Dynamic Elections and the Limits of Accountability*

John Duggan Jean Guillaume Forand

November 13, 2014

Abstract

We provide conditions under which policy outcomes are responsive to preferences of voters in a dynamic model of elections with a discrete state space and general policies and preferences. We begin with weak conditions guaranteeing that in each state, at least one type of politician can choose a policy that leads to reelection—a minimal prerequisite for electoral accountability. Strengthening our assumptions, we show that politicians whose preferences coincide with those of a fixed representative voter solve that voter's dynamic programming problem, and we show that there exist equilibria such that *all* politician types solve this problem—so that electoral accountability leads to responsive policy outcomes. Finally, under our strongest conditions, we establish an asymptotic responsiveness result to the effect that *all* equilibria approximately solve the representative dynamic programming problem as citizens become patient. When the conditions for these results are not met, examples demonstrate novel dynamic political failures.

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 variable that evolves over time and is, in principle, influenced by policy choices. The purpose of the present paper is to understand the effectiveness of incentives induced by electoral accountability in the presence of such endogenous economic

*Duggan: University of Rochester, dugg@ur.rochester.edu. Forand: University of Waterloo, jgforand@uwaterloo.ca. This paper contains results previously circulated in our working paper 'Markovian elections' (2013), available on the authors' websites. We thank seminar audiences at Concordia, Johns Hopkins, London School of Economics, Rice, Rochester, U.S.C., Warwick, and Western Ontario, and participants at the Canadian Public Economics meetings 2011, the Priorat Workshop in Theoretical Political Science 2012, CEA 2013, CETC 2013, CPEG 2011, MPSA 2013, PET 2014, SED 2014, and SITE 2014.

and political state variables, e.g., capital stock, the distribution of income, or the institutional rules governing the political system. To this end, we analyze a general dynamic model of elections and we focus on whether, and how, repeated electoral choices between incumbent office holders and challenging candidates can engender *responsive politics,* i.e., policy choices by office holders that correspond to those that would be chosen by the voters directly in a hypothetical scenario without delegation of public decisions to politicians. The key challenge to electoral accountability is that voters cannot commit in advance to electoral responses, and politicians cannot commit to future policy choices as the state evolves. In the absence of a commitment mechanism, one might expect that elections lose their sanctioning power, and that politics devolves into the pursuit of parochial objectives.

We find, however, that electoral accountability can indeed lead to different forms of responsive politics. We give weak conditions ensuring that in each state, at least one type of politician can choose a policy that leads to reelection—a minimal prerequisite for the effectiveness of elections. Under stronger assumptions, we show that in all equilibria, politicians whose preferences coincide with those of a fixed representative voter solve the voter's dynamic programming problem, i.e., they are "faithful delegates." Moreover, we show that there exist equilibria such that all politician types solve the representative voter's problem. Finally, under our most stringent conditions, we establish that the supremum of these results holds asymptotically: *all* equilibria approximately solve the representative voter's problem as this voter becomes patient. In other words, these conditions imply that equilibrium policy choices approximate the upper limit for electoral accountability, providing a strong result on policy responsiveness. In sum, our results are broadly optimistic, but we also offer cautionary tales in the form of several examples that demonstrate the possibility of novel dynamic political failures.

In line with Besley and Coate (1998), dynamic political failures have typically been identified with Pareto inefficiencies. Our normative performance criterion for elections, namely responsiveness, differs in that it compares electoral outcomes with those outcomes that would be chosen by voters directly. In a static setting, assuming policies are one-dimensional and voter preferences are single-peaked, our responsiveness criterion singles out the ideal point of the median voter as the benchmark against which electoral outcomes are compared. In a dynamic environment, however, even if we fix the identity of a representative voter (e.g., the median voter type) in all elections, the "ideal policy" of this voter would correspond to a state-contingent policy plan that is obtained as the solution to a non-trivial dynamic programming problem. In general, we therefore compare policy choices of office holders to solutions of a *representative dynamic programming problem,* so that our results are the natural dynamic analogue to the median voter theorem in the static setting.[1]

[1]Our existence and continuity results allow voters' political power to evolve over time and for the representative voter to be state-dependent. In that case, the hypothetical scenario

Model and results We construct a general dynamic model of elections. 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 over the challenger's type as well as the 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 the challenger's type is not directly observed before the election. 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 as a special case when policy is a public good level and there is a single, fixed state.[2]

We establish four main results. First, we provide conditions under which, in an arbitrary state, there exists some politician type that the representative voter in that state considers *eligible* for office, i.e., there is at least one policy that, if implemented by that politician, leads the representative voter to weakly prefer reelecting the politician to opting for the challenger. This is a minimal prerequisite of responsive politics, as in its absence no politician can obtain reelection through appropriate policy choices, so that the representative voter's preferences impose no constraints on office holders. Second, in case the representative voter is fixed across states, we show that when politicians are purely policy-motivated, or when state transitions are independent of policies, then the politician type that shares the policy preferences of the representative voter is always eligible for office in equilibrium, and furthermore such politicians implement solutions to the representative dynamic programming problem. This sharpens our first result and implies a lower bound on the representative voter's expected payoff from a challenger in equilibrium, as in each state, there is a positive probability that a challenger would, if elected, choose policies that are optimal for the voter.

The previous result is a partial responsiveness result: it guarantees that policies chosen by politicians sharing the policy preferences of the representative voter are responsive, but it imposes no restrictions on the policy choices of other politicians. To address this issue, our third result shows that when politicians are sufficiently office-motivated, there exists an equilibrium (in pure strategies) in which all politicians solve the dynamic programming problem of the representative voter. As well as establishing the possibility of total responsiveness in dynamic elections, this result provides theoretical justification for the application of the median voter theorem in models of dynamic policy choice

in which these voters choose policy directly is modeled as a dynamic "representative voting game." The correspondence between election outcomes and the equilibria of this game is the focus of a companion paper (Duggan and Forand (2013)).

[2]See also Aragones et al. (2007), in which policy is a one-dimensional ideological variable.

(e.g., Krusell et al. (1997) and Krusell and Rios-Rull (1999)). While its conditions are fairly general, however, this result is weak in that it establishes only that *some* equilibria deliver responsive policies. In our final result, we provide conditions such that *all* equilibria are arbitrarily responsive. Specifically, we establish the strong asymptotic responsiveness result that the expected discounted payoff of the representative voter at every strongly recurrent state converges to the representative voter's optimal payoff as this voter becomes patient: that is, at any state that recurs with probability one regardless of the policy choices of politicians, equilibrium policies in that state and in future states approximately solve the representative voter's dynamic programming problem.

We further delineate the limits of electoral accountability by exploring the possibility of unresponsive policy dynamics when our conditions are violated. Specifically, we can show that our results are tight, in the sense that whenever our conditions fail, we can construct simple economies exhibiting equilibria that do not meet the relevant electoral performance criterion. For example, when state transitions depend on incumbents' types (e.g., politicians vary with respect to competence, which affects the state transition), 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 that admit an "ineligibility trap" such that no politician can implement any policy that leads to reelection, even if the benefit of holding office is arbitrarily large. For another example, if politicians who share the representative voter's policy preferences also value office per se, then we show the possibility of a "curse of ambition," in which these politicians are reelected only by choosing policies that are suboptimal for the representative voter. This is a stark example of a dynamic coordination failure: conditional on retaining office, a politician that shares the policy preferences of the representative voter always chooses optimal polices for this voter, but in some equilibria, "good" policies may lead to defeat whereas only "bad" policies ensure reelection. Thus, our model presents the possibility of new sources of inefficiency that can arise from the dynamic incentives of elections.[3]

Existence of equilibrium A contribution of this paper is a modeling framework that is general (amenable to a range of structure on preferences, policies, and states), viable (equilibria are guaranteed to exist and equilibrium behavior is tractable), and practically useful (we can solve special cases to generate novel insights). Existence of equilibrium is a particularly thorny issue, as it is known that the existence of Markovian equilibria in dynamic games may fail to exist (see Levy and McLennan (2014)).[4] A benefit of the discrete state space

[3]This 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)).

[4]Our model can be reformulated as a stochastic game, but in doing so the incumbent's policy choice must be included in the description of the state at the beginning of any electoral

model is its comparative tractability: under otherwise quite general assumptions on policies and payoffs, we are able to prove existence and continuity properties of a selection of Markovian equilibria, and we can solve simple examples with relative ease. Furthermore, we can give non-constructive characterizations of the behavior of voters and politicians that hold across many special cases, enhancing the applicability of our model.

In response to equilibrium existence problems, the study of electoral dynamics often relies on explicit equilibrium constructions or the analysis of necessary conditions in specific environments. For example, 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. Our framework does not generalize papers with continuous state variables (like capital stock or debt, which are commonly formulated as continuous), 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). In models of elections, the usual difficulties encountered in proving the existence of stationary Markovian equilibria 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. The cost of this parsimony is that the existence proof is involved, as voting strategies must be excluded from the domain of the fixed point argument and backed out after the fact.

Literature The analysis of endogenous state variables is an increasing focus of the political economy literature. 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,[5] 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 noncooperative model in which the state variable is a durable public good or public debt level, respectively, 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

period. Because policy choices belong to a compact metric space, results in that literature for countable state games (see Federgruen (1978)) cannot be applied.

[5] See Bertola (1993), Alesina and Rodrik (1994), and Bassetto and Benhabib (2006)).

an extension to growth economies that preserves the stationary structure of his equilibrium in a model of repeated elections with adverse selection. Forand (2014) considers elections 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, and elections are non-trivial. Azzimonti (2011) and Battaglini (2014) 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 (2014) 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.[6]

Despite technical differences in the form of linkage, however, our framework has connections to a dynamic model of elections with adverse selection considered by Duggan (2000) and Bernhardt et al. (2004).[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 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 single-state model to allow multiple states and the more complex incentives they entail: whereas reputational issues blow up and render the adverse selection model intractable, our model remains viable.

