \\ \text{for all } j = 1, \dots, m \end{array} \right\}$$

be the set of feasible paths of policies. Then define

$$p^m(s) = \sum_{\mathbf{s} \in \Psi^m(s)} \min \left\{ \prod_{j=1}^m p(s_j | s_{j-1}, x_{j-1}) \mid \mathbf{x} \in \Phi^m(\mathbf{s}) \right\}$$

as the minimum probability that s is realized within m periods of s being previously realized. Finally, strong recurrence of s means that $\lim_{m \rightarrow \infty} p^m(s) = 1$.

Theorem 6. *Consider any strongly recurrent state s . In addition to (D1)–(D3) and (D4'), assume:*

- *policy motivation, i.e., $b_t = 0$ for all t ,*
- $\min_{t,x} q_t(\kappa | s, x) > 0$.

Let $\delta_\kappa = \delta \rightarrow 1$, and let $\{\sigma^\delta\}$ be corresponding simple Markov electoral equilibria. Then for all t ,

$$\lim_{\delta \rightarrow 1} \frac{V_\kappa^{F,\delta}(s, t)}{V_\kappa^{*,\delta}(s, \kappa)} \geq 1.$$

Proof. Fix δ , let $\sigma^\delta = (\pi^\delta, \rho^\delta)$ be a simple Markov electoral equilibrium given δ , let $\pi_\kappa^{*,\delta}$ denote an optimal strategy for the representative voter in the dynamic programming problem, and let $\hat{\sigma}^\delta = (\hat{\pi}^\delta, \hat{\rho})$ denote the equilibrium σ^δ modified as follows: for all s and all $t \neq \kappa$, let $\hat{\pi}_t^\delta(\cdot | s) = \pi_t^\delta(\cdot | s)$, let $\hat{\pi}_\kappa^\delta(\cdot | s) = \pi_\kappa^{*,\delta}(s)$, and specify $\hat{\rho}(s, t, x) = 1$ if $t = \kappa$ and $\hat{\rho}(s, t, x) = 0$ otherwise. Letting $\hat{V}_\kappa^{F,\delta}(s, t)$ denote the expected discounted payoff to the type κ voter from electing a free type t politician in state s given strategy profile $\hat{\sigma}^\delta$, payoff achieves the voter's optimum: $\hat{V}_\kappa^{F,\delta}(s, \kappa) = V_\kappa^{*,\delta}(s, \kappa)$.

The profile $\hat{\sigma}^\delta$ may not itself be an equilibrium, but it must be the case that for all states s , $V_\kappa^{F,\delta}(s, \kappa) \geq \hat{V}_\kappa^{F,\delta}(s, \kappa)$. Otherwise, because the payoffs of the representative voter κ and the type κ politicians are aligned, a joint deviation to their strategies under $\hat{\sigma}$ is mutually profitable. In turn, by the principle of optimality, this contradicts either the optimality of reelection choices for the

representative voter, from Lemma 2, or the policy choices of politicians of type κ under σ^δ . Indeed, suppose the inequality is violated at some state. Note that since payoffs are aligned between the type κ voter and politicians, we have

$$W_\kappa^{B,\delta}(s, x) = u_\kappa(s, x) + \delta[\rho(s, \kappa, x)V_\kappa^{B,\delta}(s, \kappa, x) + (1 - \rho(s, \kappa, x))V_\kappa^{C,\delta}(s, \kappa, x)]$$

and $W_\kappa^{C,\delta}(s, x) = V_\kappa^{C,\delta}(s, \kappa, x)$ for all s and all x , given equilibrium σ^δ . By supposition, σ^δ is not jointly optimal for the type κ voter and politicians, so there exist s' and x' such that either

$$\begin{aligned} & u_\kappa(s', x') + \delta[\rho(s', \kappa, x')V_\kappa^{B,\delta}(s', \kappa, x') + (1 - \rho(s', \kappa, x'))V_\kappa^{C,\delta}(s', \kappa, x')] \\ & > \int_x \left[u_\kappa(s', x) + \delta[\rho(s', \kappa, x)V_\kappa^{B,\delta}(s', \kappa, x) \right. \\ & \quad \left. + (1 - \rho(s', \kappa, x))V_\kappa^{C,\delta}(s', \kappa, x)] \right] \pi_\kappa^\delta(dx|s'), \end{aligned}$$

or $V_\kappa^{C,\delta}(s', \kappa, x') > V_\kappa^{B,\delta}(s', \kappa, x')$. But in the first case, we have

$$W_\kappa^{B,\delta}(s', x') > \max_{x \in X_\kappa(s')} \rho(s', \kappa, x)W_\kappa^{B,\delta}(s', x) + (1 - \rho(s', \kappa, x))W_\kappa^{C,\delta}(s', x),$$

contradicting optimality of π_κ^δ for the politician, and in the second case, we have

$$\begin{aligned} \bar{V}_\kappa^{B,\delta}(s', \kappa, x') &= \rho(s', \kappa, x')V_\kappa^{B,\delta}(s', \kappa, x') + (1 - \rho(s', \kappa, x'))V_\kappa^{C,\delta}(s', \kappa, x') \\ &< V_\kappa^{C,\delta}(s', \kappa, x') \\ &\leq \bar{V}_\kappa^{C,\delta}(s', \kappa, x'), \end{aligned}$$

where the equality follows from Lemma 2, a contradiction. This establishes the claim.

We have shown that for all s , $V_\kappa^{F,\delta}(s, \kappa) \geq \hat{V}^{F,\delta}(s, \kappa) = V_\kappa^{*,\delta}(s, \kappa)$. Given any strongly recurrent state s , let $\alpha = \min_{t,x} q_t(\kappa|s, x) > 0$, and note that regardless of policy choices, the probability that a type κ politician is drawn within m periods is at least $1 - (1 - \alpha)^m$. Then, given equilibrium σ^δ and using the normalization $u_\kappa \geq 0$, the representative voter's expected discounted utility from electing a free type t politician satisfies

$$V_\kappa^{F,\delta}(s, t) \geq p^m(s)(1 - (1 - \alpha)^m)\delta^m V_\kappa^{*,\delta}(s, \kappa),$$

or equivalently,

$$\frac{V_\kappa^{F,\delta}(s, t)}{V_\kappa^{*,\delta}(s, \kappa)} \geq p^m(s)(1 - (1 - \alpha)^m)\delta^m$$

for all m . Given $\epsilon > 0$, we can choose $m(\epsilon)$ sufficiently high that $p^{m(\epsilon)}(s)(1 - (1 - \alpha)^{m(\epsilon)}) \geq 1 - \epsilon$. Then, taking limits, we have

$$\sup_m p^m(s)(1 - (1 - \alpha)^m)\delta^m \geq (1 - \epsilon)\delta^{m(\epsilon)} \rightarrow 1 - \epsilon$$

as $\delta \rightarrow 1$. Since ϵ was arbitrary, the desired inequality follows. \square

The previous result relies on politicians of type κ implementing optimal policies in the representative voter's dynamic programming problem. Perhaps surprisingly, difficulties arise if we allow mixed motivation: the following example of the “curse of ambition” demonstrates how the result of Theorem 6 can fail if type κ politicians value office. This may be unintuitive, because these politicians share the representative voter's policy preferences and, in light of the benefits of office, they will be all the more willing to compromise to ensure reelection if needed. In some equilibria, however, retaining office may require implementing policies detrimental to both politicians and voters of type κ . One may think that restricting attention to deferential equilibria would allow for a weaker version of Theorem 6, since in such equilibria type κ office holders can ensure reelection in any state by Corollary 1. However, since the following example exhibits a deferential equilibrium in which type κ voters have a strict incentive to replace type κ politicians following κ -optimal actions in some states, it also shows that such a weaker results fails.

Example 4 (Curse of ambition). Let the state space be $S = \{s_1, s_{-1}, \underline{s}\}$ and the type space be $T = \{1, \kappa\}$, with type κ voters representative in all states. Feasible policies are independent of states and politicians' types and are given by $Y = \{x_1, x_{-1}\}$. Transition probabilities are independent of politicians' types and are given by $p(s_{-1}|s_{-1}, x_{-1}) = p(s_{-1}|\underline{s}, x_{-1}) = p(s_{-1}|s_1, x_{-1}) = p(s_1|s_{-1}, x_1) = p(s_1|\underline{s}, x_1) = 1$, $p(s_1|s_1, x_1) = 1 - p$ and $p(\underline{s}|s_1, x_1) = p \in (0, 1)$. Challenger selection probabilities are independent of states, policies, and incumbents' types, and they are given by $q(\kappa) = p$ and $q(1) = 1 - p$. The stage utilities of type κ voters are independent of states and politicians' types and are such that $u_\kappa(x_1) > u_\kappa(x_{-1})$. Politicians have mixed motivations with type-independent office benefit $b > 0$. The stage utilities of type 1 politicians are such that $u_1(s_1, x_1) = u_1(\underline{s}, x_1) = u_1(s_{-1}, x_{-1}) > u_1(s_1, x_{-1}) = u_1(\underline{s}, x_{-1}) = u_1(s_{-1}, x_1)$. Note that, in this example, all the conditions of Theorem 6 are respected other than the requirement that $b = 0$. Furthermore, for all states s , policy x_1 is optimal for the representative voter:

$$V_\kappa^*(s) = \frac{u_\kappa(x_1)}{1 - \delta_\kappa}.$$

We show that given any δ_κ , there exists $b^* > 0$ such that for all $b \geq b^*$, there exists a deferential simple Markov equilibrium in which type κ politicians implement policy x_{-1} in all states, type 1 politicians implement their stage-ideal policy in all states and all politicians are reelected in all states on the equilibrium path. Note that in this equilibrium, for any state s ,

$$V_\kappa^F(s, \kappa) = \frac{u_\kappa(x_{-1})}{1 - \delta_\kappa} < V_\kappa^*(s).$$

Implementing policy x_{-1} in any state leads to a transition to state s_{-1} with probability 1. In this state, both types of politicians implement policy x_{-1} , ensuring that the state remains in s_{-1} hereafter. Hence, following policy x_{-1}




Figure 3: Curse of ambition

in any state, type κ voters are indifferent between reelecting the incumbent and opting for the challenger and since the equilibrium is deferential, they opt for the incumbent. See Figure 3 for the equilibrium diagram, where darker arrows represent type 1 and lighter represent type κ .

