Electoral Accountability in Markovian Elections

John Duggan* Jean Guillaume Forand†

June 13, 2022

Abstract

We study electoral accountability in a dynamic environment with complete information. As our normative benchmark, we take the solution of the dynamic programming problem facing the representative voter as if he chose policy directly. There always exist equilibria in which the politician type corresponding to the voter is accountable, in the sense that these politicians achieve the idealized benchmark; and 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 two examples of novel political failures in a model of dynamic public investment. We do not allow the voter to commit to an optimal re-election rule; nevertheless, we identify a class of responsive voting equilibria such that voter welfare converges to the idealized benchmark as the voter becomes patient for every selection of such equilibria.

Keywords: Electoral Accountability; Dynamic Games; Representative 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

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

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

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 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 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)). Dynamic models are particularly important in applications, but they raise difficult issues surrounding the existence and characterization of Markovian equilibria, as 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 complexity of dynamic electoral incentives, for an office holder typically needs to choose between policies that satisfy the voters' standard for re-election on the one hand, and policies that sacrifice office for short-run gains on the other, which introduces a non-convexity into her optimization problem. As a consequence, equilibrium existence can depend on mixing between "shirking" (choosing the politician's best policy and being removed from office) or "compromising" policy to meet the bar for re-election. Furthermore, the calculus of both voters and politicians must be forward-looking and take into account the impact of current policy choices on future policy. In sum, dynamic models of elections present a steep analytical trade off between tractability and generality of the model.

We analyze a dynamic framework for elections that develops the standard citizen-candidate approach in a parsimonious yet general way: in each period, a political/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 citizen-candidate models, 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, any continuous stage utilities, and any continuous transition probabilities on the state and the challenger's type. To examine the effectiveness of elections when the linkage across periods is through the political/economic state variable, we abstract from private information about office holders' types: in contrast to much work on electoral accountability, we assume that a politician's type is publicly observed once she takes office. This shuts down signaling incentives, and it allows us to focus on dynamic incentives of politicians, who can use current policy to influence the calculus of voters and the choices of future politicians—our main channel of interest. We do allow, however, for ex ante uncertainty about the challenger's type, so that voters may have less information about a challenger relative to the known incumbent prior to an election.

We begin with the observation that in general dynamic environments, the representative voter's policy preferences are state dependent, so that if policy choices affect the evolution of the state, then the appropriate normative benchmark is not simply the voter's static ideal point. Rather, it is the value of the dynamic programming problem facing the representative voter as if he chose policy directly, and the analogue to the median ideal policy is the set of policy rules that solve this *representative dynamic programming problem.* We establish the 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.[2] We also observe that if politicians have sufficiently high office incentives, then there is an equilibrium in which all politician types are accountable. These results raise several questions: Can there be equilibria in which electoral incentives prevent congruent politicians from choosing voter-optimal policies? If so, are there conditions under which elections will discipline congruent politicians to choose optimally for the voter? And if so, does this force non-congruent politicians to also choose policies that are acceptable to the voter?

[1] Our results extend to the model with a countably infinite set of states in a straightforward way.

[2] We 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 voters, and they give sufficient conditions for the existence of a representative voter in each state.

Before addressing these questions, we emphasize that our framework stacks the deck against accountability: we do not allow politicians to commit to policies prior to taking office, and we do not allow the voter to commit to 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: this means that congruent politicians choose optimally for the voter, while politicians with different ideal points do not. But when the horizon is extended (either finite or infinite), the incumbent's policy choice in all but the final period can influence future states, and therefore the policy choices of future politicians. Moreover, because the incumbent may be more desirable to the voter in some states than in others, and because her policy choice can affect the distribution of the challenger's type, the incumbent's policy choice also influences the voter's decision. The incentives of an incumbent are complex, and it is unclear whether or when they would align with the preferences of the voter.

The dynamic linkage between periods, which arises naturally in the absence of commitment or any form of pre-electoral politics, can lead to new forms of political inefficiency through 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. Underlying this pathology is the dual commitment problem we highlighted at the outset, which leads to coordination failure between the voter and the congruent politician. 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. This restriction on voting strategies integrates components of retrospective and prospective voting from the literature on voting behavior, and we term it *k-responsive voting.* We establish that in a k-responsive voting equilibrium, the voter and congruent 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: given any sequence of k-responsive voting equilibria as the voter becomes patient, the expected (normalized) 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 non-congruent 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 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 incumbent chooses policies that are undesirable for the voter, but is nevertheless re-elected, because the next state is likely to be one in which other non-congruent types are even worse in expectation. 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.[3]

We illustrate our positive results and the possibility of political failures in a model of public investment that serves as a vehicle for examples throughout the paper. 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. We interpret the state as the level of a discrete stock of a durable public good that depreciates stochastically, and we interpret the policy choice as a level of public investment

[3]We 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.

that shifts probability toward states in which the stock of public good is high. For tractability, we assume that stage utilities are quadratic, that there are just two states, and that the transition probability is linear in the level of investment. 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 our functional form assumptions allow us to solve the representative dynamic programming problem analytically. 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, and highlighting the importance of a normative benchmark consistent with the dynamic structure of the environment.