2 Dynamic Electoral Framework

In this section, we describe the electoral model, including the timing of moves and information available to voters and politicians, transition probabilities on

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

[7]See 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.

states and challengers, and payoffs of all actors.

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

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

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

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. If the incumbent seeks reelection, then voters simultaneously cast ballots for the incumbent or challenger. An *electoral outcome* for a type t incumbent in state s is $e \in \{0,1\}$, where $e = 1$ indicates that the incumbent seeks reelection ($x \in Y$) and is reelected, while $e = 0$ indicates that either the incumbent does not seek reelection ($x \in Z$) or that she is defeated by the challenger. This is determined by a set $\mathcal{D}_t(s) \subseteq 2^T \backslash \{\varnothing\}$ of *decisive* coalitions of types: if the coalition of voter types τ who vote for the incumbent belongs to $\mathcal{D}_t(s)$, then the incumbent retains office ($e = 1$) in the following period.[10] We assume that $\mathcal{D}_t(s)$ is monotonic, i.e., $C \in \mathcal{D}_t(s)$ and $C \subseteq C'$ imply $C' \in \mathcal{D}_t(s)$. Otherwise, if $x \in Z$ or if the set of voter types voting for the incumbent is not decisive, then the challenger assumes office in the following period ($e = 0$). Our formulation of the electoral rule is quite general and admits a wide variety of special cases as applications: weighed majority rule, more stringent quota rules, complex electoral systems such as the US Electoral College or domestic politics within an autocratic regime.

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 $p_t\colon S \times S \times X \times \{0,1\} \to [0,1]$ is continuous and independent of the incumbent's campaign decision, i.e., $p_t(s'|s,y,0) = p_t(s'|s,\xi(y),0)$ for all $y \in Y$.

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

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

of states, office holder names, types of office holders, policy choices, and electoral outcomes. A *partial finite public history* of length $m+1$ is a complete finite public history of length m concatenated with a triple $(s_{m+1}, j_{m+1}, t_{m+1})$ representing the state and the office holder's name and type in period $m+1$,

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

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

prior to choice of a policy. An *infinite public history* is an infinite list $h \in (S \times M \times T \times X \times \{0,1\})^\infty$.

State by state commitment We assume that if an office holder chooses a policy x in a state s, and if she is subsequently reelected, then she is committed to her policy choice if the state remains s in the following period. By implication, she remains committed in successive periods in which she is reelected and the state remains s. 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 partial finite public history (h^m, s, j, t) such that $e_m = 1$, $j = j_m$, and $t = t_m$, the action set available to the office holder is $\{x_m\}$ if $s = s_m$ (she is bound to her previous choice) and is $X_t(s)$ if $s \neq s_m$ (she is free). In a remark below, we note that as a special case we obtain the model in which commitment is probabilistic, so that an office holder may be free to choose new policies following some successive realizations of the state. Also, in a remark in Section 3, we discuss the importance of state-by-state commitment for generating a dynamic linkage between current and future policy choices in the equilibria of an important special case of the model with policy-independent state and challenger transitions.

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

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

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

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.

Remark We can capture a number of interesting features by suitable specialization of the model.[11] For example, we can obtain a flexible form of commitment in which, conditional on the politician choosing x in state s and the state recurring, the politician is committed with some probability $\gamma_t(s,x)$; in particular, the case of no commitment, $\gamma_t(s,x) \equiv 0$, is allowed. We do so by doubling the states, i.e., $\tilde{S} = (S \times \{1\}) \cup (S \times \{-1\})$, and defining a transition probability $\tilde{p}_t((s',k)|(s,k),x,e) = p_t(s'|s,x,e)\gamma_t(s,x)$ and $\tilde{p}_t((s',-k)|(s,k),x,e) = p_t(s'|s,x,e)(1-\gamma_t(s,x))$, so that one state (s,k) transitions to its twin $(s,-k)$ and releases the office holder from her commitment with appropriate probability. Similarly, in addition to the several standard formulations of payoffs described above, we can capture models of rent-seeking and political agency, and we can allow voters to have preferences over office holders' types. Using the dependence of the voting rule $\mathcal{D}_t(s)$ on the state, we can accommodate term limits for incumbent politicians. Finally, while it may seem that the model restricts attention policy choices generated by sequences of individual office holders, the model also admits competition between two long-lived parties by associating each party with a politician type, exploiting dependence of the challenger transition $q_t(t'|s,x)$ on the incumbent's type, and letting the parties' preferences depend on the current office holder's type.

[11] For additional special cases and further details, see Duggan and Forand (2013).

3 Markov Electoral Equilibria

In this section, we describe strategies, continuation values, and our solution concept of Markov electoral equilibrium. The latter adapts the notion of stationary Markov perfect equilibrium to the dynamic electoral framework. It is flexible enough to admit general equilibrium existence and continuity results (Theorems 1 and 2, below), yet sufficiently restrictive to permit characterization results (Theorems 3–6, in Section 4) in terms of responsiveness of equilibrium policy choices.

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

We adopt a parsimonious view of voting strategies, letting $\rho(h_m, s, j, t, x)$ be the probability that politician j is reelected after after public history h_m, the realization of state s, being type t, and choosing policy $x \in Y$.[12] As with policy strategies, we need only consider mixed voting upon the initial transition to a state s and policy choice x: if $s = s_m$, $j = j_m$, $t = t_m$, and $x = x_m$, then $\rho(h_m, s, j, t, x) = e^m$. Also consistent with our formulation of policy strategies, we restrict attention to strategies that are stationary with respect to the state and policy choice of the preceding period and the incumbent's type: thus, we write simply $\rho(s, t, x)$ for the probability that a type t office holder is reelected following policy choice x in state s.[13] 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 x in state s is again reelected after choosing x in state s, this is not a constraint imposed on voters; rather, by stationarity of the decision problem of the electorate, it will be consistent with the incentives of voters in equilibrium. When N is finite, the probability of reelection may be decomposed

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

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

into mixed voting strategies of individual voters. When the electorate N is infinite, individual uncertainty generated by mixed voting strategies washes out due to the law of large numbers, in which case we may interpret reelection probabilities as the result of conditioning on a public randomization device; we are agnostic as to interpretation. We refer to ρ as a *Markov voting strategy,* and to $\sigma = (\pi, \rho)$ as a *Markov electoral strategy profile.*

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

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

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

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

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

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

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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

where $\rho(s,t,x)$ is unrestricted if $x \in R(s,t) \setminus P(s,t)$. 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 $\rho(s,t,x) = 0$ for all $x \in Z$.

Remark State-by-state commitment is useful for generating interesting dynamic incentives in the equilibria of an important special case of the model with policy-independent transitions (i.e., when both $q_t(\cdot|s,x)$ and $p_t(\cdot|s,x,e)$ are independent of x for all s, t and e). In this environment, the absence of commitment would imply that the continuation payoff $V_\tau^B(s,t,x)$ to voter τ from reelecting an incumbent of type t in state s, as well as the payoff $V_\tau^C(s,x)$ from opting for the challenger, are independent of the policy x implemented by the incumbent. Hence, the reelection decision of voter τ would be independent of the incumbent's policy choice (with the possible exception of the case in which voter τ is indifferent between the incumbent and the challenger following this policy), so that in equilibrium incumbents would simply implement policies that maximize their stage utilities in each state.[14]

General properties of equilibria The starting point of our analysis is the next theorem, which provides a foundation for the model by establishing existence of equilibrium under the general conditions of our framework.[15]

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

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

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

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. The following result establishes upper hemi-continuity of the equilibrium payoff correspondence.

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

Illustration of equilibrium In the following example, we illustrate the logic of equilibrium in a simple two-state model with three feasible policies and state-contingent preferences. The example embeds a finite-policy version of the single-state model of Banks and Duggan (2008), and it highlights the critical role of state transitions for the characteristics of Markov electoral equilibria: equilibrium play in any one state depends on anticipated equilibrium behavior in other

[14] For a formal argument, see Duggan and Forand (2013).
[15] Appendix A contains formal proofs of existence and continuity.

states that may be reached in future periods. To illustrate these incentives, we introduce an equilibrium diagram that is also used in Examples 3 and 4.[16]

Example 1. Let the state space be $S = \{\hat{s}, \mathring{s}\}$ 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 $p(\mathring{s}|\mathring{s}) = \mathring{p}$ and $p(\hat{s}|\hat{s}) = \hat{p}$. Challenger selection probabilities are independent of states, policies, and incumbents' types, and they are such that, for all types t, $q(t) = \frac{1}{3}$. Voters have state-independent ideal policies, with a voter of type τ having ideal policy $\hat{x}_\tau$. 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 $\hat{x}_\kappa$ is a Condorcet winner in the stage game. In state $\mathring{s}$, voters' preferences induce a Condorcet cycle, with

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

For all voter types and all states, let $\hat{u}$ denote a voter's payoff from his ideal policy, u denote his payoff from his middle-ranked policy, and $\check{u}$ denote his payoff from his 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 $b \geqslant 0$, and all types have common discount factor δ.

We assume that the single-peaked state $\hat{s}$ is absorbing, i.e., that $\hat{p} = 1$.[17] This implies that 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 $\hat{x}_\kappa$. 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 $\hat{x}_\kappa$. 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 $\hat{x}_\ell$. Hence, $\hat{x}_\kappa$ 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 $\hat{x}_\kappa$ 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 $u - \check{u} > \hat{u} - u$ ensures that voters of type $\tau \in \{\ell, r\}$ support politicians implementing

[16]Detailed arguments for this and other examples are contained in Appendix C, which is not intended for publication.

[17]In Example 2, we let $\hat{p} < 1$ and specify that $\mathring{s}$ is absorbing, i.e., $\mathring{p} = 1$.

