ORIGINAL PAPER

Representative Voting Games

John Duggan[1] · Jean Guillaume Forand[2]

Abstract

We propose the stationary Markov perfect equilibria of *representative voting games* as a benchmark to evaluate the outcomes of dynamic elections, in which the evolution of voters' political power is endogenous. We show that the equilibria of dynamic elections can achieve this benchmark if politicians are sufficiently office motivated. For arbitrary equilibria of the electoral model, we characterize the faithfulness of politicians' choices to the policy objectives of representative voters through a *delegated best-response property*. Finally, we provide conditions under which general dynamic electoral environments admit representative voters in each state.

1 Introduction

To analyze the effect of delegation in political systems, it is important to understand the outcomes that would obtain in an idealized environment in which voters retain policy-making power. These outcomes may have interest as a normative benchmark, and to the extent that they match equilibria of the political system, they can serve as an analytical shortcut. In static models with single-peaked preferences, the usual benchmark is the preferred policy of the median voter. As is well known, the median ideal policy is a Condorcet winner, distinguishing its normative status, and it is the unique equilibrium outcome of Downsian elections when candidates can commit to policy platforms before an election. This confluence holds when candidates have a range of objectives from pure office seeking to pure policy motivation, and thus, under broad conditions, Downsian competition among two candidates is consistent with the idealized benchmark. In this paper, we examine a dynamic analogue of the

✉ Jean Guillaume Forand
   jgforand@uwaterloo.ca

   John Duggan
   dugg@ur.rochester.edu

[1] University of Rochester, Rochester, USA

[2] University of Waterloo, Waterloo, Canada

static model, in which a state variable follows a controlled Markov process, and the identity of a representative voter can depend on the state and evolves stochastically over time. Within this framework, we compare the direct choices of voters to the policy outcomes of a dynamic electoral model, in which voters delegate policy-making power to political representatives, whose choices are a product of ideological and office-holding incentives. Specifically, we study the conditions under which the policy choices of politicians, who are held accountable to different voters over time through elections, conform to the idealized benchmark.

Given this dynamic environment, we take as our benchmark the *representative voting game*, in which successive representative voters exercise their political power not through elections, but by implementing policies directly. In this game, each period starts with a state; the representative voter in that state chooses a feasible policy; and a new state is realized, determining a new representative voter and a new set of feasible policies, and so on. Representative voting games are stochastic games with a finite number of players (the possible voter types), discrete states, compact action spaces (equal to the policy space), and continuous stage utilities and transition probabilities. Analogous to the median voter's ideal policy, which would be this voter's optimal choice in a static setting, our point of comparison is the state-contingent distribution over policies generated by the stationary Markov perfect equilibria of the representative voting game.

To this environment, we append a model of delegated policy-making through dynamic elections, in which potential politicians have policy preferences corresponding to voter types but also value holding office per se. As in the idealized representative voting game, each period starts with a state, but now we assume an incumbent politician holds office and (rather than the voter) chooses a feasible policy for that period; a challenger type is then drawn according to a state- and policy-dependent transition probability; the representative voter reelects the incumbent or opts for the challenger; a new state is drawn, and the process repeats. We focus on the *Markov electoral equilibria* of the electoral model, in which representative voters only control electoral outcomes in their own states, and their choices anticipate both the future electoral decisions of representative voters and the policy choices of politicians.

Preview of results Our first goal is to investigate the possibility that policy choices made by politicians in the dynamic electoral model correspond to the choices made directly by the voters in the benchmark representative voting game. In Theorem 1, we show that if politicians are sufficiently office motivated, then every stationary Markov perfect equilibrium of a representative voting game can be replicated by a Markov electoral equilibrium of the associated electoral game, in the sense that every type of politician chooses policy according to the representative voter's equilibrium strategies in each state. Interestingly, the result holds even for mixed strategy equilibria of the representative voting game, despite the fact that a politician may not be indifferent over the policies in the support of the representative voter's mixed strategy. We address this by adjusting the probability of reelection and using the promise of future office benefits to equalize the politician's payoff across policies in the support of the voter's mixed strategy. By establishing the possibility that the representative voters' control over elections extends to policy choices, this result gives

conditions under which delegation entails no loss of control by voters, and it provides foundations for using the representative voting game to study dynamic elections. In turn, this is of potential use in applications, as stationary Markov perfect equilibria of the representative voting game can be characterized with less difficulty than Markov electoral equilibria in many environments of interest.

While the equilibria of representative voting games can be supported by policy outcomes of elections under broad conditions, it may be that electoral incentives create multiple Markov electoral equilibria, including some in which politicians' policy choices bear little relation to voters' choices in the benchmark. Correspondingly, our second goal is to identify a class of Markov electoral equilibria satisfying a *delegated best-response property:* in every state, all politician types choose policies that the representative voter in that state would choose in their place, given the expected future choices of politicians. In contrast to the representative voting game, the delegated best-response property is defined in terms of optimal policies for representative voters in the electoral model, rather than the idealized benchmark, taking as given future choices of elected politicians. In Theorem 2, we establish that a Markov electoral equilibrium satisfies the delegated best-response property, if it is *convergent*, in the sense that all politician types use the same policy strategy, as well as *reelection-balanced*, in the sense that voters in all states coordinate on reelection standards that determine the same ex ante probability of reelection for all politician types. Moreover, we show by example that imbalances in a politician's electoral prospects across future states can weaken her incentives to choose policies that are optimal for the current representative voter, so that the delegated best-response property can fail in the absence of reelection-balancedness.

Theorem 2 relies on the fact that our electoral model allows for a weak form of commitment by office holders: if the politician chooses x in state s, she can also commit to choose x again if the state remains s. This type of ex post commitment, which can capture the presence of politician-specific transition costs or institutional stickiness, is weaker than the commitment assumed in Downsian models: first, it requires the politician to first make the choice of x, "putting her money where her mouth is," rather than assuming binding commitments to ex ante promises; and second, we do not assume that a politician is necessarily committed in this sense, only that she has the option to generate policy inertia across periods in which the state remains the same. In contrast to the static Downsian model, commitment plays no role in our Theorem 1: in our dynamic setting, future office benefits are sufficient to provide incentives for politicians to implement voters' optimal policies when these policies are expected of all politicians in equilibrium. However, to rule out Markov electoral equilibria in which some politicians choose policies that are not best responses for some representative voters in some states, politicians must have the means to provide voters with incentives to reelect them if instead they make optimal choices. In our model, politicians can do this by committing to some persistence in their policy choices over time.

Under the delegated best-response property, the preferences of state-contingent representative voters can provide a convenient tool for describing equilibrium behavior in elections. Our third and final goal is to address some remaining foundational issues: namely, when do Markov electoral equilibria admit representative

voters in all states, and when they do, can the identity of the representative voter in some state be easily recovered from the primitives of the underlying model of political institutions? In Theorem 3, we provide sufficient conditions for the existence of a representative voter in each state. Because elections involve a choice between distributions over streams of policies across time, the usual single-crossing condition is not adequate for this purpose, but roughly, it is enough that voters discount the future at a common rate, and that utility differences are affine linear in a parameter that varies across voters. The latter is satisfied if, for example, the policy space is one dimensional and policy utility is quadratic, with the state entering as a shift parameter on citizen ideal points. Furthermore, given any state, the representative voter in the dynamic game is the voter type that is decisive in the stage game determined by that state.