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 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 formulating policymaking as a game between parties that cannot commit to policy instruments prior to an election;[4] 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 results in Theorems 4.1 and 4.2. 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

[4]Persson 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.

voters, irrespective of the strength of office motivation, 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 as a consequence. Persson and Svensson (1989) and Alesina and Tabellini (1990) show that incumbents may distort spending in order to "tie the hands" of potential successors with different preferences. In contrast to the latter papers, Aghion and Bolton (1990) feature endogenous elections, and they show that when default on the public debt is costless, if the right-wing party is in power in the first period, then it can sometimes use public debt to change the election outcome in its favor; this is closely connected to our study of political failures due to state manipulation. Besley and Coate (1998) find that politicians may fail to implement Pareto-improving investments if they anticipate that future policy-makers will not choose to reap 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 prospect of government turnover can lead to inefficiencies in either private or public capital accumulation, and Battaglini (2014) gives a sharp characterization of equilibria in a rich model of elections, assuming parties are vote-maximizing and voting is subject to uncertainty. He shows that when voters' preferences are time-varying, temporarily powerful districts can attract inefficiently high levels of government spending. In a finite-horizon model, Callander and Raiha (2017) admit the possibility that policies can affect the outcome of elections, as in Aghion and Bolton (1990) and our model, and they show that politicians may distort investments in a durable public good to compel voters to re-elect them.[5] Finally, Bai and Lagunoff (2011) and Duggan and Forand (2021) show that these concerns are magnified when the identity of future representative voters is determined by current policy choices, introducing additional distortions into policy and voting decisions.[6]

[5] The 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.

[6] Anticipated 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

More broadly, our paper is related to the literature in economics and political science that studies the possibility and limits of electoral accountability. Much of the existing work in this literature extends Barro (1973) by adding private information, in the form of adverse selection or moral hazard or both, and it imposes significant structure on the electoral environment to achieve tractability. Duggan (2000) and Bernhardt et al. (2004) study models of pure adverse selection with no state variable or with the term of office as state variable, respectively.[7] Ferejohn (1986) considers a model of pure moral hazard, and Fearon (1999), Besley and Smart (2007), Acemoglu et al. (2103), and Duggan and Martinelli (2020) analyze two-period models with adverse selection and moral hazard. In comparison, we abstract from private information and learning (except for the possibility of uncertainty about the challenger's type), and we consider general dynamic environments. If we take Barro (1973) as our starting point, then our direction of departure is to add a linkage across periods in the form of an economic state variable, like the literature cited in the preceding paragraph, and to consider the effectiveness of elections when politicians can, in principle, influence election outcomes and the policy choices of future politicians.

2 Model

Political environment A representative voter decides between an incumbent politician and a challenger in an infinite sequence of elections. The voter is assigned a fixed type k from the finite type 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 state and winner's type are publicly observed; and the process repeats.[8] We assume

et al. (2011)).

[7]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.

[8]We can 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 paper, we

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.

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 $\beta \geqslant 0$. We assume that $u_t\colon S \times X \to \Re$ is continuous (and therefore bounded). Let $\underline{u}, \overline{u} \in \Re$ be bounds such that for all s, all x, and all t, we have $\underline{u} \leqslant u_t(s,x) \leqslant \overline{u}$. For convenience, we normalize these bounds so that $\underline{u} = 0$ and $\overline{u} = 1$. Each type t citizen discounts flows of payoffs by the factor $\delta \in [0,1)$. Thus, given a sequence $(s_1, x_1, s_2, x_2, \ldots)$ 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 $p\colon S \times S \times X \to [0,1]$ is continuous. Our asymptotic accountability results require some notion of recurrence for states, and for expositional simplicity, we assume that the state transition places positive probability on all states, i.e., $p(s'|s,x) > 0$ for all states $s', s \in S$ and policies $x \in X$. The dependence of future states on current policy choices underpins an incumbent's incentives to manipulate the state to her electoral advantage, and Examples 1 and 2 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. We take a reduced form approach to challenger selection, which 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, $q_t\colon T \times S \times X \to [0,1]$, is continuous for each type t. We also assume that the challenger is congruent with positive probability in each state and

omit it for simplicity.

following any policy choice by any incumbent: $q_t(k|s,x) > 0$ for all t, s and x. The possibility of eventually drawing a type k challenger plays an important role in our asymptotic accountability results. The dependence of future challengers on current policy choices has the potential to bias incumbents in favor of policies that engender weaker opponents. Although we do not focus on such "challenger manipulation" in this paper, our results allow for its possibility.

Representative dynamic programming problem 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 representative 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,

$$V_k^*(s) \quad = \quad \max_{x \in X(s)} u_k(s,x) + \delta \sum_{s'} p(s'|s,x) V_k^*(s'). \tag{1}$$

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

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

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

Strategies A *policy strategy* for a type t politician is a mapping $\pi_t\colon S \to \Delta(X)$, 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 $\rho\colon S \times T \times X \to [0,1]$, where $\rho(s,t,x)$ is 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, the discounted expected utility of a type t citizen from re-electing a type t' incumbent who chooses policy x in state s satisfies: for all $x \in X(s)$,

$$V_t^I(s,t',x) \quad = \quad \sum_{s'} p(s'|s,x) V_t(s',t'), \tag{2}$$

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

$$V_t^C(s,t',x) \quad = \quad \sum_{t''} q_{t'}(t''|s,x) \sum_{s'} p(s'|s,x) V_t(s',t''). \tag{3}$$

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

$$\begin{aligned} V_t(s,t') \quad = \quad \int_x \Big[& u_t(s,x) + \delta[\rho(s,t',x) V_t^I(s,t',x) \\ & + (1-\rho(s,t',x)) V_t^C(s,t',x)] \Big] \pi_{t'}(dx|s), \end{aligned} \tag{4}$$

reflecting that the office holder chooses a policy x using to the policy strategy $\pi_{t'}(\cdot|s)$ and is either re-elected or replaced by a challenger, accordingly.

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,

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

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.

Equilibrium 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

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

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

$$\rho(s,t,x) \quad = \quad \begin{cases} 1 & \text{if } V_k^I(s,t,x) > V_k^C(s,t,x)\}, \\ 0 & \text{if } V_k^I(s,t,x) < V_k^C(s,t,x)\}, \end{cases}$$

where $\rho(s,t,x)$ is unrestricted if $V_k^I(s,t,x) = V_k^C(s,t,x)$.

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) $\int_x \rho(s,t,x)\pi_t(dx|s) = 1$.