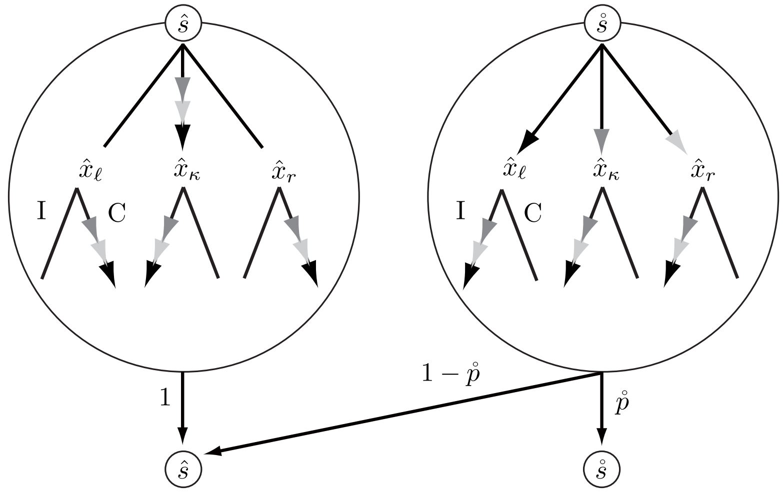

Figure 1: Equilibrium with state-dependent preferences and high b

their second-ranked policy in this equilibrium, while the assumption that $\hat{u}-u > \frac{1}{2}(u-\check{u})$ 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 rather than implement their second-ranked policy $\hat{x}_\kappa$ and retain office. See Figure 1 for a diagram of the equilibrium with high office benefit. Here, arrows emanating from a given state indicate policy choices of different types, where dark arrows represent type ℓ, medium represent type κ, and light 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 medium arrow pointing toward C from $\hat{x}_\ell$ indicates that when a type κ office holder chooses $\hat{x}_\ell$, voters replace the politician with a challenger.

Equilibrium play in state $\mathring{s}$ depends on whether non-κ politicians compromise or shirk in state $\hat{s}$. When b is high and all politicians compromise in state $\hat{s}$, as in Figure 1, then there exists a Markov electoral equilibrium in which all politicians implement their ideal policies in state $\mathring{s}$ 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 $\mathring{s}$ depend only on the policy she implements while the state remains $\mathring{s}$, 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 $\mathring{s}$, all voter types support incumbents having implemented their second-ranked policy (types ℓ support $\hat{x}_\kappa$, types κ support $\hat{x}_r$ and types r support $\hat{x}_\ell$).

When b is low and all politicians shirk in state $\hat{s}$, then there exists $\bar{p}$ such that when $\mathring{p} \leqslant \bar{p}$, there exists an equilibrium in which all politicians implement their ideal policy in state $\mathring{s}$ and only politicians of type κ are reelected in that state. Indeed, if $\mathring{p}$ is low, then in state $\mathring{s}$ 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 $\mathring{s}$, 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 $\hat{x}_\kappa$. 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 $\mathring{s}$, so the remaining types simply choose their ideal policies before being removed from office. □

4 Accountability in Dynamic Elections

The electoral accountability process has the potential, by subjecting incumbents to periodic review by voters, to discipline office holders and bring policy choices in line with voter preferences. This is so even if politicians' preferences are not aligned with the preferences of voters, so long as the value of holding office provides a sufficient incentive for office holders to put aside their personal policy preferences and to compete with the option of a challenger. Our dynamic environment presents three distinctive challenges to the efficacy of electoral accountability, however, all stemming from the absence of a commitment mechanism. First, candidates are not able to make credible promises about their policy choices in future states, so that even a candidate who would be willing to bind herself to popular policies in order to gain reelection has no way of doing so. Second, voters also have no way of committing to a reelection rule, so they cannot incentivize politicians by offering reelection in exchange for desirable policies. Third, a voter's political influence may vary across states, since membership in decisive coalitions can evolve with the state. Hence, a voter who has political power in the current state, but who anticipates losing this power in future states, faces an additional barrier to holding politicians accountable: if at all, voters can reward or punish politicians only in those states in which they have influence.

Nevertheless, we find various sets of conditions under which electoral accountability leads to different levels of policy responsiveness. Our first main result establishes the minimal precondition for policy responsiveness that in each state, there is at least one politician type that is eligible for reelection, in the sense that her weak reelection set is nonempty. Under stronger assumptions, our second result shows that, in particular, politicians with policy preferences coinciding with those of a fixed representative voter type are eligible, and the policy choices of these politicians in fact solve the dynamic programming problem of the representative voter type, so that such politicians are faithful delegates. Adding the assumption that politicians are largely office motivated, our third result establishes existence of a Markov electoral equilibrium such that all types of politician solve the representative dynamic programming problem, providing a full policy responsiveness result for at least one equilibrium. Finally, our fourth result shows that all Markov electoral equilibria approximately solve this program as voters become arbitrarily patient; the logic for this result differs from the latter and does not assume large office benefit. We start by identifying those voters who have political influence in a given state.

Representative voters In the remainder of the analysis, in order to provide the sharpest normative benchmark, we consider Markov electoral equilibria such that in all states, electoral outcomes are consistent with the preferences of a representative voter type. Formally, given a Markov electoral strategy profile σ, a type τ voter is *representative* in state s if for all t, $P(s,t) = P_\tau(s,t)$ and $R(s,t) = R_\tau(s,t)$.[18] When there is a voter type κ that is representative in all states, we can consider the hypothetical scenario in which that voter chooses policy directly in each state, and we can use this benchmark to evaluate the performance of elections. That is, we consider the *representative dynamic programming problem* with Bellman equation

$$V^*_\kappa(s) \quad = \quad \max_{x \in Y_\kappa(s)} u_\kappa(s,x) + \sum_{s'} p_\kappa(s'|s,x) V^*_\kappa(s'),$$

where we assume the representative voter is restricted to choose from policies feasible for the type κ politicians in each state.

A special case of the model is that in which there is a single state, policies are one-dimensional, stage utilities of voters are quadratic, and the type κ is the median voter: in this case, the solution of the representative dynamic programming problem is simply the choice of the median ideal policy in each period. In the classical Downsian model of elections, the well-known median voter theorem establishes that office-motivated candidates will commit to the median ideal policy, which in this special case is the unique Condorcet winner. Paralleling the role of the median ideal policy, solutions to the representative dynamic programming problem select policies on the basis of voter preferences alone, without reference to the electoral process through which policy decisions are delegated to politicians, and thus they are the natural dynamic analogue of the median in static models.

We say policy choices of office holders are *responsive* to the extent that they solve the above program, and we seek conditions under which, in the spirit of the median voter theorem, electoral accountability leads to responsive policy.

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 $\mathring{s}$ when the equilibrium in state $\hat{s}$

[18]In Duggan and Forand (2013), we provide sufficient conditions that ensure that given any Markov electoral strategy profile, a representative voters exists in all states.

called for shirking, whereas no voter type was representative in state $\mathring{s}$ when the equilibrium in state $\hat{s}$ called for compromise. In the former case, because the state transition is independent of policies, the representative dynamic programming problem is straightforward: the optimal policy is $\hat{x}_\kappa$ in both states, so responsiveness simply reduces to choosing the ideal policy of the representative voter in both states. That voter type κ need not be representative in state $\mathring{s}$ in all equilibria is natural since policy $\hat{x}_\kappa$ 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 ($\mathring{p} = 1$). If the probability $\hat{p}$ of remaining in the state with a Condorcet winner is sufficiently low, then there exists a 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 $\mathring{s}$, even if these politicians implement their third-ranked policy in state $\hat{s}$. In particular, voters of type ℓ and κ support policy $\hat{x}_\kappa$ 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 $\hat{x}_\ell$ when implemented by a politician of type ℓ. When transitions away from state $\hat{s}$ are likely, even though $\hat{x}_\kappa$ is a Condorcet winner in state $\hat{s}$, voters support incumbents based on the policies they will implement in state $\mathring{s}$. The absence of a Condorcet winner in state $\mathring{s}$ is carried into state $\hat{s}$ through dynamic incentives. □

Eligible politicians A precondition for a politician to be held accountable by a representative voter is that there exist circumstances in which the voter is willing to reelect the politician. Specifically, given a Markov electoral strategy profile σ, a politician type t is *eligible* in state s if the weak re-election set $R(s, t)$ of this politician at s contains at least one policy. Returning to Example 1, suppose that both b and $\mathring{p}$ are low, so that in equilibrium all politicians implement their ideal policies in both states $\hat{s}$ and $\mathring{s}$. Note that a politician of type ℓ, while not elected in equilibrium in state $\hat{s}$ after implementing policy $\hat{x}_\ell$, is nevertheless eligible for the representative voter κ, since implementing policy $\hat{x}_\kappa$ would lead to reelection. In state $\mathring{s}$, however, a politician of type ℓ is not eligible for voter κ, as this voter strictly prefers voting for the challenger over a politician of type ℓ following any policy by this politician.

The absence of this condition—if there is a state such that no politician type is eligible—clearly indicates a political failure, for then no type of politician can conceivably be reelected in that state, and thus each type trivially shirks with probability one. This is perhaps more paradoxical when there is a voter type that is representative in this state and politicians have mixed motives, as even a politician who shares the policy preferences of the representative voter cannot deliver a sufficiently high payoff to achieve reelection. Specifically, the payoff to a representative voter in state s from electing the challenger is an average of this voter's equilibrium payoff over all possible office holder types, and some type must be no worse than average, so one might conclude that some type of politician should be able to implement some policy that the voter prefers to the lottery over challengers.[19] In our general environment, however, the state evolves endogenously and politicians respond heterogeneously to changes in the state. This leaves the possibility that by implementing a policy that is aligned with the representative voter's preferences in state s, a politician can increase the likelihood of a transition to some state s' in which, say, the voter expects that politician to implement policies that are not aligned with his preferences. Thus, in principle, the anticipated loss of control experienced by the representative voter in state s when the state transitions to s' could lead to a breakdown in electoral accountability.