Literature review In the standard static model of collective decision-making, an odd number of voters have single-peaked preferences over a one-dimensional policy space. As a Condorcet winner, the median voter's ideal policy has both positive and normative appeal as a benchmark outcome. The question of the existence of electoral institutions generating policies in line with this benchmark was first addressed by Downs (1957): if two office-motivated candidates simultaneously commit to platforms, then this game has a unique equilibrium in which each candidate promises, and if elected implements, the ideal policy of the median voter. This result is robust in some respects, but not in others: for example, it persists if candidates are policy motivated or have mixed motivations (Calvert 1985); but equilibria with non-median policies cannot be ruled out if politicians cannot commit to policies, as in citizen-candidate models (Osborne and Slivinski 1996; Besley and Coate 1997). In our dynamic model, the absence of binding campaign promises can also undercut the possibility of a tight linkage of representative voters' preferences and office holders' policy choices. This discrepancy is rectified by Banks and Duggan (2008), who establish, in a model closely related to the single-state version of our model,[1] that when players are sufficiently patient, or when office benefits are sufficiently high, the policies chosen by office holders of all types will converge to the ideal point of the median type.[2]

In the single-state model, the representative voting game benchmark remains simple: it calls for the median voter's ideal policy to be implemented in every period. If the state evolves endogenously through policy choices, however, then another challenge to a dynamic median voter result is that, even if a single voter type is representative across all states, the representative voter need not have a fixed ideal point; rather, the optimal policy choices of the voter will be state-dependent and should be obtained as the solution to a hypothetical dynamic programming problem in which this voter can choose policies directly. In Duggan and Forand (2019), we consider

---

[1] With a single state, our model mimics the equilibrium outcomes of dynamic model of elections with adverse selection following Duggan (2000) and Bernhardt et al. (2004): while we assume that a politician's type is observed after she takes office to avoid complex updating of voters' beliefs across states, the ex post commitment that we allow proxies for these beliefs in the absence of a state variable.

[2] Forand (2014) and Van Weelden (2013) establish related results in different models of dynamic elections with a fixed median voter and underlying collective decision problem.

the special case of our model with a single representative voter type and study the relationship between this voter's dynamic programming problem and the set of Markov electoral equilibria. An important insight of that paper is that the scope for politicians to manipulate the state is a powerful source of equilibrium multiplicity, so that stringent conditions are required to rule out dynamic political failures. In fact, Duggan and Forand (2019) show it is possible that politicians who share the representative voter's policy preferences implement policy plans that are suboptimal for the voter, even if they are highly office motivated.

However, the model with a single representative voter, which Bai and Lagunoff (2011) refer to as the "permanent authority" benchmark, is not appropriate when the influence of voters varies over time, so that the identity of the representative voter depends on the state. Then the hypothetical scenario is not as simple as solving a dynamic programming problem, motivating our focus on representative voting games. In fact, the equilibrium outcomes of these games mimic those of the model in Bai and Lagunoff (2011). In contrast, we use these outcomes in our setting both as a tool to characterize the set of electoral equilibria and as a benchmark against which the equilibrium policy choices of politicians can be compared. To the possibility of delegation failure between a fixed representative voter and politicians, representative voting games add the potential for coordination failure between the various representative voters. Moreover, with a single representative voter, there is no loss from restricting attention to solutions of this voter's dynamic programming problem in which he uses pure strategies. In contrast, an additional complication in the model with multiple representative voters is that the possibility of mixed strategy equilibria in representative voting games cannot be sidestepped: as we detail below, mixing can introduce a wedge between equilibrium outcomes in the benchmark and those that can be supported by equilibria of dynamic elections.

2 Representative voters and dynamic elections

Representative voting games A *representative voting game* is described by an octuple $\mathcal{R} = (S, T, \kappa(\cdot), Y, Y(\cdot), p(\cdot), (u_t)_{t \in T}, \delta)$ such that S is a countable set of states; T is a finite set of voter types; $\kappa : S \to T$ is a mapping such that $\kappa(s)$ is the representative voter type in state s; Y is a metric space of policies and $Y(s) \subseteq Y$ is a nonempty, compact subset of feasible policies in state s; $p : S \times Y \times S \to [0, 1]$ is a state transition function such that $p(s'|s, y)$ is the probability of s' given policy choice y in state s; each $u_t : S \times Y \to \mathfrak{R}$ is a bounded, continuous stage utility function; and $\delta \in [0, 1)$ is voters' common discount factor.[3] We make the additional assumption that all states have a positive probability of recurring following all policy choices: $p(s|s, y) > 0$ for all s and y. The importance of this assumption will become clear when we introduce the electoral model below, and for now we only note that our

---

[3] We assume a type-independent discount factor to simplify notation. As can be verified from their proofs, this assumption is not needed for Theorems 1 or 2. On the other hand, the result of Theorem 3 does not hold when different voter types have different discount factors.

results go through even if these transition probabilities are arbitrarily small. Policies are chosen in an infinite sequence of periods. In each period, a state s is given, and representative voter $\kappa(s)$ chooses any policy $y \in Y(s)$, utilities $u_t(s, y)$ accrue to each voter type t, and next period's state s' is drawn from $p(\cdot|s, y)$. Given a stream $(s_1, x_1, s_2, x_2, \ldots)$ of state-policy pairs, the discounted payoff of a type t voter is

$$\sum_{\ell=1}^{\infty} \delta^{\ell-1} u_t\big(s_\ell, x_\ell\big),$$

and payoffs extend to probability distributions over such streams via expected utility.

In the representative voting game, there are no elections and voters govern directly. Correspondingly, a stationary strategy for a type t voter is a mapping $\tilde{\pi}_t : \kappa^{-1}(t) \to \Delta(Y)$, where $\kappa^{-1}(t)$ is the set of states in which t is the representative voter type and $\Delta(Y)$ is the set of Borel probability measures on Y. Let $\tilde{\pi}_{\kappa(s)}(\cdot|s)$ represent the mixture over policies used by the representative voter $\kappa(s)$ in state s, and let $\tilde{\pi} = \big(\tilde{\pi}_t\big)_t$ denote a profile of such strategies. Because the representative voting game is a well-behaved stochastic game with a finite set of players and a countable set of states, the existence of stationary Markov perfect equilibria is known from Federgruen (1978). Note that, as a standard stochastic game, the representative voting game will not admit a unique equilibrium in general.

Dynamic elections A *dynamic election* is a triple $\mathscr{E} = (\mathscr{R}, q(\cdot), b)$ such that $\mathscr{R}$ is a representative voting game; $q : S \times Y \times T \to [0, 1]$ is a continuous challenger transition probability; and b is politicians' common office benefit. Here, an infinite pool of politicians is partitioned into voter types, each period begins with some state s and an incumbent politician of some type t, and the incumbent chooses any feasible policy $y \in Y(s)$ and whether to run for reelection. A challenger is then drawn from the challenger transition, so that she is type t' with probability $q(t'|s, y)$, and the representative voter $\kappa(s)$ decides between the incumbent and challenger. Politician types are initially private information, but the type of the winning politician (i.e., the incumbent) is publicly observed, so that elections pit a known incumbent against a potentially unknown challenger. In the next period, after the election, a new state s' is realized, and the process is repeated. That is, we effectively superimpose on the representative voting game $\mathscr{R}$ an electoral system in which policies are chosen by political agents, who intervene between the representative voter in any given state and the choice of policy in that state; and instead of choosing policy directly, the representative voter in state s chooses the winner of elections in that state.