3 Failures of Accountability

In Section 4, we define a class of equilibria in which congruent politicians are accountable, and which in turn yield approximately optimal payoffs to a patient representative voter. We also show that for a more restrictive class of equilibria, non-congruent types are also (at least approximately) accountable. In this section, we motivate these results by showing how incumbents' incentives to manipulate future states allows for equilibria with accountability failures. In doing so, we introduce a simple special case of our model applied to public investment in a durable public good. We assume that the stock of public goods can be either high or low, and that this stock evolves stochastically, with higher investment increasing the probability that it is high.

To capture this in our setting, assume that the policy space is $X = [0,1]$, where policy x represents the level of investment in the public good, and that the state space is $S = \{s^1, s^2\}$, where s^2 represents the high stock. Stage utilities are quadratic in investment and additive in the benefit of the public good: $u_t(s^j, x) = -(x - \hat{x}_t^j)^2 + \Gamma^j$ for each type t and each state s^j. Here, $\hat{x}_t^j$ is a the state-dependent ideal investment for the type t politician, reflecting the possibility that types differ in their willingness to invest, and $\Gamma^2 > \Gamma^1$ are the state-dependent benefits from the public good. For simplicity, we assume that the probability of a transition to the high state increases linearly with investment and is independent of the current state, i.e., $p(s^2|x) = x$, and that ideal investment levels are ordered by type, i.e., $\hat{x}_1^j \leqslant \hat{x}_2^j \leqslant \cdots \leqslant \hat{x}_n^j$ for each $j = 1, 2$. Finally, to capture the idea that the opportunity cost of investment is lower when the public good stock is high, we assume that ideal investment levels are increasing in the state: $\hat{x}_t^1 \leqslant \hat{x}_t^2$ for all t.

With this structure in place, it is straightforward to solve the representative dynamic programming problem. The representative voter's value in state s^j, which we denote by V_k^{j*}, satisfies

$$V_k^{j*} \;=\; \max_{x\in[0,1]} -(x-\hat{x}_k^j)^2 + \Gamma^j + \delta[xV_k^{2*} + (1-x)V_k^{1*}]. \tag{5}$$

The first order condition of the Bellman equation yields $x_k^{j*} = \hat{x}_k^j + \frac{\delta}{2}[V_k^{2*} - V_k^{1*}]$, which is a linear function of the difference $V_k^{2*} - V_k^{1*}$ in continuation values. We can substitute this expression for x_k^{j*} into (5) for $j = 1, 2$, and then solve for the difference directly, and we arrive at the voter-optimal investment level

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

in each state s^j, $j = 1, 2$, where $\Gamma = \Gamma^2 - \Gamma^1$ is the difference in public good benefits. Unless the voter is myopic, the optimal investment in any state is higher than her ideal investment level, so $\hat{x}_k^j$ is *not* the relevant normative benchmark. Rather, reflecting the greater benefit of the public good in the high state, the voter would prefer to invest more than her static ideal to increase the probability of the high state.

In our model, there always exist equilibria in which congruent politicians invest optimally for the voter, and if politicians are sufficiently office-motivated, then there always exist equilibria in which non-congruent types are also accountable. Therefore, the problem of electoral accountability reduces to the possibility of equilibrium multiplicity: do there exist equilibria in which politicians do not respond to electoral incentives by solving the problem of the representative voter? It is clear that if politician types $t \neq k$ do not value office, then equilibria in which they are not accountable can exist widely. Perhaps more surprisingly, 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 Example 1, below, 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 1 (Curse of ambition). Assume that there are two citizen types, $T = \{1, 2\}$, with $k = 2$, that the challenger is type k with probability $q \in (0, 1)$, and that the type $1 \neq k$ politician's ideal investment is state-independent and always lower than the representative voter's ideal invesment, so that $\hat{x}_1^1 = \hat{x}_1^2 \equiv \hat{x}_1 < \hat{x}_k^1 < \hat{x}_k^2$. Assume further that

$$\Gamma - (x_1^* - \hat{x}_k^2)^2 \;<\; -(x_1^* - \hat{x}_k^1)^2. \tag{6}$$

It can be verified that the type 1 politician's optimal optimal policy rule is to choose $x_1^* = \hat{x}_1 + \frac{\delta\Gamma}{2}$ is each state. Thus, even if type 1 always underinvests from the perspective of the voter, (6) says that this distortion is worse for the voter when the stock of public good is high: the voter receives a greater payoff from the investment of type 1 politicians in the low state.

Given any δ and sufficiently high β, 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^*, and hence underinvests, in all states. For her part, the type k politician overinvests in all states. We present the detailed construction of equilibrium policy choices in the Appendix, but their key feature is that underinvestment of type 1 hurts the voter more when the opportunity cost of investment is low (in the high state), whereas the overinvestment of type k hurts the voter more when this cost is high (in the low state), i.e.,

$$V_k(s^2, k) \;>\; V_k(s^2, 1) \quad \text{and} \quad V_k(s^1, 1) \;>\; V_k(s^1, k). \tag{7}$$

Because investment increases the likelihood of the high state, the voter uses 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 invests more than $\tilde{x}$, and he re-elects type 1 if and only if she invests less than $\tilde{x}$. If the type k politician is sufficiently office motivated, then her equilibrium policy in the low state is $x_k^1 = \tilde{x} > x_k^{1*}$: she overinvests just enough to leave the voter indifferent between re-electing her and opting for the challenger. As the type k politician wants to invest more in the high state, she faces no policy-office tradeoff there: the voter strictly prefers to retain her in s^2 following her equilibrium policy $x_k^2 > \tilde{x}$. Finally, because type 1 invests $x_1^* < \tilde{x}$ in all states, the voter strictly prefers to retain her. □

The curse of ambition stems from a coordination failure between the voter and the type k politician: because the voter prefers to have this politician in office when the stock of public goods is high, the incumbent is only rewarded when she overinvests. Furthermore, the type k politician would be punished in the low state if she tried to reduce her investment: although this would benefit both the voter and the politician, it would increase the probability of the low state, leading the voter to prefer a low-spending type 1 politician in office. The manipulation of the future stock of public goods through current investment is key to the curse of ambition: it cannot arise if this stock is constant (i.e., if S is a singleton), or if it is independent of current investment (i.e., if p is independent of x). The negative finding of the example survives even if politicians are more willing to compromise in order to secure re-election, through increases in either 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, 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.

