

Accountability in Markovian Elections

John Duggan*

Jean Guillaume Forand†

March 16, 2021

Abstract

We study the possibility that elections achieve electoral accountability by bringing the incentives of politicians in line with those of a representative voter, in the context of a general dynamic environment. As our normative benchmark, we take the solution of the dynamic programming problem facing the voter, as if she chose policy directly. We show that there always exist equilibria in which the politician type corresponding to the representative voter is accountable, in the sense that she implements a voter-optimal policy rule and is rewarded by re-election. Furthermore, when politicians are highly office motivated, there exist equilibria in which all politician types are accountable. We demonstrate that challenges to electoral accountability stem from multiple equilibria with undesirable normative properties, and we give examples of novel political failures in a model of dynamic public investment. Finally, we identify a class of responsive voting equilibria in which the representative voter succeeds in leveraging electoral power to achieve asymptotic accountability: for any selection of such equilibria, policies implemented by politicians converge to voter-optimal policies as the voter becomes patient.

Keywords: Electoral Accountability; Dynamic Models; Median Voter

1 Introduction

The electoral process has the potential, by subjecting incumbents to periodic review by voters, to discipline office holders and bring policy choices in line with voters' preferences. This is so even if politicians do not share these preferences, so long as the value of holding office provides a sufficient incentive for incumbents to put aside their own policy preferences and to compete with the option of a challenger. When elections take place in a dynamic environment, two distinctive challenges to the efficacy of elections present themselves, both stemming from the absence of intertemporal commitment. First, in any given election, candidates may find it difficult to make credible promises about their policy choices in future environments, so that even a candidate who would be willing to bind herself to popular policies in order to gain re-election has no way of doing so. Second, voters also have no way of committing to future re-election standards, so they cannot incentivize politicians by offering re-election in exchange for desirable policies. Because office holders' expectations

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

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

of future electoral prospects drive their current actions in office, and because voters' expectations of politicians' future policy choices drive their current electoral decisions, political accountability is inherently vulnerable to this dual commitment problem.

Problems of commitment are well known in political economy and have been explored in citizen-candidate models of elections (Besley and Coate (1997); Osborne and Slivinski (1996)), the determination of macroeconomic policy (Alesina (1987)), and democratic transitions (Acemoglu and Robinson (2000)). The class of dynamic models is particularly important for applications, but they present special challenges because of the complexity of political decision-making: the choices of one actor can influence the value of a state variable, and this can in turn affect the choices of future actors. These difficulties are multiplied by the richness of dynamic electoral incentives, for in these settings, an office holder typically needs to choose between policies that satisfy voters' standard for re-election and policies that sacrifice office for short run gains. Not only must this calculus be forward-looking and take into account the future impacts of current policy choices, but the bar for re-election introduces a non-convexity into the optimization problem of an office holder. For these reasons, existence and characterization of equilibria in models of dynamic elections is difficult in general, although some progress has been made in environments with more structure, where the time horizon is short, the state space is small, persistence across time is limited, elections involve substantial uncertainty, or political institutions reduce to the decision of a single actor.

We present a dynamic framework for elections that develops the standard citizen-candidate model in a general but parsimonious way: in each period of the model, a political or economic state is given and an incumbent office holder chooses policy from a feasible set; then a challenger is drawn and an election is held; and then a new state is realized, and so on. In the spirit of the citizen-candidate approach, we preclude commitment by either politicians or voters. We allow for rich, dynamic environments with an arbitrary compact metric space of policies, an arbitrary, finite set of states,¹ an arbitrary, finite set of citizen types, and any continuous transition probabilities on the political state and the challenger's type. To focus on the effectiveness of elections when the incumbent's policy choices are linked to future policies through the state, we abstract from private information about office holders' types. In contrast much work on electoral accountability, a politician's type is observed when she takes office, which shuts down incentives to signal a persistent type and avoids difficult technical issues that would take away from our focus.² We do allow for uncertainty about the challenger's type before the election, so that voters may have less information about a challenger relative to the known incumbent.

In this setting, we establish existence of *Markov electoral equilibria*, we examine the welfare properties of these equilibria, and we apply our analysis to a new class of multi-state spatial models. We begin by observing that in general dynamic environments, the representative voter's ideal point will depend on the state, and moreover, if policy choices

¹The results of the model extend to the model with a countably infinite set of states in a straightforward way; we assume finiteness only for exposition.

²Our framework could be extended by adding a privately observed type that decays over time, as in work on political budget cycles by Rogoff (1990).

enter the transition probability on states, then the appropriate normative benchmark is not simply the voter’s static ideal point. Rather, it is the value of the representative voter’s dynamic programming problem, and the analogue to the median ideal policy is the set of policy rules that solve this *representative dynamic programming problem*. We establish the general existence of an electoral equilibrium such that the *congruent* politician type, i.e., the politician type corresponding to the voter, is *accountable*, in the sense that she uses a strategy that solves the representative dynamic programming problem and is always re-elected.³ We also observe that if politicians have sufficiently high office incentives, then there is an equilibrium in which all politician types are accountable. Thus, there always exists an equilibrium such that the congruent politician type delivers the voter’s optimal value, and when politicians are highly office motivated, there is an equilibrium that also solves the accountability problem of non-congruent politicians. These results raise two questions: Can there be equilibria in which electoral incentives prevent congruent politicians from choosing voter-optimal policies? And are there conditions under which politicians are necessarily disciplined by elections?

Before addressing these questions, we note that our framework stacks the deck against accountability: we do not allow politicians to commit to policies prior to taking office; moreover, while we assume the existence of a representative voter, we do not allow voters to coordinate on a re-election rule to optimize the performance of politicians. In the one-period model, familiar from the literature on citizen-candidates, this combination of assumptions implies that the incentives of politicians are trivial, as the elected politician simply chooses her ideal policy before the game ends. But when the horizon is extended (either finite or infinite), office holders in all but the final period must make non-trivial strategic calculations. In any such period, the incumbent’s policy choice can influence the evolution of the state, and therefore the policy choices of future politicians. Because the incumbent may be more desirable to the voter in some states than others, and also because her policy choice can affect the distribution of the challenger’s type, the voter’s calculus thus also generally depends on the incumbent’s policy choice. That is, in sharp contrast to the static model, the policy choices of an office holder can determine her re-election prospects and the identity of future office holders. We emphasize that this dynamic linkage arises naturally in the model, in the absence of commitment or any form of pre-electoral politics.

The dynamic linkage creates the possibility of electoral accountability, but it also leads to new forms of political inefficiency through the scope for state manipulation by incumbents. We show that if politicians are office motivated and the evolution of the state depends on policy choices, then there can exist equilibria in which a congruent politician chooses policies that are harmful to the voter (and herself) in order to increase her chances of retaining office, a phenomenon we refer to as the *curse of ambition*. In such examples, the politician is trapped by the expectations of the voter and forced to choose between the voter-optimal policy, which leads to removal from office, or a suboptimal policy that ensures victory. For the simplest version of the curse of ambition, assume there are two states: an absorbing state of perpetual war (with no policy choice), and a peaceful state that can be

³We assume a state-independent representative voter type for simplicity. Duggan and Forand (2021) show that equilibrium existence extends to the model with state-dependent representative voter, and they give sufficient conditions for existence of a representative voter in each state.

tipped into war by the choice of the office holder. Assume there are two types, one of which corresponds to the voter, and assume that in terms of policy payoffs, both types prefer the peaceful state, but assume also that office rents are so large as to outweigh the cost of war in the minds of the politicians. It is then easy to construct an inefficient equilibrium: we specify that in the peaceful state, the congruent politician type always chooses to go to war, and the other type opts to maintain peace; in the war state, the voter always re-elects the non-congruent politician, and he re-elects the congruent politician if and only if she chooses war; and in the war state, the voter always re-elects the incumbent, regardless of type.⁴ The voter’s strategy in the peaceful state, while unintuitive, is optimal, because the congruent politician, if re-elected after choosing peace, is expected to start a war in the next period. For the congruent politician, the strategy of starting a war is optimal, because it is the only way she can avoid removal from office.

Underlying the curse of ambition is the dual commitment problem we highlighted at the outset, which leads to coordination failure between the voter and the congruent politician. Specifically, the congruent politician’s inability to commit to peace in the future gives the voter incentives to replace her if she chooses peace today, and conversely, the voter’s inability to commit to rewarding peace with re-election gives the politician incentives to start a war. We therefore explore the welfare properties of a smaller class of equilibria in which the re-election strategy of the voter is more tightly connected to the choices of congruent politicians, in the following sense: in each state, a congruent politician is re-elected if the voter’s discounted payoff strictly exceeds a threshold, and there is at least one policy that satisfies the threshold and guarantees re-election of the incumbent.⁵ Clearly, this condition is violated by the equilibrium construction in the simple curse of ambition example, because in the peaceful state, the voter is willing to re-elect the politician after she starts a war but not after she maintains peace, although the latter choice improves the voter’s payoff. This restriction on voting strategies integrates components from notions of both retrospective and prospective voting in the literature on voting behavior, and we term it *k-responsive voting*, to indicate that in each state, the voter responds positively to at least one, sufficiently good policy choice by the corresponding politician type. We establish that in a responsive voting equilibrium, the voter and congruent politician type overcome the commitment problem, and the strategies used by these politicians are optimal policy rules for the voter. Exploiting the dynamic environment of our framework, we show that this, in turn, leads to an *asymptotic accountability* result as the voter becomes patient: given any sequence of *k*-responsive voting equilibria, the expected discounted payoffs of the voter converge to the optimal value of the representative dynamic programming problem.

For non-congruent politicians, differences in preferences can easily lead to equilibria in which the politicians do not choose optimally for the voter. In fact, the scope for state manipulation can create a *political hold-up problem*, in which incumbents of all types are re-elected, despite choosing policies that are undesirable for the voter, even if the congruent type herself is accountable. Although perhaps surprising, the intuition for this possibility

⁴In this simple example, the non-congruent politician type actually chooses optimally for the voter, but it is straightforward to augment the example with a policy choice in the peaceful state that is suboptimal for the voter, but still more desirable than going to war.

⁵This class is always nonempty, by the general equilibrium existence result stated above.

lies in the fact that once a policy is chosen in a given state, whether or not it is bad for the voter, that cost is sunk; the relevant consideration for the voter is the distribution over policies implied by that choice, and the expected performance of politicians in those states. This allows us to support an equilibrium in which a non-congruent politician chooses policies that are undesirable for the voter, but is nonetheless re-elected, because the next state is likely to be one in which other types are even worse. We then extend the notion of responsive voting to all types, and we define a *responsive voting equilibrium* as one such that the voter gives each type a threshold that is sufficient for re-election, and such that each politician type has at least one policy that achieves her threshold. When office incentives are high, we establish that the asymptotic accountability of congruent politicians serves to discipline all other types: given any sequence of responsive voting equilibria as the voter becomes patient, the policy choices of all politician types converge to solutions of the representative dynamic programming problem.^{6,7}

We highlight our asymptotic optimality result and our examples of dynamic political failures in a new environment, the *multi-state spatial model*, that is conducive to many applications of interest. This special case of our framework augments the classical spatial model with a state variable that evolves stochastically as a function of a one-dimensional policy choice. The state variable provides a dynamic linkage across periods, and it can represent the size of the economy, the level of public debt, the distribution of income or wealth, etc. In our applications, we interpret the state as the stock of a durable public good, divided into discrete categories, and the policy choice in a given period as the amount of public investment; the citizen types possess a common ordering of the set of states in our examples, but this is not generally required. For tractability, we assume that stage utilities are concave, and the transition probability is linear in policy choices. This ensures that the representative dynamic programming problem in the multi-state spatial model is well-behaved and can be solved analytically when there are two states. Because the voter's optimization problem is strictly concave, there is a unique optimal policy rule that specifies a particular level of investment in each state; and unless the voter is myopic, the optimal policy in any state strictly exceeds the voter's stage ideal point, reflecting the value of maintaining a high stock of public goods in the future. These voter-optimal policies then provide the normative benchmark against which electoral equilibria are compared.

Literature A recurring theme of the dynamic political economy literature is that commitment problems are critical for understanding policy outcomes and evaluating electoral performance. The assumption that politicians can commit to policies in one-shot elections, standard since the work of Downs (1957), has often been contested, notably in the context of citizen-candidates models (Besley and Coate (1997); Osborne and Slivinski (1996)). Extending such commitment to sequences of policy choices is even more debatable (Alesina and Rodrik (1994); Bertola (1993)), and what is now a large literature studies the dynamic policy consequences of office holders' inability to make credible campaign proposals. For example, Alesina (1987, 1988) made early contributions to the topic of political cycles by

⁶Recall that when office incentives are high, there is an equilibrium in which all politician types are accountable, so the set of responsive voting equilibria is nonempty.

⁷We also show that exact optimality holds in responsive voting equilibria, regardless of the voter's discount factor, when office benefits are high and the state transition is independent of policy.

formulating policymaking as a game between parties that cannot commit to policy instruments prior to an election;⁸ 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; and Acemoglu et al. (2008) and Yared (2010) describe the distortions in tax policies that are necessary to provide rent-seeking politicians with the incentives to limit their extractive activities. Banks and Duggan (2008) prove an asymptotic accountability result in the one-dimensional model with adverse selection, analogous to our asymptotic result in Theorem 6.1. Commitment failures are accentuated in models with term limits, e.g., Banks and Sundaram (1998), Bernhardt et al. (2004), and Besley and Case (1995), where subgame perfection directly implies that politicians choose their ideal policies in the last term of office. Duggan (2017) shows that this incentive leads to an upper bound on equilibrium payoffs of voters, irrespective of the strength of office incentives, highlighting the role of the voters’ commitment problem.

Because voters cannot commit to re-elect politicians after good policy choices, politicians may anticipate government turnover and choose poor policies, reinforcing the voters’ incentive to remove the politician. Persson and Svensson (1989) and Alesina and Tabellini (1990) show that incumbents may distort public policies in order to “tie the hands” of potential successors with different preferences. Besley and Coate (1998) find that politicians may fail to implement Pareto-improving investments if they anticipate that future policy-makers will not realize their returns. In a dynamic legislative bargaining setting, Battaglini and Coate (2008) show that legislators’ uncertainty about being included in future governing coalitions drives them to approve excessive pork barrel spending. In models of two-party competition, Azzimonti (2011, 2015) shows that the prospects of government turnover can lead to inefficiencies in either private or public capital accumulation, Battaglini (2014) shows that when voters’ preferences are time-varying, temporarily powerful districts can attract inefficiently high levels of government spending, and in a finite-horizon model, Callander and Raiha (2017) describe how parties distort investments in a durable public good to compel voters to re-elect them, which is closely connected to our study of the political failures due to state manipulation.⁹ Finally, Bai and Lagunoff (2011) and Duggan and Forand (2021) show that these concerns are magnified when future office holders are determined by current policy choices, so that they further distort policies in order to affect the identities of their successors.¹⁰

More broadly, our paper is related to the literature in economics and political science that studies the possibility and limits of electoral accountability (for comprehensive surveys, see Duggan and Martinelli (2017) and Ashworth (2012)). Existing work in this literature often focuses on the effects of politicians’ private information, either in the form of adverse selection or moral hazard, and it imposes extensive structure on the electoral environment

⁸Persson and Tabellini (2000) provide an extensive overview of the macropolitical economy literature, in which time inconsistency and lack of credible commitments plays a large role.

⁹The setting of Callander and Raiha (2017) precludes a “curse of ambition,” because neither party shares the preferences of the voter, but their main result illustrates a “political hold-up problem,” in which the initial incumbent makes investments that, from the voter’s point of view, the challenger would capitalize on poorly.

¹⁰Anticipated turnover in power has also been associated with political inefficiencies outside the realm of electoral competition, e.g., with policy “gridlock” in legislatures (Bowen et al. (2014) and Dziuda and Loeper (2016)) or with the creation of ineffective public administrations (Acemoglu et al. (2011)).

to achieve tractability. In comparison, we assume symmetric information (except for the possibility of uncertainty about the challenger’s type), but we consider very general dynamic environments. The single-state version of our model is related to the dynamic model of elections with adverse selection considered by Duggan (2000) and Bernhardt et al. (2004).¹¹ In that work, politicians are privately informed about their preferences, and there is no state variable; in equilibrium, an office holder’s initial policy choice is expected to remain the same over time if she is re-elected. There, the persistence of the information disclosed about the incumbent’s type by her initial policy acts as a commitment device, which is absent from our setting. In a more general version of our model with state-dependent representative voters, Duggan and Forand (2021) allow incumbents the option to commit to policies, thus bridging the gap with single-state models of dynamic elections with adverse selection. In this paper, the dynamic linkages across periods only occurs through the impact of current decisions on future states and challenger types.

Outline of the paper The remainder of the paper is organized as follows. Sections 2 and 4 set forth the model and describe the equilibrium concept, respectively. The intervening Section 3 describes the multi-state spatial model, which extends the classical spatial model of politics to our dynamic environment and serves as our main application throughout the paper. Section 5 establishes the existence of equilibria in which, respectively, only congruent politicians or all politician types are guaranteed to choose voter-optimal policy plans. Section 6 introduces the notion of responsive voting equilibria and details our asymptotic accountability results for equilibria in this class. Section 7 concludes, and the Appendix contains the proofs of all results.