The following result establishes conditions under which the weak reelection set is nonempty for at least one office holder type at any given state, precluding the latter possibility.[20] The proof formalizes the intuition from above: assuming that some voter is representative in s, if the best possible politician type chooses the best possible policy for the representative voter in state s, where these are identified through the voter's continuation values, then no challenger can offer a higher expected payoff.[21] This argument relies on politicians feasibly mimicking the policy choices and associated continuation payoffs of other types, and it thus requires state and challenger transitions to be independent of types. We will see, however, that the possibility of an "ineligibility trap" can be realized when the conditions of the theorem are violated.

Theorem 3. *Let σ be a Markov electoral equilibrium. Fix a state s, assume that the type κ voter is representative in s, and assume*

(C1) $Y_t(s)$ is independent of t,

(C2) $p_t(s'|s, x, e)$ is independent of t and e for all s' and x,

(C3) $q_t(t'|s, x)$ is independent of t for all t' and x.

Then there exists a type t such that $R(s, t) \neq \emptyset$.

The next example shows that the weak reelection sets of *all* politician types can be empty in some states. Specifically, we exhibit an equilibrium in which all politicians are "trapped" by voters' expectations, as any policy they may implement, whether it benefits the representative voter or not, generates transitions to states in which they are the worst possible politician for this voter.

[19] In fact, this argument ensures that the reelection set of the representative voter is nonempty in any equilibrium in the single-state model of Banks and Duggan (2008), so that all politicians can be reelected by choosing compromise policies.

[20] We can capture term limits in the electoral model by augmenting states with a counter that records the incumbent's term of office and by assuming that the collection $\mathcal{D}_t(s)$ is empty whenever an incumbent is forced out of office by the term limit. An implication of Theorem 3, under (C1)–(C3), is that in equilibrium, there cannot be a representative voter in a state at which the term limit is binding.

[21] The formal proof of this an other accountability results are contained in Appendix B.

Example 3 (**Ineligibility trap**)**.** Assume the state space is $S = \{s_1, s_{-1}, \underline{s}, s_\kappa\}$, and that the type space is $T = \{1, -1, \kappa\}$. Assume that type κ voters are representative in all states, and let t range over $\{1, -1\}$. Sets of feasible policies are independent of politicians' types and are such that $Y(s_t) = \{x_1, x_{-1}\}$, $Y(\underline{s}) = \{\underline{x}\}$ and $Y(s_\kappa) = \{x_\kappa\}$. Transition 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. States $\underline{s}$ is absorbing, while $p_t(\underline{s}|s_\kappa, x_\kappa) = p_\kappa(s_\kappa|s_\kappa, x_\kappa) = 1$ and for all policies x, $p_\kappa(s_\kappa|s_t, x) = 1$. 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) requiring that 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 $u_\kappa(x_1) > u_\kappa(x_{-1}) > u_\kappa(\underline{x}) > u_\kappa(x_\kappa)$. Politicians have mixed motivations with type-independent office benefit $b \geqslant 0$ and stage utilities such that $u_t(s_t, x_1) = u_t(s_{-t}, x_{-1}) > u_t(s_t, x_{-1}) = u_t(s_{-t}, x_1) > u_t(\underline{s}, \underline{x}) > u_t(s_\kappa, x_\kappa)$.

We claim that there exists a 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 $R(s, t) = \varnothing$. Notice that politicians of type κ induce a transition to state s_κ following any policy choice in state s_t, and that once in s_κ, the voters' preferred absorbing state $\underline{s}$ can be reached only if politicians of type 1 or -1 are in office. Hence, politicians of type κ, while never selected as challengers, would nevertheless never be reelected in state s_t.[22] 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, where dark arrows represent policy choices of type 1 politicians, medium arrows represent choices of type κ politicians, solid lines represent transition probabilities following choices of type 1 politicians, and dashed lines represent transition probabilities for type κ 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 $-t$

[22]The assumption that challengers are never of type κ is not essential, but it simplifies the presentation of the example.

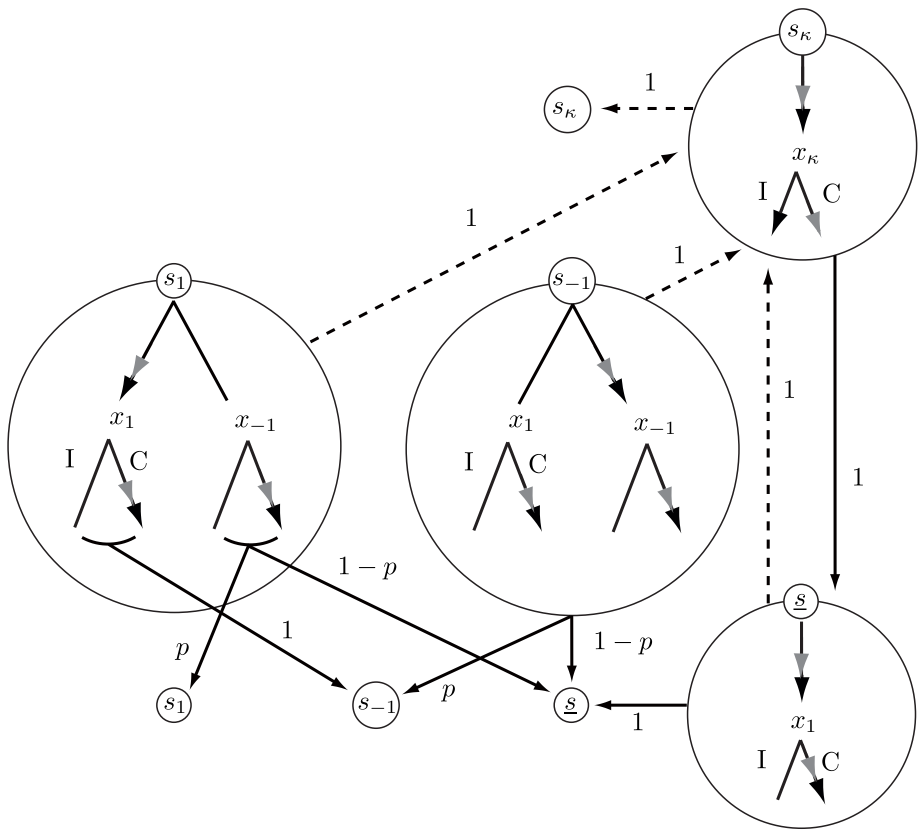

Figure 2: Ineligibility trap

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. □

Faithful delegates Theorem 3 provides conditions that guarantee the existence of a politician who is eligible for reelection, but it does not identify her type. In particular, the politician type whose weak reelection set is nonempty may conceivably vary with the Markov electoral equilibrium. This begs the question of identifying eligible politician types from the model's fundamentals, independently of the equilibrium. If we assume that some fixed voter type κ is representative in all states, and that politicians have mixed motivations, then politicians of type κ are natural candidates, as they share the representative voter's policy preferences. In particular, in the absence of electoral incentives, these politicians would choose optimal policies for the representative voter. In the following result, we strengthen the conditions of Theorem 3 to ensure that politicians who share the policy preferences of a fixed representative voter type κ have nonempty reelection sets in all equilibria. Additionally, these conditions are also sufficient to ensure that in all equilibria, politicians of type κ solve the representative voter's dynamic programming problem. We will see, however, that when the conditions of the theorem are violated, the type κ politician's desire for office may lead to the existence of states in which she is not reelected and is, in fact, ineligible for reelection.

Theorem 4. *Let σ be a Markov electoral equilibrium and assume that the type κ voter is representative in all s. Assume*

(D1) $Y_t(s)$ is independent of t for all s.

(D2) $p_t(s'|s,x,e)$ is independent of t and e for all s, s', and x,

(D3) mixed motives, i.e., $w_t(s,x) = u_t(s,x) + b_t$ for all s and all t,

and that at least one of the following hold:

1. *$\delta_\kappa = 0$,*
2. *$b_\kappa = 0$,*
3. *$p(s'|s,x)$ and $q_t(s,x)$ are independent of x for all s', s and t, and $p(s|s) > 0$ for all s.*

Then for all states s, we have $R(s,\kappa) \neq \varnothing$, and furthermore

$$V_\kappa^F(s,\kappa) \quad = \quad V_\kappa^*(s).$$

Note that the conditions of Theorem 4 pinpoint three possible specifications of the model under which the result holds: myopic citizens, pure policy motivation, and independence of the transition probabilities on states and challengers from policy choices. These specifications are extreme cases of the model, but upper hemicontinuity, from Theorem 2, implies that the expected payoff of the representative voter from the corresponding politician type will be close to optimal when any of these specifications is approximately satisfied.

The next example shows that under (D1)–(D3), if the other conditions of Theorem 3 are violated, then there may exist Markov electoral equilibria such that in some states, politicians of type κ have empty reelection sets, while in other states, they can secure reelection only by choosing policies that are suboptimal for the representative voter κ. In the latter case, conditional on being reelected, these politicians are choosing policies that they themselves would prefer not to choose. However, they are "cursed by their ambition," as they would be thrown out of office by the representative voter were they to choose good policies. These perverse incentives are the result of a coordination failure, due to the wedge driven between the congruent policy interest of voters and politicians of type κ by office benefits.

Example 4 (Curse of ambition). Assume the state space is $S = \{\underline{s}, s_1, s_{-1}, \overline{s}\}$, the type space is $T = \{1, \kappa\}$, and type κ voters are representative in all states. Let j range over $\{-1, 1\}$. Sets of feasible policies are independent of politicians' types and are such that $Y(\overline{s}) = \{\overline{x}, \underline{x}\}$, $Y(s_j) = \{\underline{x}, x_{-j}\}$ and $Y(\underline{s}) = \{\underline{x}\}$. State transition probabilities are independent of politicians' types and satisfy $p(s_{-1}|\overline{s}, \overline{x}) = p(s_1|\overline{s}, \underline{x}) = 1$, $p(s_{-j}|s_j, x_{-j}) = p(\underline{s}|s_j, \underline{x}) = 1$, and $p(\underline{s}|\underline{s}, \underline{x}) = 1$, so that all transitions are deterministic and state $\underline{s}$ is absorbing. Challenger selection probabilities are independent of states, policies, and incumbents' types and are such that $q(\kappa) = q(1) = \frac{1}{2}$. Note that all conditions (D1)–(D3) are respected. The payoffs to type κ voters are independent of states and are such that $u_\kappa(\overline{x}) > u_\kappa(x_1) = u_\kappa(x_{-1}) > u_\kappa(\underline{x})$. Politicians have mixed motivations with type-independent office benefit satisfying

$$b \quad \geqslant \quad \frac{(1 - \frac{1}{2}\delta)[u_\kappa(x_1) - u_\kappa(\underline{x})]}{\delta}.$$

Type 1 voters have state-independent stage utilities such that the inequalities $u_1(x_1) = u_1(x_{-1}) > u_1(\underline{x}) > u_1(\overline{x})$ hold. The common discount factor satisfies $\delta > 0$. Finally, note that the unique solution to the representative voter's dynamic programming problem calls for policy $\overline{x}$ in state $\overline{s}$, policy x_j in state s_{-j}, and policy $\underline{x}$ in state $\underline{s}$.