In Example 1, 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, the latter could not be re-elected. With more than two politician types, however, it is possible that the type k politician always chooses optimally for the voter, all politician types $t \neq k$ choose suboptimally for the voter, and they are *still* always re-elected. This possibility arises because a type $t \neq k$ politician's choice may lead to states where the challenger is even worse than the incumbent, and it confronts the voter with a hold-up problem, forcing him to re-elect the incumbent, despite the fact that her policy choices are suboptimal, and despite the fact that the type k politician is accountable.

Example 2 (Political hold-up problem). In the model of public investment, 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 investments 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 myopically ideal investment 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

$$q \quad < \quad \frac{(1-\delta)(1-2\hat{x})}{1-\delta(1-2\hat{x})}, \tag{8}$$

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 $\Gamma \approx 0$.

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 future stock of public goods, and each such politician type invests in a way that pushes 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 increases the probability of the low state and is preferred to the challenger, because the voter wants to avoid electing a high-investment type 3 in this case. The symmetric logic leads the voter to strictly prefer type 3 to the challenger after she overinvests in each 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 (8). Note that the right-hand side of (8) is decreasing in δ. Thus, if a type k challenger is likely, so that the left-hand side of (8) is higher, then equilibrium construction requires that the voter be relatively impatient. □

As with the curse of ambition, the political hold-up problem can only arise through state manipulation. If the stock of public good is constant or evolves exogenously, then there cannot be any equilibrium in which type k is accountable, *all* non-congruent politicians are re-elected, and some of these implement suboptimal investments. There remains the possibility that under some conditions, compelling equilibria possess positive welfare properties for the voter. Example 1 shows that any such class of equilibria must overcome the curse of ambition, so that voter behavior provides incentives for congruent politicians to choose optimal policies. However, Example 2 shows that type k accountability alone is not enough for the voter to achieve his optimal payoff in all states. To gain insight into conditions under which elections induce accountability, recall that the hold-up problem creates a political failure when inequality (8) holds, and that given any $q > 0$, the inequality fails for δ close enough to one. In fact, the construction in the example fails as $\delta \to 1$, 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. We turn to this question in the next section.

4 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 "retrospective voting" rules that condition their choices on past performance in a simple way, whether these rules are consistent with rational voting, and whether they lead to electoral accountability.[9] Theoretically, retrospective voting has been formulated in different ways,[10] but it is known that the simple conditioning of voting decisions on past outcomes 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. We maintain the equilibrium viewpoint and investigate the normative properties of a class of equilibrium voting strategies, which we call "responsive voting," that captures the intuitive features of simple retrospective rules while allowing for more general voting behavior.

4.1 Accountability of the Congruent Type

To motivate our focus on responsive voting, recall that the state manipulation illustrated in Example 1 rests on a disconnect between politicians' performance and their resulting electoral rewards. Counterintuitively, the type k incumbent is replaced even if she chooses the best possible policy for the voter, given the profile of voting and policy strategies: if the politician best responds for the voter, then she increases the probability of transitioning to a state in which she is the worst possible office holder

[9] Anderson (2007) and Healy and Malhotra (2013) survey the empirical literature on retrospective 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.

[10] For 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.

for the voter, leading the voter to elect the challenger after such choices. These considerations lead us to 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 voting used in the literature.

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 weakly exceeds the threshold and guarantees re-election. This definition is permissive in several ways: first, it only applies to type k politicians; second, the threshold in any state can be high; and 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

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

(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

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

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 Definition 1 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 subsection, we examine the welfare properties of the larger class of k-responsive equilibria.

The main result of this section establishes that a k-responsive voting equilibrium exists,[11] and 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, the state, and the incumbent type. Thus, regardless of the discount factor, the voter's payoffs in a k-responsive voting equilibrium are closely related to the voter's optimal value. The impact of this result is magnified by the fact that the discounted payoffs compared are not normalized: they are discounted sums of stage utilities that may be expected to "blow up" for discount factors close to one. A direct implication is that if for any sequence of discount factors δ such that the representative voter becomes arbitrarily patient, then for every corresponding sequence of k-responsive voting equilibria, the voter's *normalized* payoffs become close to the value of the representative dynamic programming problem. 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 4.1. *There is k-responsive voting equilibrium, and the difference between the expected discounted sum of voter payoffs in equilibrium and the value of the representative dynamic programming problem is uniformly bounded: there exists $M > 0$ such that given arbitrary discount factor δ, if σ^{δ} is a k-responsive voting equilibrium, then for all states s and all types t,*

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

Furthermore, as the voter becomes patient, the normalized equilibrium payoffs of the voter converge to the optimum: for all s and all t,

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

[11] Theorem B.1, in the Supplementary Appendix, shows that there is an equilibrium in which the type k politician is accountable, giving us a k-responsive equilibrium. Such an equilibrium is easy to construct when there are only two politician types, but the general argument leverages an existence result in Duggan and Forand (2018).

Theorem 4.1 establishes that 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 directly implies that the normalized payoffs become arbitrarily close, so that equilibria become *approximately optimal* for the representative voter as the voter becomes patient.