2 Dynamic Electoral Framework

Political environment A representative voter decides between an incumbent politician and a challenger in an infinite sequence of elections. The voter is assigned a type k from the finite set T , and we assume an infinite pool of politicians of each type, with a politician’s type typically denoted t . We sometimes refer to a politician who is type k (the same type as the voter) as *congruent*. Politician types are initially private information and are independently distributed. Each period begins with a state s and a politician who holds office, the state and the office holder’s type being publicly observed. The office holder chooses a policy x ; a challenger whose type is private information is selected; an election is held; a new state is realized, and the winner’s type, along with the state, is publicly observed; and the process repeats.¹² We assume that states belong to a finite set S ; that policies belong to a compact metric space X ; and that in every state s , the set of feasible policies is a nonempty, closed (and therefore compact) subset $X(s)$ of X .

¹¹See further applications include the analysis of competence (Meirowitz (2007)), parties (Bernhardt et al. (2009)), valence (Bernhardt et al. (2011)), and taxation (Camara (2012)). Duggan (2014) provides a folk theorem for the model when non-Markovian equilibria are permitted.

¹²It is straightforward to allow the incumbent the option of not running for re-election, albeit at the cost of additional notation (see Duggan and Forand (2021)). As this choice plays no role in our results in this paper, we omit it for simplicity.

Payoffs The stage utility of a type t citizen from policy x in state s is $u_t(s, x)$, while a politician who holds office receives an additional office benefit $\beta \geq 0$. We assume that $u_t: S \times X \rightarrow \mathfrak{R}$ is continuous and bounded, so that for all s , all x , and all t , we have $\underline{u} \leq u_t(s, x) \leq \bar{u}$. For convenience, we normalize these bounds so that $\underline{u} = 0$ and $\bar{u} = 1$. Each type t citizen discounts flows of payoffs by the factor $\delta \in [0, 1)$. Thus, given a sequence $(s_1, x_1, s_2, x_2, \dots)$ of state-policy pairs, the discounted payoffs of a type t citizen is

$$\sum_{\ell=1}^{\infty} \delta^{\ell-1} [u_t(s_\ell, x_\ell) + I_\ell \beta],$$

where I_ℓ is an indicator function taking value one if the citizen holds office in period ℓ and zero otherwise.

State transitions States are used to describe the political and/or economic environment in the current period. Given that an office holder chooses a policy x in state s , a new state s' is drawn with probability $p(s'|s, x)$: thus, states evolve according to a controlled Markov process. We assume that the transition probability $p: S \times S \times X \rightarrow [0, 1]$ is continuous and that it has full support, i.e., $p(s'|s, x) > 0$ for all states $s', s \in S$ and policies $x \in X$. This assumption is not needed for our existence results in Section 5, namely, Theorem 5.1 and Proposition 5.1, but our asymptotic accountability results in Section 6 require some notion of recurrence for states, and we assume that the state transition places positive probability on all states for expositional simplicity. The dependence of future states on current policy choices underpins an incumbent's incentives to manipulate the state to her electoral advantage, and Examples 2 and 3 illustrate how these incentives can drive failures of accountability.

Challengers After the office holder chooses policy, a challenger is drawn from the pool of politicians who have never held office, so the challenger's type is not observed by voters before the election. Challenger selection depends on the current incumbent and the previous state and policy choice: let $q_t(t'|s, x)$ denote the probability that the challenger is type t' , given that a type t incumbent chose policy x in state s . We assume that the transition probability on challenger types, $q_t: T \times S \times X \rightarrow [0, 1]$, is continuous for each type t . We assume that the challenger is congruent with positive probability in each state and following any policy choice by any incumbent: $q_t(k|s, x) > 0$ for all t, s and x . This assumption is not required for our equilibrium existence results in Section 5, but the possibility of eventually drawing a type k challenger plays an important role in our asymptotic accountability results, and we impose the assumption at the outset to simplify the statements of these results. The dependence of future challengers on current policy choices has the potential to bias incumbents in favour of policies that engender weaker opponents. Although we do not focus on such "challenger manipulation" in this paper, all our results allow for its possibility.

Representative dynamic programming problem Given our assumption of a representative voter, our normative benchmark in the analysis of accountability is the optimal value for the voter in the associated *representative dynamic programming problem*, in which the voter directly chooses any policy $x \in X(s)$ in state s and receives utility $u_k(s, x)$, the next state s' is realized from $p(\cdot|s, x)$, and so on. Under our maintained compactness and continuity conditions, this program has a unique value V_k^* , which solves the associated

Bellman equation: for all s ,

$$V_k^*(s) = \max_{x \in X(s)} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) V_k^*(s').$$

Let $\Phi^*(s)$ denote the set of voter-optimal policies in state s , i.e.,

$$\Phi^*(s) = \arg \max_{x \in X(s)} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) V_k^*(s').$$

A *policy rule* is a mapping $\phi: S \rightarrow X$ such that for all s , we have $\phi(s) \in X(s)$, i.e., ϕ assigns a feasible policy to each state. By the optimality principle, a policy rule ϕ^* is voter-optimal if and only if it selects from the correspondence of voter-optimal policies, i.e., for all s , we have $\phi^*(s) \in \Phi^*(s)$.

Remark 1

(a) Our general formulation of policies and preferences captures finite models, as well as multidimensional policy spaces with smooth, well-behaved stage utilities. The dependence of the feasible set on the state allows us to interpret s as a state of the economy, which can affect the set of feasible policies.

(b) Although we assume the state transition has full support for simplicity, this is not needed for our equilibrium existence results, and our asymptotic accountability results extend to the model with more general state transitions, including deterministic ones. It is straightforward to add one or more absorbing states to capture finite-horizon models, and we can add a “calendar” that tracks the term of office of an incumbent to incorporate term limits in the model.

(c) To capture two-party competition, we can partition the set of types into disjoint sets D and R , with $D \cup R = T$, and restrict challenger transitions as follows: for all states s and policies x , $q_t(R|s, x) = 1$ for all $t \in D$ and $q_t(D|s, x) = 1$ for all $t \in R$. That is, given an incumbent from the party in power, the challenger is always drawn from the party out of power, with the challenger transition reflecting the nomination process of that party. To apply our asymptotic accountability results, we technically need the type k politician to belong to both parties, but it is straightforward to circumvent this formality by assuming there are two politician types, say k^D and k^R , each belonging to a different party and sharing the stage utility of the voter. The results of Section 6 are proved under the assumption that $q_t(k|s, x) > 0$, but they extend to the two-party model, as long as each party contains a type that corresponds to the voter and has positive probability.

(d) The canonical model used to represent politics over a single issue dimension is the *classical spatial model*: the set $X \subseteq \mathfrak{R}$ of policy outcomes is a closed interval, and each citizen type t has continuous, strictly quasi-concave utility over policies. Competition between two candidates in a majority-rule election can be layered on this model of voting in different ways: in the Downsian approach, the candidates simultaneously commit to policy platforms to win the election; at the opposite extreme, the citizen-candidate approach assumes the candidates cannot commit to policy, so that voters anticipate that each candidate will choose her ideal policy if elected, and they vote accordingly. Our framework rules out commitment

and can be viewed as a dynamic extension of the citizen-candidate approach. We generalize the classical model to the dynamic framework in the next section. \square

3 Multi-State Spatial Model

In this section, we specialize our policy environment to a dynamic version of the classical spatial model, with multiple states and non-trivial dynamics. To begin, assume that $X \subseteq \mathfrak{R}$ is a closed interval, and that for each state s and type t , the stage utility function $u_t(s, x)$ can be decomposed as follows: there exist parameters $\omega_t, \zeta_t \in \mathfrak{R}$ with $\omega_t > 0$ for each type t and mappings $v: S \times X \rightarrow \mathfrak{R}$ and $c: S \times X \rightarrow \mathfrak{R}$ such that for all s , all t , and all x , we have

$$u_t(s, x) = \omega_t v(s, x) - c(s, x) + \zeta_t. \quad (1)$$

An advantage of this functional form is that it provides microfoundations for our assumption of a representative voter. Indeed, given a model of the voting population, where the proportion of type t voters is $\mu_t \geq 0$, if there is no subset $C \subseteq T$ such that $\sum_{t \in C} \mu_t = \frac{1}{2}$, then Duggan and Forand (2021) show that there is a type k citizen that is representative, in the sense that in every state and after every policy choice, the majority winner of the election is the preferred candidate of the type k citizen.

We define the *multi-state spatial model* by adding the assumptions that stage utility is strictly concave, and that the state transition probability is linear in policy:

- $X \subseteq \mathfrak{R}$ is a closed interval,
- for each state s and each type t , $u_t(s, x)$ has the functional form in (1), with $v(s, x)$ continuous and concave in x and $c(s, x)$ continuous and strictly convex in x ,
- for all states s and s' , the transition probability $p(s'|x)$ is linear in x .

Let \hat{x}_t^s denote the unique maximizer of $u_t(s, \cdot)$. The representative dynamic programming problem, which has Bellman equation

$$V_k(s) = \max_{x \in X} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) V_k(s'),$$

has unique solution V_k^* , and under the concavity assumptions of the multi-state spatial model, the representative voter has a unique optimal policy rule $\phi^*: S \rightarrow X$. The voter-optimal policy rule ϕ^* is the relevant normative benchmark in the multi-state spatial model, rather than the median ideal point in any state, and in fact, Theorem 3 of Duggan and Forand (2021) implies that among all policy rules, the optimal rule ϕ^* of the representative voter is a Condorcet winner; that is, it beats all other rules in pairwise majority votes, making it the natural extension of the median ideal point to the dynamic setting.

In the remainder of this section, we impose further structure on the multi-state spatial model to solve the representative dynamic programming problem analytically and to simplify the examples that follow. We assume that $X = [0, 1]$ is the unit interval, that there

are just two states, $S = \{s^1, s^2\}$, and that the state transition depends on policy but not on the current state; for simplicity, we will just assume $p(s^2|x) = x$ and $p(s^1|x) = 1 - x$. Moreover, assume that stage utility $u_t(x, s^j)$ is quadratic in x , i.e.,

$$u_t(s^j, x) = -(x - \hat{x}_t^j)^2 + \Gamma_j,$$

with ideal point \hat{x}_t^j and state-dependent constant term Γ_j , and that ideal points are ordered by type: indexing types as $T = \{1, \dots, n\}$, we assume that for each state s^j , we have $\hat{x}_1^j \leq \hat{x}_2^j \leq \dots \leq \hat{x}_n^j$. Finally, to capture the idea that the higher state is better for citizens, we assume that ideal points and constant terms are increasing in the state, i.e., for all t ,

$$\hat{x}_t^1 \leq \hat{x}_t^2 \quad \text{and} \quad \Gamma_2 \equiv \Gamma > 0 = \Gamma_1.$$

To conserve on notation, we write the voter's optimal values as a vector, $V_k^* = (V_k^{1*}, V_k^{2*})$, so that the representative voter's optimal value reduces to

$$V_k^{j*} = \max_{x \in Y} -(x - \hat{x}_k^j)^2 + \Gamma_j + \delta[xV_k^{2*} + (1-x)V_k^{1*}] \quad (2)$$

for each $j = 1, 2$.

We interpret the model as one of public investment, where the policy x in any period is an investment that confers present benefits and costs. It may be, for example, that x is a level of infrastructure spending, and that a state s^i measures the stochastically depreciating stock of public good. We allow ideal points to vary with the state to capture the idea that public investment crowds out present consumption, and that because the economy is tighter in lower states, the opportunity cost of investment is higher, leading to lower ideal levels of investment and to higher ideal investment levels in higher states. A key feature of the multi-state spatial model is that unless voters are myopic, the median ideal policy in any state is *not* the relevant normative benchmark: each voter would prefer a policy choice above his stage ideal policy, because it increases the probability of transitioning to the higher state. Figure 1 depicts the two-state model with stage utilities for three citizen types in black, along with the representative voter's induced utility, which incorporates the future benefits of transitioning to the high state, in gray. As the figure suggests, the representative voter's optimal policy choices in states s^1 and s^2 , denoted x^{1*} and x^{2*} , shift to the right by an increment that is determined by the parameters of the model and reflects the voter's willingness to trade off present consumption for future public good.

Solving the first order condition of the Bellman equation (2), the voter's optimal choice in state s^j satisfies

$$x_k^{j*} = \hat{x}_k^j + \frac{\delta}{2}[V_k^{2*} - V_k^{1*}], \quad (3)$$

and in particular, since $V_k^{2*} > V_k^{1*}$, the voter's induced ideal policy is higher than the stage ideal point in both states.¹³ Henceforth, let $\Delta = V_k^{2*} - V_k^{1*}$ denote the difference in

¹³For some parameterizations, the optimal policy will be a corner solution, i.e., $x_k^{j*} = 1$. For ease of exposition, we assume both in this section and in the examples that follow that solutions are interior and a first-order analysis can be used. This can be guaranteed if, for example, Γ is sufficiently low.




Figure 1: Induced utility in the multi-state spatial model

continuation values at the two states. Since $x_k^{j*} = \hat{x}_k^j + \frac{\delta}{2}\Delta$ solves the voter's maximization problem, we can substitute this into the Bellman equation to obtain

$$V_k^{j*} = -\left(\frac{\delta}{2}\Delta\right)^2 + \Gamma_j + x_k^{j*}\delta\Delta + \delta V_k^{1*}. \quad (4)$$

Taking the difference between the two states, we see that

$$\Delta = \Gamma + \delta(x^{2*} - x^{1*})\Delta = \Gamma + \delta(\hat{x}_k^2 - \hat{x}_k^1)\Delta,$$

and we solve for Δ to obtain the difference in continuation values,

$$\Delta = \frac{\Gamma}{1 - \delta(\hat{x}_k^2 - \hat{x}_k^1)},$$

as a function of the parameters. Substituting into (3), we obtain the voter-optimal policies,

$$x_k^{j*} = \hat{x}_k^j + \frac{\delta\Gamma}{2(1 - \delta(\hat{x}_k^2 - \hat{x}_k^1))}, \quad (5)$$

and the voter's optimal values can be solved directly using (4).

Given the optimal policy rule of the representative voter, the evolution of the state and policy is inherently dynamic, transitioning between low and high states with probability that is endogenously determined. In particular, if Γ is close to zero, then the voter's problem becomes close to the static problem of the classical spatial model, and the optimal

policies converge to the stage ideal points, i.e., $x^{j*} \rightarrow \hat{x}_k^j$. However, for a given $\Gamma > 0$, the optimal policy of the voter exceeds the stage ideal policy, reflecting the future returns to current investment. The voter invests more in the higher state, and thus the probability of transitioning to the high state is higher when beginning from this state: it is more likely to maintain the high state than it is to move to s^2 from s^1 . The influence of the future naturally increases as the voter becomes more patient, and the limiting optimal policy is

$$x_k^{j*} \rightarrow \hat{x}_k^j + \frac{\Gamma}{2(1 - (\hat{x}_k^2 - \hat{x}_k^1))}$$

as $\delta \rightarrow 1$. We use the solution of the representative dynamic programming problem as our normative criterion in the dynamic electoral framework, and we pose the problem of accountability as that of achieving the representative voter's optimal value in each state. We will see that at the heart of the problem is equilibrium multiplicity, and that the performance of elections largely hinges on whether the voter's responses effectively tie current policy choices to future electoral rewards.

4 Markov Electoral Equilibrium

Strategies A *policy strategy* for a type t politician is a mapping $\pi_t: S \rightarrow \Delta(X)$, where $\Delta(X)$ is the set of Borel probability measures on X , and $\pi_t(\cdot|s)$ represents the mixture over policies used by the type t politician in state s . Let $\pi = (\pi_t)_t$ denote a profile of policy strategies. A *voting strategy* is a Borel measurable mapping $\rho: S \times T \times X \rightarrow [0, 1]$, where $\rho(s, t, x)$ represents the probability that a type t office holder is re-elected following a policy choice of x in state s . Let $\sigma = (\pi, \rho)$ denote a profile containing both policy and voting strategies.

Continuation values Given a strategy profile σ , we can define continuation values for a type t citizen. The discounted expected policy utility of the citizen from electing a type t' incumbent who chooses policy x in state s satisfies: for all $x \in X(s)$,

$$V_t^I(s, t', x) = \sum_{s'} p(s'|s, x) V_t(s', t'), \quad (6)$$

where $V_t(s, t')$ is the expected discounted utility to the citizen from a type t' office holder in state s , calculated before a policy is chosen. The discounted expected utility of electing a challenger following the choice of x in state s by a type t' office holder is defined by

$$V_t^C(s, t', x) = \sum_{t''} q_{t''}(t''|s, x) \sum_{s'} p(s'|s, x) V_t(s', t''). \quad (7)$$

Finally, $V_t(s, t')$ is given by

$$V_t(s, t') = \int_x \left[u_t(s, x) + \delta[\rho(s, t', x) V_t^I(s, t', x) + (1 - \rho(s, t', x)) V_t^C(s, t', x)] \right] \pi_{t'}(dx|s), \quad (8)$$

reflecting that the office holder chooses a policy x according to the policy strategy $\pi_t(\cdot|s)$, and is either re-elected or replaced by a challenger.