In addition to choosing policy, the office holder chooses whether to run for reelection; we model this by using Y to represent choices of policy and the decision to run for reelection, and using a copy of Y, denoted Y^d, to represent policy choices and the decision to drop out of politics. We maintain the convention that $Y \cap Y^d = \emptyset$; we assume a mapping $\xi : Y \cup Y^d \to Y^d$ so that for all $y \in Y$, $\xi(y) = z$ is the element of Y^d corresponding to y, and for all $z \in Y^d$, $\xi(z) = z$; and we let $Y^d(s) = \xi(Y(s))$ be the feasible policy choices for an office holder who chooses not to seek reelection in state s. We assume that the challenger and state transitions are independent of the incumbent's decision to run, i.e., $q(t'|s, x) = q(t'|s, \xi(x))$ and $p(s'|s, x) = p(s'|s, \xi(x))$ for all $x \in Y$.

Commitment power We assume that office holders who run for reelection have the option to bind themselves to policies through a weak form of commitment. Specifically, we assume that if a type t office holder implements policy x in a state s, then she can choose to commit herself to implementing policy x again if she is subsequently reelected and the next state remains s (i.e., $s' = s$). We assume that an office holder's decision to commit is public, in that the representative voter in state s observes whether the politician is *bound* to x or *free* to choose any feasible policy before making his reelection decision. The politician's commitment to x is broken when the state transitions away from s (i.e., $s' \neq s$).

This commitment differs from the usual assumption in the Downsian model, where both candidates can commit to arbitrary platforms *before* an election; here, in contrast, it is the incumbent who may be committed to a policy that has actually been implemented in a state *after* an election. This ex post form of commitment is more consonant with the citizen-candidate approach to elections, and as we also highlight in Duggan and Forand (2019), it plays a useful role in aligning the outcomes of dynamic elections with those preferred by representative voters. As noted in the Introduction, commitment power plays a role in some of our results but not others, and we will discuss this further in the text.

Analogous to our treatment of office holders' decision to drop out, we model the choice of committing to policies by making a further copy of Y, denoted Y^c, where policy choices in Y^c involve commitments, while policy choices in Y do not (this is consistent with the absence of commitment in the representative voting game). Formally, we assume a mapping $\varphi : Y \cup Y^c \rightarrow Y^c$ such that for every policy choice $y \in Y$ that is free of commitment, $\varphi(y)$ denotes same choice of policy along with commitment, and for all $y \in Y^c$, $\varphi(y) = y$; and we let $Y^c(s) = \varphi(Y(s))$ be the feasible policy choices with commitment. Let $X = Y \cup Y^c \cup Y^d$ represent the space of simultaneous policy, commitment, and campaign decisions, and let $x \in X$ denote a generic choice for the incumbent. Finally, challenger and state transitions are independent of incumbents' policy commitments, i.e., $q(t'|s, x) = q(t'|s, \varphi(x))$ and $p(s'|s, x) = p(s'|s, \varphi(x))$ for all $x \in Y$. Note that our assumption that $p(s|s, x) > 0$ for all states s and policies x implies that incumbents' option to commit to policies is always meaningful.

To fix ideas, we present a simple application of our model, to which we return throughout the text to illustrate our main results.

Example (Dynamic deficit reduction) We first specify the representative voting game. Let the state space be $S = \{\bar{s}, \underline{s}, s_0\}$, where we interpret both $\bar{s}$ and $\underline{s}$ as states in which the government is in a poor fiscal position, say due to high accumulated debt or unfunded future liabilities. Furthermore, suppose that the economy is strong in state $\bar{s}$ but weak in state $\underline{s}$. Feasible policies in both these states are $Y(\bar{s}) = Y(\underline{s}) = \{\bar{x}, \underline{x}\}$, where we interpret $\bar{x}$ as high government spending and $\underline{x}$ as implementation of austerity measures. For its part, s_0 captures the state in which the government's fiscal problems have been rectified, and for simplicity we model this as an absorbing state with a single policy: $Y(s_0) = \{x_0\}$ and $p(s_0|s_0, x_0) = 1$. In a high debt state, choosing policy $\bar{x}$ ensures that the government's fiscal problems persist, although the strength of the economy may vary: specifically, transition

probabilities are such that for all $s, s' \in \{\bar{s}, \underline{s}\}$ with $s \neq s'$, we have $p(s|s, \bar{x}) = p > 0$ and $p(s'|s, \bar{x}) = 1 - p$. On the other hand, choosing policy $\underline{x}$ in such a state can resolve the government's fiscal problems with positive probability: for simplicity, we specify transition probabilities such that $p(s|s, \underline{x}) = p > 0$ and $p(s_0|s, \underline{x}) = 1 - p$. Let the type space be $T = \{h, d\}$, where type h is a fiscal hawk and type d is a fiscal dove. Assume that hawkish voters are representative when the economy is weak (i.e., $\kappa(\underline{s}) = h$), and that dovish voters are representative when the economy is strong or fiscal problems have been resolved (i.e., $\kappa(\bar{s}) = \kappa(s_0) = d$). Suppose that doves prefer big government and do not care about the economy:

$$\hat{u} = u_d(\bar{s}, \bar{x}) = u_d(\underline{s}, \bar{x}) > \tilde{u} = u_d(s_0, x_0) > \check{u} = u_d(\bar{s}, \underline{x}) = u_d(\underline{s}, \underline{x}).$$

Suppose that hawks agree with doves that austerity measures should not be imposed when the economy is weak, but that they think spending should be reduced when the economy is strong:

$$\hat{u} = u_h(\bar{s}, \underline{x}) > \tilde{u} = u_d(\bar{s}, \bar{x}) = u_h(\underline{s}, \bar{x}) = u_h(s_0, x_0) > \check{u} = u_h(\underline{s}, \underline{x}).$$

This representative voting game has a unique stationary Markov perfect equilibrium, which is such that $\tilde{\pi}_d(\bar{x}|\bar{s}) = \tilde{\pi}_h(\bar{x}|\underline{s}) = 1$. In equilibrium, no representative voter implements austerity measures, and the state never transitions to s_0 from another state. Hawks do not want to fight deficits when they have political power, because in that case the economy is weak. They would want fiscal problems to be addressed when the economy is strong, but in that state doves control policy and choose to continue running deficits. We can append a dynamic election to the representative voting game above by specifying office benefit $b \geqslant 0$ as well as state- and policy-independent challenger transition probabilities such that $q(h) = q(d) = \frac{1}{2}$. $\qquad\square$

3 Markov electoral equilibria

Strategies A stationary *Markov policy strategy* for a type t politician is a mapping $\pi_t : S \to \Delta(X)$, where $\pi_t(\cdot|s)$ represents the mixture over policies used by the type t politician when free in state s.