The proof of the result involves two steps and uses 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 fix any k-responsive voting equilibrium, and we "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, explained in more detail below, 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. In the second step, we note that the Markov chains determined by these two profiles converge at a geometric rate to their ergodic distribution, and this implies that the difference between the voter's payoffs from the two profiles has a bound M with the properties stated in Theorem 4.1. The construction of the second strategy profile, which gives a lower bound on the voter's equilibrium payoff and has an ergodic distribution that places probability one on voter-optimal policies in all states, 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 the voter and the congruent politician type.

To describe the first step of the proof in more detail, we consider a variant of the electoral game in which the representative voter and all type k politicians act as a unitary player: this player controls electoral outcomes and, when a congruent politician holds office, policy choices as well. Given any strategy profile σ in the original game, we can formulate the best response problem of the unitary player as a dynamic programming problem, in which the policy strategies of types $t \neq k$ are fixed, and we let $\tilde{V}_k$ denote the value function for the unitary actor's problem.[12] The next

[12]In this best response problem, the unitary player decides the election outcome given (s,t,x),

lemma describes a useful necessary condition for k-responsive voting equilibria: the voter's payoff achieves the best response payoff of the unitary actor in the modified game. Part (i) of the lemma establishes that the equilibrium voting strategy solves the voter's optimal retention problem, i.e., 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. Part (ii) of the lemma states that when a congruent politician holds office, the policy choice of the incumbent 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 k-responsive voting rule, in the sense defined here.

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

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

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

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

Returning to the proof of Theorem 4.1, consider any k-responsive voting equilibrium σ. Letting ϕ^* be an optimal policy rule for the voter, an upper bound for the voter's payoff is given by the strategy profile such that each politician type chooses according to ϕ^*, and 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 chooses according to ϕ^*, while other politician types use their equilibrium policy 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. Obviously, 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

where the type t incumbent chose policy x in state s; and it chooses policy as well given (s,k), where the state is s and the incumbent is type k. The Bellman equation defining this value function is intuitively straightforward but notationally cumbersome, so we relegate it to the Appendix.

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; in particular, it is independent of the voter's discount factor and strategies used by type $t \neq k$ politicians. Taken together, these arguments deliver the result.[13]

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 4.1 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 thereafter 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, contrary to the "wait it out" strategy, the voter retains a type $t \neq k$ politician in order to avoid challenger types who are worse. The point, however, is that in all such equilibria, the voter's payoffs cannot fall below the lower bound provided by the construction.

4.2 Accountability of All Politician Types

Theorem 4.1 leverages the incentives of incumbents in k-responsive voting equilibria to deduce an accountability result for congruent politicians. It is possible that the policy choices of non-congruent politicians may also respond positively as the voter becomes patient, but they do not necessarily do so. The next example illustrates the possibility that the congruent politician type is accountable, while a second type chooses suboptimally for the voter but is re-elected in each state, and a third type shirks, choosing her myopically ideal investment level and being removed from office in all states. Here, the relatively poor policy choices of the third type introduce slack into the re-election constraint facing the second type, allowing her to distort policies in her preferred direction, while satisfying the threshold for re-election. As the

[13] Our argument uses the assumption that the challenger is congruent 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 this next best type will typically vary with the discount factor and the equilibrium.

voter becomes patient, however, this slack disappears, and the latter type is forced to choose policies arbitrarily close to voter-optimal to retain office; she is willing to make this compromise if the rewards of office are sufficiently high, giving us approximate accountability of a non-congruent type.

Example 3 (Approximate accountability). Return to the setting from Example 2, 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 investment 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 $\Gamma \approx 0$, so that optimal policy rules for all types consist of choosing their myopically ideal levels of investment: $x_t^{j*} = \hat{x}_t^j$ for all types t and all states s^j. Assuming 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 $\delta \to 1$; and in contrast, type 3 chooses her myopic ideal investment 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 investment 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 whenever he prefers her policy to x_1: for all j, $\rho(s^j, 1, x) = 1$ if and only if $x \in [x_1, 1 - x_1]$. 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 investment 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 in both states. By computation, it can be verified that the indifference condition yields

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

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

In addition to high office motivation, the key feature driving the positive result of Example 3 is that the type $1 \neq k$ politician has the option of choosing policies that are good enough for the voter that she is rewarded with re-election following such choices.

In this subsection, we analyze the welfare consequences for equilibria in which the voter is similarly responsive toward the policy choices of a subset of non-congruent politician types. We will see that when such politician types are sufficiently office motivated, they become accountable in the limit, along with the congruent type.

Definition 2. *Fix a subset $K \subseteq T$ of types with $k \in K$. The Markov electoral equilibrium σ is a* ***K-responsive voting equilibrium*** *if for each type $t \in K$, parts (i) and (ii) of Definition 1 hold for the type t politician.*

Any k-responsive equilibrium is a K-responsive equilibrium with $K = \{k\}$ and, because $k \in K$, any K-responsive equilibrium is also k-responsive. We will show that when politicians are sufficiently office motivated, equilibria satisfying the enhanced responsiveness condition exist,[14] and that the accountability of type k politicians spills over to the other politician types $t \in K$. Indeed, in such an equilibrium, a highly office motivated type $t \in K$ incumbent will always be re-elected with positive probability, i.e., she will choose policies that are acceptable to the voter. An implication of Theorem 4.1 is that as the voter becomes patient, all incumbents essentially compete against an accountable challenger, namely, the prospect that the challenger is type k. Thus, if a type $t \neq k$ politician chooses policies that are acceptable to the voter, then her policy choices must be approximately voter-optimal. Taken together, these observations imply that if office benefit is large and the voter becomes patient, then in any corresponding sequence of K-responsive voting equilibria, *all* politician types in K choose policies that are approximately optimal for the voter. This is stated in part (i) of Theorem 4.2, below.