In addition to payoffs from policies, a type t office holder evaluates future expected discounted office benefit from choosing policy x in state s , conditional on being re-elected. For all $x \in X(s)$, we define this as follows,

$$B_t(s, x) = \sum_{s'} p(s'|s, x) \int_{x'} [\beta + \delta \rho(s', t, x')] B_t(s', x') \pi_t(dx'|s),$$

reflecting the fact that the office holder receives β in the period following her re-election and, conditional on choosing policy x' in the next state s' and being re-elected again, receives $B_t(s', x')$ in the future.

Re-election sets Given a strategy profile $\sigma = (\pi, \rho)$ and policy choice x in state s by a type t incumbent, the representative 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 and all incumbent types t , the sets

$$\begin{aligned} P_k(s, t) &= \{x \in X(s) : V_k^I(s, t, x) > V_k^C(s, t, x)\} \\ R_k(s, t) &= \{x \in X(s) : V_k^I(s, t, x) \geq V_k^C(s, t, x)\} \end{aligned}$$

of policies that yield the voter an expected discounted payoff strictly and weakly greater, respectively, than the expected discounted payoff of a challenger. We refer to these as the *strict* and *weak re-election sets*, respectively. Note that choosing $x \in R_k(s, t)$ is necessary for re-election of a type t incumbent in state s , and choosing $x \in P_k(s, t)$ is sufficient.

Equilibrium concept A 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 the representative voter in all states. Formally, we require that (i) for all s and all t , $\pi_t(\cdot|s)$ puts probability one on solutions to

$$\max_{x \in X(s)} u_t(s, x) + \beta + \delta \left[\rho(s, t, x) (V_t^I(s, t, x) + B_t(s, x)) + (1 - \rho(s, t, x)) V_t^C(s, t, x) \right],$$

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

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

where $\rho(s, t, x)$ is unrestricted if $x \in R_k(s, t) \setminus P_k(s, t)$.¹⁴ Intuitively, a type t office holder maximizes her stage utility plus future expected discounted payoff, which combines policy utility and office benefit (in case the politician is re-elected) and the continuation value of a challenger (in case the politician loses).

¹⁴In the special case of the framework in which the state and challenger transition probabilities are independent of policy, the policy choice of the incumbent is not payoff-relevant for the voter at the electoral stage. Nevertheless, our concept allows the voter to condition on policy, consistent with the notion of stationary Markov perfect equilibrium from the literature on stochastic games.

Accountable politicians In the analysis of accountability, it is useful to have a designation for politician types who respond to electoral incentives by choosing policies that solve the representative dynamic programming problem and who are rewarded with electoral success by the voter. Formally, given a strategy profile σ , we say that a type t politician is *accountable* if for each state s , (i) $V_k(s, t) = V_k^*(s)$, and (ii) $\int_x \rho(s, t, x) \pi_t(dx|s) = 1$.

5 Existence of Equilibria with Accountable Politicians

In this section, we establish existence of Markov electoral equilibria in which some politician types are accountable; thus, once elected, they choose optimal policies for the representative voter and remain in office thereafter. In such an equilibrium, the optimal value of the representative dynamic programming problem may be achieved eventually—if the initial office holder is accountable, or if an incumbent is removed and replaced by an accountable challenger—but it is also possible that a politician type that is not accountable assumes office and is never removed. We present an example showing that there can exist equilibria in which no type is accountable, due to the “curse of ambition,” and we give another example of an equilibrium in which the type k politician is accountable, yet because the voter faces a “political hold-up problem,” other politician types are never removed from office, despite choosing policies that are suboptimal from the perspective of the voter. We first prove the general existence of an equilibrium in which the type k politician is accountable.

Theorem 5.1. *There is a Markov electoral equilibrium in which the type k politician is accountable.*

The proof follows from an application of the equilibrium existence result in Theorem 1 of Duggan and Forand (2018), which allows the state transition to depend on the incumbent’s type and the electoral outcome, and which does not assume the state transition places positive probability on all states.¹⁵ To apply that result, we transform our model by specifying that if a type k incumbent is removed from office, then the game moves to a bad state and remains at that state thereafter. Theorem 1 of Duggan and Forand (2018) then delivers an equilibrium $\tilde{\sigma}$ of the transformed model in which type k politicians are always re-elected, regardless of policy choice. This removes the wedge between the incentives of the type k politician and the voter, and this allows us to deduce that the equilibrium strategy $\tilde{\pi}_k$ of the type k politician is optimal for the voter. We then map this equilibrium to a strategy profile σ of the original model, maintaining policy strategies and modifying $\tilde{\rho}$ so that for all s and all x , the type k politician is re-elected with probability one if and only if $V_k^I(s, k, x) \geq V_k^C(s, t, x)$. This preserves equilibrium conditions of $\tilde{\sigma}$, and we conclude that σ is an equilibrium in which the type k politician is accountable.

In Example 1, below, we construct an equilibrium of the two-state spatial model such that the congruent politician type is accountable. That construction is simplified by the assumption that there are only two types, but when there are more than two types, existence

¹⁵The framework of Duggan and Forand (2018) does not assume a representative voter, and it allows general politician payoffs.

of equilibrium is a challenging problem. To see that the proof cannot proceed by simply constructing the desired equilibrium, suppose that we specify a strategy profile in which the type k politician implements an optimal policy rule and the voter re-elects her in all states. Clearly, when the type k politician holds office, neither the politician nor the voter can benefit from adopting any other policy or voting strategies, so the relevant strategic interaction takes place between the voter and the remaining politician types. Abstracting from voting (say each type $t \neq k$ is removed with some exogenous probability), the politician types play a dynamic game that is non-trivial, because an incumbent's choices will influence the evolution of the state, and thus the policy choices of future politicians; if some politician types are more desirable in some states than others, then this consideration will affect the policy choice of the incumbent, and of course all other politician types make similar calculations. Adding the voter, the technical challenges are multiplied, as the optimization problem of an office holder is potentially discontinuous: one policy choice may create voter indifference and lead to re-election, while nearby policies may generate a strict preference for the challenger and lead to removal from office. Such discontinuities are typical in bargaining games (Duggan, 2017), and they require delicate fixed point arguments.

Remark 2 The electoral game between politician types and the voter can be formulated as a stochastic game, in which policy choice stages alternate with voting stages. In a policy choice stage, a state s and an incumbent type t are given, and the incumbent's policy choice x leads to a choice of the voter. The state variable for the voting stage in the stochastic game must contain all relevant information for the voter, i.e., the state s , the incumbent's type t , and also the incumbent's policy choice x . The latter component is problematic for equilibrium existence: in the stochastic game formulation of the electoral model, the state variable in the voting stage contains a component, x , that depends on the politician's choice in a deterministic way. Unless the set of feasible policies is finite, this violates the condition of norm continuity used in existence of stationary Markov equilibria (Duggan (2012); Nowak and Raghavan (1992)), and it is known that otherwise well-behaved stochastic games with deterministic state transitions may fail to have a stationary Markov perfect equilibrium (Levy (2013)). In the absence of norm continuity, subgame perfect equilibria exist in general dynamic games, but existence rests on atomless moves by nature (Barelli and Duggan (2020); He and Sun (2020)), which are not present here. When nature's moves have atoms, Luttmer and Mariotti (2003) show that even continuous games with sequential moves can fail to possess a subgame perfect equilibrium. \square

Next, in the multi-state model with two states and two types, we give a constructive example of an equilibrium in which the type k politician is accountable.

Example 1 (Accountable type k politician). Consider the two-state spatial model with quadratic utility, state-independent transition probabilities, and two citizen types, $T = \{1, 2\}$, with $k = 2$. Assume that the challenger is type k with probability $q \in (0, 1)$, and that the type $1 \neq k$ politician's ideal policy is state-independent and to the left of the voter's: $\hat{x}_1^1 = \hat{x}_1^2 = \hat{x}_1 < \hat{x}_k^1 < \hat{x}_k^2$. To construct a Markov electoral equilibrium in which the congruent politician is accountable, we specify that this type implements the voter-optimal rule (x_k^{1*}, x_k^{2*}) as described by (5), and that this politician is re-elected in all states following all policies: $\rho(s^j, k, x) = 1$ for all j and all x . Clearly, this voting strategy is optimal because, given that the type k politician is choosing optimal policies in all future

states irrespective of her current policy choice, no challenger can be strictly preferred to the type k politician by the voter. It remains to determine the equilibrium policy strategy of the type 1 politician, as well as the re-election rule that the voter applies to this type. Because there are only two politician types, and type k is accountable, it follows that a type 1 incumbent can be at least as good for the voter as the challenger in any state only if she implements the voter's optimal rule in all states. Therefore, we can focus on two types of equilibria: one in which type 1 is accountable, and another in which type 1 shirks and implements her preferred policies conditional on being replaced.

An equilibrium in which type 1 is never re-elected always exists. To see this, specify the voting strategy $\rho(s^j, 1, x) = 0$ for all j and all x , and note that this is optimal for the voter irrespective of the policy choices of type 1, because type k is accountable. Using first order conditions as in (3), the equilibrium investment of type 1 in all states is

$$x_1 = \hat{x}_1 + \frac{\delta}{2} (q[V_1(s^2, k) - V_1(s^1, k)] + (1 - q)[V_1(s^2, 1) - V_1(s^1, 1)]). \quad (9)$$

Note that even if the type 1 politician is not re-elected in state s^j , her equilibrium policies must be best-responses to the policies she expects future office holders to implement. Her choices trade off two concerns. First, the type 1 politician has incentives to influence the state to provide insurance in case a type k politician comes to power. For example, if $q \approx 1$ and $V_1(s^1, k) > V_1(s^2, k)$, so that type 1 prefers the low state under the voter-optimal plan, then her equilibrium policy $x_1 < \hat{x}_1$ underinvests in all states relative to her own optimal rule, which, according to (5) applied to type 1, would be to invest $x_1^* = \hat{x}_1 + \frac{\delta\Gamma}{2}$ in all states. Second, the type 1 politician has incentives to align her choices with potential type 1 replacements. For example, if $q \approx 0$, then (9) reduces to (5), and successive type 1 politicians approximately implement their optimal rule even if they each hold office only for one term.

Whether an equilibrium with an accountable type 1 politician exists depends on this politician's willingness to compromise on policy to remain in office. To construct such an equilibrium, we specify a voting strategy such that the type 1 politician is re-elected if and only if she implements the voter-optimal rule: for all j , $\rho(s^j, 1, x_k^{j*}) = 1$ and $\rho(s^j, 1, x) = 0$ otherwise. In equilibrium, the accountable type 1 obtains the payoff $V_1(s^j, k) + \frac{\beta}{1-\delta}$ in state s^j . If she deviates, then her best option is to choose a best-response against a challenger who will implement the voter-optimal rule with probability one irrespective of her type. The optimal such policy in state s^j is given by (9) with $q = 0$. Let $V_1^D(s^j) + \beta$ denote the payoff to this best deviation in state s^j . Hence, type 1 is incentivized to be accountable in equilibrium if and only if $\frac{\delta\beta}{1-\delta} \geq \max\{V_1^D(s^2) - V_1(s^2, k), V_1^D(s^1) - V_1(s^1, k)\}$, which is satisfied if $\delta\beta$ is sufficiently high. \square

The Markov electoral equilibrium established in Theorem 5.1 exists generally, but it delivers no restrictions on the policy choices of type $t \neq k$ politicians. Because these types have different policy goals than the voter, the possibility of accountability depends on their willingness to compromise their policy choices, and this rests on their desire to stay in office. As anticipated in Example 1, the following result shows that equilibria in which all type

$t \neq k$ politicians are accountable exist whenever they value future office benefits enough to overcome policy losses from compromise.

Proposition 5.1. *If $\delta\beta$ is large, then there is an electoral equilibrium in which all politician types are accountable.*

In contrast to the equilibrium existence proof of Theorem 5.1, the equilibrium construction establishing Proposition 5.1 is straightforward. To construct the equilibrium, we fix a voter-optimal policy rule, i.e., a selection ϕ^* from Φ^* , and we specify that all politician types choose policies according to ϕ^* , and that the voter re-elects incumbents if and only if they choose according to this rule. That this strategy profile forms an equilibrium hinges on the following logic: in any state s , a type t office holder can guarantee re-election by following the rule ϕ^* , generating a stream of future office benefits with payoff $\frac{\delta\beta}{1-\delta}$; by deviating to $x \neq \phi^*(s)$, the politician can achieve a policy gain in the current period, but at the cost of being replaced by a challenger, who also follows the rule. As discussed in Example 1, because the distribution of the state can depend on x , it is possible that the deviation to x generates future policy gains—not by changing policy choices of future politicians (which are taken as given in equilibrium), but by increasing the likelihood of states that are advantageous to the current office holder. But if office incentives are sufficiently large, then it is impossible for policy gains to compensate the incumbent for sacrificing re-election.

Together, Theorem 5.1 and Proposition 5.1 show that the problem of electoral accountability reduces to the possibility of equilibrium multiplicity: these results leave open the question of whether there are other equilibria in which politicians do not respond to electoral incentives by solving the problem of the representative voter. In Example 1, we have already demonstrated that equilibria in which type $t \neq k$ politicians are not accountable will exist widely. Perhaps more surprisingly, we present an example showing that the congruent politician can also fail to be accountable. For this to happen, the shared policy goals of the voter and politician must be trumped by the latter’s office motivation. In the equilibrium in Example 2, the type k politician faces a choice between choosing policies that both she and the voter prefer but then being replaced by the challenger, or choosing suboptimal policies and staying in office. Preferring to retain office, the politician is “cursed” by her ambition and chooses the latter option.

Example 2 (Curse of ambition). Returning to the setting from Example 1, we assume further that

$$\Gamma - (x_1^* - \hat{x}_k^2)^2 < -(x_1^* - \hat{x}_k^1)^2, \quad (10)$$

recalling that $x_1^* = \hat{x}_1 + \frac{\delta\Gamma}{2}$ is the optimal investment for the type 1 politician in all states. In words, the voter prefers the type 1 politician’s optimal policy in the low state s^1 to her optimal policy in the high state: even though the voter gains Γ in the high state, type 1 underinvests too much relative to the voter’s preferences in s^2 .

We claim that given any δ and for β sufficiently high, there exists a Markov electoral equilibrium such that both politician types are re-elected in all states, but neither of them solves the representative dynamic programming problem. In the equilibrium, the type 1

politician chooses x_1^* . From the perspective of the voter, this represents underinvestment in all states. For her part, the type k politician overinvests in all states. However, as we detail below, the distortions imposed by the type 1 politician are relatively worse in the high state s^2 , while those imposed by the type k politician are worse in the low state s^1 , i.e.,

$$V_k(s^2, k) > V_k(s^2, 1) \quad \text{and} \quad V_k(s^1, 1) > V_k(s^1, k). \quad (11)$$

Because the likelihood of the high state increases in investment, the voter's re-election strategy follows a threshold rule: there exists \tilde{x} such that in any state s^j , the voter re-elects the type k politician if and only if she chooses policy $x \geq \tilde{x}$ and re-elects type 1 if and only if she implements a policy satisfying $x \leq \tilde{x}$. The curse of ambition stems from a coordination failure between the voter and the type k politician: because the type k politician is expected to overinvest in the future, the voter does not reward her with re-election when she reduces her investments. Specifically, the dilemma for type k is that the threshold investment \tilde{x} is above the voter's optimal investment x_k^{1*} in the low state, so that she is replaced in that state when she chooses policies that both she and the voter prefer. If the type k politician is sufficiently office-motivated, then her equilibrium policy in the low state is $x_k^1 = \tilde{x}$: she distorts investments upwards to increase the probability of the high state just enough to leave the voter indifferent between re-electing her and opting for the challenger. Because the type k politician wants to invest more in the high state, she faces no policy-office tradeoff in that state: the voter strictly prefers to retain her in that state following her equilibrium policy $x_k^2 > x_k^1 = \tilde{x}$. Finally, because the type 1 politician invests $x_1^* < \tilde{x}$ in all states, the voter strictly prefers to retain her.

We now describe how to construct the equilibrium policy strategy of the type k politician in more detail. Intuitively, we start from the voter-optimal rule (x_k^{1*}, x_k^{2*}) and distort low state investments upward, while allowing the type k politician to adjust her high state investment optimally against that distortion. This leads to investment rules yielding lower policy payoffs to type k . The cutoff investment \tilde{x} is then defined as the level of overinvestment in the low state that leaves the voter indifferent between having either type in office following \tilde{x} (with type 1 implementing her own optimal rule, which is not voter-optimal). We then set $x_k^1 = \tilde{x}$, as described above. We set x_k^2 to be the optimal response for k in state s^2 to setting \tilde{x} in state s^1 , which ensures the optimality of her policy strategy in s^2 .