We claim that there exists a Markov electoral equilibrium such that $\pi_\kappa(\overline{x}|\overline{s}) = \pi_\kappa(\underline{x}|s_j) = \pi_\kappa(\underline{x}|\underline{s}) = 1$ and $\pi_1(\underline{x}|\overline{s}) = \pi_1(x_{-j}|s_j) = \pi_1(\underline{x}|\underline{s}) = 1$. In this equi-

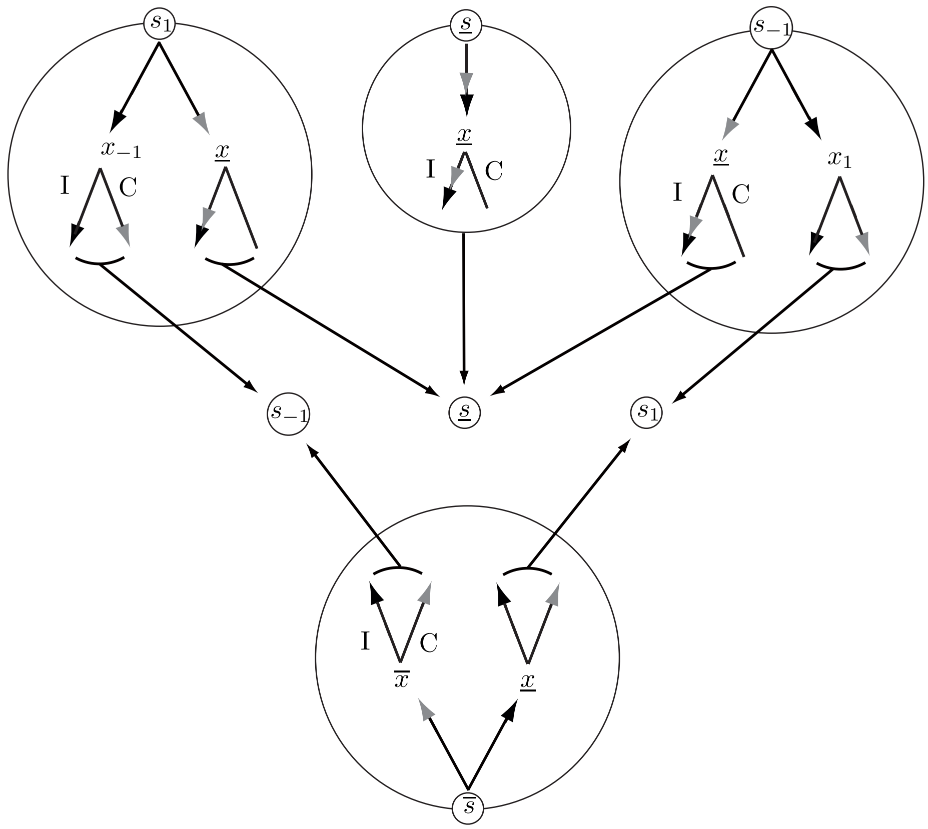

Figure 3: Curse of ambition

librium, a politician of type κ is ineligible for voter κ in state $\overline{s}$. By Theorem 3, politician 1 must therefore be eligible in that state, and in fact the representative voter strictly prefers to vote in favor of this politician following all policies in $\overline{s}$. To see this, note that voter κ expects a politician of type κ to implement policy $\underline{x}$ in all future periods following any policy choice in $\overline{s}$, while he expects a politician of type 1 to alternate between implementing x_1 and x_{-1}, which this voter strictly prefers. See Figure 3 for the equilibrium diagram, where dark arrows represent policy choices of type 1 politicians, and medium arrows represent choices of type κ politicians.

In this equilibrium, a politician of type κ is eligible in state s_j, but she is only reelected if she chooses policy $\underline{x}$, which does not correspond to the optimal policy in the representative dynamic programming problem. In fact, in that state, politicians of type 1 implement solutions to this problem (although these politicians implement suboptimal policies for κ voters in state $\overline{s}$). To see this, note that an almost identical argument to that above ensures that the representative voter supports the challenger following the choice of x_{-j} by type κ politicians in state s_j. Furthermore, the representative voter is indifferent between any incumbent and any challenger following the choice of policy $\underline{x}$ in state s_j, and, in this equilibrium, the voter always opts for the incumbent in this case, as well as in state $\underline{s}$. The office benefit is assumed to be high enough that a politician of type κ prefers retaining office in state s_j to implementing her ideal policy x_j in that state. □

Responsive politicians To this point, we have provided a partial responsiveness result to the effect that in all Markov electoral equilibria, the policy strategies of politicians who share the preferences of the representative voter solve the voter's dynamic programming problem. Given a particular equilibrium, a more demanding normative criterion for elections is that in each state, *all* politician types choose optimal policies of the representative voter. In this case, voters suffer no loss of control from the delegation of policy decisions to politicians. In the next result, we provide relatively general conditions under which there is at least one Markov electoral equilibrium that solves the representative dynamic programming problem. Moreover, we can restrict attention to equilibria in pure strategies: when there are multiple solutions to the representative dynamic programming problem, a politician's (pure) policy strategy selects the best policy for the politician from the voter's optimal policies. The key assumption is that politicians value office sufficiently, so that it is prima facie possible that they are willing to implement optimal policies of the representative voter type in equilibrium. Logically speaking, the other preconditions of the theorem are that the solutions of the representative dynamic programming problem are attainable when policies are determined by political competition, so that feasible policy sets and state transitions are independent of politicians' types, with the latter also independent of electoral outcomes.

Theorem 5. *Assume that the type κ voter is representative in all s. Assume (D1)–(D3), with $\delta_t > 0$ and b_t large for all t. Then there is a Markov electoral equilibrium $\sigma = (\pi, \rho)$ in pure strategies such that for all t, the strategy π_t of type t politicians is an optimal policy rule for the representative dynamic programming problem.*

A median voter result is established in the dynamic framework with a single state by Banks and Duggan (2008). They show that policies chosen by office holders of all types converge to the ideal policy of the representative voter when players are sufficiently patient, or when office benefits are sufficiently high; thus, the representative voter does not incur any cost from the delegation of policy choices to politicians in any equilibrium.[23] Theorem 5 shows that a form of the result extends to the general model with an endogenous state variable,

[23] Related results are obtained by Forand (2014), in a model with commitment, and by Van Weelden (2013), in a single-state model without commitment but with complete information about politicians' types.

with the dynamic analogue of the median ideal policy being the solution of the representative dynamic programming problem. The result does not rule out other equilibria in the multi-state model that fail to solve the representative dynamic programming problem, however, and Example 4, where no politician type chooses optimal policies for the representative voter in all states, demonstrates the possibility of such equilibria. Hence, in our general dynamic environment, electoral accountability cannot, on its own, ensure policy responsiveness.

Our final result strengthens the conclusion of the preceding theorem by establishing the asymptotic responsiveness for *all* Markov electoral equilibria: we provide conditions under which, as voters become patient, the representative voter's optimal payoff is approached in a strong sense. Under the conditions of Theorem 4, the policy choices of type κ politicians solve the representative voter's dynamic programming problem in all equilibria. One voting strategy, which may not be optimal for the representative voter, is to simply retain any type κ office holder and reject all other types. 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,[24] and therefore the representative voter's payoffs approach the optimal level. Two assumptions are key for the success of the previous simple voting rule at a given state. First, the state must recur with probability one 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

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

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

$$\Phi^m(\mathbf{s}) \quad = \quad \left\{ (x_0, \ldots, x_m) \middle| \begin{array}{c} x_0 \in Y(s_0) \text{ and } x_j \in Y(s_j) \\ \text{for all } j = 1, \ldots, m \end{array} \right\}$$

be the set of feasible paths of policies. Then define

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

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

[24]Duggan and Forand (2013) show that in a Markov electoral equilibrium, if a voter type is representative in all states, then the equilibrium voting strategy solves the associated optimal retention problem for the voter.

Theorem 6. *Let $\delta_\kappa = \delta \to 1$, and let $\{\sigma^\delta\}$ be corresponding Markov electoral equilibria such that given each δ, the type κ voter is representative in all s. In addition to the conditions of Theorem 4, assume that*

$$\min_{s,t,x} q_t(\kappa|s,x) > 0.$$

Then for all strongly recurrent states s and all t,

$$\lim_{\delta\to 1} \frac{V_\kappa^{F,\delta}(s,t)}{V_\kappa^{*,\delta}(s,\kappa)} \quad \geqslant \quad 1,$$

where $V_\kappa^{F,\delta}(s,t)$ denotes the expected discounted payoff to the type κ voter from electing a free type t politician in state s given strategy profile σ^δ, and $V_\kappa^{*,\delta}(s,\kappa)$ denotes this voter's optimal payoff in state s in the representative dynamic programming problem with discount factor δ.