Let $\pi = (\pi_t)_t$ denote a profile of such strategies. A *Markov voting strategy* is a Borel measurable mapping $\rho : S \times T \times X \to [0, 1]$, where $\rho(s, t, x)$ represents the probability that the representative voter in state s reelects a type t office holder following a free policy choice of x in state s. The precise form of mixed voting we use is such that mixing occurs when the incumbent is free and chooses policy x in state s; if the incumbent currently bound to x in state s (and thus reelected in the previous period after choosing x in state s), then the representative voter $\kappa(s)$ reelects the incumbent with probability one. This focus is not a constraint imposed on the voter; rather, by stationarity of the voter's decision problem, it remains optimal to reelect the incumbent again when the politician is bound to a policy that was previously sufficient for reelection. We refer $\sigma = (\pi, \rho)$ as a *Markov electoral strategy profile*.

Continuation values Given a Markov electoral strategy profile σ, we can define continuation values for a type t citizen. If $x \in Y(s) \cup Y^c(s)$, then the discounted expected policy utility of the citizen from electing a type t' incumbent who chooses policy x in state s satisfies:

$$V_t^I(s,t',x) = p(s|s,x)\left[\mathbb{1}_{x\in Y^c}\left[u_t(s,x) + \delta V_t^I(s,t',x)\right] + \mathbb{1}_{x\in Y}V_t^F(s,t')\right]$$
$$+ \sum_{s'\neq s} p(s'|s,x)V_t^F(s',t'),$$

where $V_t^F(s,t')$ is the expected discounted utility to the citizen from a type t' office holder who is free in state s, calculated before a policy is chosen. In words, if the incumbent is reelected, then with probability $p(s|s,x)$, the state remains s, and in this case, either the politician has committed to choose x again and will be reelected; or the incumbent has opted to be free in s. In all other states $s' \neq s$, the incumbent is free. When an office holder chooses $x \in Y^d(s)$ and thus not to stand for reelection, we have $V_t^I(s,t',x) = V_t^C(s,x)$, where $V_t^C(s,x)$ is the expected discounted utility of electing a challenger following the choice of x in state s and is defined by

$$V_t^C(s,x) = \sum_{t'} q(t'|s,x) \sum_{s'} p(s'|s,x)V_t^F(s',t').$$

That is, when a challenger is elected, the new office holder is free for every realization of next period's state. Finally, $V_k^F(s,t')$ is given by

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

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

In addition to payoffs from policies, a type t office holder must evaluate future expected discounted office benefit from choosing policy x in state s, conditional on being reelected, defined as follows: for all $x \in Y(s) \cup Y^c(s)$,

$$B_t(s,x) = p(s|s,x)\left[\mathbb{1}_{x\in Y^c}\left[b + \delta B_t(s,x)\right] + \mathbb{1}_{x\in Y}B_t^F(s)\right]$$
$$+ \sum_{s'\neq s} p(s'|s,x)B_t^F(s'),$$

where the expected discounted office benefit for a type t office holder who is free in state s is

$$B_t^F(s) = \int_{x'} \left[b + \delta\rho(s,t,x')B_t(s,x')\right]\pi_t(dx'|s),$$

reflecting the fact that the office holder receives b in the current period and, conditional on choosing policy x' and being reelected, receives $B_t(s, x')$ in the future. For all $x \in Y^d(s)$, set $B_t(s, x) = 0$.

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

$$
\begin{aligned}
P_{\kappa(s)}(s, t) &= \{x \in Y(s) \cup Y^c(s) : V^I_{\kappa(s)}(s, t, x) > V^C_{\kappa(s)}(s, x)\} \\
R_{\kappa(s)}(s, t) &= \{x \in Y(s) \cup Y^c(s) : V^I_{\kappa(s)}(s, t, x) \geqslant V^C_{\kappa(s)}(s, x)\}
\end{aligned}
\tag{1}
$$

of policies that yield the type $\kappa(s)$ voter an expected discounted utility strictly and weakly greater, respectively, than the expected discounted utility of a challenger. We refer to these as the *strict* and *weak reelection sets,* respectively.

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

$$
\max_{x \in X(s)} u_t(s, x) + b + \delta \left[\rho(s, x, t) \left[V^I_t(s, t, x) + B_t(s, x) \right] + (1 - \rho(s, x, t)) V^C_t(s, x) \right],
$$

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

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

where $\rho(s, t, x)$ is unrestricted if $x \in R_{\kappa(s)}(s, t) \setminus P_{\kappa(s)}(s, t)$. Intuitively, a type t office holder maximizes current period utility plus future expected discounted payoff, which combines policy utility and office benefit (in case the politician is reelected) and the continuation value of a challenger (in case the politician loses). Duggan and Forand (2018) establish existence of Markov electoral equilibria in a more general framework that does not assume the existence of representative voters and that allows general politician payoffs.

Special classes of equilibria Our goal in this paper is to relate the policy outcomes of representative voting games, which are generated by a stationary Markov perfect equilibrium $\tilde{\pi}$, to those of dynamic elections, which are generated by a Markov electoral equilibrium (π, ρ). To do this, we focus on restricted classes of Markov electoral equilibria. First, because policy in a given state is set by a single voter in the representative voting game but by many potential office holders in the electoral game, dynamic elections cannot mimic the choices of representative voters if different politician types implement different policies. Therefore, we say that a Markov electoral equilibrium $\sigma = (\pi, \rho)$ is *convergent* if $\pi_t(\cdot|s) = \pi_{t'}(\cdot|s)$ for all states s and types t and t'. In convergent Markov electoral equilibria, representative voters in all states are indifferent between all types

of office holders. In this case, voters may nevertheless apply different reelection standards to different incumbent types, and in turn, as we illustrate in our running example below, these heterogeneous reelection incentives can generate a gap between the policies chosen by politicians and those preferred by representative voters. Therefore, our second equilibrium restriction imposes some uniformity across states in the treatment of different politician types: we say that a Markov electoral equilibrium $\sigma = (\pi, \rho)$ is *reelection-balanced* if there exists $R^* \in [0, 1]$ such that $\int_x \rho(s, t, x) \pi_t(dx|s) = R^*$ for all states s and all types t. In words, while different policy choices may lead to different reelection probabilities, all incumbents ex ante expect to be reelected with the same probability in all states in a reelection-balanced electoral equilibrium.

To be clear, we make no claim that all compelling Markov electoral equilibria must be convergent and reelection-balanced. Rather, we focus on this class because our results indicate that equilibria in which politicians adopt divergent policies or face different reelection rates will not, in general, produce outcomes that are consonant with direct policy-making by representative voters. As an analogy, static models of elections that can generate non-median equilibrium outcomes are important and useful in applications. However, this does not reduce the value of using the median voter's preferred policy as an idealized benchmark, or of understanding the conditions that yield median convergence as an electoral outcome.