To formalize the procedure from the previous paragraph, fix $z^1 \in [x_k^{1*}, 1]$ and consider the problem in which a type k politician is free to choose any policy $x^2 \in [0, 1]$ in the high state but is constrained in the low state: she is forced to overinvest by choosing some policy $x^1 \geq z^1 \geq x_k^{1*}$. The policy payoffs to the type k politician in this problem solve the value functions: for $j = 1, 2$,

$$\tilde{V}_k(s^j, z^1) = \max_{x^j \in [I_{j=1} z^1, 1]} -(x^j - \hat{x}_k^j)^2 + \Gamma_j + \delta \left[x^j \tilde{V}_k(s^2, z^1) + (1 - x^j) \tilde{V}_k(s^1, z^1) \right], \quad (12)$$

where $I_{j=1}$ is an indicator function taking the value of 1 in the low state, in which the type k politician is constrained to overinvest. A first note is that if $z^1 > x_k^{1*}$, then the solution to (12) for $j = 1$ has $x^1 = z^1$: recalling (3), interior solutions to (12) for $j = 1, 2$ imply that $\tilde{V}_k(s^j, z^1) = V_k^*(s^j)$ for all j , a contradiction. From this, it also follows that type k 's policy payoff decreases as the constraint on investment in the low state becomes more stringent:

$z^1 > z^{1'}$, then $\tilde{V}_k(s^1, z^1) < \tilde{V}_k(s^1, z^{1'})$. Applying the envelope theorem to (12) for $j = 2$ yields

$$\frac{\partial}{\partial z^1} \tilde{V}_k(s^2, z^1) = \frac{\delta(1-x^2)}{1-\delta x^2} \frac{\partial}{\partial z^1} \tilde{V}_k(s^1, z^1) < 0, \quad (13)$$

where the inequality follows because, from above, $\frac{\partial}{\partial z^1} \tilde{V}_k(s^1, z^1) < 0$. We then define \tilde{x} as level of the low state investment constraint (and hence, by above, of low-state investment by type k) that makes the voter indifferent between having a type k or a type 1 in the continuation game, conditional on an incumbent having invested \tilde{x} , i.e., as the unique solution to

$$\tilde{x} \tilde{V}_k(s^2, \tilde{x}) + (1 - \tilde{x}) \tilde{V}_k(s^1, \tilde{x}) = \tilde{x} V_k(s^2, 1) + (1 - \tilde{x}) V_k(s^1, 1). \quad (14)$$

A final task is to verify that our construction satisfies inequalities (11). To see that it does, notice that

$$\begin{aligned} \tilde{V}_k(s^2, x_k^{1*}) &= V_k^{2*} \\ &> V_k^{1*} \\ &= \tilde{V}_k(s^1, x_k^{1*}) \\ &> V_k(s^1, 1) \\ &> V_k(s^2, 1), \end{aligned}$$

where both equalities follow from the construction of \tilde{V}_k , the first inequality follows because $\Gamma > 0$, the second inequality follows because the type 1 politician does not implement a voter-optimal rule, and the final inequality follows from (10). Therefore, because $\frac{\partial}{\partial z^1} [\tilde{V}_k(s^2, z^1) - \tilde{V}_k(s^1, z^1)] > 0$ by (13), it must be that $V_k(s^2, k) = \tilde{V}_k(s^2, \tilde{x}) > \tilde{V}_k(s^1, \tilde{x}) = V_k(s^1, k)$, so that (11) must hold by the indifference condition (14) defining \tilde{x} . \square

Example 2 shows that by providing incentives for state manipulation, equilibrium expectations can drive striking failures of accountability. The example is robust to politicians' willingness to compromise in order to secure re-election, through either increases in office benefit β or the discount factor δ , both of which promote accountability in other models of dynamic elections (Banks and Duggan (2008); Forand (2014); Van Weelden (2013)). In fact, the political failure in the example stems from the opposite consideration: the congruent politician type chooses suboptimal policies precisely to stay in office in the long run, not to realize short-term gains. Note that in Example 2, the type 1 politician is re-elected because the voter does not have a better option: if the type k politician were accountable while the type 1 politician was not, then the latter could not be re-elected. With more than two politician types, however, it is in fact possible that politician types other than k always chooses suboptimally for the voter, yet are re-elected—even if the type k politician is accountable. This possibility arises because a type $t \neq k$ politician's policy choice may lead to states where the challenger may be even worse.

The next example turns to the multi-state spatial model with three types, and it constructs an equilibrium in which both politician types other than k choose a policy that

threatens states in which the voter prefers the incumbent to the expected challenger. This confronts the voter with a hold-up problem, forcing him to re-elect the incumbent, despite the fact that her policy choices are suboptimal. Thus, even if the type k politician is accountable, if the voter is not too patient, then other politician types can manipulate the state to their advantage and secure re-election.

Example 3 (Political hold-up problem). Return to the two-state spatial model, but now assume three citizen types: $T = \{1, 2, 3\}$, with $k = 2$. A challenger is type k with probability q , and of each of the remaining types have probability $\frac{1-q}{2}$. Fix $0 < \hat{x} < \frac{1}{2}$, and assume that the ideal points of types 1 and 3 are state independent: $\hat{x}_1^j = \hat{x}$ and $\hat{x}_3^j = 1 - \hat{x}$ for all s^j . Assume that the type k politician has state-dependent preferences, which are such that her stage-ideal policy agrees with type 1 in state s^1 and with type 3 in state s^2 : $\hat{x}_k^1 = \hat{x}$ and $\hat{x}_k^2 = 1 - \hat{x}$. Assume that q and δ jointly satisfy

$$q < \frac{(1 - \delta)(1 - 2\hat{x})}{1 - \delta(1 - 2\hat{x})}, \quad (15)$$

so that the voter's incentives to dismiss an incumbent in the hope of drawing a type k challenger are not too strong. Finally, for simplicity, assume $\Gamma \approx 0$.

We claim that there exists a Markov electoral equilibrium in which all politician types implement their optimal policy rules and are always re-elected. Thus, the type k politician is accountable, but because the type 1 politician underinvests in state s^2 and the type 3 politician overinvests in state s^1 , no other type is accountable. Clearly, in this equilibrium, the voter strictly prefers a type k incumbent to all other politician types in all states: $V_k(s^j, k) > V_k(s^j, t)$ for all j and all $t = 1, 3$. But then why does the voter retain all type $t \neq k$ politicians? Because the voter's ranking of these types depends on the state, and each such politician type chooses policies that push transitions towards the state in which they are ranked higher. That is, it can be computed that $V_k(s^1, 1) - V_k(s^1, 3) = V_k(s^2, 3) - V_k(s^2, 1) = (1 - 2\hat{x})^2 > 0$. Therefore, by underinvesting, the type 1 politician favors transitions to state s^1 and is preferred to the challenger, because the voter wants to avoid electing a type 3 when anticipating a transition to that state. The symmetric logic leads the voter to strictly prefer type 3 to the challenger after she overinvests in any state.

The condition $V_k^I(s^j, 1, \hat{x}) > V_k^C(s^j, 1, \hat{x})$, which ensures that the voter strictly prefers the type 1 politician to a challenger following \hat{x} in any state s^j (and by symmetry, the condition that the voter strictly prefers type 3 following $1 - \hat{x}$), reduces to (15). Note that the right-hand side of (15) is decreasing in δ . Thus, if a type k challenger is likely, so that the left-hand side of (15) is higher, then equilibrium incentives require that the voter be relatively impatient. \square

There remains the possibility that under certain conditions, equilibria possess positive welfare properties for the voter. In Example 3, the hold-up problem created a political failure when the inequality (15) holds. However, given any $q > 0$, the inequality fails for δ close enough to one. In fact, the construction in the example fails as $\delta \rightarrow 1$, for the risk to the voter of drawing a challenger who chooses suboptimally in all future periods becomes too great: even if a worse challenger might be drawn, that temporary setback is

outweighed by the chance of installing a type k politician who chooses optimally for the voter. This observation raises the possibility that prospects for accountability improve as the voter becomes patient, a topic to which we turn in the next section.

6 Accountability and Responsive Voting

A substantial literature in political science has examined the behavior of voters in light of information available to them prior to an election. Empirically, some key issues are whether voters use simple “retrospective voting” rules that condition their choices on past performance, whether these rules are consistent with voters responding rationally to politicians’ behavior, and whether such conditioning leads to electoral accountability.¹⁶ Theoretically, retrospective voting has been formulated in different ways,¹⁷ but it is known that the conditioning of voting decisions on past outcomes in a simple way is consistent with “prospective voting,” as assumed in equilibrium modeling: Ferejohn (1986), Fearon (1999) and Duggan (2000) restrict attention to equilibria with a utility threshold that is necessary and sufficient for re-election; and Ashworth et al. (2017), and Duggan and Martinelli (2020) assume a belief threshold.¹⁸ We maintain the equilibrium viewpoint, and investigate the normative properties of a class of equilibrium voting strategies, which we call “responsive voting,” that retains the intuitive features of simple retrospective rules while allowing for more general voting behavior.

To motivate our restriction, recall the state manipulation illustrated in Example 2. That example has the counterintuitive feature that the type k incumbent is replaced even if she chooses the best possible policy for the voter, conditional on the profile of voting and policy strategies: roughly, if the politician best responds for the voter, then she increases the probability of a transition to the state in which she is the worst possible office holder for the voter, leading the voter to elect the challenger after such choices. This disconnect between performance and rewards creates the possibility of a political failure in which the voter has a strict preference to replace the type k incumbent, even though both the voter and politician would prefer to coordinate their choices on optimal policies. While the voter indeed conditions on policy choices in a simple way in the equilibrium of Example 2, the voting strategy does not adequately tie electoral success of the incumbent to current policy outcomes.

This leads us to focus on a class of equilibria in which the voter responds positively to sufficiently good policy choices by the congruent politician type. To strengthen our accountability results, we define the class broadly and isolate it using a weak notion that reflects the dynamic structure of our framework and captures of the spirit of retrospective

¹⁶Anderson (2007) and Healy and Malhotra (2013) survey the empirical literature on retrospective voting, with the former focused on economic voting. Huber et al. (2012), Kayser and Peress (2012), and Woon (2012) study the rationality of retrospective voting rules, while Healy and Malhotra (2009) focus on how politicians respond to such rules.

¹⁷For example, Ashworth and de Mesquita (2014) interpret retrospective voting as not observing policy choices, but only the voter’s level of welfare, and Esponda and Pouzo (2019) model retrospective voting as boundedly rational updating of beliefs.

¹⁸In the latter papers, observed policy outcomes are used to update beliefs, so the belief threshold is equivalent to a policy cutoff.

voting used in the literature. Specifically, we focus on equilibria such that in every state, there is a threshold that is, in a sense, sufficient for the voter to re-elect a type k incumbent: she is re-elected if the voter's discounted payoff strictly exceeds the threshold, and there is at least one policy that meets the threshold and guarantees re-election of the incumbent. This definition is permissive in several ways. First, it only applies to type k politicians. Second, the threshold in any state can be high, which decreases the restrictiveness of our definition. Third, it is not sufficient that the threshold is met with exact equality, as we allow for some policy choices to meet the threshold exactly yet fail to secure victory for the incumbent. We do, however, require that there is at least one policy that equals or exceeds the threshold and does lead to re-election, so the threshold cannot be so high as to be infeasible; and if there are no policies that strictly exceed the threshold, then there must be one that meets it exactly and leads to re-election.

Definition 1. *A Markov electoral equilibrium σ is a k -responsive voting equilibrium if for each state s , there exists a threshold $u_s \in \Re$ such that:*

- (i) *for all policies $x \in X(s)$, if x strictly exceeds the threshold, then a type k incumbent is re-elected with probability one: $\rho(s, k, x) = 1$ if*

$$u_k(x, s) + \delta[\rho(s, t, x)V_k^I(s, t, x) + (1 - \rho(s, t, x))V_k^C(s, t, x)] > u_s,$$

- (ii) *there is at least one policy $x \in X(s)$ that satisfies the threshold and that secures re-election for the type k incumbent with probability one: $\rho(s, k, x) = 1$ and*

$$u_k(x, s) + \delta[\rho(s, t, x)V_k^I(s, t, x) + (1 - \rho(s, t, x))V_k^C(s, t, x)] \geq u_s.$$

Note that if a Markov electoral equilibrium σ is such that the type k politician is accountable, then it is a k -responsive voting equilibrium. Indeed, for each state s , we can set $u_s = V_k^*(s)$ equal to the voter's optimal value in state s . Then condition (i) in the definition of k -responsive voting is satisfied vacuously, because the voter's optimal value cannot be strictly exceeded. By accountability, the type k politician places probability one on optimal policies and is re-elected with probability one in each state s , and thus there is at least one feasible policy $x \in X(s)$ that is both optimal for the voter and secures re-election with probability one, fulfilling condition (ii) in the definition. In this section, we examine the welfare properties of the larger class of k -responsive equilibria. By Theorem 5.1, there is an equilibrium in which the type k politician is accountable, and this immediately implies that k -responsive voting equilibria exist in general, so that our analysis is non-vacuous.

The main result of this section establishes that the shortfall of the voter's discounted payoff in any k -responsive voting equilibrium, relative to the optimal value, is uniformly bounded by a constant that is independent of the discount factor δ , the state s , and the incumbent type t . This bound implies that regardless of the rate of time discounting, the voter's payoffs in a k -responsive voting equilibrium has some connection to the voter's optimal value, and the implications are reinforced by the fact that the discounted payoffs compared are not normalized: they are discounted sums of stage utilities that should be expected to "blow up" for discount factors close to one, giving the bound on the absolute difference between these sums greater impact. A direct implication is that if we consider

a sequence of discount factors δ such that the representative voter becomes arbitrarily patient, there is a corresponding sequence of k -responsive voting equilibria (by Theorem 5.1), and for every such sequence, the voter's normalized payoffs become close to the value of the representative dynamic programming problem. To state the result formally, given a discount factor δ and a strategy profile σ , let $V_k^\delta(s, t)$ denote the expected discounted payoff to the voter from electing a type t politician in state s , and let $V_k^{*,\delta}(s)$ denote the voter's optimal value in state s .

Theorem 6.1. *There exists $M > 0$ such that the difference between the expected discounted sum of voter payoffs in equilibrium and the value of the representative dynamic programming problem is uniformly bounded: given arbitrary discount factor δ , if σ^δ is a k -responsive voting equilibrium, then for all states s and all types t ,*

$$V_k^{*,\delta}(s) - V_k^\delta(s, t) \leq M.$$

In particular, for all δ , there exists a k -responsive voting equilibrium σ^δ , and as the voter becomes patient, the normalized equilibrium payoffs of the voter converge to the optimum: for all s and all t ,

$$\lim_{\delta \rightarrow 1} (1 - \delta)V_k^{*,\delta}(s) - (1 - \delta)V_k^\delta(s, t) = 0.$$

Thus, the gap between the voter's optimal value and his discounted payoff in any k -representative equilibrium becomes negligible as the voter becomes patient, and the rate of convergence is fast: because the discounted payoffs $V_k^{*,\delta}(s)$ and $V_k^\delta(s, t)$ are not normalized, these discounted sums of stage utilities will typically diverge to infinity as δ goes to one, so the conclusion that the difference between them is bounded is quite strong: in particular, it trivially implies that the normalized payoffs become arbitrarily close, so that equilibria become *approximately optimal* for the representative voter as the voter becomes patient.¹⁹ The proof of the result involves two steps, starting from the fact that because we assume the state transition is positive on every state, every strategy profile determines a unique ergodic distribution on state-policy pairs.

In the first step, we show that given any k -responsive voting equilibrium, we can “sandwich” the voter's equilibrium payoffs between the payoffs from two particular strategy profiles. In one profile, all politician types choose optimally for the voter, achieving the value of the representative dynamic programming problem and providing an upper bound for his equilibrium payoff. The other profile is such that the choices of type $t \neq k$ politicians may not be optimal, but it determines the same ergodic distribution as the first. The second step relies on the fact that the Markov chains determined by these two profiles converge at a geometric rate to the ergodic distribution, which implies (after some algebra) that the difference between the voter's payoffs from the two profiles is bounded by M with the properties stated in Theorem 6.1. Thus, the voter's equilibrium payoffs converge to the optimum in the strong sense used there. The existence of the second profile, which gives a lower bound on the voter's equilibrium payoff and has ergodic distribution that places probability one

¹⁹Examples 2 and 4 (below) show, respectively, that k -responsive voting equilibria and voter patience cannot be dropped, so Theorem 6.1 is tight in this sense. It is straightforward to construct examples showing that the conditions are not generally necessary.

on optimal state-policy pairs, does not follow immediately from the structure of the model or the equilibrium concept, but we derive it from dynamic programming arguments that leverage equilibrium incentives of both the voter and the congruent politician type. We discuss these steps in reverse order.