Taken together, Theorems 3–6 inform us of the limits of electoral accountability in delivering responsive policy choices by political office holders. Our results provide conditions under which equilibrium policies accord with the preferences of a representative voter, and insofar as they establish the possibility of policy responsiveness, the results are broadly optimistic. To obtain the strongest concordance in Theorem 6, we impose the conditions of Theorem 4 and consider the effects of patience; in particular, in case the representative voter is not myopic and the corresponding politician type values office, the theorem assumes that the state transition probability is independent of policy choice. This still allows non-trivial dynamics, as the representative voter must anticipate policy choices of the incumbent and challenger in future states, but it implies that an optimal policy rule simply maximizes the representative voter's stage utility in each state. We give cautionary examples of political failures demonstrating that electoral accountability does not generally lead to responsiveness outside the conditions we provide. In connection to Theorem 6, observe that in Example 4, if $b \geqslant u_\kappa(x_1) - u_\kappa(\underline{x})$, then the selected Markov electoral equilibrium persists when δ is arbitrarily close to one, with the implication that the curse of ambition is resistant to the effects of patience, and that the conclusion of the theorem does not extend to the general model with positive office benefit and policy dependent state transitions. Thus, the identification of further positive results will rely on the imposition of additional structure within the general framework.

A Proof of Existence and Continuity

This appendix consists of the proof of Theorems 1 and 2. The fixed point argument 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.[25] Normalize utilities so that the images of $\frac{u_t}{1-\delta_t}$ and $\frac{w_t}{1-\delta_t}$ lie in $[0, \overline{u}]$ for all types. Recall that $v_{s,t,\tau}$ and $w_{s,t}$ denote expected discounted payoffs to voters and office holders, so we can assume $v_{s,t,\tau}, w_{s,t} \in [0, \overline{u}]$ for all s and all t. Let $\pi_{s,t}$ represents the mixture over policies played by an office holder who is free at s upon the initial transition to that state, i.e., given an equilibrium σ, $\pi_{s,t}$ corresponds to $\pi_t(\cdot|s)$. Then $\pi = (\pi_{s,t}) \in \Delta(X)^{S\times T}$ is the vector mixing probabilities, where $\Delta(\cdot)$ denotes the space of Borel probability measures endowed with the weak* topology.

Define the nonempty, convex product space

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

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

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

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

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

[25] This proof follows the lines of Duggan (2011), which establishes existence of equilibrium in a model of dynamic bargaining. See the latter paper for an informal discussion of the proof technique.

or equivalently,

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

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:

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

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 $x \in Y$,

$$\begin{aligned} W_t^B(s,x,\theta^+) \quad &= \quad w_t(s,x,\gamma) + \delta_t p_t(s|s,x,1,\gamma)W_t^B(s,x,\theta^+) \\ &\quad + \delta_t \sum_{s'\neq s} p_t(s'|s,x,1,\gamma)w_{s',t}, \end{aligned} \tag{5}$$

or equivalently,

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

and for all $x \in Z$,

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

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

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

and for each coalition C, define the correspondences

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

and as well define the correspondences

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

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 hemi-continuous, since it has open graph. Then Aliprantis and Border's (2006) Lemma 17.29 implies that the extended real-valued function

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

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

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

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

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

is lower semi-continuous. Now define

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

so, assuming $W^B_t(s,y,\theta^+) > W^B_\tau(s,\xi(y),\theta^+)$, we must have

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

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

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

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

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

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

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 $Y \cap Z = \varnothing$). Given s and θ^+, the correspondence $\mathcal{R}_t(s, \cdot, \theta^+)$ gives the reelection probabilities, as a function of the policy choice in s, that are consistent with the office holder's payoff $w_{s,t}$ in θ^+, but note that these reelection probabilities will not generally satisfy the conditions required in a Markov electoral equilibrium: it may be that $\hat{\rho}_\tau(s, x, \theta^+) < 1$ for some $x \in P_t(s, \theta^+)$, and it may be that $\hat{\rho}_\tau(s, x, \theta^+) > 0$ for some $x \in Y \backslash R_t(s, \theta^+)$. This discrepancy will be resolved after the fixed point argument. In any case, a voter's or politician's induced expected discounted utilities will be determined by the precise way that reelection probabilities depend on policies, i.e., by a selection from $\mathcal{R}_t(s, \cdot, \theta^+)$.

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

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

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

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

for all $x \in X$, all $y = (w, v, \gamma) \in Y$, and all $r \in [0, 1]$.[26] Let the correspondence

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

F consist of integrals of f with respect to $\mu = \pi_{s,t}$, i.e.,

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

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

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

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

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

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

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

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 $\operatorname{supp}(\pi_{s,t}) \subseteq \hat{P}_t(s,\theta^+)$, and therefore $f_t(s,\theta^+) \leqslant g_t(s,\theta^+)$. It follows that $w_{s,t} \in \hat{\mathcal{W}}_t(s,\theta^+) = [f_t(s,\theta^+), g_t(s,\theta^+)]$. In particular, this implies that for all $x \in \operatorname{supp}(\pi_{s,t})$, we have $W_t^B(s,x,\theta^+) \geqslant w_{s,t} \geqslant f_t(s,\theta^+)$. Let $\hat{\rho}_t(s,\cdot,\theta^+)$ be the selection of reelection probabilities such that for all s, all t, and all τ,

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

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

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

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 (6) reduces to $W_t^B(s,x,\theta^+)$. We have noted that $W_t^B(s,x,\theta^+) \geqslant w_{s,t} \geqslant f_t(s,\theta^+)$, and furthermore, $f_t(s,\theta^+) \geqslant \overline{Z}_t(s,\theta^+) \geqslant W_t^B(s,x,\theta^+)$. Combining these two inequalities, we have $w_{s,t} = W_t^B(s,x,\theta^+)$, as claimed.

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

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

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

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 choose $\xi(x)$ instead of x.

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

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

To define policy choice strategies, consider any state s and office holder of type t. We specify that the politician mixes according to $\pi_{s,t}$, modified to correct the discrepancy in Case 1 above. For x in the support of $\pi_{s,t}$ with $x \in P_t(s,\theta^+)$, so that $W_t^B(s,x,\theta^+) = W_t^B(s,\xi(x),\theta^+)$, 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 $A \subseteq X$,

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

and

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

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

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

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

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

By construction, and using the expression in (7), the values $V_\tau^B(\cdot,\theta^+)$, $V_\tau^C(\cdot,\theta^+)$, 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 (8) into (5), the values $W_t^B(\cdot,\theta^+)$ fulfill the recursive condition (4) that uniquely defines $W_t^B(\cdot;\sigma^*)$. Therefore,

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

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

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

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

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

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

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

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

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

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

as required.

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

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

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

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

for all Borel measurable $A \subseteq Y$.

To establish that $(\pi,w,v) \in \mathcal{F}(\pi,w,v,\gamma)$, define $\hat{\rho}_t\colon S \times X \to [0,1]$ as follows. Fix a state s. First, we specify that $\hat{\rho}_t(s,x) = \rho^*(s,t,x) = 0$ for all $x \in Z$. Second, for $x \in Y$ such that $W_t^B(s,x;\sigma^*) = W_t^B(s,\xi(x);\sigma)$, 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 $w_{s,t} \geqslant W^B(s,x;\sigma^*)$, 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

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

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 Proofs of Accountability Results

Proof of Theorem 3. Fix state s, and suppose, in order to deduce a contradiction that for all types t, $R(s,t) = \varnothing$. Thus, for all types t and all policies x, $\rho(s,t,x) = 0$ and

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

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

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

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

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

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

Proof of Theorem 4. Fix a Markov electoral equilibrium $\sigma = (\pi, \rho)$. A first remark is that under $(D1)$ and $(D2)$, if it is the case that $V_\kappa^F(s,\kappa) = V_\kappa^*(s)$ for all states s, then we must also have $R(s,\kappa) \neq \varnothing$. To see this, note that for all $x' \in \operatorname{supp}(\pi_\kappa(\cdot|s))$, we have

$$x' \in \operatorname{argmax}_{x \in Y(s)} \;\; u_\kappa(s,x) + \delta_\kappa \sum_{s'} p(s'|s,x) V_\kappa^*(s'),$$

and that correspondingly

$$\begin{aligned} V_\kappa^B(s,\kappa,x') &= \frac{p(s|s,x')u_\kappa(s,x') + \sum_{s'\neq s} p(s'|s,x)V_\kappa^F(s',\kappa)}{1-\delta_\kappa p(s|s,x')} \\ &= \sum_{s'} p(s'|s,x)V_\kappa^*(s') \\ &\geqslant V_\kappa^C(s,\kappa,x'), \end{aligned}$$

as desired, where the inequality follows since $V_\kappa^F(s,t) \leqslant V_\kappa^*(s)$ for all politician types t.

To prove part 1, suppose that $\delta_\kappa = 0$. It follows that

$$x' \in \text{argmax}_{x \in Y(s)} \quad u_\kappa(s,x) + \delta_\kappa \sum_{s'} p(s'|s,x)V_\kappa^*(s'),$$

and furthermore

$$x'' \in \text{argmax}_{x \in Y(s)} \quad \rho(s,\kappa,x)W_\kappa^B(s,x) + (1-\rho(s,\kappa,x))W_\kappa^C(s,x),$$

if and only if

$$x', x'' \in \text{argmax}_{x \in Y(s)} \quad u_\kappa(s,x),$$

so that $V_\kappa^F(s,\kappa) = V_\kappa^*(s)$, as desired.