Example (Continued) To reinforce this last point, we return to our example and construct a Markov electoral equilibrium that is neither convergent nor reelection-balanced. Specifically, consider a Markov policy strategy profile in which fiscally dovish politicians always run deficits (i.e., $\pi_d(\bar{x}|s) = 1$ for all $s = \bar{s}, \underline{s}$), and fiscal hawks impose austerity if and only if the economy is strong (i.e., $\pi_h(\underline{x}|\bar{s}) = \pi_h(\bar{x}|\underline{s}) = 1$). Suppose that politicians do not commit to strategies in any state, although this is irrelevant to our results in this example. Furthermore, consider a Markov voting strategy such that the dovish representative voter reelects an incumbent when the economy is strong if and only if she is also dovish (i.e., $\rho(\bar{s}, d, x) = 1$ and $\rho(\bar{s}, h, x) = 0$ for all x); the hawkish representative voter reelects an incumbent when the economy is weak if and only if she is also hawkish (i.e., $\rho(\underline{s}, h, x) = 1$ and $\rho(\underline{s}, d, x) = 0$ for all x); and all politicians are reelected in state s_0 (i.e., $\rho(s_0, t, x_0) = 1$ for all t). If both the office benefit b and the state persistence probability p are low, then the profile $\sigma = (\pi, \rho)$ is a Markov electoral equilibrium.

Notice that this Markov electoral equilibrium generates policy outcomes that differ from the representative voting game's unique equilibrium: a hawkish politician introduces fiscal reforms when the economy is strong even if the dovish representative voter in that state would prefer to run a deficit. Furthermore, this politician would not be reelected in that state, even if she committed to expansionary fiscal policy. The reason for this is that the voter would anticipate that this politician would stay in office if the economy became weak, until eventually the economy became strong again, in which case this hawk would return to implementing fiscal reforms. Therefore, the dovish voter opts for the challenger in the hope of securing a dovish incumbent in the continuation game. A similar logic explains why a dovish politician cannot be reelected when the economy is weak, even if in that state both

politician types run deficits: the hawkish representative voter prefers that policies be set by policies set by hawks rather than doves, if the economy becomes strong.   □

4 Delegation and representative voting games

If the equilibrium outcomes of representative voting games are to be used as a benchmark to evaluate electoral performance, then we need to determine whether elections can ever achieve this benchmark. Put differently, can the equilibrium policy choices of voters in the representative voting game be delegated to politicians? We answer this in the affirmative: if $\tilde{\pi}$ is an equilibrium of the representative voting game, and if politicians place sufficient value on holding office in the future, then we can construct a convergent and reelection-balanced Markov electoral equilibrium $\sigma = (\pi, \rho)$ such that in each state s, office holders of all types t will use the mixed strategy $\tilde{\pi}_{\kappa(s)}(\cdot|s)$ of the representative voter in s.

Theorem 1 *Assume that δb is large, and let $\tilde{\pi}$ be a stationary Markov perfect equilibrium of the representative voting game. Then, given any $R < 1$, there exists a convergent and reelection-balanced Markov electoral equilibrium σ with ex ante reelection probability $R^* \geqslant R$ in which politicians implement the equilibrium from the voting game: for all s and all t, $\pi_t(\cdot|s) = \tilde{\pi}_{\kappa(s)}(\cdot|s)$. If $\tilde{\pi}$ is in pure strategies, then such a Markov electoral equilibrium exists for $R^* = 1$.*

If the equilibrium $\tilde{\pi}$ from the representative voting game is pure, then proving the result is simple, as we can explain by returning to our example.

Example (Continued) Recall that the unique stationary Markov perfect equilibrium of the representative voting game has all representative voters run deficits when the economy is both strong and weak. Therefore, to construct the Markov electoral equilibrium from Theorem 1, we specify that all politicians choose to run deficits as well (i.e., $\pi_t(\overline{x}|s) = 1$ for all t and all $s = \overline{s}, \underline{s}$), and that furthermore they are reelected if and only if they run deficits (i.e., $\rho(s, t, x) = 1$ when $s = \overline{s}, \underline{s}$ if and only if $x = \overline{x}$, along with $\rho(s_0, t, x_0) = 1$). Because representative voters expect all politicians to choose the same policies in future states that other voters would have chosen in the representative voting game, no politician can improve the current representative voter's payoff by implementing fiscal reforms, so that all voters' reelection decisions are optimal. Finally, high office motivation gives politicians the incentives to run deficits whether or not this agrees with their policy preferences, and hence the equilibrium is reelection-balanced with $R^* = 1$.   □

On the other hand, if the equilibrium $\tilde{\pi}$ is mixed, then because a politician of arbitrary type t may have very different preferences than the representative voter, our result may seem surprising. To prove it, we must induce a politician of type t to mix over policies in state s according to $\tilde{\pi}_{\kappa(s)}(\cdot|s)$, and we use mixed voting strategies,

along with the assumption that politicians are sufficiently office motivated, to accomplish this. The preferences of a type t politician over all policies in the support of $\tilde{\pi}_{\kappa(s)}(\cdot|s)$, together with the requirement that she be indifferent between all these policies in equilibrium, pin down the relative magnitudes of politicians' associated reelection probabilities. To ensure that the electoral equilibrium is reelection-balanced, we use a fixed point argument to align the ex ante reelection probabilities of all politician types across all states. Note also that the associated reelection-balanced equilibrium can have an ex ante reelection probability R^* that is arbitrarily close to one, but if some politician is not indifferent over all policies in the support of $\tilde{\pi}$ in some state, then it must be that $R^* < 1$. In our construction, a higher ex ante reelection probability entails a higher threshold that the discounted office benefit δb must exceed in order to support the equilibrium.

We make two further remarks on Theorem 1. First, the result would be easier to prove if we did not insist on constructing a Markov electoral equilibrium that is reelection-balanced (hence avoiding the fixed point argument described above). However, Theorem 2 below, which rules out Markov electoral equilibria in which politicians choose suboptimal policies for representative voters, depends critically on reelection-balancedness. Thus, inclusion of this restriction in Theorem 1 reinforces Theorem 2 by ensuring it is non-vacuous when politicians are highly office motivated. Second, politicians' commitment power plays no role in Theorem 1. More precisely, in the equilibrium we construct politicians never choose to commit to policies. Again, commitment will be critical for Theorem 2, and we will discuss this further below.

By demonstrating that the prospect of retention provides sufficient incentives for office-motivated politicians to reproduce the equilibrium policy choices of representative voting games, Theorem 1 provides evidence of the latter's validity as a benchmark. Nevertheless, the usefulness of this benchmark is increased if we can delimit the Markov electoral equilibria that cannot be replicated in the benchmark game among voters. To evaluate the constraints that representative voters' preferences impose on politicians' choices, we begin with a criterion that compares equilibrium policies in the electoral model to those that representative voters would direct politicians to choose if they could. We say that a Markov electoral equilibrium $\sigma = (\pi, \rho)$ satisfies the *delegated best-response property* if all politician types choose optimal policies for all representative voters in all states: for all s and t, $\pi_t(\cdot|s)$ puts probability one on solutions to

$$\max_{x \in X(s)} u_{\kappa(s)}(s, x) + \delta \left[\rho(s, t, x) V^I_{\kappa(s)}(s, t, x) + (1 - \rho(s, t, x)) V^C_{\kappa(s)}(s, t, x) \right].$$

The hypothetical scenario facing a representative voter in the definition of the delegated best-response property is similar to his best response problem in the representative voting game. The key distinction is that in the representative voting game, only voters choose policies, whereas in a Markov electoral equilibrium satisfying the delegated best-response property, it is as though the representative voter in state s chooses policies in that state, but anticipates that future policies will again be delegated to politicians by other representative voters.