In no state can politicians of type κ be reelected by implementing policy x_1 . In states s_{-1} and \underline{s} , implementing policy x_1 leads to a transition to state s_1 with probability 1. However, in that state, politicians of type κ are expected to implement policy x_{-1} , whereas politicians of type 1 are expected to implement policy x_1 , leading to strict incentives for voters of type κ to opt for the challenger. In both states s_2 and \underline{s} , an incumbent of type κ is replaced if she chooses policy x_1 , because she cannot commit to implementing her ideal policy in the future, because transitions occur with probability one. In state s_1 , policy

x_1 leads to s_1 or \underline{s} . However, since in state \underline{s} a type κ politician implements policy x_{-1} while a type 1 politician implements policy x_1 , voters of type κ have a strict incentive to replace politicians of type κ who have implemented policy x_1 in state s_1 . Finally, type κ politicians will implement policy x_{-1} in all states in order to be reelected if they value office enough. \square

The previous strong core convergence result was asymptotic in nature, showing that the representative voter's optimal payoff is nearly achieved when the voter is sufficiently patient. We state a final strong core convergence theorem, focusing on deferential equilibria, that holds independently of the players' discount factors. To obtain the stronger equivalence, we assume large office benefit, ensuring that all politicians choose to compromise, and, importantly, we add the assumption that the state transition is independent of policy choice. Also, since the previous result provided only a lower bound on the representative voter's equilibrium payoffs, it left open the possibility that politicians of type other than κ could implement policies yielding payoffs higher than those generated by solutions to the representative dynamic programming problem. Here, we assume that feasible policies are independent of politicians' types and give an exact result.

Define a state s as *non-trivial* if there exist s' , t and x such that $p_t(s|s', x) > 0$. Our next result is restricted to non-trivial states. However, this condition is essentially unrestrictive, as a state s that fails to satisfy it can occur only if it is the initial state of the game, after which the state transitions away from s with probability one and visits only non-trivial states thereafter.

Theorem 7. *Let σ be a deferential simple Markov electoral equilibrium. In addition to (D1)–(D3) and (D4'), assume*

- b_t large for all t ,
- $Y_t(s)$ is independent of t for all s ,
- $p(s'|s, x)$ is independent of x for all s and all s' ,
- $\inf_{s,t,x} q_t(\kappa|s, x) > 0$.

Then $V_\kappa^F(s, t) = V_\kappa^*(s)$ for all non-trivial s and all t .

Proof. Since the equilibrium is deferential and b_t is large for all t , Corollary 1 implies that for all s , all t , and all $x \in \text{supp}(\pi_t(\cdot|s))$, we have $\pi_t(R(s, t)|s) = 1$ and $\rho(s, t, x) = 1$. Furthermore, we claim that $V_\kappa^F(s, \kappa) = V_\kappa^*(s)$ for all s . First, note that the Bellman equation for the representative voter simplifies to

$$V_\kappa^*(s) = \max_{x \in Y(s)} u_\kappa(s, x) + \delta_\kappa \sum_{s'} p(s'|s) V_\kappa^*(s'),$$

so that π^* is optimal for the representative voter's optimization problem if and only if $\pi^*(\arg \max_{x \in Y(s)} u_\kappa(s, x)|s) = 1$. Second, note that for all $x \in \text{supp}(\pi_\kappa(\cdot|s))$,

$$W_\kappa^B(s, x) = \frac{\max_{x' \in Y(s)} u_\kappa(s, x') + b_\kappa + \delta_\kappa \sum_{s' \neq s} p(s'|s) \int_{x''} W^B(s', x'') \pi_\kappa(dx''|s')}{1 - \delta_\kappa p(s|s)},$$

where we use the fact that $\rho(s, \kappa, x) = 1$. Hence, $\pi_\kappa(\arg \max_{x \in Y(s)} u_\kappa(s, x)|s) = 1$, and $V_\kappa^F(s, \kappa) = V_\kappa^*(s)$, as claimed.

Next, note that for all states s and policies $x \in \text{supp}(\pi_\kappa(\cdot|s))$, we have $x \in \text{supp}(\pi^*(\cdot|s))$ and $\rho(s, \kappa, x) = 1$, and it follows that $u_\kappa(s, x) + \delta_\kappa V_\kappa^B(s, \kappa, x) = V_\kappa^*(s)$. Hence, for all $x \in \text{supp}(\pi_\kappa(\cdot|s))$,

$$\begin{aligned} V_\kappa^B(s, \kappa, x) - \sum_{s'} p(s'|s) V_\kappa^F(s', \kappa) &= \\ &= p(s|s) [u_\kappa(s, x) + \delta_\kappa V_\kappa^B(s, \kappa, x) - V_\kappa^F(s, \kappa)] \\ &= p(s|s) [V_\kappa^*(s) - V_\kappa^F(s, \kappa)] \\ &= 0. \end{aligned} \tag{8}$$

Now, given any state s , suppose there exist politician type \tilde{t} and policy $\tilde{x} \in \text{supp}(\pi_{\tilde{t}}(\cdot|s))$ such that

$$\begin{aligned} V_\kappa^B(s, \tilde{t}, \tilde{x}) &= \min \left\{ V_\kappa^B(s, t, x) : (t, x) : t \in T, x \in \text{supp}(\pi_t(\cdot|s)) \right\} \\ &< \sum_{s'} p(s'|s) V_\kappa^F(s', \kappa). \end{aligned} \tag{9}$$

Then, using the fact that $\rho(s, t, x) = 1$ for all t' and all $x \in \text{supp}(\pi_{t'}(\cdot|s))$, we have

$$\begin{aligned} V_\kappa^C(s, \tilde{t}, \tilde{x}) &= \sum_{t'} q_{\tilde{t}}(t'|s, \tilde{x}) \left[p(s|s) V_\kappa^F(s, t') + \sum_{s' \neq s} p(s'|s) V_\kappa^F(s', t') \right] \\ &= \sum_{t'} q_{\tilde{t}}(t'|s, \tilde{x}) \left[p(s|s) \int [u_\kappa(s, t', x) + \delta V_\kappa^B(s, t', x)] \pi_{t'}(dx|s) \right. \\ &\quad \left. + \sum_{s' \neq s} p(s'|s) V_\kappa^F(s', t') \right] \\ &= \sum_{t'} q_{\tilde{t}}(t'|s, \tilde{x}) \int \left[p(s|s) [u_\kappa(s, t', x) + \delta V_\kappa^B(s, t', x)] \right. \\ &\quad \left. + \sum_{s' \neq s} p(s'|s) V_\kappa^F(s', t') \right] \pi_{t'}(dx|s) \end{aligned}$$

$$\begin{aligned}
&= \sum_{t'} q_{\tilde{t}}(t'|s, \tilde{x}) \int V_{\kappa}^B(s, t', x) \pi_{t'}(dx|s) \\
&> V_{\kappa}^B(s, \tilde{t}, \tilde{x}),
\end{aligned}$$

contradicting $\rho(s, \tilde{t}, \tilde{x}) = 1$, with the inequality following from (9) and $q_{\tilde{t}}(\kappa|s, \tilde{x}) > 0$. Therefore, it must be that for all states s , types t , and policies $x \in \text{supp}(\pi_{\kappa}(\cdot|s))$ and $x' \in \text{supp}(\pi_t(\cdot|s))$, we have

$$V_{\kappa}^B(s, t, x') = \sum_{s'} p(s'|s) V_{\kappa}^F(s', \kappa) = V_{\kappa}^B(s, \kappa, x), \quad (10)$$

where the second equality follows from (8).

Finally, we claim that $V_{\kappa}^F(s, t) = V_{\kappa}^*(s)$ for all s and t . Indeed, note that $V_{\kappa}^F(s, t) \leq V_{\kappa}^*(s) = V_{\kappa}^F(s, \kappa)$. Suppose, toward a contradiction, that $V_{\kappa}^F(s, t) < V_{\kappa}^F(s, \kappa)$, and, since s is non-trivial, fix \tilde{s} such that $p(s|\tilde{s}) > 0$. Then, for all $x \in \text{supp}(\pi_{\kappa}(\cdot|\tilde{s}))$,

$$\begin{aligned}
&V_{\kappa}^B(\tilde{s}, \kappa, x) \\
&= \sum_{s'} p(s'|\tilde{s}) V_{\kappa}^F(s', \kappa) \\
&> \sum_{s'} p(s'|\tilde{s}) V_{\kappa}^F(s', t) \\
&= p(\tilde{s}|\tilde{s}) \int [u_{\kappa}(\tilde{s}, x') + \delta_{\kappa} V_{\kappa}^B(\tilde{s}, t, x')] \pi_t(dx'|\tilde{s}) + \sum_{s' \neq \tilde{s}} p(s'|\tilde{s}) V_{\kappa}^F(s', t) \\
&= \int V_{\kappa}^B(\tilde{s}, t, x') \pi_t(dx'|\tilde{s}),
\end{aligned}$$

where the first equality follows from (8) and the inequality by supposition. Therefore, there exists $x' \in \text{supp}(\pi_t(\cdot|\tilde{s}))$ such that $V_{\kappa}^B(\tilde{s}, \kappa, x) > V_{\kappa}^B(\tilde{s}, t, x')$, contradicting (10). \square

A significant limitation of the preceding result is that it assumes the state transition is independent of policy choices, precluding many interesting forms of dynamic linkage across periods. The next example illustrates that when the state transition depends on policy, simple Markov equilibria may fail to implement the optimal policies of a voter type that is representative in all states, even when politicians are arbitrarily office motivated: when the players are relatively impatient (so that Theorem 6 does not apply), it is possible that there exist deferential equilibria in which some politician types choose extreme policies (bounded away from the representative voter's optimum) yet are continually reelected. In particular, if an extreme type of politician holds office in the first period, then she will remain in office and choose an extreme policy in all future states, despite the fact that the representative voter could elect a challenger, who is type κ with positive probability. Thus, our model can explain a new type of political inefficiency that arises from the dynamic incentives of elections.