To prove part 2, suppose that $b_\kappa = 0$. We claim that the equilibrium profile $\sigma = (\pi, \rho)$ must maximize the joint payoffs of type κ voters and politicians, which implies that $V_\kappa^F(s,\kappa) = V_\kappa^*(s)$ for all s. Suppose, towards a contradiction, that an alternative Markov strategy profile achieved a higher joint payoff for these types than profile σ. A first possibility is that, under σ, voters of type κ make reelection decisions about politicians of type κ that are not optimal for these politicians. In that case, by the principle of optimality and recalling (4), there exists a state s and a policy x such that either

(a) $\rho(s,\kappa,x) < 1$ and

$$\frac{W_\kappa^B(s,x) - w_\kappa(s,x)}{\delta_\kappa} > \frac{W_\kappa^C(s,x) - w_\kappa(s,x)}{\delta_\kappa}, \text{ or}$$

(b) $\rho(s,\kappa,x) > 0$ and

$$\frac{W_\kappa^B(s,x) - w_\kappa(s,x)}{\delta_\kappa} < \frac{W_\kappa^C(s,x) - w_\kappa(s,x)}{\delta_\kappa},$$

where, given part 1, we impose that $\delta_\kappa > 0$. Note that, since $b_\kappa = 0$, we have

$$W_\kappa^B(s,x) = u_\kappa(s,x) + \delta_\kappa \left[\rho(s,\kappa,x)V_\kappa^B(s,\kappa,x) + (1-\rho(s,\kappa,x)V_\kappa^C(s,\kappa,x))\right]$$

and

$$W_\kappa^C(s,x) \quad = \quad u_\kappa(s,x) + \delta_\kappa V_\kappa^C(s,\kappa,x).$$

Hence, if case (a) obtains, then it follows that

$$V_\kappa^B(s,\kappa,x) > V_\kappa^C(s,\kappa,x),$$

which in turn implies that $\rho(s,\kappa,x) = 1$, yielding the desired contradiction. Similarly, if (b) obtains it must be that $V_\kappa^C(s,\kappa,x) > V_\kappa^B(s,\kappa,x)$ and that $\rho(s,\kappa,x) = 0$, yielding the desired contradiction.

The second possibility is that politicians of type κ make policy decisions under σ that are not optimal for the representative voter type. In that case, there are a state s and policies $x \in Y(s)$ and $x' \in \text{supp}(\pi_\kappa(\cdot|s))$ such that

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

But then

$$\begin{aligned} & \rho(s,\kappa,x)W_\kappa^B(s,x) + (1-\rho(s,\kappa,x))W_\kappa^C(s,x) \\ > & \, \rho(s,\kappa,x')W_\kappa^B(s,x') + (1-\rho(s,\kappa,x'))W_\kappa^C(s,x'), \end{aligned}$$

yielding the desired contradiction.

To prove part 3, first note that if $p(s'|s,x)$ is independent of x for all s' and s, we have that

$$x' \in \text{argmax}_{x \in Y(s)} \;\; u_\kappa(s,x) + \delta_\kappa \sum_{s'} p(s'|s)V_\kappa^*(s'),$$

if and only if

$$x' \in \text{argmax}_{x \in Y(s)} \; u_\kappa(s,x),$$

Fix some state s and suppose that $R(s,\kappa) = \varnothing$, so that $\rho(s,\kappa,x) = 0$ for all x. Hence, we have

$$x' \in \text{argmax}_{x \in Y(s)} \;\; w_\kappa(s,x) + \delta_\kappa V_\kappa^C(s',\kappa),$$

if and only if

$$x' \in \text{argmax}_{x \in Y(s)} \;\; u_\kappa(s,x),$$

where we use the fact that $V_\kappa^C(s',\kappa,x)$ is independent of x when $q_t(t'|s,x)$ is independent of x. Now suppose there exists $x \in R(s,\kappa)$ such that $x \notin$

$\text{argmax}_{x' \in Y(s)}\ u_\kappa(s,x')$. In that case, for $\tilde{x} \in \text{argmax}_{x' \in Y(s)}\ u_\kappa(s,x')$, using $p(s|s) > 0$, we have that

$$\begin{aligned} V_\kappa^B(s,\kappa,\tilde{x}) &= \frac{p(s|s)u_\kappa(s,\tilde{x}) + \delta_\kappa \sum_{s' \neq s} p(s'|s) V_\kappa^F(s',\kappa)}{1 - \delta_\kappa p(s|s)} \\ &> V_\kappa^B(s,\kappa,x) \\ &\geqslant V_\kappa^C(s,\kappa), \end{aligned}$$

so that $\rho(s,\kappa,\tilde{x}) = 1$. Hence, it follows that $\pi\left(\arg\max_{x \in Y(s)} u_\kappa(s,x)|s\right) = 1$, so that $V_\kappa^F(s,\kappa) = V_\kappa^*(s)$, as desired. □

Proof of Theorem 5. Fix a state s, and 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 Markov strategy profile in pure strategies. Note that $V_\kappa^F(s,t) = V_\kappa^*(s)$. Second, note that given state s, politician type t and a choice of $x \in Y(s)$, the payoff to player κ is $u_\kappa(s,x) + \delta_\kappa \tilde{V}_\kappa(s,x)$, where $\tilde{V}_\kappa$ satisfies

$$\tilde{V}_\kappa(s,x) = \sum_{s'} p(s'|s,x) \int_{x'} [u_\kappa(s',x') + \delta_\kappa \tilde{V}_\kappa(s',x')] \pi_{s',t}^*(dx'),$$

and that, given any t, $V_\kappa^B(s,t,x) \leqslant \tilde{V}_\kappa(s,x)$ for all x, with equality if and only if $x \in X^*(s)$. Then for all s, all t, and all x,

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

so that $V^B_\kappa(s,t,x) \leqslant V^C_\kappa(s,t,x)$, with equality if and only if $x \in X^*(s)$. Thus, the voting rule ρ satisfies the conditions for equilibrium. 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 $0 \leqslant u_t \leqslant \overline{u}$, assuming b_t sufficiently large, and using $\delta_t > 0$, we have

$$W^B_t(s, x^*_t(s)) \;\geqslant\; \frac{b_t}{1-\delta_t} \;\geqslant\; b_t + \frac{\overline{u}}{1-\delta_t} \;\geqslant\; W^C_t(s,x),$$

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

Proof of Theorem 6. Fix δ, and let $\sigma^\delta = (\pi^\delta, \rho^\delta)$ be a Markov electoral equilibrium given δ. From Theorem 4, it follows that if a politician of type κ is in office in any state s, then $V^{F,\delta}_\kappa(s,\kappa) = V^{*,\delta}_\kappa(s,\kappa)$. Let $\hat{\rho}$ denote the voting strategy ρ^δ modified as follows: for all s and x, $\hat{\rho}(s,t,x) = 1$ if $t = \kappa$ and $\hat{\rho}(s,t,x) = 0$ otherwise. Letting $\hat{V}^{F,\delta}_\kappa(s,t)$ denote the expected discounted payoff to the type κ voter from electing a free type t politician in state s given strategy profile $\hat{\sigma}^\delta = (\pi^\delta, \hat{\rho})$, we have that $\hat{V}^{F,\delta}_\kappa(s,\kappa) = V^{*,\delta}_\kappa(s,\kappa)$. The profile $\hat{\sigma}^\delta$ may not itself be an equilibrium, but because the equilibrium voting strategy solves the representative voter type's optimal retention problem, it must be the case that for all s and all t, we have $V^{F,\delta}_\kappa(s,t) \geqslant \hat{V}^{F,\delta}_\kappa(s,t)$.

For all s, let $\alpha = \min_{s,t,x} q_t(\kappa|s,x) > 0$, 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 $u_\kappa \geqslant 0$, the representative voter's expected discounted utility from electing a free type t politician satisfies

$$V^{F,\delta}_\kappa(s,t) \;\geqslant\; \hat{V}^{F,\delta}_\kappa(s,t) \;\geqslant\; p^m(s)(1-(1-\alpha)^m)\delta^m V^{*,\delta}_\kappa(s,\kappa)$$

for all m, and this implies

$$\frac{V^{F,\delta}_\kappa(s,t)}{V^{*,\delta}_\kappa(s,\kappa)} \;\geqslant\; p^m(s)(1-(1-\alpha)^m)\delta^m$$

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

$$\sup_m p^m(s)(1-(1-\alpha)^m)\delta^m \;\geqslant\; (1-\epsilon)\delta^{m(\epsilon)} \;\to\; 1-\epsilon$$

as $\delta \to 1$. Since ϵ was arbitrary, the desired inequality follows. □

C Proofs for Examples (Not for Publication)

Proofs for Example 1. In an equilibrium with compromise in state $\hat{s}$, all politician types implement policy $\hat{x}_\kappa$ and are reelected. In such an equilibrium, the payoff to a politician of type $t \in \{\ell, r\}$ from implementing policy $\hat{x}_\kappa$ is $\frac{u+b}{1-\delta}$, while her 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

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

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 $\hat{x}_\kappa$ 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 $\hat{u} + b + \frac{1}{3}\delta\left[V_t^F(\hat{s}, \ell) + V_t^F(\hat{s}, \kappa) + V_t^F(\hat{s}, r)\right]$. 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 $\hat{x}_\kappa$ 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 $\hat{x}_\kappa$ is $V_t^F(\hat{s}, \kappa) + \frac{b}{1-\delta}$, an equilibrium with shirking in state $\hat{s}$ exists if and only if

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

which, since

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

holds if and only if

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

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

$$\frac{\delta b}{1-\delta} \leqslant 2(\hat{u}-u) - (u-\check{u}).$$

Since $\hat{u} - u > \frac{1}{2}(u - \check{u})$ 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 $\hat{x}_\kappa$ is $V_t^F(\hat{s}, \kappa)$, while his payoff from the challenger is $\frac{1}{3}\left[V_\tau^F(\hat{s}, \ell) + V_\tau^F(\hat{s}, r) + V_\tau^F(\hat{s}, \kappa)\right]$. Hence, type τ voters support the incumbent if and only if $\hat{u} - u \leqslant u - \check{u}$, which holds strictly by assumption.