We now show that in every convergent and reelection-balanced Markov electoral equilibrium, including the construction used to prove Theorem 1, politicians choose the best responses of representative voters.

Theorem 2 *If a Markov electoral equilibrium $\sigma = (\pi, \rho)$ is convergent and reelection-balanced, then it satisfies the delegated best-response property.*

Driving Theorem 2 is the fact that in any convergent and reelection-balanced Markov electoral equilibrium, politicians of type $\kappa(s)$, i.e., who are the same type as the representative voter in state s, must choose policies that are best responses for the representative voter $\kappa(s)$ in that state: in such an equilibrium, if politicians of type $\kappa(s)$ instead choose policies that are not optimal for representative voter $\kappa(s)$, then we show that these politicians can profitably deviate to a policy x that is preferred by this voter and, furthermore, must be rewarded with reelection. Finally, because the equilibrium is convergent, this extends to all politician types other than $\kappa(s)$, establishing the delegated best-response property. There are two steps in the argument above, the first is to establish that the representative voter $\kappa(s)$ has incentives to reelect the politician if he deviates to the preferred policy x in state s, and the second is to show that the politician of type $\kappa(s)$ has incentives to deviate to x in the first place. The key to the first step is politicians' commitment power, and the key to the second is our restriction to reelection-balanced equilibria, and we address each of these in turn.

An important point is that even if the representative voter in state s strictly benefits if some politician chooses the policy x, this fact on its own does not ensure that the voter retains the politician. The issue is that the politician's deviation to x is not necessarily a credible indication that she will implement x in future occurrences of state s. If the incumbent has no commitment power, then he is expected to return to equilibrium policy choices in case s recurs. Because the equilibrium under consideration is convergent, voter $\kappa(s)$ is indifferent between all politician types following the choice of x in s. But if the incumbent has commitment power, then there is a positive probability that the voter's gain from x in s also accrues in the next period, so that she has a strict incentive to reelect him. Notice that Theorem 2 does not require that politicians actually commit to policies in equilibrium (or even choose to run for reelection); rather, it says that the option to commit is incompatible with policy choices by politicians in state s that are suboptimal for the representative voter in that state.

Because Theorem 2 depends on the policy preferences of representative voter $\kappa(s)$ and politicians of this type being aligned, it does not require that politicians place a high value on holding office in the future. This does not mean, however, that the wedge between voters and politicians of type $\kappa(s)$ introduced by office motivation is unimportant, only that in reelection-balanced equilibria, it is inoperative. The deviation described above by a politician of type $\kappa(s)$ to the policy x will improve her policy payoffs and, as argued above, it will also lead to reelection with probability one in state s. However, this policy choice could fail to improve her overall payoffs if it generated transitions to states in which she is less

likely to be reelected. This concern is taken care of by reelection-balancedness, because a deviation by a type $\kappa(s)$ politician to policy x in state s has no effect on her reelection probability in other states $s' \neq s$.[4]

Given the similarity between the hypothetical scenario in the definition of the delegated best-response property and the best response problem of the representative voter in the benchmark, Theorem 2 provides insight into conditions under which Markov electoral equilibria must correspond to equilibria of the representative voting game. Specifically, given a convergent and reelection-balanced Markov electoral equilibrium $\sigma = (\pi, \rho)$, we consider whether the induced strategy profile $\tilde{\pi}$ in the representative voting game is a stationary Markov perfect equilibrium. Formally, we define the induced profile by $\tilde{\pi}_{\kappa(s)}(A|s) = \pi_t(A \cup \xi(A) \cup \varphi(A)|s)$ for all s, arbitrary t, and all open $A \subseteq Y(s)$, taking the marginal on policy choices across the politicians' decisions to commit, drop out, or neither. We provide conditions for $\tilde{\pi}$ to be an equilibrium of the representative voting game in Corollary 1 below, but the correspondence does not hold in general. This is due to the difficulty of replicating the distribution over policy sequences generated by σ in the electoral model through $\tilde{\pi}$ in the representative voting game. In particular, suppose that the policy strategies π involve mixing, and that an incumbent politician chooses some policy x in state s, is reelected, and that s recurs. If the incumbent has chosen to commit to x, then she implements x again. In the representative voting game, on the other hand, the representative voter randomizes according to $\tilde{\pi}_{\kappa(s)}(\cdot|s)$ after successive realizations of s. Put differently, incumbents' policy commitments in the game with politicians, when viewed in the context of the representative voting game, generate non-stationary policy outcomes.[5]

As anticipated above, there are two cases in which the outcomes of a convergent and reelection-balanced Markov electoral equilibrium can be replicated by equilibrium play in the representative voting game: when policy strategies are pure and when incumbents do not exercise commitment power in equilibrium, so that electoral turnover does not interact with policy choices. In these cases, under the conditions of Theorem 2, Markov electoral equilibria replicate the policy outcomes of equilibria of the representative voting game.

Corollary 1 *Consider a convergent and reelection-balanced Markov electoral equilibrium $\sigma = (\pi, \rho)$. Define the strategy profile $\tilde{\pi}$ in the representative voting game by $\tilde{\pi}_{\kappa(s)}(A|s) = \pi_t(A \cup \varphi(A) \cup \xi(A)|s)$ for all s, arbitrary t, and all open $A \subseteq Y(s)$, and suppose that either*

---

[4] Theorem 2 is related to Proposition 4.2 in Duggan and Forand (2019), where we focus on ruling out the implementation of policy rules that are not solutions to the representative voter's dynamic programming problem. There, the representative voter is fixed, and equilibrium coordination between representative voters is not an issue, so we can rely on a refinement of voting strategies which assumes only that, in all states, all politician types have available some policy which leads to reelection.

[5] If we allowed for history-dependent persistence of policy choices in the representative voting game, it would be the case that the outcomes of any Markov electoral equilibrium satisfying the conditions of Theorem 2 could be replicated by a (nonstationary) subgame perfect equilibrium of the representative voting game.

1. *the policy profile π is pure, or*
2. *politicians do not use commitment, i.e., $\pi_t(Y^c(s)|s) = 0$ for all s and t.*

Then $\tilde{\pi}$ is a stationary Markov perfect equilibrium of the representative voting game.

The restriction to reelection-balanced equilibria in Theorem 2 is strong, but by returning to our example we can illustrate a fundamental insight regarding the limits of policy control by competing representative voters: if representative voters in future states fail to coordinate on balanced reelection standards, then politicians may not have incentives to choose policies that are optimal for the representative voter in the current state.

Example (Continued) If the state persistence probability p is low and the office benefit b is high, then there exists a Markov electoral equilibrium such that all politicians run deficits when the economy is weak but implement fiscal reforms when the economy is strong (i.e., $\pi_t(\bar{x}|\underline{s}) = \pi_t(\underline{x}|\bar{s}) = 1$ for all types t). All politicians are reelected following all policy choices when the economy is strong or when the budget problem has been resolved (i.e., $\rho(s, t, x) = 1$ for all t and x if $s = \bar{s}, s_0$). Meanwhile, when the economy is weak, only hawkish politicians are ever reelected, and they are reelected following all policy choices (i.e., $\rho(\underline{s}, h, x) = 1$ and $\rho(\underline{s}, d, x) = 0$ for all x). Because dovish politicians expect to be reelected with probability one when the economy is strong but with probability zero when the economy is weak, this equilibrium is not reelection-balanced. Note, however, that it is convergent.