To describe the second step above, we define the dynamic programming problem of a unitary type k actor who controls electoral outcomes and, when a congruent politician holds office, policy choices as well. This actor makes decisions in two kinds of situations: when a type k politician holds office in a state s , the actor chooses policy $x \in X(s)$ and also makes a retention decision $r \in \{0, 1\}$; and when a type $t \neq k$ politician holds office in state s and chooses x , the type k actor makes only the retention decision. Thus, a “state” in the unitary actor’s problem has the form (s, k) or the form (s, t, x) with $t \neq k$. Given an equilibrium σ , let \tilde{V}_k denote the value function for the problem, and note that by the optimality principle, we have: for all s ,

$$\begin{aligned} \tilde{V}_k(s, k) = & \max_{(x,r) \in X(s) \times \{0,1\}} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) \left(r \tilde{V}_k(s', k) + (1-r)[q_k(k|s, x) \tilde{V}_k(s', k) \right. \\ & \left. + \sum_{t' \neq k} q_k(t'|s, x) \int_{x'} \tilde{V}_k(s', t', x') \pi_{t'}(dx'|s') \right), \end{aligned}$$

and for all s , all $t \neq k$, and all x ,

$$\begin{aligned} \tilde{V}_k(s, t, x) = & \max_{r \in \{0,1\}} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) \left(r \int_{x'} \tilde{V}_k(s', t, x') \pi_t(dx'|s') \right. \\ & \left. + (1-r)[q_t(k|s, x) \tilde{V}_k(s', k) + \sum_{t' \neq k} q_t(t'|s, x) \int_{x'} \tilde{V}_k(s', t', x') \pi_{t'}(dx'|s')] \right). \end{aligned}$$

Parsing the Bellman equation for the unitary actor, note that if the type t incumbent is retained in state s after choosing policy x , then the state transitions to a new state s' , and the incumbent’s type remains t ; and if the challenger is elected, then the new office holder’s type is drawn from $q_t(\cdot|s, x)$. When the “state” has the form (s, k) , the actor chooses policy as well as the election outcome, and when it has the form (s, t, x) with $t \neq k$, the policy choice is determined by the politician’s equilibrium strategy $\pi_t(\cdot|s)$. In particular, if the challenger is elected and is of type $t' \neq k$, then the continuation payoffs $\tilde{V}_k(s', t', x')$ are integrated over policy choices using $\pi_{t'}(\cdot|s')$.

The next lemma describes a useful necessary condition for k -responsive voting equilibria: the value of the unitary actor problem is achieved by the type k citizen in all “states.” Part (i) of the lemma establishes that the equilibrium voting strategy solves the voter’s optimal retention problem, an observation that is critical for the analysis of equilibria. In the definition of equilibrium, the voter does not choose a dynamic plan governing electoral outcomes in all future elections, but rather the voter plays a limited role: she can check office holders retroactively, but the voter takes future policy choices and electoral outcomes as given. Nevertheless, the Unitary Actor Lemma shows that in equilibrium, it is as if the voter chooses an optimal re-election rule that, in addition to determining the winner of

the current election, also dictates electoral outcomes in all future states, for all incumbent types, and after all policy choices.^{20,21}

Part (ii) of the lemma states that when a congruent politician holds office, the policy choice of the politician and retention decision of the voter jointly maximize the expected discounted payoff of the voter. The voter and type k politicians are separate players who can have distinct incentives, but those incentives are brought into alignment when the voter uses a responsive voting rule, in the weak sense defined here. Taken together, the two parts of the lemma mean that in a k -responsive voting equilibrium, we can treat the voter and type k politicians as a unitary actor who controls electoral outcomes after all triples (s, t, x) , and who, when a type k politician holds office, chooses policy as well. The lemma does not restrict the policy choices of other politician types, which are reflected in the terms $V_k(s, t)$ with $t \neq k$, but it implies that in such an equilibrium, the type k policy strategy and the voting strategy (π_k, ρ) form a best response for the unitary actor to the strategies of other politician types.

Unitary Actor Lemma. *In any k -responsive voting equilibrium, we have:*

(i) *for each state s , each type $t \neq k$, and each policy $x \in X(s)$,*

$$u_k(s, x) + \delta[\rho(s, t, x)V_k^I(s, t, x) + (1 - \rho(s, t, x))V_k^C(s, t, x)] = \tilde{V}_k(s, t, x),$$

(ii) *for each state s , $V_k(s, k) = \tilde{V}_k(s, k)$.*

With the lemma in place, we can return to the first step above to complete the logic of the proof of Theorem 6.1. Letting ϕ^* be any optimal policy rule for the voter, we can construct the upper bound by specifying that each politician type chooses according to ϕ^* , and that the voter removes the incumbent until a congruent politician type is drawn, after which the incumbent is retained thereafter. Denote this profile by $\tilde{\sigma}$. For the lower bound, we specify that the type k politician choose according to ϕ^* , while other politician types use their equilibrium strategies; and again, the voter removes the incumbent until a type k politician is drawn, after which the incumbent is retained thereafter. Denoting this profile by $\hat{\sigma}$, we claim that the voter's equilibrium payoff lies between the payoffs determined by the two profiles. Indeed, it is obvious that the equilibrium payoff cannot exceed the optimal value. To see that $\hat{\sigma}$ provides a lower bound, note that the unitary actor has the option of removing all type $t \neq k$ incumbents continually in each state, until a type k candidate is selected, and then using ϕ^* thereafter. By the Unitary Actor Lemma, it follows that the voter's equilibrium payoffs weakly exceed those of $\hat{\sigma}$, as claimed. The ergodic distributions over (s, x) pairs determined by $\tilde{\sigma}_k$ and $\hat{\sigma}$ coincide, and finally, we exploit the geometric convergence of Markov chains to show that the voter's discounted payoff from $\hat{\sigma}$ is within M of the optimal value, where the bound depends only on the state transition and, in

²⁰Note that the unitary actor does not commit to a re-election rule; rather, the actor's problem takes policy strategies as given for types $t \neq k$. Despite the assumption of a representative voter, the solution with commitment would require significant coordination on the part of voters, and it would rely on history dependence of strategies to make commitment credible.

²¹Part (i) of the Unitary Actor Lemma is also true for $t = k$ and holds in all equilibria, and it does not depend on the selection of a k -responsive voting equilibrium.

particular, is independent of the voter’s discount factor and strategies used by type $t \neq k$ politicians. Taken together, these arguments deliver the result.²²

We have already discussed the implications of the bound for k -responsive voting equilibria as the voter becomes patient. We must emphasize, however, that while the proof of Theorem 6.1 uses a particular voting strategy to establish a lower bound on the voter’s equilibrium payoff—the voter essentially waits for a congruent politician and continually retains the first one realized—there should be no expectation that this strategy is in fact used in equilibrium. For a given discount factor, there may well be equilibria in which the voter retains a type $t \neq k$ politician in order to avoid challenger types who are worse. That incumbent may choose policies that are close enough to optimal that the voter has a strict preference to re-elect the politician, or it may be that the voter is indifferent between the incumbent and a randomly drawn challenger. In the latter case, it may be that the re-election incentives discipline the type t politician to compromise her policy choices, choosing policies that are better for the voter than she otherwise would in the absence of the electoral mechanism. Proposition 5.1 shows that when politicians are sufficiently office motivated, there are always equilibria of this sort.²³ The “wait it out” construction is a useful construct, but the equilibrium voting strategy may be distinct and involve re-election of other politician types. The point, of course, is that in all such equilibria, the payoffs cannot fall below the lower bound provided by the construction.

Example 4 (Approximate accountability). Return to the setting from Example 3 where, given $0 < \hat{x} < \frac{1}{2}$, we have $\hat{x}_1^j = \hat{x}$ and $\hat{x}_3^j = 1 - \hat{x}$ for all s^j . Now, however, assume that the ideal point of type k is fixed at $\hat{x}_k^j = \frac{1}{2}$ for all states s^j . It is enough for our purposes to assume that $\Gamma \approx 0$, so that optimal policy rules for all types consist of choosing stage-optimal levels of investment: $x_t^{j*} = \hat{x}_t^j$ for all types t and all states s^j . Assuming that β is high enough, we claim that for all δ , there exists a Markov electoral equilibrium in which type k is accountable, type 1 underinvests, is re-elected in all states, and is approximately accountable as $\delta \rightarrow 1$; in contrast, type 3 chooses her stage-optimal policy in all states and is never re-elected. In these equilibria, the threat of selecting a type 3 challenger allows a type 1 politician to distort policies downward, but the threat diminishes as the voter becomes patient, and the extent of distortion goes to zero.

To construct the equilibrium, let $x_1 < \frac{1}{2}$ denote the equilibrium investment of type 1 in all states. In any state, the voter re-elects type 1 if and only if she chooses x_1 : for all j , $\rho(s^j, 1, x) = 1$ if and only if $x = x_1$. Type 1 chooses to compromise as long as the office benefit β is high enough. Type 3 chooses policy $1 - \hat{x}$ in all states and is not re-elected following any policy: for all j , $\rho(s^j, 3, x) = 0$ for all x . The equilibrium policy of type 1 must be such that in all states, the voter is indifferent between her and the challenger: $V_k^I(s^j, 1, x_1) = V_k^C(s^j, 1, x_1)$ for all j , where in this setting, these expressions are identical

²²Our argument uses the assumption that the challenger is type k with positive probability. Otherwise, the Unitary Actor Lemma implies that the voter’s equilibrium payoff is bounded below by the policy strategy of the “next best” type, but the identity and strategy of the next best type will typically vary with the discount factor.

²³Indeed, in an equilibrium in which all politician types choose optimal policies for the voter, it may be that type k incumbents are removed from office, thereby failing our definition of accountability.

in both states. By computation, it can be verified that the indifference condition yields

$$\left(\frac{1}{2} - x_1\right)^2 = 2(1 - \delta) \frac{\frac{1}{2} - \hat{x}^2(1 - q)}{1 - \frac{\delta(1-q)}{2}(1 + q(4 - q))}.$$

Taking limits, the right-hand side goes to zero as $\delta \rightarrow 1$, and thus the type 1 politician's policy converges to the voter optimum, i.e., $x_1 \rightarrow \frac{1}{2}$, as desired. \square

We have expressed Theorem 6.1 in terms of voter payoffs, but it has direct implications for the optimality of policy choices themselves. Namely, we consider the distribution of policies in an equilibrium with k -responsive voting, which is averaged over the policy choices of all incumbents and the voter's re-election decisions, and we show that this distribution converges to the optimal policies of the representative voter as he becomes patient. To state this result precisely, we need to work with the stochastic process over future state-policy pairs (s', x') generated by an equilibrium σ given a type t incumbent in state s ; this is a probability measure $\mu_{s,t}$ over infinite sequences $(s^m, x^m)_{m=1}^\infty$, where $s^1 = s$. Let $\mu_{s,t}^m$ denote the marginal distribution on state-policy pairs (s^m, x^m) in period m , where $s^1 = s$. We aggregate these marginals across time by geometric discounting to define the probability measure

$$\mu_{s,t}^\delta = (1 - \delta) \sum_{m=1}^{\infty} \delta^{m-1} \mu_{s,t}^m, \quad (16)$$

which depends on the initial state and politician type, and where we highlight the dependence on δ for future use. Our summary statistic for the equilibrium policies in any state s is then the conditional $\mu_{s,t}^\delta(\cdot|s)$ of the aggregate measure $\mu_{s,t}^\delta$, which is well-defined since $\mu_{s,t}^\delta(X \times \{s\}) > 0$. Thus, for a Borel subset $A \subseteq X$ of policies, $\mu_{s,t}^\delta(A|s)$ measures the probability, given initial state s and politician type t , that future policy choices, conditional on being in state s , belong to A . We use this measure to aggregate across periods in a way that reflects the time preferences of the voter.

The next corollary shows that for every state s and every politician type t , the equilibrium probability distribution $\mu_{s,t}^\delta$ on policy choices at s allocates probability mass close to the optimal policies $\Phi^{*,\delta}(s)$ as the voter becomes patient, where $\Phi^{*,\delta}(s)$ denotes the voter's optimal policies in state s given discount factor δ .

Corollary 6.1. *Given any δ , let σ^δ be a k -responsive voting equilibrium. Then for all initial states s and types t , and for all $\epsilon > 0$, the probability that policies are within ϵ of optimal under the aggregate measure $\mu_{s,t}^\delta$ converges to one as the voter becomes patient:*

$$\lim_{\delta \rightarrow 1} \mu_{s,t}^\delta(B_\epsilon(\Phi^{*,\delta}(s))|s) = 1,$$

where $B_\epsilon(\Phi^{*,\delta}(s))$ is the open ball of radius ϵ around the set $\Phi^{*,\delta}(s)$ of optimal policies.

Theorem 6.1 focuses on responsive voting when the incumbent is congruent, and it leverages the incentives of this politician to deduce an accountability result. As noted above, the policy choices of type non-congruent politicians may also respond positively as the voter

becomes patient, but they do not necessarily do so. In the remainder of this section, we consider the stronger consequences for such politicians when there is a threshold sufficient for their re-election as well. We will see that when type $t \neq k$ politicians are sufficiently office motivated, they also become accountable in the limit; and if the state transition is policy independent, then all politician types are exactly accountable, even if the voter is not arbitrarily patient. Next, we extend the definition of responsive voting to incumbents of all types.

Definition 2. *The Markov electoral equilibrium σ is a **responsive voting equilibrium** if for each type t , parts (i) and (ii) of Definition 1 hold for the type t incumbent.*

Extending the observation for k -responsive equilibria, if all politician types are accountable in an equilibrium σ , then σ is a responsive voting equilibrium; moreover, Proposition 5.1 implies that responsive voting equilibria generally exist when office incentives are sufficiently high. To understand the implications of the responsive voting condition, note that an implication of Corollary 6.1 is that when the voter is patient, all incumbents are essentially competing against an accountable challenger, namely, the potential of a type k politician. Thus, if a type $t \neq k$ politician chooses acceptable policies in all states, then her policy choices must also be approximately optimal. Moreover, if the voter uses a responsive voting strategy and such a politician is sufficiently office motivated, then she will always be re-elected with positive probability, i.e., she will choose policies that are acceptable to the voter. Taken together, these conditions imply that the accountability of type k politicians spills over to the other politician types: if office benefit is large and the voter becomes patient, then in any corresponding sequence of responsive voting equilibria, *all* politician types choose policies that are approximately optimal for voter. This is stated in part (i) of Corollary 6.2, below.

These conclusions can be sharpened when the policy transition is independent of the state, i.e., the state encapsulates shocks that are exogenous to policy choices.²⁴ In this case, it is impossible for a politician to manipulate the state by choice of policy, and competition with accountable type k politicians disciplines all other politician types: part (ii) of Corollary 6.2 establishes that if office incentives are large, then in any responsive voting equilibrium, all politician types choose policies that are exactly, rather than approximately, optimal. Note that if the state transition depends on policy, then politicians can manipulate the state, and the exact accountability result from part (ii) does not hold. This was illustrated in Example 3: in the equilibrium from that example, which satisfies responsive voting, all types are re-elected but only the type k politician is accountable.

Corollary 6.2. *Given any δ , assume that $\delta\beta$ is large, and let σ^δ be a responsive voting equilibrium. Then:*

- (i) *for all states s , all types t , and all $\epsilon > 0$, the probability that policy choices of the type t office holder are within ϵ of optimal converges to one as the voter becomes patient:*

$$\lim_{\delta \rightarrow 1} \pi_t^\delta(B_\epsilon(\Phi^{*,\delta}(s))|s) = 1,$$

²⁴For example, shocks may capture climate change, public health crises, technological innovations, or the term of office of the incumbent.

(ii) if $p(\cdot|s, x)$ is policy-independent for each state s , then each type t politician is accountable.

The proof of Corollary 6.2 relies on the assumption that type $t \neq k$ politician are sufficiently office motivated, so that they are willing to compromise their policy choices to improve their chances for re-election, and also on the fact that in a responsive voting equilibrium, the politician has the opportunity to choose a policy that guarantees re-election with probability one. There is some subtlety to the proof because the combination of these assumptions does not imply that the type $t \neq k$ incumbent is re-elected with probability one; rather, we can infer that with probability one, the politician chooses a policy x such that conditional on x , the voter re-elects the incumbent with positive probability. This implies that politicians always choose policies that are acceptable to the voter. Part (i) of the corollary then follows from Theorem 6.1: because the challenger is type k with positive probability, and because all other politician types are acceptable, the expected challenger becomes close to optimal, implying that all politician types must be choosing policies that approximate solutions to the representative programming problem.

The argument for part (ii) of the corollary continues with the observation that in an equilibrium in which all politicians choose policies that are acceptable to the voter, we must have:

$$\text{for all } s, \text{ all } t, \text{ and } \pi_t(\cdot|s)\text{-almost all } x, V_k^I(s, t, x) \geq V_k^C(s, t, x). \quad (17)$$

There can exist equilibria satisfying (17) in which a politician is re-elected after policies that are suboptimal for the voter, if the politician manipulates the state: she may choose suboptimal policies that steer the state transition toward states in which the incumbent is preferable to an untried challenger, inducing the voter to re-elect. But when the state transition is policy independent, it is impossible to manipulate the state in this way, and the value of retaining an incumbent is also independent of policy:

$$V_k^I(s, t) = \sum_{s'} p(s'|s) V_t(s', t),$$

where we remove the notational dependence of V_k^I and p on policy. Moreover, $V_k^C(s, t, x)$ is just an ‘‘average’’ payoff from a new office holder following policy x in state s , i.e.,

$$V_k^C(s, t, x) = \sum_{t'} q_t(t'|s, x) \sum_{s'} p(s'|s) V_k(s', t') = \sum_{t'} q_t(t'|s, x) V_k^I(s, t'),$$

even if the distribution over these challengers depends on both the incumbent’s type and on policy. Given any state s , choose $\tilde{t} \in \arg \min_t V_k^I(s, t)$, and consider any \tilde{x} in the support of $\pi_{\tilde{t}}(\cdot|s)$, and note that (17) implies

$$V_k^I(s, \tilde{t}) \geq V_k^C(s, \tilde{t}, \tilde{x}) = \sum_t q_{\tilde{t}}(t|s, \tilde{x}) V_k^I(s, t) \geq V_k^I(s, \tilde{t}),$$

which, because $q_{\tilde{t}}(k|s, \tilde{x}) > 0$, ensures that $V_k^I(s, \tilde{t}) = V_k^I(s, k)$. But then, by the Unitary Actor Lemma, congruent politicians achieve the voter’s optimal value, so that for all t , we have

$$V_k^*(s) \geq V_k^I(s, t) \geq V_k^I(s, \tilde{t}) = V_k^*(s).$$

That is, all politician types achieve the optimal value.