Example 5 (Dynamic policy extremism). Let the state space be $S = \{s_{-1}, s_1\}$ and the type space be $T = \{-1, \kappa, 1\}$, with type κ voters representative in all states, and let t range over $\{-1, 1\}$. Feasible policies are independent of states and politicians' types and are given by $\{\hat{x}_{-1}, \hat{x}_1, \hat{x}_\kappa\}$. Transition probabilities are independent of incumbents' types and are such that $p(s_t|s_t, \hat{x}_t) = p(s_t|s_t, \hat{x}_\kappa) = p(s_{-t}|s_t, \hat{x}_{-t}) = 1$. Challenger selection probabilities are independent of states, policies and incumbent types and are such that $q(\kappa) = q \in (0, 1)$ and $q(t) = \frac{1}{2}(1 - q)$. The stage utilities to voter κ satisfy $u_\kappa(s_t, \hat{x}_\kappa) > u_\kappa(s_t, \hat{x}_t) > u_\kappa(s_t, \hat{x}_{-t})$. Politicians have mixed motivations with type-independent office benefit $b \geq 0$, and the stage utilities of t -types are independent of states and satisfy $u_t(\hat{x}_t) > u_t(\hat{x}_\kappa) > u_t(\hat{x}_{-t})$. Note that in this example, all conditions of Theorem 7 are respected, other than the requirement that state transitions be independent of policies. Furthermore, the representative voter's stage ideal policy is optimal for the voter in both states:

$$V_\kappa^*(s_t) = \frac{u_\kappa(\hat{x}_\kappa)}{1 - \delta_\kappa}.$$

We assume that q and δ_κ jointly satisfy

$$\begin{aligned} q &\leq \frac{(1 - \delta_\kappa) [u_\kappa(s_t, \hat{x}_t) - u_\kappa(s_t, \hat{x}_{-t})]}{2u_\kappa(s_t, \hat{x}_\kappa) - (1 + \delta_\kappa)u_\kappa(s_t, \hat{x}_t) + (1 - \delta_\kappa)u_\kappa(s_t, \hat{x}_{-t})} \\ &\in (0, 1). \end{aligned} \quad (11)$$

We show that there exists a simple Markov electoral equilibrium in which all politicians choose their stage-ideal policy and are always reelected. Note that in this equilibrium,

$$V_\kappa^F(s_t, t) = \frac{u_\kappa(s_t, \hat{x}_t)}{1 - \delta_\kappa} < V_\kappa^*(s_t).$$

Voters of type κ support incumbents of type $t \in \{-1, 1\}$ in state s_t following policy \hat{x}_t , since the state persists with probability one and opting for the challenger entails a sufficiently high probability of electing a politician of type $-t$, who implements policy \hat{x}_{-t} (and is subsequently reelected, since the state transitions to s_{-t}). See Figure 4 for the equilibrium diagram, where dark arrows represent type -1 , medium represent type κ , and light represent type 1 .

An equilibrium of this type exists as long as q , the probability that a challenger is of type κ , or discount factor δ_κ are sufficiently small (i.e., (11) is fulfilled). To reconcile this example with Theorem 6, suppose that $b = 0$. Suppose further that

$$\frac{1 - \delta_t}{\delta_t} [u_t(\hat{x}_t) - u_t(\hat{x}_\kappa)] > [u_t(\hat{x}_\kappa) - u_t(\hat{x}_{-t})],$$

so that in every equilibrium, t -type politicians choose policy \hat{x}_t in all states. If (11) holds, then the profile described above is an equilibrium. However, if (11)




Figure 4: Dynamic policy extremism

fails, then it is not optimal in state s_t for the type κ voter to reelect type t politicians following \hat{x}_t or type $-t$ politicians following \hat{x}_{-t} . As shown in the proof of Theorem 6, when $b = 0$, type κ politicians choose policy \hat{x}_κ in all states in any equilibrium. Therefore, under these conditions there exists a unique equilibrium in which the payoffs to the type κ voter satisfies

$$\begin{aligned}
& \lim_{\delta_\kappa \rightarrow 1} \frac{V_\kappa^{F, \delta_\kappa}(s_t, t)}{V^*, \delta_\kappa(s_t)} \\
&= \lim_{\delta_\kappa \rightarrow 1} \frac{[u_\kappa(s_t, \hat{x}_t) + \delta_\kappa V_\kappa^C(s_t, \hat{x}_t)]}{V^*, \delta_\kappa(s_t)} \\
&= \lim_{\delta_\kappa \rightarrow 1} \frac{[u_\kappa(s_t, \hat{x}_t) + \frac{\delta_\kappa}{1 - \delta_\kappa(1-q)} \left[\frac{qu_\kappa(s_t, \hat{x}_\kappa)}{1 - \delta_\kappa} + \frac{1-q}{2} [u_\kappa(s_t, \hat{x}_t) + u_\kappa(s_t, \hat{x}_{-t})] \right]]}{\frac{u_\kappa(s_t, \hat{x}_\kappa)}{1 - \delta_\kappa}} \\
&= 1,
\end{aligned}$$

in conformity with the conclusion of Theorem 6. Hence, although politicians of type t implement their ideal policy, they are successively replaced until a politician of type κ attains and retains office. \square

A Existence and Continuity Proof

This appendix consists of the proof of Theorems 1 and 2. The fixed point argument, which follows Duggan (2011), will take place in a space consisting of

the product of policy strategies, ex ante expected payoffs for voters, and ex ante expected payoffs for politicians. Normalize utilities so that the images of $\frac{u_t}{1-\delta_t}$ and $\frac{w_t}{1-\delta_t}$ lie in $[0, \bar{u}]$ 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 $v_{s,t,\tau}, w_{s,t} \in [0, \bar{u}]$ 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 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 probability measures endowed with the weak* topology.

Define the nonempty, convex product space

$$\Theta = (\Delta(X)^{S \times T}) \times ([0, \bar{u}]^{S \times T}) \times ([0, \bar{u}]^{S \times T \times T}),$$

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 $[0, \bar{u}]$ in the real line with the Euclidean topology. Then the product topology on $(\mathcal{M}^{S \times T}) \times (\mathfrak{R}^{S \times T}) \times (\mathfrak{R}^{S \times T \times T})$ 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 $\mathcal{F}: \Theta^+ \rightrightarrows \Theta$ 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 simple Markov electoral equilibrium in the model parameterized by γ ; and conversely, each simple 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 $x \in Y$,

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

or equivalently,

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

and for all $x \in Z$, we adopt the convention that $V_\tau^B(s, t, x, \theta^+) = V_\tau^C(s, t, x, \theta^+)$. 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:

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

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 $x \in Y$,

$$\begin{aligned} W_t^B(s, x, \theta^+) &= w_t(s, x, \gamma) + \delta_t p_t(s|s, x, 1, \gamma) W_t^B(s, x, \theta^+) \\ &\quad + \delta_t \sum_{s' \neq s} p_t(s'|s, x, 1, \gamma) w_{s', t}, \end{aligned} \quad (12)$$

or equivalently,

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

and for all $x \in Z$,

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

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

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

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

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

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

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

is well-defined by nonemptiness and compactness of $Z_t(s)$ and continuity of $W_t^B(s, \cdot, \theta^+)$; 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

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

is lower semi-continuous. Now define

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

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 $W_t^B(s, \cdot)$ and lower semi-continuity of f_t , $\hat{\mathcal{P}}_t(s, \cdot)$ has closed graph in $\Theta^+ \times X$. Define $\mathcal{P}: \Theta^+ \rightrightarrows \Delta(X)^{S \times T}$ by

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

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

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

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

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

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

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

is compact. To see that $\hat{\Theta}_t(s) \neq \emptyset$, choose any $\theta^+ = (\pi, w, v, \gamma)$ such that $\pi_{s,t}$ puts probability one on an outcome that maximizes $W_t^B(s, x, \theta^+)$ 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 $\mathcal{W}: \Theta^+ \rightrightarrows [0, \bar{u}]^{S \times T}$ by

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

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 $y \in R_t(s, \theta^+) \setminus P_t(s, \theta^+)$, for then the equilibrium conditions on voting strategies impose no restrictions on the probability of reelection. Specifically, we use the observation that if $y \in \text{supp}(\pi_{s,t})$, 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 $y \in \text{supp}(\pi_{s,t})$,

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

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

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

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

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

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

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

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

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., $W_t^B(s, x, \theta^+) \neq W_t^B(s, \xi(x), \theta^+)$. It is unrestricted if she chooses a policy $x \in Y$ such that she is indifferent between winning or losing following x , e.g., $W_t^B(s, x, \theta^+) = W_t^B(s, \xi(x), \theta^+)$.