This electoral equilibrium does not satisfy the delegated best-response property: a dovish politician implements fiscal reforms when the economy is strong, even if running deficits would be optimal for the representative voter in this state (because all politicians run deficits when the economy is weak). However, the fact that a dovish politician places a high value on office creates a wedge between her preferences and those of dovish voters, and it leads to her to choose suboptimal policies: a dovish politician could secure reelection by running a deficit when the economy is strong, but she forecasts that once the state transitions to $\underline{s}$ she will not be reelected, although she also chooses her preferred policy, which is also the preferred policy of dovish voters, in that state. Instead, facing an imbalance in equilibrium reelection probabilities across states $\bar{s}$ and $\underline{s}$, a dovish politician sacrifices policy payoffs by implementing austerity measures when the economy is strong, in order to maximize her long-run office benefits. □

5 Existence of representative voters

Our model of representative voting games rests on the assumption that there exists a representative voter in each state. A more micro-founded modelling approach would allow for a richer description of political interactions in all states and characterize those institutional arrangements whose electoral outcomes can be described through the preferences of state-dependent representative voters. To that end, for

each state s, fix a set $\mathscr{D}(s) \subseteq 2^T \setminus \{\emptyset\}$ of *decisive* coalitions of types: the interpretation is that if the coalition of voter types who vote for the incumbent belongs to $\mathscr{D}(s)$, then the incumbent retains office. Electoral outcomes must now be defined through the preferences of decisive coalitions of voters. Fix a Markov electoral equilibrium $\sigma = (\pi, \rho),$[6] and given any state s and any incumbent type t, let $P_\tau(s, t)$ and $R_\tau(s, t)$ denote the strong and weak reelection sets of a voter of type τ, which are defined as in (1). For all coalitions $C \subseteq T$, define

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

and let the *strict* and *weak reelection sets* for incumbent type t in state s be denoted by

$$P(s, t) = \bigcup \{P_C(s, t) : C \in \mathscr{D}(s)\} \text{ and } R(s, t) = \bigcup \{R_C(s, t) : C \in \mathscr{D}(s)\},$$

respectively.

A type $\kappa(s)$ voter is *representative* in state s if $P(s, t) = P_{\kappa(s)}(s, t)$ and $R(s, t) = R_{\kappa(s)}(s, t)$. In words, the type $\kappa(s)$ voter strictly prefers to elect one candidate over the other if and only if a decisive coalition of types in state s shares this preference. Note that the property of being representative, as it depends on voters' reelection sets, is endogenously determined within a specific equilibrium. This raises two important problems: we seek conditions that ensure the existence of representative voters in all equilibria; and want to identify representative voters from the model's fundamentals. In particular, can representative voters be identified through their stage utilities, which are primitives, without reference to their continuation values, which are endogenous? We address these issues in Theorem 3 below.

The result relies on assumptions on both voters' preferences and on the game's decisive coalitions. First, say stage utilities are *ordered by type* if there exist parameters $\omega_\tau, \zeta_\tau \in \mathfrak{R}$ for each voter type τ and mappings $v : S \times X \to \mathfrak{R}$ and $c : S \times X \to \mathfrak{R}$ such that for all τ, all s, and all x, we have

$$u_\tau(s, x) = \omega_\tau v(s, x) - c(s, x) + \zeta_\tau.$$

Note that if Y and S are one-dimensional and utility is quadratic, with the state entering as a shift parameter on ideal points, $\hat{x}_\tau + s$, then stage utilities are ordered by type. Indeed, write

$$
\begin{aligned}
u_\tau(x, s) &= -(\hat{x}_\tau + s - x)^2 \\
&= -\hat{x}_\tau^2 - s^2 - x^2 - 2\hat{x}_\tau s + 2\hat{x}_\tau x + 2sx \\
&= 2\hat{x}_\tau(s + x) - (s^2 + x^2 - 2sx) - \hat{x}_\tau^2.
\end{aligned}
$$

This has the required form, if we set

---

[6] This augmented game is a special case of the model from Duggan and Forand (2018), where we prove general existence of Markov electoral equilibria.

$$\omega_\tau = 2\hat{x}_\tau, \quad v(s,x) = s + x, \quad c(s,x) = s^2 + x^2 - 2sx, \quad \zeta_\tau = -\hat{x}_\tau^2.$$

Second, given a state s, say $\mathscr{D}(s)$ is a *weighted majority rule* if there exist weights $n_\tau(s) \geqslant 0$ for each voter type τ with $\sum_{\tau \in T} n_\tau(s) = 1$ such that $\mathscr{D}(s) = \{C : \sum_{\tau \in C} n_\tau(s) > \frac{1}{2}\}$. Furthermore, say $\mathscr{D}(s)$ is *strong* if every blocking coalition is decisive, i.e., there is no coalition C with $\sum_{\tau \in C} n_\tau(s) = \frac{1}{2}$. Assuming stage utilities are ordered by type, and given a weighted majority rule, we say κ is a *weighted median* type at s if

$$\sum \{n_\tau(s) : \omega_\tau < \omega_k\} \leqslant \frac{1}{2} \text{ and } \sum \{n_\tau(s) : \omega_\tau > \omega_k\} \leqslant \frac{1}{2}.$$

If the weighted majority rule is strong, as is the case for generic weights, then there is a unique weighted median type.

Next, we establish that for a rich class of dynamic electoral environments, the weighted median type is a representative voter in each state.

Theorem 3 *Let σ be a Markov electoral equilibrium, and let s be any state. Suppose that stage utilities are ordered by type and that $\mathscr{D}(s)$ is a strong weighted majority rule. Then there exists a representative voter type $\kappa(s)$ in s, and furthermore $\kappa(s)$ is the weighed median type at s.*