7 Concluding Discussion

The results of this paper inform us of the possibilities for—and limits of—electoral accountability in disciplining policy choices by politicians. We propose a general electoral framework that extends the standard citizen-candidate model by adding a finite set of states and, consequently, the possibility of non-trivial dynamics, and we establish general existence of equilibrium. Our results support the view that elections can be an effective mechanism for holding politicians accountable, for if office incentives are strong, then there exists an equilibrium in which all politician types choose optimal policies for the voter. However, this positive result is attenuated by the possibility that there may also exist equilibria in which an office holder retains office by manipulating the state, despite choosing policies that are suboptimal for the voter; this is true even for the congruent politician type, who may suffer from a curse of ambition. Such equilibria require voting behavior that violates a responsive voting condition, and when attention is focused on responsive equilibria, the negative welfare consequences of state manipulation are diminished: the type k politician always chooses optimally for the voter, and as citizens become patient, the voter’s payoff in all equilibria approaches the optimal value in a strong sense.

Our dynamic elections framework can be augmented with greater structure, and a number of directions are worthy of attention. For example, we have abstracted away from electoral uncertainty (beyond the possibility of voter mixing) to maximize parsimony, but the model could incorporate unobserved valence shocks to generate meaningful uncertainty. Adding “large,” uniform shocks in the multi-state spatial model would concavify the optimization problem of an office holder, but a joint restriction on preferences and uncertainty used by Duggan and Martinelli (2020) would allow the introduction of “small” shocks and permit a key trade off that can arise in our model: a politician may have to choose between compromising in order to ensure reelection or shirking to obtain desirable policy today at the cost of re-election. Another direction of interest is the endogenizing of challenger entry, which could provide microfoundations for our assumption that congruent challengers have positive probability, and would bring our framework even closer to classic citizen-candidate models. The addition of adverse selection in the form of persistent, privately observed types is an interesting question, but this would create signaling incentives and difficult technical issues. These are discussed in detail in Duggan and Martinelli (2017); we suspect that meaningful characterizations would require a careful selection of equilibrium, and that the general solution of the existence problem would rely on careful specification of off-path beliefs and history-dependent equilibrium strategies.

We have illustrated the applicability of our framework and motivated examples by introducing the multi-state spatial model, with the interpretation of public investment, where the state represents discrete categories of a stock of public goods. For ease of exposition, we left the connection between policy and citizen utility abstract, but the framework can accommodate economic structure that permits numerous interpretations, including the choice of macroeconomic policy instruments. This creates the possibility of addressing traditional topics (such as political cycles, the size of government, rent seeking, and redistributive politics) in an economy with an evolving state variable.

To close the paper, we provide an example of a simplified model a macroeconomic policy making, which also serves to describe the problems created by endogenous economic outcomes in our framework. Suppose that an office holder chooses a level of inflation, $x \in [0, 1]$, and this determines employment y according to a linear Philips curve, $y(s, x, x^e) = s + x - x^e$, where s is the state of the economy (e.g., high s indicates that the economy is utilizing resources more efficiently), and x^e is expected inflation. In the spirit of Persson and Tabellini (1990), we assume that each citizen type t has an inflation target $\hat{x}_t(s)$ and quadratic loss from inflation around this target, but here we permit the target to depend on the state, as in $u_t(s, x) = -(x - \hat{x}_t(s))^2$, and we assume utility is linear in employment. This gives us stage utility

$$u_t(s, x, x^e) = u_t(s, x) + \lambda_t(s + x - x^e),$$

where $\lambda_t \geq 0$ is a type-dependent weight on employment.²⁵ In the usual ideological interpretation, more liberal politician types will have greater coefficients λ_t .

A strategy profile in the model augmented with this macroeconomic structure consists of three mappings, π_t , ρ , and ξ^e , where π_t and ρ are policy and voting strategies, and $\xi^e: S \times T \rightarrow [0, 1]$ gives the voter's expectation $\xi^e(s, t)$ of inflation x in each state s when the incumbent is type t . The equilibrium conditions on policy and voting strategies are as in our framework, and we add the condition that for each s ,

$$\xi^e(s, t) = E[\pi_t(\cdot|s)].$$

Thus, in equilibrium, inflation is fully anticipated, and because utility is linear in inflation, continuation values have the same form as in (6)–(8); in particular, that part of employment y due to unexpected inflation, $x - x^e$, drops out of the expressions for the continuation values, leaving only the contribution of the strength of the economy, captured by s . In equilibrium, $\pi_t(\cdot|s)$ places probability one on maximizers of

$$u_t(s, x) + \lambda_t(s + x - E[\pi_t(\cdot|s)]) + \delta \left[\rho(s, t, x) [V_t^I(s, t, x) + B_t(s, x)] + (1 - \rho(s, t, x)) V_t^C(s, t, x) \right].$$

Since stage utility is linear, the voter's expected inflation $E[\pi_t(\cdot|s)]$ does not affect the best response policies of the politician: equilibrium policies solve

$$\max_{x \in [0, 1]} \tilde{u}_t(s, x) + \delta \left[\rho(s, t, x) [V_t^I(s, t, x) + B_t(s, x)] + (1 - \rho(s, t, x)) V_t^C(s, t, x) \right],$$

where the office holder's stage utility is

$$\tilde{u}_t(s, x) = u_t(s, x) + \lambda_t(s + x).$$

Thus, endogenous economic outcomes introduce a wedge, $\lambda_k x$, between the incentives of the voter and the type k politician in the current period. This is the well-known problem of time inconsistency, cast in the context of the dynamic electoral framework with an economic state variable.

²⁵To simplify the discussion, we assume the state transition $p(s'|s, x)$ depends only on the state and inflation. Similar remarks hold if, instead, it depends linearly on surprise inflation, as in $p(s'|s, x - x^e)$.

Because the type k politician’s objective function evaluates future policies using the same stage utility as the voter but evaluates current policy via $\tilde{u}_t(s, x)$, her policy choices may appear to exhibit “myopia” or non-standard intertemporal preferences (e.g., quasi-geometric discounting), and this effect will be greater when the voter is more liberal, i.e., λ_k is higher. We conjecture that our general results on equilibrium existence and on accountability with high enough office benefit carry over to the framework with an endogenous economic outcome. However, the discrepancy between the preferences of the voter and type k politician pose a challenge for the Unitary Actor Lemma and our positive results on accountability as $\delta \rightarrow 1$. The Markovian elections framework provides a platform for evaluating the effect of this incongruence and the extent to which it varies with the economic state.

More generally, the incorporation of non-trivial dynamics gives us access to new structure and expands the set of questions we can address: How do political cycles depend on the state of the economy? Does the time-inconsistency problem advantage conservative politician types over liberal ones, whose lack of credibility leads to higher inflation choices? Under what conditions do there exist equilibria in which liberal politicians mix over policies, generating uncertainty and real effects on employment? How does the ergodic distribution of economic states and policies vary with respect to parameters, and what can we say about equilibrium paths of employment and inflation? Further work is needed to develop analytically and numerically tractable versions of the model with enhanced institutional and economic structure to address these and other questions of interest.

A Proofs of Results

Proof of Theorem 5.1. As explained in the text, we embed our model in the more general framework of Duggan and Forand (2018), modifying our model by adding a bad absorbing state $s_b \notin S$, and then applying Theorem 1 of that paper. The augmented set of states is $\tilde{S} = S \cup \{s_b\}$. We then specify the state transition such that for all $s, s' \in S$, all $t \in T$, all $x \in X(s)$, and all electoral outcomes $e \in \{0, 1\}$ (with $e = 1$ when the incumbent is re-elected),

$$\tilde{p}_t(s'|s, x, e) = \begin{cases} 1 & \text{if } s' = s_b, t = k, \text{ and } e = 0, \\ & \text{or if } s = s_b, \\ p(s'|s, x) & \text{else.} \end{cases}$$

That is, the state transition is otherwise the same as in our model, but if a type k politician is removed from office, then the state transitions to the absorbing bad state. We define stage utility functions as in our model, but we assign a bad payoff to the voter in the bad state:

$$\tilde{u}_k(s, x) = \begin{cases} -2 & \text{if } s = s_b, \\ u_k(s, x) & \text{else,} \end{cases}$$

where we recall that stage utility u_k is bounded between $\underline{u} = 0$ and $\bar{u} = 1$. With this specification, Theorem 1 of Duggan and Forand (2018) yields a Markov electoral equilibrium $\tilde{\sigma} = (\tilde{\pi}, \tilde{\rho})$.

We claim that the type k politician is re-elected in every state $s \neq s_b$ following every policy choice. Intuitively, this follows because the voter strictly prefers to avoid reaching the bad state s_b from any state $s \neq s_b$, so that he re-elects the type k politician following any policy choice: for all $x \in X(s)$, we have $\tilde{\rho}(s, k, x) = 1$. Indeed, consider any policy choice $x \in X(s)$ by a type k incumbent at state $s \neq s_b$. The voter's discounted payoff from re-electing the incumbent is at least equal to $\frac{-2\delta}{1-\delta}$, and this is strictly greater than the payoff of electing a challenger, which is $\frac{-2}{1-\delta}$. By part (ii) of the definition of equilibrium, the voter re-elects the incumbent, establishing the claim.

Next, we establish that the type k politician's policy strategy π_k solves the representative dynamic programming problem: $\tilde{V}_k(s, k) = V_k^*(s)$ for all $s \neq s_b$. Intuitively, this follows because in the equilibrium $\tilde{\sigma}$ of the augmented model, the type k politician's expected discounted office benefit at every state $s \neq s_b$ for every policy choice $x \in X(s)$ is $\frac{\beta}{1-\delta}$. Therefore, since office benefits are constant with respect to her policy choice, the type k politician implements the policies that the voter would choose in her place. Formally, because she is always re-elected, the type k politician solves

$$\max_{x \in X(s)} \tilde{u}_k(s, x) + \delta \tilde{V}_k^I(s, t, x) + \frac{\beta}{1-\delta},$$

in each state $s \neq s_b$. Equivalently, for each $s \neq s_b$, $\pi_k(\cdot|s)$ places probability one on solutions to

$$\max_{x \in X(s)} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) \tilde{V}_k(s, k),$$

and thus

$$\tilde{V}_k(s, k) = \max_{x \in X(s)} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) \tilde{V}_k(s, k).$$

We conclude that $\tilde{V}_k(\cdot, k)$ solves the Bellman equation for the voter, and thus $\tilde{V}_k(s, k) = V_k^*(s)$ for all $s \neq s_b$.

We map $\tilde{\sigma}$ to a strategy profile σ in our model by defining π_t as the restriction of $\tilde{\pi}_t$ to S , and we define ρ so that for all $s \in S$ and all $x \in X(s)$,

$$\rho(s, t, x) = \begin{cases} \tilde{\rho}(s, t, x) & \text{if } t \neq k, \\ 1 & \text{if } t = k \text{ and } \tilde{V}_k^I(s, k, x) \geq \tilde{V}_k^C(s, x), \\ 0 & \text{else.} \end{cases}$$

Let $V_t(s, t')$, $V_t^I(s, t', x)$, and $V_t^C(s, t', x)$ denote the continuation values generated by σ in our model. In particular, the type k politician chooses optimally for the voter in σ , and thus for all $s \in S$ and all $x \in \text{supp}(\pi_k(\cdot|s))$, we have

$$V_k^I(s, k, x) = \sum_{s'} p(s'|s, x) V_k(s', k) = \sum_{s'} p(s'|s, x) V_k^*(s') \geq V_k^C(s', x).$$

Then the type k politician is re-elected with probability one at all states $s \in S$, which implies that the continuation values V_t , V_t^I , and V_t^C agree with \tilde{V}_t , \tilde{V}_t^I , and \tilde{V}_t^C at all $s \neq s_b$. Finally, we conclude that σ is a Markov electoral equilibrium of our model, and that the type k politician is accountable, as required. \square

Proof of Proposition 5.1. Let ϕ be an optimal policy rule for the voter. Define policy strategies such that for all s and all t , $\pi_t(\{\phi(s)\}|s) = 1$, and define the voting strategy such that for all s , all t , and all $x \in X(s)$, we have $\rho(s, t, x) = 1$ if $x = \phi(s)$, and $\rho(s, t, x) = 0$ otherwise. Obviously, all politician types are accountable under strategy profile $\sigma = (\pi, \rho)$. Because the voter is indifferent between the incumbent and the challenger following any policy choice in any state, it follows that ρ satisfies the conditions for equilibrium: because $V_k(s, t) = V_k^*(s)$ for all s and all t , then for every policy $x \in X(s)$ we have

$$V_k^I(s, t, x) - V_k^C(s, x) = \sum_{s'} p(s'|s, x) [V_k(s', t) - V_k^*(s')] = 0.$$

To verify that policy strategies π_t are optimal for all politicians in all states, let $V_t^\phi(s)$ denote the expected policy utility to the type t office holder from following the rule ϕ in state s and thereafter:

$$V_t^\phi(s) = u_t(s, \phi(s)) + \delta \sum_{s'} p(s'|s, \phi(s)) V_t^\phi(s').$$

Then the total expected payoff from following the rule ϕ in state s , and holding office in perpetuity, is $V_t^\phi(s) + \frac{\beta}{1-\delta}$. The expected payoff from deviating to $x \neq \phi(s)$ in state s , and then being replaced by a challenger, is no more than $\frac{\bar{u}}{1-\delta} + \beta$. Therefore, using the normalization $\bar{u} = 1$ and $\underline{u} = 0$, a sufficient condition for office holder t to follow the policy rule ϕ in all states s is $\frac{\delta\beta}{1-\delta} \geq 1$, as required. \square

Proof of the Unitary Actor Lemma. Consider any k -responsive equilibrium σ . We first prove part (ii). Clearly, we have $\hat{V}_k(s, k) \geq V_k(s, k)$, as the unitary actor has the option of using the equilibrium strategies of the voter and type k politician. To prove the opposite inequality, we show that in a k -responsive equilibrium, the type k politician achieves the voter's optimal value. For each state s , define

$$\hat{v}_k(s) = \max_{x \in X(s)} u_k(s, x) + \delta [\rho(s, k, x) V_k^I(s, k, x) + (1 - \rho(s, k, x)) V_k^C(s, k, x)].$$

With compactness of $X(s)$ and continuity of the above objective function, the definition of k -responsiveness implies that for each state s , there is a maximizing policy $\hat{x}(s) \in X(s)$ such that $\rho(s, k, \hat{x}(s)) = 1$ and

$$\hat{v}_k(s) = u_k(s, \hat{x}(s)) + \delta [\rho(s, k, \hat{x}(s)) V_k^I(s, k, \hat{x}(s)) + (1 - \rho(s, k, \hat{x}(s))) V_k^C(s, k, \hat{x}(s))].$$

For future reference, note that since $\hat{x}(s)$ leads to re-election with probability one, we have

$$\hat{v}_k(s) = u_k(s, \hat{x}(s)) + \delta V_k^I(s, k, \hat{x}(s)) \tag{18}$$

for each state s .

Let $\hat{\pi}_k$ denote the policy strategy such that $\hat{\pi}(\{\hat{x}(s)\}|s) = 1$ for all s , and let $\hat{V}_k(s)$ denote the expected payoff to the voter from using policy rule $\hat{\pi}_k$ in state s , i.e., for each s , we have

$$\hat{V}_k(s) = u_k(s, \hat{x}(s)) + \delta \sum_{s'} p(s'|s, \hat{x}(s)) \hat{V}_k(s').$$

Here, $\hat{v}_k(s)$ is the best policy payoff the type k politician can achieve by deviating from π_k to $\hat{\pi}_k$ for one period, under the assumption that the better candidate for the voter (incumbent or challenger) is elected subsequently. In contrast, $\hat{V}_k(s)$ is the policy payoff that would be achieved by following $\hat{\pi}_k$ in all states and continually being re-elected. It is not obvious that these are the same, but our arguments imply that this is indeed the case.