Moreover, \mathcal{R}_t has closed graph because ρ_t and W_t^B are continuous (using the convention that $Y \cap Z = \emptyset$). 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 simple 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 $x \in Y \setminus R_t(s, \theta^+)$. 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 $v' = (v'_{s,t,\tau})_\tau \in [0, \bar{u}]^T$ of induced expected discounted utilities defined by

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

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 $\Theta^+ \times [0, \bar{u}]^T$ 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 $([0, \bar{u}]^{S \times T}) \times ([0, \bar{u}]^{S \times T \times T}) \times \Gamma$, 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

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

for all $x \in X$, all $y = (w, v, \gamma) \in Y$, and all $r \in [0, 1]$.¹⁵ Let the correspondence F consist of integrals of f with respect to $\mu = \pi_{s,t}$, i.e.,

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

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 $\mathcal{V}: \Theta^+ \rightrightarrows [0, \bar{u}]^{S \times T \times T}$ by

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

¹⁵Note 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 .

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 simple 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 $[0, \bar{u}]^{S \times T} \times [0, \bar{u}]^{S \times T \times T}$. 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 3. *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 simple Markov electoral equilibrium $\sigma^* = (\pi^*, \rho^*)$ of the game parameterized by γ such that for all s and all t ,*

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

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

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 $\text{supp}(\pi_{s,t}) \subseteq \hat{P}_t(s, \theta^+)$, and therefore $f_t(s, \theta^+) \leq g_t(s, \theta^+)$. 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 $x \in \text{supp}(\pi_{s,t})$, we have $W_t^B(s, x, \theta^+) \geq w_{s,t} \geq f_t(s, \theta^+)$. 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 :

$$w_{s,t} = \hat{\rho}_t(s, x, \theta^+) W_t^B(s, x, \theta^+) + (1 - \hat{\rho}_t(s, x, \theta^+)) W_t^B(s, \xi(x), \theta^+). \quad (13)$$

Indeed, consider $x \in \text{supp}(\pi_{s,t})$. If $x \in Y$ and $W_t^B(s, x, \theta^+) \neq W_t^B(s, \xi(x), \theta^+)$, 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 $W_t^B(s, x, \theta^+) = W_t^B(s, \xi(x), \theta^+)$, 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 (13) reduces to $W_t^B(s, x, \theta^+)$. We have noted that $W_t^B(s, x, \theta^+) \geq w_{s,t} \geq f_t(s, \theta^+)$, and furthermore, $f_t(s, \theta^+) \geq \bar{Z}_t(s, \theta^+) \geq W_t^B(s, x, \theta^+)$. Combining these two inequalities, we have $w_{s,t} = W_t^B(s, x, \theta^+)$, as claimed.

To construct a simple 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 $x \in Y \setminus R_t(s, \theta^+)$. 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 simple 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 other special conditions. Since we consider $x \in P_t(s, \theta^+)$, we have $f_t(s, \theta^+) \geq \bar{W}_t(s, \theta^+) \geq W_t^B(s, x, \theta^+)$. But if $x \in \text{supp}(\pi_{s,t})$, then we have noted that $W_t^B(s, x, \theta^+) \geq w_{s,t} \geq f_t(s, \theta^+)$. Combining these inequalities, we have $W_t^B(s, x, \theta^+) = w_{s,t}$. Thus, $W_t^B(s, x, \theta^+) > W_t^B(s, \xi(x), \theta^+)$ 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 $W_t^B(s, x, \theta^+) \leq W_t^B(s, \xi(x), \theta^+)$, and since $x \in \text{supp}(\pi_{s,t})$, we also have

$$W_t^B(s, x, \theta^+) \geq f_t(s, \theta^+) \geq \bar{Z}_t(s, \theta^+) \geq W_t^B(s, \xi(x), \theta^+).$$

Combining these inequalities, we see that the problem described above can only arise if $W_t^B(s, x, \theta^+) = W_t^B(s, \xi(x), \theta^+)$, 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 chooses $\xi(x)$ instead of x .

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

Case 3: The policy choice belongs to $X \setminus R_t(s, \theta^+)$. We specify $\rho^*(s, t, x) = 0$. It is possible that $\hat{\rho}_t(s, x, \theta^+) > 0$ for $x \in Y \setminus R_t(s, \theta^+)$, 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 $\pi_{s,t}(Y \setminus R_t(s, \theta^+)) = 0$, 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 $W_t^B(s, x, \theta^+) = W_t^B(s, \xi(x), \theta^+)$, 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

$A \subseteq X$,

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

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

$$\begin{aligned} v_{s,t,\tau} &= \int_x [\rho^*(s, t, x) \hat{V}_\tau(s, t, x, \theta^*) \\ &\quad + (1 - \rho^*(s, t, x)) \hat{V}_\tau(s, t, \xi(x), \theta^*)] \pi_t^*(dx|s) \end{aligned} \quad (14)$$

$$w_{s,t} = \rho^*(s, t, x) W_t^B(s, x, \theta^+) + (1 - \rho^*(s, t, x)) W_t^B(s, \xi(x), \theta^+) \quad (15)$$

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

By construction, and using the expression in (14), the values $V_\tau^B(\cdot, \theta^+)$, $V_\tau^C(\cdot, \theta^+)$, and $\{v_{s,t,\tau}\}_{s,t}$ fulfill the recursive conditions (1)–(3), that uniquely define $V_\tau^B(\cdot; \sigma^*)$, $V_\tau^C(\cdot; \sigma^*)$, and $V_\tau^F(\cdot; \sigma^*)$ in the model parameterized by γ . Furthermore, substituting (15) into (12), the values $W_t^B(\cdot, \theta^+)$ fulfill the recursive condition (4) that uniquely defines $W_t^B(\cdot; \sigma^*)$. Therefore,

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

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

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

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

Next, we argue that the simple Markov strategy profile $\sigma^* = (\pi^*, \rho^*)$ satisfies the conditions for equilibrium. Indeed, ρ^* clearly satisfies condition (ii) in the definition of simple 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 $x \in X$,

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

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

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

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

$$w_{s,t} \geq f_t(s, \theta^+) \geq \bar{Z}_t(s, \theta^+) \geq W_t^B(s, \xi(x); \sigma^*),$$

as required.

For the “if” direction of the lemma, consider a simple Markov electoral equilibrium σ^* of the game parameterized by γ satisfying conditions of Lemma 3, so that for all s and all t , we have

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

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 $w_{s,t} \geq W_t^B(s, z; \sigma^*)$ 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 $x \in Y \setminus R(s, t; \sigma^*)$, i.e., we specify that

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

for all Borel measurable $A \subseteq Y$.

To establish that $(\pi, w, v) \in \mathcal{F}(\pi, w, v, \gamma)$, define $\hat{\rho}_t: S \times X \rightarrow [0, 1]$ 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 $W_t^B(s, x; \sigma^*) = W_t^B(s, \xi(x); \sigma)$, we specify that $\hat{\rho}_t(s, x) = \rho^*(s, t, x)$. Third, for $x \in Y$ such that $W_t^B(s, x; \sigma^*) > W_t^B(s, \xi(x); \sigma)$, we require: (i) if $w_{s,t} \geq W_t^B(s, x; \sigma^*)$, then $\hat{\rho}_t(s, x) = 1$, (ii) if $w_{s,t} = W_t^B(s, \xi(x); \sigma^*)$, then $\hat{\rho}_t(s, x) = 0$, and (iii) if $W_t^B(s, x; \sigma^*) > w_{s,t} > W_t^B(s, \xi(x); \sigma^*)$, then the politician's expected discounted utility is exactly $w_{s,t}$, i.e.,

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

Fourth, for $x \in Y$ such that $W_t^B(s, x; \sigma^*) < W_t^B(s, \xi(x); \sigma)$, 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 $x \in \text{supp}(\pi_t^*(\cdot|s))$, 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

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

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.

B Proof of Representative Voter Theorem

In this appendix, we provide the proof of Theorem 3. Let $\text{ca}(X)$ denote the Banach space of signed Borel measures on X , endowed with the total variation norm, $\|\cdot\|_v$. Let $\Xi = [\text{ca}(X)^3]^{S \times T \times X}$ be the space of mappings from triples (s, t, x) to triples of Borel measures on X , so we may write an element $\xi \in \Xi$ as a triple $\xi = (\lambda, \mu, \nu)$ of component functions. Endow Ξ with the sup norm, i.e., for all $\xi, \xi' \in \Xi$, we have

$$\|\xi - \xi'\| = \sup \left\{ \begin{array}{l} \|\lambda(s, t, x) - \lambda'(s, t, x)\|_v, \\ \|\mu(s, t, x) - \mu'(s, t, x)\|_v, \\ \|\nu(s, t, x) - \nu'(s, t, x)\|_v \end{array} \mid (s, t, x) \in S \times T \times X \right\}.$$

As such, Ξ is itself a Banach space; in particular, it is a complete metric space. For all s and all x , let $\Delta(s, x)$ be the unit mass on (s, x) . Given the simple Markov electoral equilibrium σ , and using the assumption of a common discount factor from (B3), we define the mapping $\Phi: \Xi \rightarrow \Xi$ by

$$\begin{aligned} \Phi^F(\xi)(s, t, x) &= \int_{x'} \left[(1 - \delta)\Delta(s, x') + \delta[\rho(s, t, x')\mu(s, t, x') \right. \\ &\quad \left. + (1 - \rho(s, t, x'))\nu(s, t, x')] \right] \pi_t(dx'|s) \\ \Phi^B(\xi)(s, t, x) &= p_t(s|s, x, 1)[(1 - \delta)\Delta(s, x) + \delta\mu(s, t, x)] \\ &\quad + \sum_{s' \neq s} p_t(s'|s, x, 1)\Phi^F(\xi)(s', t, x) \\ \Phi^C(\xi)(s, t, x) &= \sum_{t'} q_t(t'|s, x) \sum_{s'} p_t(s'|s, x, 0)\Phi^F(\xi)(s', t', x), \end{aligned}$$

where we write the values of Φ as $\Phi(\xi) = (\Phi^F(\xi), \Phi^B(\xi), \Phi^C(\xi))$. Note that the first component Φ^F is in fact constant in x .