These conclusions can be sharpened when the policy transition is independent of the state. In this case, it is impossible for a politician to manipulate the state by choice of policy and, as long as all types face a responsive voting rule (i.e., $K = T$), then competition with accountable type k politicians disciplines all other politician types: part (ii) of Theorem 4.2 establishes that if office incentives are large, then in any T-responsive voting equilibrium, all politician types choose policies that are exactly, rather than approximately, optimal. Recall from Example 2 that if the state transition depends on policy, then the exact accountability result from part (ii) does

[14] It is straightforward to show, as stated in Theorem B.2 in the Supplementary Appendix, that when $\delta\beta$ is high, there is an equilibrium in which all politician types are accountable, giving us a K-responsive equilibrium.

not hold: in the equilibrium from that example, which satisfies T-responsive voting, all types are re-elected, but only the type k politician is accountable.

Theorem 4.2. *Given any δ, assume that $\delta\beta$ is large, and fix $K \subseteq T$. Then there is a K-responsive voting equilibrium, and:*

(i) for all states s, all types $t \in K$, 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:

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

where $\{\sigma^{\delta}\}$ is any selection of K-responsive voting equilibria, and $B_{\epsilon}(\Phi^{*,\delta}(s))$ is the open ball of radius ϵ around the set $\Phi^{*,\delta}(s)$ of optimal policies,

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

The proof of Theorem 4.2 relies on the assumption that non-congruent type $t \in K$ politicians are sufficiently office motivated, so they are willing to compromise their policy choices to gain re-election, and also on the fact that in a K-responsive voting equilibrium, these politicians have 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 Theorem then follows from Theorem 4.1: because the challenger is type k with positive probability, then the expected challenger becomes close to optimal, implying that all politician types $t \in K$ must choose policies that approximate solutions to the representative programming problem.

To manipulate the state and be reelected, a politician must choose suboptimal policies that steer the state transition toward states in which the incumbent is preferable to an untried challenger. But this is impossible when the state transition is policy independent, in which case the value of retaining an incumbent is independent of policy. But then if all incumbent types are at least as good as a challenger for the voter, and a challenger is just a lottery over incumbent types, it follows that all politician types must deliver the same payoffs to the voter in all states. Therefore, because the type k is accountable, all politicians $t \neq k$ must also be accountable, as stated in part (ii) of the theorem.

5 Conclusion

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. Our results support the view that elections can be an effective mechanism for holding politicians accountable, but that conclusion is attenuated by the possibility that there may 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 congruent politicians, who may suffer from a curse of ambition. Our positive results are predicated on equilibria in which voters respond positively to sufficiently good policy choices, breaking the curse and delivering accountability of congruent politicians; in turn, this spills over to other politician types, who choose policies that are approximately voter optimal. In addition to parametric conditions on office motivation (so incumbents forego short-run gains from shirking) and voter patience (so voters are willing to reject a less than ideal incumbent), our analysis highlights the role of voters, through the conditioning of election outcomes on incumbent performance, in the well functioning of democratic elections.

A Appendix: Proofs of Results

Completing the construction of the curse of ambition example. Here, we 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 $x^1 \geqslant z^1 \geqslant x_k^{1*}$. The policy payoffs to the type k politician in this problem solve the value functions: for $j = 1, 2$,

$$\tilde{V}_k(s^j, z^1) \;=\; \max_{x^j \in [I_{j=1} z^1, 1]} -(x^j - \hat{x}_k^j)^2 + \Gamma^j + \delta \left[x^j \tilde{V}_k(s^2, z^1) + (1 - x^j)\tilde{V}_k(s^1, z^1) \right], \quad (9)$$

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 (9) for $j = 1$ has $x^1 = z^1$: interior solutions to (9) 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: $z^1 > z^{1\prime}$, then $\tilde{V}_k(s^1, z^1) < \tilde{V}_k(s^1, z^{1\prime})$. Applying the envelope theorem to (9) for $j = 2$ yields

$$\frac{\partial}{\partial z^1}\tilde{V}_k(s^2, z^1) \;=\; \frac{\delta(1 - x^2)}{1 - \delta x^2}\frac{\partial}{\partial z^1}\tilde{V}_k(s^1, z^1) \;<\; 0, \quad (10)$$

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

$$\tilde{x}\tilde{V}_k(s^2, \tilde{x}) + (1 - \tilde{x})\tilde{V}_k(s^1, \tilde{x}) \;=\; \tilde{x}V_k(s^2, 1) + (1 - \tilde{x})V_k(s^1, 1). \quad (11)$$

A final task is to verity that our construction satisfies inequalities (7). To see that it does, note that

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

$$
\begin{aligned}
&> V_k(s^1, 1)\\
&> V_k(s^2, 1),
\end{aligned}
$$

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 (6). Therefore, because $\frac{\partial}{\partial z^1}[\tilde{V}_k(s^2, z^1) - \tilde{V}_k(s^1, z^1)] > 0$ by (10), 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 (7) must hold by the indifference condition (11) defining $\tilde{x}$. □

Bellman equation for the unitary actor. The unitary actor 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 this 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 unitary actor's problem. For all s, the value function $\tilde{V}_k$ satisfies

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

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

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

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 $\pi_{t'}(\cdot|s')$. $\square$

Proof of the Unitary Actor Lemma. Consider any k-responsive equilibrium σ. We first prove part (ii). Clearly, we have $\tilde{V}_k(s,k) \geqslant V_k(s,k)$, 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

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

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

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

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

$$\hat{v}_k(s) \quad = \quad u_k(s,\hat{x}(s)) + \delta V_k^I(s,k,\hat{x}(s)) \tag{12}$$

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

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

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 $\hat{V}_k(s) \geqslant V_k(s,k)$ 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

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

by (4) and the definition of $\hat{v}_k$. In particular, $\hat{v}_k(s) \geqslant V_k(s,k)$.