The first step in the proof of (ii) is to show that $\hat{V}_k(s) \geq V_k(s, k)$ for all s . Suppose the type k politician deviates to the non-stationary strategy such that in the first term of office, she chooses according to $\hat{\pi}_k$, and she reverts to π_k thereafter. The politician's policy payoff to this deviation is $\hat{V}_k^1(s) \equiv \hat{v}_k(s)$, and for each s , this satisfies

$$\begin{aligned} V_k(s, k) &= \int_x \left(u_k(s, x) + \delta[\rho(s, k, x)V_k^I(s, k, x) + (1 - \rho(s, k, x))V_k^C(s, k, x)] \right) \pi_k(dx|s) \\ &\leq \hat{v}_k(s), \end{aligned}$$

by (8) and the definition of \hat{v}_k . In particular, $\hat{v}_k(s) \geq V_k(s, k)$.

Next, suppose the type k politician chooses according to $\hat{\pi}_k$ in the first two terms of office and reverts to π_k thereafter. Letting $\hat{V}_k^2(s)$ denote the politician's policy payoff from this deviation in state s , note that

$$\begin{aligned} \hat{V}_k^2(s) &= u_k(s, \hat{x}(s)) + \delta \sum_{s'} p(s'|s, \hat{x}(s)) \hat{V}_k^1(s) \\ &\geq u_k(s, \hat{x}(s)) + \delta \sum_{s'} p(s'|s, \hat{x}(s)) V_k(s, k) \\ &= u_k(s, \hat{x}(s)) + \delta V_k^I(s, k, \hat{x}(s)) \\ &= \hat{v}_k(s), \end{aligned}$$

where the first equality uses $\rho(s, k, \hat{x}(s)) = 1$, the inequality follows from the above argument, the second equality follows from (6), and the last equality uses (18). We conclude that $\hat{V}_k^2(s) \geq \hat{v}_k(s) \geq V_k(s, k)$.

Continuing recursively, we construct a sequence of non-stationary strategies, indexed by m , for the type k politician such that she chooses according to $\hat{\pi}_k$ in the first m terms of office and reverts to π_k thereafter, along with a sequence $\{\hat{V}_k^m\}$ of policy payoffs from deviations of duration m , evaluated at the first term of office, as $m \rightarrow \infty$. This sequence of policy payoffs satisfies $\hat{V}_k^m(s) \geq \hat{v}_k(s) \geq V_k(s, k)$ for all s and all m , and it converges pointwise to the function \hat{V}_k , which is the policy payoff to the type k politician from using the policy strategy $\hat{\pi}_k$. Thus, for each s , we have $\hat{V}_k^m(s) \rightarrow \hat{V}_k(s)$, and since $\hat{V}_k^m(s) \geq \hat{v}_k(s) \geq V_k(s, k)$ holds in each state s for every m , we conclude that

$$\hat{V}_k(s) \geq V_k(s, k) \tag{19}$$

for each state s . This completes the first step.

The second step is to show that, in fact, (19) holds with equality. Note that if the type k politician deviates from π_k to $\hat{\pi}_k$, then she wins with probability one in every state. Thus, the politician's payoff from deviating in state s is $\hat{V}_k(s) + \frac{\beta}{1-\delta}$. Since σ is an equilibrium,

the payoff from deviating to $\bar{\pi}_k$ cannot exceed the politician's equilibrium payoff, it follows that for each s ,

$$V_k(s, k) + B_k(s) \geq \hat{V}_k(s) + \frac{\beta}{1 - \delta}.$$

This implies that $V_k(s, k) \geq \hat{V}_k(s)$ for each s , and by (19), we conclude that

$$\hat{V}_k(s) = \hat{v}_k(s) = V_k(s, k) \quad (20)$$

for each state s , completing the second step.

Returning to part (ii) of the lemma, it follows from (20) that for each s and each x , we have

$$V_k^I(s, k, x) = \sum_{s'} p(s'|s, x) V_k(s', k) = \sum_{s'} p(s'|s, x) \hat{V}_k(s). \quad (21)$$

Next, we claim that for each state s , the policy $\hat{\pi}_k(s)$ solves

$$\max_{x \in X(s)} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) \hat{V}_k(s).$$

Indeed, given any state s , let x' solve the above maximization problem, and note that

$$\begin{aligned} & u_k(s, \hat{x}(s)) + \delta \sum_{s'} p(s'|s, \hat{x}(s)) \hat{V}_k(s) \\ &= u_k(s, \hat{x}_k(s)) + \delta V_k^I(s, k, \hat{x}(s)) \\ &= u_k(s, \hat{x}_k(s)) + \delta [\rho(s, k, \hat{x}(s)) V_k^I(s, k, \hat{x}(s)) + (1 - \rho(s, k, \hat{x}(s))) V_k^C(s, k, \hat{x}(s))] \\ &\geq u_k(s, x') + \delta [\rho(s, k, x') V_k^I(s, k, x') + (1 - \rho(s, k, x')) V_k^C(s, k, x')] \\ &\geq u_k(s, x') + \delta V_k^I(s, k, x') \\ &= u_k(s, x') + \delta \sum_{s'} p(s'|s, x') \hat{V}_k(s), \end{aligned}$$

where the first equality follows from (21), second equality uses $\rho(s, k, \hat{x}(s)) = 1$, the first inequality follows from construction of $\hat{x}(s)$, the second inequality follows from part (ii) of the definition of equilibrium, and the last equality follows from (21). This establishes the claim, and we conclude that \hat{V}_k solves the voter's Bellman equation, i.e., it is the optimal value for the voter. Finally, since $V_k(s, k) = \hat{V}_k(s)$, it follows that $V_k(s, k) \geq \tilde{V}_k(s, k)$ for all s , completing the proof of part (ii).

To prove part (i) of the lemma, we need to show that the function

$$\phi(s, t, x) \equiv u_k(s, x) + \delta [\rho(s, t, x) V_k^I(s, t, x) + (1 - \rho(s, t, x)) V_k^C(s, t, x)]$$

satisfies the recursion

$$\begin{aligned} \phi(s, t, x) &= \max_{r \in \{0, 1\}} u_k(s, x) + \delta \sum_{s'} p(s'|s, x) \left(r \int_{x'} \phi(s', t, x') \pi_t(dx'|s') \right. \\ &\quad \left. + (1 - r) \left[q_t(k|s, x) V_k(s', k) + \sum_{t' \neq k} q_t(t'|s, x) \int_{x'} \phi(s', t', x') \pi_{t'}(dx'|s) \right] \right), \end{aligned}$$

where we use part (ii) to substitute $V_k(s', k)$ for $\tilde{V}_k(s', k)$. Note that by (8), we have

$$V_k(s, t) = \int_x \phi(s, t, x) \pi_t(dx|s),$$

and thus by (6), it follows that

$$V_k^I(s, t, x) = \sum_{s'} p(s'|s, x) \int_{x'} \phi(s', t, x') \pi_t(dx'|s').$$

As well, by (7), we have

$$V_k^C(s, t, x) = \sum_{s'} p(s'|s, x) \left[q_t(k|s, x) V_k(s', k) + \sum_{t' \neq k} q_t(t'|s, x) \int_{x'} \phi(s', t', x') \pi_{t'}(dx'|s) \right].$$

Thus, the recursion reduces to

$$\phi(s, t, x) = \max_{r \in \{0,1\}} u_k(s, x) + \delta [r V_k^I(s, t, x) + (1-r) V_k^C(s, t, x)],$$

which holds by part (ii) of the definition of equilibrium, completing the proof of part (i), as required. \square

Proof of Theorem 6.1. Fix δ , let $\sigma^\delta = (\pi^\delta, \rho^\delta)$ be a k -responsive equilibrium given δ , and let $\phi^*: S \rightarrow X$ be an optimal policy rule for the voter. For comparison with equilibrium dynamics, we let $\tilde{\sigma}$ be a strategy profile in which the type k politician is always re-elected, each type $t \neq k$ politician is always removed from office, and every politician type t chooses according to ϕ^* , i.e., for all s , $\tilde{\pi}_t(\{\phi^*(s)\}|s) = 1$. This determines a Markov chain \tilde{P} on state-type pairs as follows:

$$\tilde{P}((s', t')|(s, t)) = \begin{cases} p(s'|s, \phi^*(s)) & \text{if } t = k = t', \\ 0 & \text{if } t = k \neq t', \\ p(s'|s, \phi^*(s)) q_t(t'|s, \phi^*(s)) & \text{else.} \end{cases}$$

The voter's payoff from σ can be expressed in terms of the Markov chain \tilde{P} as follows: for each s and t ,

$$\tilde{V}_k(s, t) = \sum_{m=1}^{\infty} \delta^{m-1} \sum_{(s', t')} u_k(s', \phi^*(s')) \tilde{P}^m((s', t')|(s, t)),$$

where \tilde{P}^m is the m th product of \tilde{P} . Of course, this is just the value of the representative dynamic programming problem, $V_k^{*, \delta}(s) = \tilde{V}_k(s, t)$. Because $\tilde{P}((s', k)|(s, t)) > 0$ for all s , all s' , and all t , the chain is irreducible, and it possesses a unique ergodic distribution, denoted \tilde{p} , and this places probability one on $S \times \{k\}$. In particular, for all s and all t , we have $\tilde{P}^m(\cdot|s, t) \rightarrow \tilde{p}$. For later use, let

$$\tilde{V} = \frac{1}{1-\delta} \sum_{(s', t')} u_k(s', \phi^*(s')) \tilde{p}((s', t'))$$

denote the voter's expected payoff from the ergodic distribution \tilde{p} , divided by $1 - \delta$.

Now, specify the strategy profile $\hat{\sigma}$ such that: (i) the voter always re-elects the type k politician and rejects all other types, i.e., for each s , each t , and each x ,

$$\hat{\rho}(s, t, x) = \begin{cases} 1 & \text{if } t = k, \\ 0 & \text{else.} \end{cases}$$

(ii) the type k politician chooses according to the optimal policy rule ϕ^* , and (iii) for all $t \neq k$, $\hat{\pi}_t = \pi_t^\delta$. It may be that $\hat{\sigma}$ is not an equilibrium, but by the Unitary Actor Lemma, the voter's payoffs in the equilibrium σ^δ must be at least as great as from the profile $\hat{\sigma}$. The strategy profile σ determines a Markov chain \hat{P} on state-type pairs as follows:

$$\hat{P}((s', t')|(s, t)) = \begin{cases} p(s'|s, \phi^*(s)) & \text{if } t = k = t', \\ 0 & \text{if } t = k \neq t', \\ \int_x p(s'|s, x) q_t(t'|s, x) \hat{\pi}_t(dx|s) & \text{else.} \end{cases}$$

The voter's payoff from $\hat{\sigma}$ is

$$\hat{V}_k(s, t) = \sum_{m=1}^{\infty} \delta^{m-1} \sum_{(s', t')} \left[\int_x u_k(s', x) \hat{\pi}_{t'}(dx|s') \right] \hat{P}^m((s', t')|(s, t)), \quad (22)$$

and as mentioned, the Unitary Actor Lemma implies that $V_k^\delta(s, t) \geq \hat{V}_k(s, t)$ for all s and all t . Again, because $\hat{P}((s', k)|(s, t)) > 0$, this chain is irreducible, and it possesses a unique ergodic distribution, \hat{p} , and this places probability one on $S \times \{k\}$. In particular, for all s and all t , $\hat{P}^m(\cdot|s, t) \rightarrow \hat{p}$. For later use, let

$$\hat{V} = \frac{1}{1 - \delta} \sum_{(s', t')} \left[\int_x u_k(s', x) \hat{\pi}_{t'}(dx|s') \right] \hat{p}((s', t')) \quad (23)$$

denote the voter's expected payoff from the ergodic distribution \hat{p} , divided by $1 - \delta$.

We claim that $\tilde{p} = \hat{p}$. By continuity of q_t and compactness of $X(s)$, the probability of a type k challenger has a positive lower bound,

$$\alpha = \min_{s \in S, t \in T, x \in X(s)} q_t(k|s, x) > 0.$$

Thus, given any s and $t \neq k$, the probability that a type k challenger is drawn over m periods is at least $1 - (1 - \alpha)^m$. According to $\hat{\sigma}$, once a type k politician is drawn to replace the incumbent, she remains in office thereafter, and the Markov chain is identical to \tilde{P} , i.e.,

$$\hat{P}((s', k)|(s, k)) = \tilde{P}((s', k)|(s, k)) \quad \text{and} \quad \hat{P}((s', t')|(s, k)) = \tilde{P}((s', t')|(s, k)) = 0$$

for all s , all s' , and all $t' \neq k$. Together these observations imply that given any initial pair (s, t) , we can write

$$\hat{P}^m(\cdot|(s, t)) = (1 - \alpha)^m \hat{p}^m(s, t) + (1 - (1 - \alpha)^m) \tilde{P}^m(\cdot|(s, t)),$$

where $\hat{p}^m(s, t)$ is a distribution determined by strategies of types $t \neq k$. Taking limits as $m \rightarrow \infty$, we have $\hat{p} = \tilde{p}$, as claimed. Note that since the type k politician uses ϕ^* in $\hat{\sigma}$, the equivalence of \hat{p} and \tilde{p} implies $\tilde{V} = \hat{V}$.

Next, we establish that convergence to the ergodic distributions is geometric, with parameters that depend on primitives of the model, but not the strategies of the players. By continuity of p and compactness of $X(s)$, the probability of transitioning from one state to any other has a positive lower bound,

$$\gamma = \min_{s, s' \in S, x \in X(s)} p(s'|s, x) > 0.$$

Thus, for all s , all s' , and all t , we have

$$\min \{ \hat{P}((s', k)|(s, t)), \tilde{P}((s', k)|(s, t)) \} \geq \alpha\gamma > 0.$$

Standard convergence results for Markov chains (cf. Doob (1953), Case (b), p.173) then imply that

$$|\hat{P}^m((s', k)|(s, t)) - \hat{p}((s', k))| \leq (1 - |S|\alpha\gamma)^{m-1}$$

for all m , and likewise

$$|\tilde{P}^m((s', k)|(s, t)) - \tilde{p}((s', k))| \leq (1 - |S|\alpha\gamma)^{m-1}$$

for all m .

Next, we claim that there exists $\hat{M} > 0$ that is independent of δ and such that for all (s, t) ,

$$|\hat{V}_k(s, t) - \hat{V}| \leq \hat{M}. \tag{24}$$

Indeed,

$$\begin{aligned} |\hat{V}_k(s, t) - \hat{V}| &= \sum_{m=1}^{\infty} \sum_{(s', t')} \delta^{m-1} \left| \hat{P}^m((s', t')|(s, t)) - \hat{p}((s', t')) \right| \int_x u_k(s', x) \hat{\pi}_{t'}(dx|s') \\ &\leq \sum_{m=1}^{\infty} \sum_{(s', t')} \delta^{m-1} \left| \hat{P}^m((s', t')|(s, t)) - \hat{p}((s', t')) \right| \bar{u} \\ &\leq \sum_{m=1}^{\infty} |S||T| \delta^{m-1} (1 - |S|\alpha\gamma)^{m-1} \bar{u} \\ &= \bar{u}|S||T| \sum_{m=1}^{\infty} (\delta - \delta|S|\alpha\gamma)^{m-1} \\ &= \frac{\bar{u}|S||T|}{1 - \delta + \delta|S|\alpha\gamma} \\ &\leq \frac{\bar{u}|S||T|}{|S|\alpha\gamma} \\ &= \frac{\bar{u}|T|}{\alpha\gamma}, \end{aligned}$$

where the first equality follows from (22) and (23), the first inequality from our bound on stage utilities, the second inequality from geometric convergence, and the last inequality from $\alpha\gamma \leq 1$. Setting $\hat{M} = \frac{\bar{u}|T|}{\alpha\gamma}$, the claim is proved.