We claim that Φ is a contraction mapping with modulus δ . To see this, consider any $\xi, \xi' \in \Xi$ and any $s, t, x \in S \times T \times X$, and first note that

$$\begin{aligned} &\|\Phi^F(\xi)(s, t, x) - \Phi^F(\xi')(s, t, x)\|_v \\ &\leq \delta \int_{x'} \left[\rho(s, t, x') \|\mu(s, t, x') - \mu'(s, t, x')\|_v \right. \\ &\quad \left. + (1 - \rho(s, t, x')) \|\nu(s, t, x') - \nu'(s, t, x')\|_v \right] \pi_t(dx'|s) \\ &\leq \delta \int_{x'} \max \left\{ \begin{array}{l} \|\mu(s, t, x') - \mu'(s, t, x')\|_v, \\ \|\nu(s, t, x') - \nu'(s, t, x')\|_v \end{array} \right\} \pi_t(dx'|s) \\ &\leq \delta \|\xi - \xi'\|. \end{aligned}$$

Furthermore,

$$\|\Phi^B(\xi)(s, t, x) - \Phi^B(\xi')(s, t, x)\|_v$$

$$\begin{aligned}
&\leq \delta p_t(s|s, x, 1) \|\mu(s, t, x) - \mu'(s, t, x)\|_v \\
&\quad + \sum_{s' \neq s} p_t(s'|s, x, 1) \|\Phi^F(\xi)(s', t, x) - \Phi^F(\xi')(s', t, x)\|_v \\
&\leq \delta \|\xi - \xi'\|,
\end{aligned}$$

and

$$\begin{aligned}
&\|\Phi^C(\xi)(s, t, x) - \Phi^C(\xi')(s, t, x)\|_v \\
&\leq \sum_{t'} q_t(t'|s, x) \sum_{s'} p_t(s'|s, x, 0) \|\Phi^F(\xi)(s', t', x) - \Phi^F(\xi')(s', t', x)\|_v \\
&\leq \delta \|\xi - \xi'\|.
\end{aligned}$$

Therefore, $\|\Phi(\xi) - \Phi(\xi')\| \leq \delta \|\xi - \xi'\|$, as claimed.

By the contraction mapping theorem, the mapping Φ has a unique fixed point, which we denote $\xi^* = (\lambda^*, \mu^*, \nu^*)$. Accordingly, we have

$$\begin{aligned}
\lambda^*(s, t) &= \int_{x'} \left[(1 - \delta) \Delta(s, x') + \delta [\rho(s, t, x') \mu^*(s, t, x') \right. \\
&\quad \left. + (1 - \rho(s, t, x')) \nu^*(s, t, x')] \right] \pi_t(dx'|s) \\
\mu^*(s, t, x) &= p_t(s|s, x, 1) [(1 - \delta) \Delta(s, x) + \delta \mu^*(s, t, x)] \\
&\quad + \sum_{s' \neq s} p_t(s'|s, x, 1) \lambda^*(s', t) \\
\nu^*(s, t, x) &= \sum_{t'} q_t(t'|s, x) \sum_{s'} p_t(s'|s, x, 0) \lambda^*(s', t'),
\end{aligned}$$

where we omit the argument x from λ^* , which is constant in x . Let λ_X^* , μ_X^* , and ν_X^* denote the marginals of these probability measures on the policy space. For each τ , define the mappings $U_\tau^F: S \times T \rightarrow \mathfrak{R}$ and $U_\tau^B, U_\tau^C: S \times T \times X \rightarrow \mathfrak{R}$ by

$$\begin{aligned}
U_\tau^F(s, t) &= \int_{x'} u_\tau(x') \lambda_X^*(s, t)(dx') \\
U_\tau^B(s, t, x) &= \int_{x'} u_\tau(x') \mu_X^*(s, t, x)(dx') \\
U_\tau^C(s, t, x) &= \int_{x'} u_\tau(x') \nu_X^*(s, t, x)(dx').
\end{aligned}$$

for all s , all t , and all x .

Note that these functions solve the system of equations

$$\begin{aligned}
U_\tau^F(s, t) &= \int_{x'} \left[(1 - \delta) u_\tau(x') + \delta [\rho(s, t, x') U_\tau^B(s, t, x') \right. \\
&\quad \left. + (1 - \rho(s, t, x')) U_\tau^C(s, t, x')] \right] \pi_t(dx'|s)
\end{aligned}$$

$$\begin{aligned}
U_\tau^B(s, t, x) &= p_t(s|s, x, 1)[(1 - \delta)u_\tau(x) + \delta U_\tau^B(s, t, x)] \\
&\quad + \sum_{s' \neq s} p_t(s'|s, x, 1)U_\tau^F(s', t) \\
U_\tau^C(s, t, x) &= \sum_{t'} q_t(t'|s, x) \sum_{s'} p_t(s'|s, x, 0)U_\tau^F(s', t),
\end{aligned}$$

for all s , all t , and all x , and therefore we have

$$\begin{aligned}
\frac{U_\tau^F(s, t)}{1 - \delta} &= \int_{x'} \left[u_\tau(x') + \delta [\rho(s, t, x') \frac{U_\tau^B(s, t, x')}{1 - \delta} \right. \\
&\quad \left. + (1 - \rho(s, t, x')) \frac{U_\tau^C(s, t, x')}{1 - \delta} \right] \pi_t(dx'|s) \\
\frac{U_\tau^B(s, t, x)}{1 - \delta} &= \frac{p_t(s|s, x, 1)u_\tau(x) + \sum_{s' \neq s} p_t(s'|s, x, 1) \frac{U_\tau^F(s', t)}{1 - \delta}}{1 - p_t(s|s, x, 1)\delta} \\
\frac{U_\tau^C(s, t, x)}{1 - \delta} &= \sum_{t'} q_t(t'|s, x) \sum_{s'} p_t(s'|s, x, 0) \frac{U_\tau^F(s', t)}{1 - \delta}
\end{aligned}$$

for all s , all t , and all x . That is, after removing the normalization by $1 - \delta$, these functions satisfy conditions (1)–(3), which uniquely characterize the continuation values V^F , V^B , and V^C .

Therefore, we can write these continuation values as integrals,

$$\begin{aligned}
V_\tau^F(s, t) &= \frac{1}{1 - \delta} \int_{x'} u_\tau(x') \lambda_X^*(s, t, x)(dx') \\
V_\tau^B(s, t, x) &= \frac{1}{1 - \delta} \int_{x'} u_\tau(x') \mu_X^*(s, t, x)(dx') \\
V_\tau^C(s, t, x) &= \frac{1}{1 - \delta} \int_{x'} u_\tau(x') \nu_X^*(s, t, x)(dx'),
\end{aligned}$$

with respect to a collection of probability measures that varies with s , t , and x . For state s and incumbent type t , let κ be the total median type, as in (B5). Using (B1), (B2), and (B4), part (iii) of the lemma of Banks and Duggan (2006) implies that for all x , we have:

$$V_\kappa^B(s, t, x) > V_\kappa^C(s, t, x) \Leftrightarrow \exists C \in \mathcal{D}_t(s) : \forall \tau \in C : V_\tau^B(s, t, x) > V_\tau^C(s, t, x),$$

and

$$V_\kappa^B(s, t, x) \geq V_\kappa^C(s, t, x) \Leftrightarrow \exists C \in \mathcal{D}_t(s) : \forall \tau \in C : V_\tau^B(s, t, x) \geq V_\tau^C(s, t, x),$$

Thus, type κ is representative in s for t , as required.

C Proofs for Examples

Proofs for Example 1. In an equilibrium with compromise in state \hat{s} , all politician types implement policy \hat{x}_κ and are reelected. In such an equilibrium, the

payoff to a politician of type $t \in \{\ell, r\}$ from implementing policy \hat{x}_κ is $\frac{u+b}{1-\delta}$, while its payoff to implementing policy \hat{x}_t is $\hat{u} + b + \frac{\delta u}{1-\delta}$. Hence, an equilibrium with compromise in state \hat{s} exists if and only if

$$\frac{\delta b}{1-\delta} \geq \hat{u} - u.$$

In such an equilibrium, we have that for all politician types t and voter types τ , $V_\tau^B(\hat{s}, t, \hat{x}_\kappa) = V_\tau^C(\hat{s})$, so that reelecting all politicians implementing policy \hat{x}_κ is optimal for all voter types.