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

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

where the first equality uses $\rho(s,k,\hat{x}(s)) = 1$, the inequality follows from the above argument, the second equality follows from (2), and the last equality uses (12). We conclude that $\hat{V}_k^2(s) \geqslant \hat{v}_k(s) \geqslant V_k(s,k)$.

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 $m \to \infty$. This sequence of policy payoffs satisfies $\hat{V}_k^m(s) \geqslant \hat{v}_k(s) \geqslant V_k(s,k)$ 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 $\hat{V}_k^m(s) \to \hat{V}_k(s)$, and since $\hat{V}_k^m(s) \geqslant \hat{v}_k(s) \geqslant V_k(s,k)$ holds in each state s for every m, we conclude that

$$\hat{V}_k(s) \geqslant V_k(s,k) \qquad (13)$$

for each state s. This completes the first step.

The second step is to show that, in fact, (13) 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 $\overline{\pi}_k$ cannot exceed the politician's equilibrium payoff, it follows that for each s,

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

This implies that $V_k(s,k) \geqslant \hat{V}_k(s)$ for each s, and by (13), we conclude that

$$\hat{V}_k(s) \;\;=\;\; \hat{v}_k(s) \;\;=\;\; V_k(s,k) \tag{14}$$

for each state s, completing the second step.

Returning to part (ii) of the lemma, it follows from (14) that for each s and each x, we have

$$V_k^I(s,k,x) \;\;=\;\; \sum_{s'} p(s'|s,x)V_k(s',k) \;\;=\;\; \sum_{s'} p(s'|s,x)\hat{V}_k(s). \tag{15}$$

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

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

where the first equality follows from (15), 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 (15). 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 $V_k(s,k) \geqslant \tilde{V}_k(s,k)$ 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

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

where we use part (ii) to substitute $V_k(s',k)$ for $\tilde{V}_k(s',k)$. Note that by (4), we have

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

and thus by (2), 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 (3), we have

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

Thus, the recursion reduces to

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

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

Proof of Theorem 4.1. Theorem B.1, in the Supplementary Appendix, shows that there is an equilibrium in which the type k politician is accountable, which implies that k-responsive voting equilibria also exist. Here, we establish the uniform bounded stated in Theorem 4.1. Fix δ, let $\sigma^\delta = (\pi^\delta, \rho^\delta)$ be a k-responsive equilibrium given δ, and let $\phi^*: S \to X$ 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 mth 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 $\tilde{P}^m(\cdot|s,t) \to \tilde{p}$. 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

$$\hat{V}_k(s,t) \;=\; \sum_{m=1}^{\infty} \delta^{m-1} \sum_{(s',t')} \left[\int_x u_k(s',x)\hat{\pi}_{t'}(dx|s') \right] \hat{P}^m((s',t')|(s,t)), \qquad (16)$$

and as mentioned, the Unitary Actor Lemma implies that $V_k^\delta(s,t) \geqslant \hat{V}_k(s,t)$ 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, $\hat{P}^m(\cdot|s,t) \to \hat{p}$. For later use, let

$$\hat{V} \;=\; \frac{1}{1-\delta} \sum_{(s',t')} \left[\int_x u_k(s',x)\hat{\pi}_{t'}(dx|s') \right] \hat{p}((s',t')) \qquad (17)$$

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 $m \to \infty$, 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

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

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

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

for all m, and likewise

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

for all m.

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

$$|\hat{V}_k(s,t) - \hat{V}| \;\leqslant\; \hat{M}. \tag{18}$$

Indeed,

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

where the first equality follows from (16) and (17), the first inequality from our bound on stage utilities, the second inequality from geometric convergence, and the last inequality from $\alpha\gamma \leqslant 1$. Setting $\hat{M} = \frac{\overline{u}|T|}{\alpha\gamma}$, 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), we have

$$|\tilde{V}_k(s,t) - \tilde{V}| \;\; \leqslant \;\; \tilde{M}. \tag{19}$$

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

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

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 (18) and (19). Finally, note that for all s and all t, we have the inequalities

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

and we conclude that for all s and all t,

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

as required. □

Proof of Theorem 4.2. Theorem B.2, in the Supplementary Appendix, shows that when $\delta\beta$ is large, there is an equilibrium in which all politician types are accountable, implying that K-responsive voting equilibria exist. Here, we focus on parts (i) and (ii) of Theorem 4.2. 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 K-responsive voting equilibrium. Then, for each states s and each type $t \in K$, we have $\rho(s,t,x) > 0$ for almost all $x \in \text{supp}(\pi_t^{\delta}(\cdot|s))$. In turn, this implies that

$$\text{for all } s, \text{ all } t \in K, \text{ and } \pi_t(\cdot|s)\text{-almost all } x,\; V_k^I(s,t,x) \geqslant V_k^C(s,t,x). \tag{20}$$

To see this, fix state s and type $t \in K$, 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 K-responsive voting equilibrium. Recalling the normalization $\overline{u} = 1$ 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.

To prove part (i) of the corollary, we need to translate the conclusion of Theorem 4.1, which is stated in terms of voter payoffs, into its implications for the optimality of policy choices themselves. Namely, we consider the distribution of policies in an equilibrium with K-responsive voting (and hence also 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. In turn, because every type $t \in K$ chooses policies that lead to reelection when they value office, each of their policy choices must also converge to optimal policies.

More 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 $\mu^m_{s,t}$ 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

$$\mu^{\delta}_{s,t} \quad = \quad (1-\delta) \sum_{m=1}^{\infty} \delta^{m-1} \mu^m_{s,t}, \tag{21}$$

which depends on the initial state and politician type, and where we highlight the dependence on δ. Our summary statistic for the equilibrium policies in any state s is then the conditional $\mu^{\delta}_{s,t}(\cdot|s)$ of the aggregate measure $\mu^{\delta}_{s,t}$, which is well-defined since $\mu^{\delta}_{s,t}(X \times \{s\}) > 0$. Thus, for a Borel subset $A \subseteq X$ of policies, $\mu^{\delta}_{s,t}(A|s)$ 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.