We now turn to state $\mathring{s}$. 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 $\mathring{s}$ 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 his middle-ranked policy x. If the state transitions, then all politician types are expected compromise at policy $\hat{x}_\kappa$. Hence, we have that $V_\tau^B(\mathring{s}, t, x) - V_\tau^C(\mathring{s}) \geqslant 0$ if and only if $\hat{u} - u \leqslant u - \check{u}$, 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 $\mathring{p} \leqslant \bar{p}$, there exists an equilibrium in which all politicians implement their ideal policy in state $\mathring{s}$ 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 $\mathring{s}$ are given by

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

This implies that, for any transition probability $\mathring{p}$, we have that $V_\tau^C(\mathring{s}) = V_\tau^C(\hat{s})$, since $V_\tau^B(\mathring{s}, \kappa, \hat{x}_\kappa) = V_\tau^B(\hat{s}, \kappa, \hat{x}_\kappa)$. 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 $\mathring{s}$ is given by

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

We show that if $\mathring{p}$ 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 $\tau \neq t$, $\lim_{\mathring{p} \to 0} V_\tau^B(\mathring{s}, t, \hat{x}_\tau) = u_\tau(\hat{s}, \hat{x}_t) + \delta V_\tau^C(\hat{s})$ 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 $\mathring{p} \leqslant \bar{p}$, voters of type $\tau \neq t \in \{\ell, r\}$ support the challenger against an incumbent of type t that has implemented policy $\hat{x}_\tau$. Since, for any policy x, $V_\tau^B(\mathring{s}, t, x) \leqslant V_\tau^B(\mathring{s}, t, \hat{x}_\tau)$, then, for any $\mathring{p} \leqslant \bar{p}$, 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 $\hat{x}_\kappa$. It is sufficient to show that these politicians always obtain the support of voters of type ℓ. This follow by the assumption that $\hat{u} - u < u - \check{u}$ since, as noted above, $V_\ell^B(\mathring{s}, t, \hat{x}_\kappa) = V_\ell^B(\hat{s}, t, \hat{x}_\kappa)$ and $V_\ell^C(\mathring{s}) = V_\ell^C(\hat{s})$. □

Proofs for Example 2. In the absorbing state $\mathring{s}$, 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 his second-ranked policy x implemented by a politician of type t, we have that

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

while his payoff to a challenger is

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

Hence, we have that $V_\tau^B(\mathring{s}, t, x) \geqslant V_\tau^C(\mathring{s})$ 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

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

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

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

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

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

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

A simple computation verifies that $\bar{p}_1 < 1$. Since the denominator of (9) 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

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

while

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

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

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

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 $\bar{p} = \min\{\bar{p}_1, \bar{p}_2\}$. □

Proofs for Example 3. The aim is to show that there exists a 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 $s \in \{s_1, s_{-1}\}$, $R(s,t) = \varnothing$. 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

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

and

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

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

$$\begin{aligned}
& V_\kappa^F(s_{-t}, -t) - V_\kappa^F(s_{-t}, t) && (9)\\
&= u_\kappa(x_1) - u_\kappa(x_{-1}) + \delta_\kappa\left[V_\kappa^C(s_{-t}, -t, x_1) - V_\kappa^C(s_{-t}, t, x_{-1})\right] \\
&= u_\kappa(x_1) - u_\kappa(x_{-1}) + \delta_\kappa(1-p)\Bigg[\frac{1}{2}V_\kappa^F(s_{-t}, -t) + \frac{1}{2}V_\kappa^F(s_{-t}, t) \\
&\qquad - \frac{1}{1-\delta_\kappa}u_\kappa(\underline{x})\Bigg] \\
&> 0. && (10)
\end{aligned}$$

Furthermore, we assume that p satisfies

$$0 \; < \; p \; < \; 1 - \frac{u_\kappa(x_1) - u_\kappa(x_{-1})}{\delta_\kappa\left[\frac{1}{2}u_\kappa(x_1) + \frac{1}{2}u_\kappa(x_{-1}) - u_\kappa(\underline{x})\right]}, \tag{11}$$

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 $V_\kappa^B(s_t, t, x_1) = V_\kappa^F(s_{-t}, t)$,

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

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 $V_\kappa^B(s_t, t, x_{-1}) = (1-p)\frac{u_\kappa(\underline{x})}{1-\delta_k} + pV_\kappa^F(s_t, -t)$, we have that

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

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

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

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

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 (10), this condition is equivalent to

$$
\begin{aligned}
&u_\kappa(x_1) - u_\kappa(x_{-1}) && (12) \\
&\quad < \delta_\kappa(1-p)\left[\frac{1}{2}V_\kappa^F(s_{-t},-t) + \frac{1}{2}V_\kappa^F(s_{-t},t) - \frac{1}{1-\delta_\kappa}u_\kappa(\underline{x})\right]. && (13)
\end{aligned}
$$

Since $V_\kappa^F(s_{-t},-t) \geqslant u_\kappa(x_1) + \frac{\delta_\kappa}{1-\delta_\kappa}u_\kappa(\underline{x})$ and $V_\kappa^F(s_{-t},t) \geqslant u_\kappa(x_{-1}) + \frac{\delta_\kappa}{1-\delta_\kappa}u_\kappa(\underline{x})$, a sufficient condition for (13) is the assumption (11) 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

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

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

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

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. □

Proofs for Example 4. We start by verifying that voting strategies for representative voter type κ are optimal. First, given any policy $x \in Y(\overline{s})$, we have that

$$\begin{aligned} V_\kappa^B(\overline{s}, 1, x) &= V_\kappa^F(s_t, 1) \\ &= \frac{u_\kappa(x_1)}{1-\delta} \\ &> \frac{u_\kappa(\underline{x})}{1-\delta} \\ &= V_\kappa^F(s_t, \kappa) \\ &= V_\kappa^B(\overline{s}, \kappa, x), \end{aligned}$$

which, since $V_\kappa^C(\overline{s}, x) = \frac{1}{2}\left[V_\kappa^F(s_t, \kappa) + V_\kappa^F(s_t, 1)\right]$, implies both that voting in favor of type 1 politicians in state $\overline{s}$ is optimal for voter κ, and that type κ politicians are ineligible in that state.

Second, we have that

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

$$
\begin{aligned}
&= \frac{u_\kappa(x_1)}{1-\delta} \\
&> \frac{u_\kappa(\underline{x})}{1-\delta} \\
&= V_\kappa^F(s_{-t}, \kappa) \\
&= V_\kappa^B(s_t, \kappa, x_{-t}),
\end{aligned}
$$

which, since $V_\kappa^C(s_t, x_{-t}) = \frac{1}{2}\left[V_\kappa^F(s_{-t}, \kappa) + V_\kappa^F(s_{-t}, 1)\right]$, implies that voting in favor of type 1 politicians and against type κ politicians following policy x_{-t} in state s_t is optimal for voter κ.

Third, for any type $\tau \in \{\kappa, 1\}$, we have that

$$
\begin{aligned}
V_\kappa^B(s_t, \tau, \underline{x}) &= V_\kappa^F(\underline{s}, \tau) \\
&= \frac{u_\kappa(\underline{x})}{1-\delta},
\end{aligned}
$$

which, since $V_\kappa^C(s_t, \underline{x}) = \frac{1}{2}\left[V_\kappa^F(\underline{s}, \kappa) + V_\kappa^F(\underline{s}, 1)\right]$, implies that voting in favor of all politician types following policy $\underline{x}$ in state s_t is optimal for voter κ.

Fourth, since all politicians are trivially eligible in state $\underline{s}$, supporting incumbents of all types is optimal for voter κ in that state.

Next, we verify that policy strategies are optimal. First, for any policy $x \in Y(\overline{s})$, the payoff to a politician of type κ in state $\overline{s}$ is

$$
W_\kappa^C(\overline{s}, x) = u_\kappa(x) + b + \delta V_\kappa^C(\overline{s}, \overline{x}),
$$

so that implementing her ideal policy $\overline{x}$ in state $\overline{s}$ is optimal for that politician. Similarly, the payoff from policy x to a politician of type 1 in state $\overline{s}$ is

$$
W_1^B(\overline{s}, x) = u_1(x) + b + \delta W_1^B(s_t, x_t),
$$

so that implementing her ideal policy $\underline{x}$ in state $\overline{s}$ is optimal for that politician.

Second, the payoff to a politician of type 1 in state s_t from implementing policy x_{-t} is

$$
\frac{u_1(x_1) + b}{1-\delta},
$$

while her payoff from implementing policy $\underline{x}$ is

$$
\frac{u_1(\underline{x}) + b}{1-\delta},
$$

so that implementing policy x_t is optimal for this politician in that state.

Third, the payoff to a politician of type κ in state s_t from implementing policy $\underline{x}$ is

$$
\frac{u_\kappa(\underline{x}) + b}{1-\delta},
$$

while her payoff from implementing policy x_{-t} is

$$u_\kappa(x_1) + b + \frac{\delta}{2(1-\delta)}\left[u_\kappa(x_1) + u_\kappa(\underline{x})\right],$$

so that implementing policy $\underline{x}$ is optimal for this politician in that state since

$$b \geqslant \frac{(1-\frac{1}{2}\delta)[u_\kappa(x_1) - u_\kappa(\underline{x})]}{\delta}.$$

Finally, since $b > 0$ and incumbents that implement policy $\underline{x}$ in state $\underline{s}$ are reelected, running for reelection is optimal for all politician types in that state.

□

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., 2008. A dynamic model of democratic elections in multidimensional policy spaces. Quarterly Journal of Political Science 3, 269–299.

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

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., 2014. A dynamic theory of electoral competition. Theoretical Economics 9 (2), 515–554.

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., Forand, J. G., 2013. Markovian elections, unpublished paper.

Duggan, J., Kalandrakis, T., 2012. Dynamic legislative policy making. Journal of Economic Theory 147 (5), 1653–1688.

Federgruen, A., 1978. On n-person stochastic games with denumerable state space. Advances in Applied Probability 10, 452–471.

Forand, J. G., 2014. Two-party competition with persistent policies. Journal of Economic Theory 152, 64–91.

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., McLennan, A., 2014. Corrigendum to: Discounted stochastic games with no stationary nash equilibrium: Two examples. Econometrica forthcoming.

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.

Van Weelden, R., 2013. Candidates, credibility, and re-election incentives. The Review of Economic Studies 80 (4), 1622–1651.

Yared, P., 2010. Politicians, taxes and debt. Review of Economic Studies 77, 806–840.