In an equilibrium with shirking in state \hat{s} , all politician types implement their ideal policies, with only politicians of type κ being reelected. In such an equilibrium, the payoff to a politician of type $t \in \{\ell, r\}$ from implementing policy \hat{x}_t is $\hat{u} + b + \frac{1}{3}\delta [V_t^F(\hat{s}, \ell) + V_t^F(\hat{s}, \kappa) + V_t^F(\hat{s}, r)]$. Since, for any voter type τ , $V_\tau(\hat{s}, t, \hat{x}_\kappa) = V_\tau(\hat{s}, \kappa, \hat{x}_\kappa)$, then a politician of type t would be reelected if it instead implemented policy \hat{x}_κ since, as shown below, voters have a strict incentive to reelect politicians of type κ following this policy. Since the payoff of a politician of type t following a deviation to policy \hat{x}_κ is $V_t^F(\hat{s}, \kappa) + \frac{b}{1-\delta}$, an equilibrium with shirking in state \hat{s} exists if and only if

$$\frac{\delta b}{1-\delta} \leq \hat{u} - u + \frac{1}{3}\delta [V_t^F(\hat{s}, \ell) + V_t^F(\hat{s}, r) - 2V_t^F(\hat{s}, \kappa)],$$

which, since

$$V_t^F(\hat{s}, \ell) + V_t^F(\hat{s}, r) - 2V_t^F(\hat{s}, \kappa) = \frac{1}{1-\frac{2}{3}\delta} [\hat{u} + \tilde{u} - 2u],$$

holds if and only if

$$\frac{\delta b}{1-\delta} \leq \frac{1}{1-\frac{2}{3}\delta} \left[\left(1 - \frac{1}{3}\delta\right) (\hat{u} - u) - \frac{1}{3}\delta(u - \tilde{u}) \right].$$

Given $\hat{u} - u < u - \tilde{u}$, the righthand side of the above inequality is decreasing in δ , so the inequality is satisfied whenever

$$\frac{\delta b}{1-\delta} \leq 2(\hat{u} - u) - (u - \tilde{u}).$$

Since $\hat{u} - u > \frac{1}{2}(u - \tilde{u})$ by assumption, the above inequality is satisfied whenever office benefits are sufficiently low.

In an equilibrium with shirking, the payoff to a voter of type $\tau \in \{\ell, r\}$ from reelecting a politician of any type having implemented policy \hat{x}_κ is $V_\tau^F(\hat{s}, \kappa)$, while its payoff from the challenger is $\frac{1}{3} [V_\tau^F(\hat{s}, \ell) + V_\tau^F(\hat{s}, r) + V_\tau^F(\hat{s}, \kappa)]$. Hence, type τ voters support the incumbent if and only if $\hat{u} - u \leq u - \tilde{u}$, which holds strictly by assumption.

We now turn to state \hat{s} . First, suppose that the equilibrium involves compromise in state \hat{s} . Then there exists an equilibrium in which all politicians

implement their ideal policies in state \hat{s} and are reelected. In this equilibrium, voters will vote against an incumbent that has implemented their third-ranked policy and in favor of an incumbent that has implemented their first-ranked policy. Hence, we only need to consider the incentives of a voter of type τ facing a politician of type t having implemented its middle-ranked policy x . If the state transitions, then all politician types are expected compromise at policy \hat{x}_κ . Hence, we have that $V_\tau^B(\hat{s}, t, x) - V_\tau^C(\hat{s}) \geq 0$ if and only if $\hat{u} - u \leq u - \check{u}$, which holds strictly by assumption. Note that this is the same condition as in the compromise equilibrium in state \hat{s} . The argument for politicians' incentives is similar.

Now suppose that the equilibrium involves shirking in state \hat{s} . Then there exists \bar{p} such that, whenever $\hat{p} \leq \bar{p}$, there exists an equilibrium in which all politicians implement their ideal policy in state \hat{s} and only politicians of type κ are reelected in that state. In this equilibrium, the payoffs to a voter of type τ from a challenger in state \hat{s} are given by

$$V_\tau^C(\hat{s}) = \frac{1}{3}\hat{p}[\hat{u} + u + \check{u} + \delta V_\tau^B(\hat{s}, \kappa, \hat{x}_\kappa) + 2\delta V_\tau^C(\hat{s})] + (1 - \hat{p})V_\tau^C(\hat{s}).$$

This implies that, for any transition probability \hat{p} , we have that $V_\tau^C(\hat{s}) = V_\tau^C(\hat{s})$, since $V_\tau^B(\hat{s}, \kappa, \hat{x}_\kappa) = V_\tau^B(\hat{s}, \kappa, \hat{x}_\kappa)$. Also, for any politician type $t \in \{\ell, r\}$ and any policy x , the payoffs to a voter of type τ from an incumbent in state \hat{s} is given by

$$V_\tau^B(\hat{s}, t, x) = \hat{p}[u_\tau(\hat{s}, x) + \delta V_\tau^B(\hat{s}, t, x)] + (1 - \hat{p})[u_\tau(\hat{s}, \hat{x}_t) + \delta V_\tau^C(\hat{s})],$$

We show that if \hat{p} is sufficiently low, then no politician of type $t \in \{\ell, r\}$ can be reelected in equilibrium following any policy choice. This implies that for such politicians, implementing their ideal policy is optimal. Note that, for any voter of type $\tau \neq t$, $\lim_{\hat{p} \rightarrow 0} V_\tau^B(\hat{s}, t, \hat{x}_\tau) = u_\tau(\hat{s}, \hat{x}_t) + \delta V_\tau^C(\hat{s})$ and $V_\tau^C(\hat{s}) > \frac{u_\tau(\hat{s}, \hat{x}_t)}{1 - \delta}$. This implies that there exists \bar{p} such that, for all $\hat{p} \leq \bar{p}$, voters of type $\tau \neq t \in \{\ell, r\}$ support the challenger against an incumbent of type t that has implemented policy \hat{x}_τ . Since, for any policy x , $V_\tau^B(\hat{s}, t, x) \leq V_\tau^B(\hat{s}, t, \hat{x}_\tau)$, then, for any $\hat{p} \leq \bar{p}$, voters of type τ never support any incumbent of type other than κ .

It remains only to show that politicians of type κ are reelected following policy \hat{x}_κ . It is sufficient to show that these politicians always obtain the support of voters of type ℓ . This follows by the assumption that $\hat{u} - u < u - \check{u}$ since, as noted above, $V_\ell^B(\hat{s}, t, \hat{x}_\kappa) = V_\ell^B(\hat{s}, t, \hat{x}_\kappa)$ and $V_\ell^C(\hat{s}) = V_\ell^C(\hat{s})$. \square

Proofs for Example 2. In the absorbing state \hat{s} , there exists an equilibrium in which politicians of all types implement their ideal policy and are reelected in that state. To see this, note that the payoff to a voter of type τ that votes in favor of its second-ranked policy x implemented by a politician of type t , we have that

$$V_\tau^B(\hat{s}, t, x) = \frac{u}{1 - \delta},$$

while its payoff to a challenger is

$$V_\tau^C(\hat{s}) = \frac{1}{3} \frac{1}{1-\delta} [u + \hat{u} + \check{u}].$$

Hence, we have that $V_\tau^B(\hat{s}, t, x) \geq V_\tau^C(\hat{s})$ since, by assumption, $u - \check{u} > \hat{u} - u$.

If \hat{p} is sufficiently low, then there exists an equilibrium such that all types of politicians implement their ideal policy and are reelected in state \hat{s} . Note that for all voter types $\tau \in \{\ell, r\}$, we have that

$$V_\tau^C(\hat{s}) = V_\tau^C(\hat{s}).$$

Hence, since we also have that $V_\ell^B(\hat{s}, \kappa, \hat{x}_\kappa) = V_\ell^B(\hat{s}, \kappa, \hat{x}_\kappa)$, then it is optimal for voters of type ℓ to vote in favor of policy \hat{x}_κ when proposed by a politician of type κ .

To see that voters of type r vote in favor of policy \hat{x}_ℓ when proposed by a politician of type ℓ , note that

$$V_r^B(\hat{s}, \ell, \hat{x}_\ell) = \hat{p} [\check{u} + \delta V_r^B(\hat{s}, \ell, \hat{x}_\ell)] + (1 - \hat{p}) V_r^F(\hat{s}, \ell).$$

Computation yields that $V_r^B(\hat{s}, \ell, \hat{x}_\ell) \geq V_r^C(\hat{s})$ if and only if

$$\hat{p} \leq \frac{u - \check{u} - (\hat{u} - u)}{(3 - \delta)u - (3 - 2\delta)\check{u} - \delta\hat{u}} \equiv \bar{p}_1.$$

A simple computation verifies that $\bar{p}_1 < 1$. Since the denominator of (16) is linear in δ , it attains a minimum at $\delta \in \{0, 1\}$ and for both these values, we have that this denominator is positive. Hence, $\bar{p}_1 > 0$.

To see that voter κ votes in favor of policy \hat{x}_r when proposed by a politician of type r , note that

$$V_\kappa^B(\hat{s}, r, \hat{x}_r) = \hat{p} [u + \delta V_\kappa^B(\hat{s}, r, \hat{x}_r)] + (1 - \hat{p}) V_\kappa^F(\hat{s}, r) = V_\kappa^F(\hat{s}, r),$$

while

$$V_\kappa^C(\hat{s}) = \frac{1}{3} [V_\kappa^F(\hat{s}, \kappa) + V_\kappa(\hat{s}, r) + V_\kappa^F(\hat{s}, \ell)],$$

where $V_\kappa^F(\hat{s}, \ell) = \frac{1}{1-\delta\bar{p}} [\hat{p}u + (1 - \hat{p})V_\kappa^F(\hat{s}, \ell)]$. Hence, computation yields that $V_\kappa^B(\hat{s}, r, \hat{x}_r) \geq V_\kappa^C(\hat{s})$ if and only if

$$\hat{p} \leq \frac{u - \check{u} - (\hat{u} - u)}{u - \check{u} - \delta(\hat{u} - u)} \equiv \bar{p}_2.$$

Given our assumption that $u - \check{u} - (\hat{u} - u) > 0$, we have that $\bar{p}_2 \in (0, 1)$. Hence, let \bar{p} in the text be such that $\bar{p} = \min\{\bar{p}_1, \bar{p}_2\}$. \square