Returning to the proof of part (i) of the corollary, fix state s, type $t \in K$, and discount factor δ. Analogously to $\mu^{\delta}_{s,t}$, let $\tilde{\mu}^{\delta}_{s,t}$ denote the discounted and time-aggregated marginal distribution over state-policy pairs (s', x') generated by the equilibrium policy distribution π^{δ}_{t} of type t incumbents starting from state s, *conditional* on this incumbent staying in office in all future periods. 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 $V^{I}_{k}(s', t, x) \geqslant V^{C}_{k}(s', t, x)$. 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 $(\pi^{\delta}, \tilde{\rho}^{\delta})$ as under $(\pi^{\delta}, \rho^{\delta})$. Notice also that by property (20), the voter retains the type t politician with probability one in all states under strategy $\tilde{\rho}^{\delta}$. 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

$$\begin{aligned}(1-\delta)V^{\delta}_{k}(s,t) &= (1-\delta)\sum_{m=1}^{\infty}\delta^{m-1}\int_{(s',x')} u_k(s',x')\mu^{m,\delta}_{s,t}(d(s',x')) \\ &= \int_{(s',x')} u_k(s',x')\mu^{\delta}_{s,t}(d(s',x')) \\ &= \int_{(s',x')} u_k(s',x')\tilde{\mu}^{\delta}_{s,t}(d(s',x')). \qquad (22)\end{aligned}$$

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

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

Furthermore, because $\tilde{\mu}^{\delta}_{s,t}$ belongs to the set $\Delta(Y)$, which is compact with the weak* topology, we can go to a further subsequence (if needed) such that $\tilde{\mu}^{\delta}_{s,t}$ converges weak* to a limit $\tilde{\mu}$. Using Theorem 4.1 and (22), we have

$$\overline{v}_s = \lim_{\delta\to 1}(1-\delta)V^{*,\delta}_{k}(s) = \lim_{\delta\to 1}(1-\delta)V^{\delta}_{k}(s,t) = \int_{(s',x')} u_k(s',x')\tilde{\mu}(d(s',x')) \qquad (23)$$

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

Now, suppose toward a contradiction that result of part (i) of Corollary 4.2 does not hold, so that for some $\epsilon > 0$, we have

$$\liminf_{\delta \to 1} \pi_t^\delta(B_\epsilon(\overline{\Phi}^*(s))|s) \quad < \quad 1.$$

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

$$\liminf_{\delta \to 1} \tilde{\mu}_{s,t}^\delta(B_\epsilon(\overline{\Phi}^*(s))|s) \quad < \quad 1. \tag{24}$$

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

$$\tilde{\mu}(X(s) \backslash B_\epsilon(\overline{\Phi}^*(s))|s) \quad > \quad 0,$$

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

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

Since s has positive marginal probability under $\tilde{\mu}$, this implies

$$\int_{(s',x')} u_k(s',x')\tilde{\mu}(d(s',x')) \quad < \quad \max_{x' \in X(s)} \sum_{s'} p(s'|s,x')\overline{v}_{s'}. \tag{25}$$

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

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

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

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

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

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

where the first equality follows from (23). This contradicts (25), however, and we conclude that $\tilde{\mu}_{s,t}^\delta(B_\epsilon(\overline{\Phi}^*(s))|s) \to 1$, as required.

To prove part (ii) of the corollary, first note that when the state transition is policy independent, 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 (20) implies

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

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

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

That is, all politician types achieve the optimal value. □

B Supplementary Appendix: Existence of Equilibria with Accountable Politicians

We first prove the general existence of an equilibrium in which the type k politician is accountable.

Theorem B.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 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. 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 it 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 $V_k^I(s,k,x) \geqslant V_k^C(s,t,x)$. This preserves equilibrium conditions of $\tilde{\sigma}$, and we conclude that σ is an equilibrium in which the type k politician is accountable.

Proof of Theorem B.1. To 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),

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

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:

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

where we recall that stage utility u_k is bounded between $\underline{u} = 0$ and $\overline{u} = 1$. 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

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

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 $x \in X(s)$,

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

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

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

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. □

The Markov electoral equilibrium established in Theorem B.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.

Theorem B.2. *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 B.1, the proof of Theorem B.2 proceeds by a straightforward equilibrium construction.

Proof of Theorem B.2. 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

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

To verify that policy 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 $\frac{\overline{u}}{1-\delta} + \beta$. Therefore, using the normalization $\overline{u} = 1$ and $\underline{u} = 0$, a sufficient condition for office holder t to follow the policy rule ϕ in all states s is $\frac{\delta\beta}{1-\delta} \geqslant 1$, as required. □

References

Acemoglu, D., G. Egorov, and K. Sonin (2103). A political theory of populism. *Quarterly Journal of Economics 128*, 771–805.

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.

Aghion, P. and P. Bolton (1990). Government domestic debt and the risk of default: A political-economic model of the strategic role of debt. In R. Dornbusch and M. Draghi (Eds.), *Public Debt Management: Theory and History*. Cambridge University Press.

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. 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.

Barro, R. (1973). The control of politicians: An economic model. *Public Choice 14*, 19–42.

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.

Besley, T. and M. Smart (2007). Fiscal restraint and voter welfare. *Journal of Public Economics 91*, 755–773.

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. (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 (2020). Electoral accountability and responsive democracy. *The Economic Journal 130*, 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, J. (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.* 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.

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.

Meirowitz, A. (2007). Probabilistic voting and accountability in elections with uncertain policy constraints. *Journal of Public Economic Theory 9*, 41–68.

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 (2000). *Political Economics: Explaining Economic Policy*. Cambridge, MA: MIT Press.

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.