By an analogous argument, there exists $\tilde{M} > 0$ that is independent of δ and such that for all (s, t) ,

$$|\tilde{V}_k(s, t) - \tilde{V}| \leq \tilde{M}. \quad (25)$$

To finish the proof, set $M = \hat{M} + \tilde{M}$. For all (s, t) , we have

$$\begin{aligned} V_k^{*,\delta}(s) - \hat{V}_k(s, t) &= |\tilde{V}_k(s, t) - \hat{V}_k(s, t)| \\ &\leq |\tilde{V}_k(s, t) - \tilde{V} + \tilde{V} - \hat{V}_k(s, t)| \\ &= |\tilde{V}_k(s, t) - \tilde{V} + \hat{V} - \hat{V}_k(s, t)| \\ &\leq |\tilde{V}_k(s, t) - \tilde{V}| + |\hat{V} - \hat{V}_k(s, t)| \\ &\leq \tilde{M} + \hat{M} \\ &= M, \end{aligned}$$

where the first equality follows from the fact that $\tilde{\sigma}$ achieves the voter's optimal value, the second equality follows from $\tilde{p} = \hat{p}$ (which implies $\tilde{V} = V^*$), and the last inequality follows from (24) and (25). Finally, note that for all s and all t , we have the inequalities

$$V_k^{*,\delta}(s) \geq V_k^\delta(s, t) \geq \hat{V}_k(s, t),$$

and we conclude that for all s and all t ,

$$V_k^{*,\delta}(s) - V_k^\delta(s, t) \leq M,$$

as required. \square

Proof of Corollary 6.1. Note that the normalized values $(1 - \delta)V_k^{*,\delta}(s')$ belong to the compact interval $[0, 1]$, and thus we can without loss of generality consider a subsequence $v^\delta = ((1 - \delta)V_k^{*,\delta}(s'))_{s' \in S}$ with pointwise limit \bar{v} . Let $\bar{\Phi}^*(s)$ denote the voter's optimal policies in state s given these limiting values \bar{v} , i.e.,

$$\bar{\Phi}^*(s) = \arg \max_{x \in X(s)} \sum_{s'} p(s'|s, x) \bar{v}_{s'}.$$

Fix any state s and any type t . Because payoffs are additively separable across time, we can write the representative voter's (normalized) expected discounted payoff from (s, t) , given discount factor δ , as

$$\begin{aligned} (1 - \delta)V_k^\delta(s, t) &= (1 - \delta) \sum_{m=1}^{\infty} \delta^{m-1} \int_{(s', x')} u_k(s', x') \mu_{s,t}^{m,\delta}(d(s', x')) \\ &= \int_{(s', x')} u_k(s', x') \mu_{s,t}^\delta(d(s', x')). \end{aligned} \quad (26)$$

Because $\mu_{s,t}^\delta$ belongs to the set $\Delta(Y)$, which is compact with the weak* topology, we can go to a further subsequence (if needed) such that $\mu_{s,t}^\delta$ converges weak* to a limit μ . Using Theorem 6.1, we have

$$\bar{v}_s = \lim_{\delta \rightarrow 1} (1 - \delta)V_k^{*,\delta}(s) = \lim_{\delta \rightarrow 1} (1 - \delta)V_k^\delta(s, t) = \int_{(s', x')} u_k(s', x') \mu(d(s', x')) \quad (27)$$

By our full support assumption on state transitions, the marginal probability of $\mu_{s,t}^\delta$ and μ on s is positive, and thus, because S is finite, the conditional measures $\mu_{s,t}^\delta(\cdot|s)$ converge weak* to $\mu(\cdot|s)$.

Now, suppose toward a contradiction that the limit does not hold, so that for some $\epsilon > 0$, we have

$$\liminf_{\delta \rightarrow 1} \mu_{s,t}^\delta(B_\epsilon(\bar{\Phi}^*(s))|s) < 1. \quad (28)$$

Since $X(s) \setminus B_\epsilon(\bar{\Phi}^*(s))$ is compact, weak* convergence implies that

$$\mu(X(s) \setminus B_\epsilon(\bar{\Phi}^*(s))|s) > 0,$$

and thus there exists $\eta > 0$ and a Borel measurable subset $Y \subseteq X(s) \setminus B_\epsilon(\bar{\Phi}^*(s))$ such that $\mu(Y|s) > 0$ and for all $x \in Y$, we have

$$\sum_{s'} p(s'|s, x) \bar{v}_{s'} + \eta \leq \max_{x' \in X(s)} \sum_{s'} p(s'|s, x') \bar{v}_{s'}.$$

Since s has positive marginal probability under μ , this implies

$$\int_{(s', x')} u_k(s', x') \mu(d(s', x')) < \max_{x' \in X(s)} \sum_{s'} p(s'|s, x') \bar{v}_{s'}. \quad (29)$$

Define the mapping $U_s: Y \times [0, 1] \times [0, 1]^S \rightarrow \mathbb{R}$ by

$$U_s(x|\delta, v) = (1 - \delta)u_k(s, x) + \delta \sum_{s'} p(s'|s, x)v_{s'},$$

and note that it is jointly continuous in (x, δ, v) . By definition of optimal value, we have

$$v_s^\delta = \max_{x \in X(s)} U(x|\delta, v^\delta).$$

By the theorem of the maximum, this maximized value is continuous, and taking $\delta \rightarrow 1$, we have

$$\int_{(s', x')} u_k(s', x') \mu(d(s', x')) = \bar{v}_s = \max_{x \in X(s)} U(x|1, \bar{v}) = \max_{x \in X(s)} \sum_{s'} p(s'|s, x) \bar{v}_{s'},$$

where the first equality follows from (27). This contradicts (29), however, and we conclude that $\mu_{s,t}^\delta(B_\epsilon(\bar{\Phi}^*(s))|s) \rightarrow 1$, as required. \square

Proof of Corollary 6.2. Both parts of this result rely on the following claim, discussed in the text: given any δ , assume that $\delta\beta$ is large, and let σ^δ be a responsive voting equilibrium. Then, for each states s and each type t , we have $\rho(s, t, x) > 0$ for almost all $x \in \text{supp}(\pi_t^\delta(\cdot|s))$. To see this, fix state s and type t , and consider a policy strategy π_t' for t in which she is re-elected with probability one in all states: $\int \rho(s', t, x)\pi_t'(dx|s') = 1$ for all s' . Such a policy strategy exists by part (ii) of the definition of responsive voting equilibrium. Recalling the normalization $\bar{u} = 1$ and $\underline{u} = 0$, the payoff to type t from strategy π_t' in state s is no less than $\frac{\beta}{1-\delta}$. Now consider an open set of policies $A \subseteq X(s)$ following which type t is not re-elected in s , i.e., $\rho(s, t, x) = 0$ for all $x \in A$, and let π_t'' be any policy strategy for type t that puts positive probability on policies in A , i.e., $\pi_t''(A|s) > 0$. The payoff to type t in s from π_t'' is at most $\beta + \frac{1}{1-\delta}$. If $\delta\beta > 1$, then the payoff to t in s from π_t'' is strictly lower than her payoff from π_t' , and hence π_t'' cannot be optimal.

The claim established above ensures that property (17) holds, and the remainder of part (ii) of Corollary 6.2 is proved in the text. The proof of part (i) follows from steps closely related to those in the proof of Corollary 6.1. Instead of repeating that proof, we explain how to modify it to prove our result. Fix state s , type t , and discount factor δ , and recall from the preface to Corollary 6.1 that $\mu_{s,t}^\delta$ denotes the discounted and time-aggregated marginal distribution over state-policy pairs (s', x') generated by the equilibrium profile σ^δ , with μ denoting its limiting distribution as $\delta \rightarrow 1$. Analogously, let $\tilde{\mu}_{s,t}^\delta$ denote the discounted and time-aggregated marginal distribution over state-policy pairs (s', x') generated by the equilibrium policy distribution π_t^δ of type t incumbents starting from state s , conditional on this incumbent staying in office in all future periods, with $\tilde{\mu}$ denoting its limit as $\delta \rightarrow 1$. Consider a voting strategy, $\tilde{\rho}^\delta$, identical to the equilibrium voting strategy ρ^δ except that in all states s' , $\rho(s', t, x) = 1$ whenever $V_k^I(s', t, x) \geq V_k^C(s', t, x)$. By construction, $\tilde{\rho}^\delta$ must also solve the voter's optimal retention problem (part (i) of the Unitary Actor Lemma), so that the voter achieves the same payoffs under $(\pi_t^\delta, \tilde{\rho}^\delta)$ as under $(\pi_t^\delta, \rho^\delta)$. Notice also that by property (17), the voter retains the type t politician with probability one in all states under strategy $\tilde{\rho}^\delta$. Therefore, recalling (26) from the proof of Corollary 6.1, we have

$$\begin{aligned} (1 - \delta)V_k^\delta(s, t) &= \int_{(s', x')} u_k(s', x')\mu_{s,t}^\delta(d(s', x')) \\ &= \int_{(s', x')} u_k(s', x')\tilde{\mu}_{s,t}^\delta(d(s', x')). \end{aligned} \quad (30)$$

Now, suppose toward a contradiction that result of part (i) of Corollary 6.2 does not hold, so that for some $\epsilon > 0$, we have

$$\liminf_{\delta \rightarrow 1} \pi_t^\delta(B_\epsilon(\bar{\Phi}^*(s))|s) < 1.$$

By the construction of the distribution $\tilde{\mu}_{s,t}^\delta$ from π_t^δ , this implies that

$$\liminf_{\delta \rightarrow 1} \tilde{\mu}_{s,t}^\delta(B_\epsilon(\bar{\Phi}^*(s))|s) < 1. \quad (31)$$

Using (31) instead of (28) to mimic the steps in the proof of Corollary 6.1, and exploiting

(30), leads to the following version of (29):

$$\int_{(s',x')} u_k(s',x')\mu(d(s',x')) = \int_{(s',x')} u_k(s',x')\tilde{\mu}(d(s',x')) < \max_{x' \in X(s)} \sum_{s'} p(s'|s,x')\bar{v}_{s'},$$

after which the remainder of the proof of Corollary 6.1 can be applied. \square

References

- Acemoglu, D., M. Golosov, and A. Tsyvinski (2008). Political economy of mechanisms. *Econometrica* 76, 619–641.
- Acemoglu, D. and J. A. Robinson (2000). Why did the west extend the franchise? democracy, inequality, and growth in historical perspective. *Quarterly Journal of Economics* 115(4), 1167–1199.
- Acemoglu, D., D. Ticchi, and A. Vindigni (2011). Emergence and persistence of inefficient states. *Journal of the European Economic Association* 9(2), 177–208.
- Alesina, A. (1987). Macroeconomic policy in a two-party system as a repeated game. *Quarterly Journal of Economics* 102, 651–678.
- Alesina, A. (1988). Credibility and policy convergence in a two-party system with rational voters. *American Economic Review* 78, 796–806.
- Alesina, A. and D. Rodrik (1994). Distributive politics and economic growth. *Quarterly Journal of Economics* 109, 465–490.
- Alesina, A. and G. Tabellini (1990). A positive theory of budget deficits and government debt. *Review of Economic Studies* 57, 403–414.
- Anderson, C. J. (2007). The end of economic voting? contingency dilemmas and the limits of democratic accountability. *Annual Review of Political Science* 10, 271–296.
- Ashworth, S. (2012). Electoral accountability: Recent theoretical and empirical work. *Annual Review of Political Science* 15, 183–201.
- Ashworth, S., E. Bueno de Mesquita, and A. Friedenbergh (2017). Accountability and information in elections. *American Economic Journal: Microeconomics* 9(2), 95–138.
- Ashworth, S. and E. B. de Mesquita (2014). Is voter competence good for voters?: Information, rationality, and democratic performance. *American Political Science Review*, 565–587.
- Azzimonti, M. (2011). Barriers to investment in polarized societies. *American Economic Review* 101, 2182–2204.
- Azzimonti, M. (2015). The dynamics of public investment under persistent electoral advantage. *Review of Economic Dynamics* 18(3), 653–678.

- Bai, J. H. and R. Lagunoff (2011). On the faustian dynamics of policy and political power. *Review of Economic Studies* 78(1), 17–48.
- Banks, J. and J. Duggan (2008). A dynamic model of democratic elections in multidimensional policy spaces. *Quarterly Journal of Political Science* 3, 269–299.
- Banks, J. and R. Sundaram (1998). Optimal retention in agency problems. *Journal of Economic Theory* 82, 293–323.
- Barelli, P. and J. Duggan (2020). Subgame-perfect equilibrium in games with almost perfect information: Dispensing with public randomization. *Theoretical Economics forthcoming*.
- Battaglini, M. (2014). A dynamic theory of electoral competition. *Theoretical Economics* 9(2), 515–554.
- Battaglini, M. and S. Coate (2008). A dynamic theory of public spending, taxation, and debt. *American Economic Review* 98, 201–236.
- Bernhardt, D., O. Camara, and F. Squintani (2011). Competence and ideology. *Review of Economic Studies* 78(487-522).
- Bernhardt, D., L. Campuzano, F. Squintani, and O. Camara (2009). On the benefits of party competition. *Games and Economic Behavior* 66, 685–707.
- Bernhardt, D., S. Dubey, and E. Hughson (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. and A. Case (1995). Does electoral accountability affect economic policy choices? evidence from gubernatorial term limits. *Quarterly Journal of Economics* 110(3), 769–798.
- Besley, T. and S. Coate (1997). An economic model of representative democracy. *Quarterly Journal of Economics* 112, 85–114.
- Besley, T. and S. Coate (1998). Sources of inefficiency in a representative democracy: A dynamic analysis. *American Economic Review* 88, 139–156.
- Bowen, T. R., Y. Chen, and H. Eraslan (2014). Mandatory versus discretionary spending: The status quo effect. *American Economic Review* 104(10), 2941–74.
- Callander, S. and D. Raiha (2017). Durable policy, political accountability, and active waste. *Quarterly Journal of Political Science* 12(1), 59–97.
- Camara, O. (2012). Economic policies of heterogeneous politicians.
- Doob, J. (1953). *Stochastic Processes*. New York: John Wiley & Sons.
- Downs, A. (1957). *An Economic Theory of Democracy*. New York: Harper and Bros.

- Duggan, J. (2000). Repeated elections with asymmetric information. *Economics and Politics* 12, 109–136.
- Duggan, J. (2012). Noisy stochastic games. *Econometrica* 80, 2017–2045.
- Duggan, J. (2014). A folk theorem for repeated elections with adverse selection. *Political Science Research and Methods* 2, 213–242.
- Duggan, J. (2017). Term limits and bounds on policy responsiveness in dynamic elections. *Journal of Economic Theory* 170, 426–463.
- Duggan, J. and J. G. Forand (2018). Existence of markov electoral equilibria. Unpublished paper.
- Duggan, J. and J. G. Forand (2021). Representative voting games. *Social Choice and Welfare*, forthcoming.
- Duggan, J. and C. Martinelli (2017). The political economy of dynamic elections: Accountability, commitment, and responsiveness. *Journal of Economic Literature* 55(3), 916–84.
- Duggan, J. and C. Martinelli (2020). Electoral accountability and responsive democracy. *The Economic Journal* 130(627), 675–715.
- Dziuda, W. and A. Loeper (2016). Dynamic collective choice with endogenous status quo. *Journal of Political Economy* 124(4), 1148–1186.
- Esponda, I. and D. Pouzo (2019). Retrospective voting and party polarization. *International Economic Review* 60(1), 157–186.
- Fearon, James, D. (1999). Electoral accountability and the control of politicians: Selecting good types versus sanctioning poor performance. In A. Przeworski, S. Stokes, and B. Manin (Eds.), *Democracy, Accountability, and Representation*. New York: Cambridge University Press.
- Ferejohn, J. (1986). Incumbent performance and electoral control. *Public Choice* 50, 5–25.
- Forand, J. G. (2014). Two-party competition with persistent policies. *Journal of Economic Theory* 152, 64–91.
- He, W. and Y. Sun (2020). Dynamic games with (almost) perfect information. *Theoretical Economics* 15, 811–859.
- Healy, A. and N. Malhotra (2009). Myopic voters and natural disaster policy. *American Political Science Review*, 387–406.
- Healy, A. and N. Malhotra (2013). Retrospective voting reconsidered. *Annual Review of Political Science* 16, 285–306.
- Huber, G. A., S. J. Hill, and G. S. Lenz (2012). Sources of bias in retrospective decision making: Experimental evidence on voters’ limitations in controlling incumbents. *American Political Science Review*, 720–741.

- Kayser, M. A. and M. Peress (2012). Benchmarking across borders: electoral accountability and the necessity of comparison. *American Political Science Review*, 661–684.
- Krusell, P., V. Quadrini, and J.-V. Rios-Rull (1997). Politico-economic equilibrium and economic growth. *Journal of Economic Dynamics and Control* 21, 243–272.
- Krusell, P. and J.-V. Rios-Rull (1999). On the size of the us government: Political economy in the neoclassical growth model. *American Economic Review* 89, 1156–1181.
- Levy, Y. (2013). Discounted stochastic games with no stationary nash equilibrium: Two examples. *Econometrica* 81, 1973–2007.
- Luttmer, E. and T. Mariotti (2003). The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A comment. *Econometrica* 71, 1909–1911.
- Meirowitz, A. (2007). Probabilistic voting and accountability in elections with uncertain policy constraints. *Journal of Public Economic Theory* 9, 41–68.
- Nowak, A. and T. Raghavan (1992). Existence of stationary correlated equilibria with symmetric information for discounted stochastic games. *Mathematics of Operations Research* 17, 519–526.
- Osborne, M. and A. Slivinski (1996). A model of competition with citizen candidates. *Quarterly Journal of Economics* 111, 65–96.
- Persson, T. and L. Svensson (1989). Why a stubborn conservative would run a deficit: Policy with time-inconsistent preferences. *Quarterly Journal of Economics* 104, 325–345.
- Persson, T. and G. Tabellini (1990). *Macroeconomic Policy, Credibility and Politics*. Amsterdam: Harwood Academic Publishers.
- Persson, T. and G. Tabellini (2000). *Political Economics: Explaining Economic Policy*. Cambridge, MA: MIT Press.
- Rogoff, K. (1990). Equilibrium political budget cycles. *American Economic Review* 80, 21–36.
- Van Weelden, R. (2013). Candidates, credibility, and re-election incentives. *Review of Economic Studies* 80(4), 1622–1651.
- Woon, J. (2012). Democratic accountability and retrospective voting: A laboratory experiment. *American Journal of Political Science* 56(4), 913–930.
- Yared, P. (2010). Politicians, taxes and debt. *Review of Economic Studies* 77, 806–840.