Proofs for Example 3. The aim is to show that there exists a simple Markov electoral equilibrium in which all t -type politicians implement policy x_1 in state s_t and policy x_{-1} in state s_{-t} and in which, for all states $s \in \{s_1, s_{-1}\}$, $R(s, t) = \emptyset$. We start by deriving the equilibrium voting strategies of κ -type voters. Many of the computations will depend on the difference $V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t)$, the increment in the payoffs of voter κ in state s_{-t} from having an incumbent of type $-t$ rather than of type t . Note that, since

$$V_\kappa^C(s_{-t}, -t, x_1) = \frac{1}{2} [V_\kappa^F(s_t, -t) + V_\kappa^F(s_t, t)],$$

and

$$V_\kappa^C(s_{-t}, t, x_{-1}) = (1-p) \frac{u_\kappa(\underline{x})}{1-\delta_\kappa} + \frac{1}{2} p [V_\kappa^F(s_{-t}, -t) + V_\kappa^F(s_{-t}, t)],$$

and that, by symmetry, $V_\kappa^F(s_t, -t) = V_\kappa^F(s_{-t}, t)$ and $V_\kappa^F(s_t, t) = V_\kappa^F(s_{-t}, -t)$, we have that

$$\begin{aligned} & V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t) & (16) \\ & = u_\kappa(x_1) - u_\kappa(x_{-1}) + \delta_\kappa [V_\kappa^C(s_{-t}, -t, x_1) - V_\kappa^C(s_{-t}, t, x_{-1})] \\ & = u_\kappa(x_1) - u_\kappa(x_{-1}) + \delta_\kappa (1-p) \left[\frac{1}{2} V_\kappa^F(s_{-t}, -t) + \frac{1}{2} V_\kappa^F(s_{-t}, t) \right. \\ & \quad \left. - \frac{1}{1-\delta_\kappa} u_\kappa(\underline{x}) \right] \\ & > 0. & (17) \end{aligned}$$

Furthermore, we assume that p satisfies

$$0 < p < 1 - \frac{u_\kappa(x_1) - u_\kappa(x_{-1})}{\delta_\kappa \left[\frac{1}{2} u_\kappa(x_1) + \frac{1}{2} u_\kappa(x_{-1}) - u_\kappa(\underline{x}) \right]}, \quad (18)$$

which can be can always hold as long as $u_\kappa(\underline{x})$ is sufficiently small.

First, to verify that opting for the challenger in state s_t against a t -type incumbent following policy x_1 is uniquely optimal, note that, since $V_\kappa^B(s_t, t, x_1) = V_\kappa^F(s_{-t}, t)$,

$$\begin{aligned} V_\kappa^C(s_t, t, x_1) - V_\kappa^B(s_t, t, x_1) & = \frac{1}{2} V_\kappa^F(s_{-t}, t) + \frac{1}{2} V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t) \\ & = \frac{1}{2} [V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t)] \\ & > 0. \end{aligned}$$

Second, to verify that opting for the challenger in state s_t against a t -type incumbent following policy x_{-1} is uniquely optimal, note that, since $V_\kappa^B(s_t, t, x_{-1}) = (1-p) \frac{u_\kappa(\underline{x})}{1-\delta_\kappa} + p V_\kappa^F(s_t, -t)$, we have that

$$V_\kappa^C(s_t, t, x_{-1}) - V_\kappa^B(s_t, t, x_{-1})$$

$$\begin{aligned}
&= \left[(1-p) \frac{u_\kappa(\underline{x})}{1-\delta_\kappa} + p \left[\frac{1}{2} V_\kappa^F(s_t, t) + \frac{1}{2} V_\kappa^F(s_t, -t) \right] \right. \\
&\quad \left. - \left[(1-p) \frac{u_\kappa(\underline{x})}{1-\delta_\kappa} + p V_\kappa^F(s_t, -t) \right] \right] \\
&= \frac{1}{2} p [V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t)] \\
&> 0.
\end{aligned}$$

Third, to verify that opting for the challenger in state s_{-t} against a t -type incumbent following policy x_1 is uniquely optimal, note that $V_\kappa^C(s_{-t}, t, x_1) - V_\kappa^B(s_t, t, x_1) > 0$ if and only if

$$\begin{aligned}
(1-p) \frac{u_\kappa(\underline{x})}{1-\delta_\kappa} + \frac{1}{2} p [V_\kappa^F(s_{-t}, -t) + V_\kappa^F(s_{-t}, t)] \\
- \left[(1-p) \frac{u_\kappa(\underline{x})}{1-\delta_\kappa} + p [u_\kappa(x_1) + \delta_\kappa V_\kappa^C(s_{-t}, t, x_1)] \right] \\
> 0,
\end{aligned}$$

which, since $u_\kappa(x_1) + \delta_\kappa V_\kappa^C(s_{-t}, t, x_1) = u_\kappa(x_1) - u_\kappa(x_{-1}) + V_\kappa^F(s_{-t}, t)$ and $p > 0$, holds if and only if

$$\frac{u_\kappa(x_1) - u_\kappa(x_{-1})}{V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t)} < \frac{1}{2}.$$

Using our expression (17), this condition is equivalent to

$$\begin{aligned}
&u_\kappa(x_1) - u_\kappa(x_{-1}) \tag{19} \\
&< \delta_\kappa (1-p) \left[\frac{1}{2} V_\kappa^F(s_{-t}, -t) + \frac{1}{2} V_\kappa^F(s_{-t}, t) - \frac{1}{1-\delta_\kappa} u_\kappa(\underline{x}) \right]. \tag{20}
\end{aligned}$$

Since $V_\kappa^F(s_{-t}, -t) \geq u_\kappa(x_1) + \frac{\delta_\kappa}{1-\delta_\kappa} u_\kappa(\underline{x})$ and $V_\kappa^F(s_{-t}, t) \geq u_\kappa(x_{-1}) + \frac{\delta_\kappa}{1-\delta_\kappa} u_\kappa(\underline{x})$, a sufficient condition for (20) is the assumption (18) that p is sufficiently small. Fourth, to verify that opting for the challenger in state s_{-t} against a t -type incumbent following policy x_2 is uniquely optimal, note that

$$\begin{aligned}
&V_\kappa^C(s_t, t, x_{-1}) - V_\kappa^B(s_t, t, x_{-1}) \\
&= p \left[\frac{1}{2} V_\kappa^F(s_{-t}, t) + \frac{1}{2} V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t) \right] \\
&= \frac{1}{2} p [V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t)] \\
&> 0.
\end{aligned}$$

Finally, we verify that politicians' proposal strategies are optimal. To verify that implementing policy x_1 is optimal for t -type politicians in state s_t , note that their payoffs in that state to implementing policy x_1 are

$$b + V_t^F(s_t, t) = b + u_t(s_t, x_1) + \frac{1}{2} \delta_t [V_t^F(s_{-t}, t) + V_t^F(s_{-t}, -t)],$$

while, by the symmetry of transition probabilities and of type t 's payoffs, their payoffs to implementing policy x_{-t} are

$$\begin{aligned} b + V_t^F(s_t, -t) &= b + u_t(s_t, x_{-1}) + \delta_t(1-p) \frac{u_t(\underline{s}, \underline{x})}{1-\delta_t} \\ &\quad + \delta_t p \left[\frac{1}{2} V_t^F(s_t, t) + \frac{1}{2} V_t^F(s_t, -t) \right]. \end{aligned}$$

Finally,

$$\begin{aligned} &V_t^F(s_t, t) - V_t^F(s_t, -t) \\ &= u_t(s_t, x_1) - u_t(s_t, x_{-1}) + \delta_t(1-p) \left[\frac{1}{2} V_t^F(s_t, t) + \frac{1}{2} V_t^F(s_t, -t) - \frac{u_t(\underline{s}, \underline{x})}{1-\delta_t} \right] \\ &> 0. \end{aligned}$$

To verify that implementing policy x_{-1} is optimal for t -type politicians in state s_{-t} , note that their payoffs in that state are

$$b + u_t(s_{-t}, x_{-1}) + V_t^C(s_{-t}, t, x_{-1}),$$

while their payoff to implementing policy x_1 is

$$b + u_t(s_{-t}, x_1) + V_t^C(s_{-t}, t, x_1).$$

The desired inequality follows since $V_t^C(s_{-t}, t, x_{-1}) = V_t^C(s_{-t}, t, x_1)$. Note that office benefits b are irrelevant for t -type politicians' decisions in states s_t and s_{-t} since no policy leads to reelection. \square

Proofs for Example 4. We assume that politicians' office benefit b satisfies

$$\begin{aligned} b &\geq b^* \\ &\equiv \frac{(2 - \delta_\kappa - p)u_\kappa(x_1) - (1 - p)u_\kappa(x_{-1})}{\delta_\kappa} \\ &> 0. \end{aligned} \tag{21}$$

We start by verifying that the voting strategies of κ -type voters are optimal and satisfy the requirements of differential equilibria. First, to verify that reelecting any politician type $t \in \{1, \kappa\}$ following policy x_{-1} in any state s is optimal for κ -type voters, note that

$$V_\kappa^B(s, t, x_{-1}) = \frac{u_\kappa(x_{-1})}{1 - \delta_\kappa} = V_\kappa^C(s, x_{-1}).$$

Second, to verify that (i) reelecting type 1 politicians following policy x_1 is optimal for κ -type voters in states s_1 and \underline{s} and that (ii) opting for the challenger

over a κ -type politician following policy x_1 is optimal in states s_{-1} and \underline{s} , note that

$$\begin{aligned}
V_\kappa^B(s_1, 1, x_1) &= V_\kappa^B(\underline{s}, 1, x_1) \\
&= \frac{u_\kappa(x_1)}{1 - \delta_k} \\
&> p \frac{u_\kappa(x_{-1})}{1 - \delta_k} + (1 - p) \frac{u_\kappa(x_1)}{1 - \delta_k} \\
&= V_\kappa^C(s_1, x_1) = V_\kappa^C(\underline{s}, x_1) = V_\kappa^C(s_{-1}, x_1) \\
&> \frac{u_\kappa(x_{-1})}{1 - \delta_k} \\
&= V_\kappa^B(s_{-1}, \kappa, x_1) = V_\kappa^B(\underline{s}, \kappa, x_1).
\end{aligned}$$

Third, to verify that opting for the challenger over a κ -type politician following policy x_1 is optimal for κ -type voters in state s_1 , note that since

$$V_\kappa^B(s_1, \kappa, x_1) = \frac{1}{1 - \delta_k(1 - p)} \left[(1 - p)u_\kappa(x_1) + p \frac{u_\kappa(x_{-1})}{1 - \delta_k} \right],$$

we have that

$$\begin{aligned}
V_\kappa^B(s_1, \kappa, x_1) - V_\kappa^C(s_1, x_1) &= \frac{p(1 - p)\delta_k}{(1 - \delta_k)(1 - \delta_k(1 - p))} [u_\kappa(x_{-1}) - u_\kappa(x_1)] \\
&< 0.
\end{aligned}$$

Note that κ -type voters' strategies are consistent with a deferential equilibrium in that they elect challengers only when they are strictly preferred to the incumbent. Note also that κ -type voters' incentives are independent of δ_κ .

We now verify that politicians' proposal strategies are optimal. Since type 1 politicians obtain their maximal payoff in all states and are always reelected, we focus on the incentives of κ -type politicians. In all states, the equilibrium payoff to κ -type politicians is

$$\frac{u_\kappa(x_{-1}) + b}{1 - \delta_k},$$

while they payoff to deviating to policy x_1 is

$$u_\kappa(x_1) + b + \frac{\delta_\kappa}{1 - \delta_\kappa} [(1 - p)u_\kappa(x_1) + pu_\kappa(x_{-1})].$$

Hence, that κ -type politicians' proposals are optimal follows from assumption (21). \square

Proofs for Example 5. We show that there exists a simple Markov equilibrium in which all politicians implement their stage-ideal policy and are always reelected. Since these policy strategies are optimal for politicians for any b , we

need only verify the optimality of κ -type voters' reelection strategies. To verify that reelecting t -type politicians following policy \hat{x}_t in state s_t is optimal, note that

$$V_{\kappa}^B(s_t, t, \hat{x}_t) = \frac{u_{\kappa}(s_t, \hat{x}_t)}{1 - \delta_{\kappa}},$$

while

$$\begin{aligned} V_{\kappa}^C(s_t, \hat{x}_t) &= \frac{1}{2}(1 - q) \frac{u_{\kappa}(s_t, \hat{x}_t)}{1 - \delta_{\kappa}} + \frac{1}{2}(1 - q) \left[u_{\kappa}(s_t, \hat{x}_{-t}) + \delta_{\kappa} \frac{u_{\kappa}(s_t, \hat{x}_t)}{1 - \delta_{\kappa}} \right] \\ &\quad + q \frac{u_{\kappa}(s_t, \hat{x}_{\kappa})}{1 - \delta_{\kappa}}, \end{aligned}$$

where we exploit the symmetry of κ voter's payoffs with respect to t . Thus, by (11), we have that

$$V_{\kappa}^B(s_t, t, \hat{x}_t) \geq V_{\kappa}^C(s_t, \hat{x}_t).$$

Since $V_{\kappa}^B(s_{-t}, t, \hat{x}_t) = V_{\kappa}^B(s_t, t, \hat{x}_t)$ and $V_{\kappa}^C(s_{-t}, \hat{x}_t) = V_{\kappa}^C(s_t, \hat{x}_t)$, it follows that reelecting t -type politicians following policy \hat{x}_t in state s_{-t} is also optimal for κ -type voters. \square

References

- Acemoglu, D., Egorov, G., Sonin, K., 2012. Dynamics and stability of constitutions, coalitions, and clubs. *American Economic Review* 102 (4), 1446–76.
- Aghion, P., Bolton, P., 1990. Government domestic debt and the risk of default: A political-economic model of the strategic role of debt. In: Dornbusch, R., Draghi, M. (Eds.), *Public Debt Management: Theory and History*. Cambridge: Cambridge University Press.
- Alesina, A., Rodrik, D., 1994. Distributive politics and economic growth. *Quarterly Journal of Economics* 109, 465–490.
- Alesina, A., Tabellini, G., 1990. A positive theory of budget deficits and government debt. *Review of Economic Studies* 57, 403–414.
- Aliprantis, C., Border, K., 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer: New York, NY.
- Aragones, E., Palfrey, T., Postlewaite, A., 2007. Political reputations and campaign promises. *Journal of the European Economic Association* 5, 846–884.
- Azzimonti, M., 2011. Barriers to investment in polarized societies. *American Economic Review* 101, 2182–2204.

- Bai, J. H., Lagunoff, R., 2011. On the faustian dynamics of policy and political power. *The Review of Economic Studies* 78 (1), 17–48.
- Banks, J., Duggan, J., 2001. A multidimensional model of repeated elections, wallis Institute of Political Economy Working Paper 24.
- Banks, J., Duggan, J., 2006. A social choice lemma on voting over lotteries with applications to a class of dynamic games. *Social Choice and Welfare* 26, 285–304.
- Banks, J., Duggan, J., 2008. A dynamic model of democratic elections in multi-dimensional policy spaces. *Quarterly Journal of Political Science* 3, 269–299.
- 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.
- Banks, J., Sundaram, R., 1993. Adverse selection and moral hazard in a repeated election model. In: William Barnett et al. (Ed.), *Political Economy: Institutions, Information, Competition, and Representation*. Cambridge University Press.
- Banks, J., Sundaram, R., 1998. Optimal retention in agency problems. *Journal of Economic Theory* 82, 293–323.
- Barro, R., 1973. The control of politicians: An economic model. *Public Choice* 14, 19–42.
- Bassetto, M., Benhabib, J., 2006. Redistribution, taxes, and the median voter. *Review of Economic Dynamics* 9, 211–223.
- Battaglini, M., 2011. A dynamic theory of electoral competition.
- Battaglini, M., Coate, S., 2007. Inefficiency in legislative policy-making: A dynamic analysis. *American Economic Review* 97, 118–149.
- Battaglini, M., Coate, S., 2008. A dynamic theory of public spending, taxation, and debt. *American Economic Review* 98, 201–236.
- Bernhardt, D., Camara, O., Squintani, F., 2011. Competence and ideology. *Review of Economic Studies* 78 (487–522).
- Bernhardt, D., Campuzano, L., Squintani, F., Camara, O., 2009. On the benefits of party competition. *Games and Economic Behavior* 66, 685–707.
- Bernhardt, D., Dubey, S., Hughson, E., 2004. Term limits and pork barrel politics. *Journal of Public Economics* 88, 2383–2422.
- Bertola, G., 1993. Factor shares and savings in endogenous growth. *American Economic Review* 83, 1184–1210.

- Besley, T., Coate, S., 1998. Sources of inefficiency in a representative democracy: A dynamic analysis. *American Economic Review* 88, 139–156.
- Camara, O., 2012. Economic policies of heterogeneous politicians.
- Canes-Wrone, B., Herron, M., Shotts, K., 2001. Leadership and pandering: A theory of executive policymaking. *American Journal of Political Science* 45, 532–550.
- Coate, S., Morris, S., 1995. On the form of transfers to special interests. *Journal of Political Economy* 103, 1210–1235.
- Duggan, J., 2000. Repeated elections with asymmetric information. *Economics and Politics* 12, 109–136.
- Duggan, J., 2011. Coalitional bargaining equilibria, unpublished paper.
- Duggan, J., 2012. Noisy stochastic games. *Econometrica* 80, 2017–2045.
- Duggan, J., 2013. A folk theorem for repeated elections with adverse selection, unpublished paper.
- Duggan, J., Kalandrakis, T., 2012. Dynamic legislative policy making. *Journal of Economic Theory* forthcoming.
- Forand, J. G., 2012. Two-party competition with persistent policies.
- Klein, P., Krusell, P., Rios-Rull, J.-V., 2008. Time-consistent public policy. *The Review of Economic Studies* 75 (3), 789–808.
- Krusell, P., Quadrini, V., Rios-Rull, J.-V., 1997. Politico-economic equilibrium and economic growth. *Journal of Economic Dynamics and Control* 21, 243–272.
- Krusell, P., Rios-Rull, J.-V., 1999. On the size of the us government: Political economy in the neoclassical growth model. *American Economic Review* 89, 1156–1181.
- Levy, Y., 2012a. A discounted stochastic game with no stationary equilibria: The case of absolutely continuous transitions. Tech. rep., The Center for the Study of Rationality, Hebrew University, Jerusalem.
- Levy, Y., 2012b. A discounted stochastic game with no stationary equilibria: The case of deterministic transitions. Tech. rep., The Center for the Study of Rationality, Hebrew University, Jerusalem.
- Maskin, E., Tirole, J., 2004. The politician and the judge: Accountability in government. *American Economic Review* 94, 1034–1054.
- Meirowitz, A., 2007. Probabilistic voting and accountability in elections with uncertain policy constraints. *Journal of Public Economic Theory* 9, 41–68.

- Persson, T., Svensson, L., 1989. Why a stubborn conservative would run a deficit: Policy with time-inconsistent preferences. *Quarterly Journal of Economics* 104, 325–345.
- Rogoff, K., Sibert, A., 1988. Elections and macroeconomic policy cycles. *Review of Economic Studies* 55, 1–16.
- Yared, P., 2010. Politicians, taxes and debt. *Review of Economic Studies* 77, 806–840.