To prove the theorem, let σ be a Markov electoral equilibrium, and let s be any state and t any politician type. To prove that $P(s,t) = P_\kappa(s,t)$ and $R(s,t) = R_\kappa(s,t)$, note that if an incumbent chooses policy x and is reelected, then a probability distribution over future sequences of state-policy pairs, $\{(x_r, s_r)\}_{r=1}^\infty$, is determined. Let $\mu_{s,t,x}^r$ denote the marginal on state-policy pairs r periods hence, and define

$$\mu_{s,t,x} = (1 - \delta) \sum_{r=1}^\infty \delta^{r-1} \mu_{s,t,x}^r$$

as the probability measure that aggregates over these marginals according to the discounted sum. Because all voter types τ share the same discount factor δ, we have

$$V_\tau^I(s,t,x) = \int_{(s',x')} u_\tau(s',x') \mu_{s,t,x}(d(s',t')).$$

Similarly, let $v_{s,t,x}^r$ denote the marginal on state-policy pairs r periods hence if a challenger is elected instead, and define

$$v_{s,t,x} = (1 - \delta) \sum_{r=1}^\infty \delta^{r-1} v_{s,t,x}^r,$$

so that

$$V_\tau^C(s,t,x) = \int_{(s',x')} u_\tau(s',x') v_{s,t,x}(d(s',t')).$$

Using type-ordered utilities and strong weighed majority rule, the corollary to Proposition 3 of Duggan ([2014]) implies that the weighted median type $\kappa(s)$ voter is decisive over lotteries, and thus a weighted majority of voters strictly prefer $\mu_{s,t,x}$ to $v_{s,t,x}$ if and only if the type $\kappa(s)$ voter strictly prefers $\mu_{s,t,x}$ to $v_{s,t,x}$, and it follows that $P(s,t) = P_{\kappa(s)}(s,t)$. As well, a weighted majority of voters weakly prefer $\mu_{s,t,x}$ to $v_{s,t,x}$ if and only if the type $\kappa(s)$ voter weakly prefers $\mu_{s,t,x}$ to $v_{s,t,x}$, and thus $R(s,t) = R_{\kappa(s)}(s,t)$, completing the proof.

6 Conclusion

In this paper, we propose the outcomes of representative voting games as a benchmark to evaluate the outcomes of dynamic elections in which the voters' political power evolves over time. We show that this benchmark is well-founded, in that the existence of representative voters can be guaranteed for a rich class of dynamic electoral environments. Our main results establish the relevance of this benchmark, in that equilibria of the representative voting game can be supported by equilibria of the dynamic electoral model if politicians are sufficiently office motivated. Moreover, we clarify when the preferences of representative voters constrain politicians' choices across all electoral equilibria, in that office holders choose best response policies for the representative voter in each state: the delegated best-response property holds if politicians policy choices are convergent, and representative voters coordinate on electoral standards across states (reelection-balanced equilibria). To understand when equilibria of the electoral game match those of the benchmark, we show that every convergent and reelection-balanced Markov electoral equilibrium corresponds to an equilibrium of the representative voting game, unless the equilibrium involves mixing or the use of commitment. Perhaps surprisingly, our results also show that the connection between the electoral model and the benchmark is delicate, and that when delegation to politicians relies on commitment and mixing, it may introduce a wedge between electoral equilibrium outcomes and the direct choices of representative voters.

Appendix

Proof of Theorem 1 Let $\tilde{\pi}$ be a strategy profile in the representative voting game. For all states s, types t and policies x, let $\tilde{V}_t(s,x)$ denote the payoff from policy x in state s to a representative voter of type t, which solves the recursive equation

$$\tilde{V}_t(s,x) = u_t(s,x) + \delta p(s|s,x)\tilde{V}_t(s,x) + \delta \sum_{s' \neq s} p(s'|s,x) \int_{x'} \tilde{V}_t(s',x')\tilde{\pi}_{\kappa(s')}(dx'|s'),$$

(2)

and let $\tilde{V}_t(s) = \int_x \tilde{V}_t(s,x)\tilde{\pi}_{\kappa(s)}(dx|s)$. If furthermore the profile $\tilde{\pi}$ is a stationary Markov perfect equilibrium of the representative voting game, then, for all states s, $\tilde{\pi}_{\kappa(s)}(\cdot|s)$ puts probability one on solutions to

$$\max_{x \in Y(s)} \tilde{V}_{\kappa(s)}(s, x), \tag{3}$$

so that

$$\tilde{V}_{\kappa(s)}(s) \geqslant \tilde{V}_{\kappa(s)}(s, x) \tag{4}$$

for all policies x, with equality if and only if x is a best-response for $\kappa(s)$ against $\tilde{\pi}$ in the representative voting game. To define the Markov electoral strategy $\sigma = (\pi, \rho)$, we specify that for all s and all t, $\pi_t(\cdot|s) = \tilde{\pi}_{\kappa(s)}(\cdot|s)$. Therefore, because $\tilde{\pi}_{\kappa(s)}(Y^c(s)|s) = 0$ for all s, it follows that $V_t^F(s, t') = \tilde{V}_t(s)$ and that $V_t^B(s, t', x) = V_t^C(s, t', x)$ for all states s, types t and t' and policies x. In particular, any voting strategy ρ is optimal.

Assume further that the equilibrium profile $\tilde{\pi}$ is in pure strategies. We specify the voting strategy such that, for all states s and types t, $\rho(s, t, x) = 1$ if and only if $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$. To verify the optimality of the voting strategy, fix any state s, incumbent type t and policy x. We have that

$$\begin{aligned} V_{\kappa(s)}^I(s, t, x) - V_{\kappa(s)}^C(s, t, x) &= p(s|s, x)\Big[u_{\kappa(s)}(s, x) + \delta V_{\kappa(s)}^I(s, t, x) - V_{\kappa(s)}^F(s, t)\Big] \\ &= p(s|s, x)\Big[\tilde{V}_{\kappa(s)}(s, x) - \tilde{V}_{\kappa(s)}(s)\Big] \\ &\leqslant 0, \end{aligned}$$

with equality whenever $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, as desired. The second equality follows from (8) and the inequality follows from (4). To verify the optimality of policy strategies, normalize stage utilities such that $\underline{u} \leqslant u_t(s, x) \leqslant \overline{u}$ for all states s, types t and policies x, and assume that $\delta b > \overline{u} - \underline{u}$. Fix any state s and type t. An office holder of type t obtains a payoff of at least $\frac{\underline{u}+b}{1-\delta}$ if she implements policy $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, while her payoff to implementing any policy $x' \notin \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$ is at most $\frac{\overline{u}}{1-\delta} + b$, so that choosing policy x is optimal, as desired.

Now assume that the equilibrium profile $\tilde{\pi}$ is in mixed strategies. For all states s, types t and policies $x \notin \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, we specify the voting strategy such that $\rho(s, t, x) = 0$. In particular, note that because $\tilde{\pi}_{\kappa(s)}(Y(s)|s) = 1$ for all s, we have that, for any s and t, $\rho(s, t, x) = 0$ for all $x \in Y^c(s) \cup Y^d(s)$. To construct the voting strategy for policies $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, we assume that δb is large enough that

$$\frac{\overline{u} - \underline{u}}{1 - \delta} < \frac{\delta b}{1 - \delta \epsilon} \min\{\epsilon, 1 - \epsilon\}, \tag{5}$$

where we use our normalization of stage utilities. Note that given any $0 < \epsilon < 1$, we can set b large enough that (5) is satisfied. Our goal is to define a continuous mapping $f : [\epsilon, 1] \to [\epsilon, 1]$, the fixed point R^* of which will allow us to construct the voting strategy $\rho(s, t, x)$ for $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$ such that, in each state, the ex ante probability of reelection of all politician types is R^*, that is, for which $\int_x \rho(s, t, x)\tilde{\pi}_{\kappa(s)}(dx|s) = R^*$ for all s and t. To this end, let $R \in [\epsilon, 1]$, and fix state s and type t.

Given any $r \in [0, 1]$, define $\tilde{w}_{s,t}(x, r)$ as the payoff to a type t politician in s if she chooses policy x and is reelected with probability r, given that she will be reelected

with ex ante probability R in all future periods (because no politician commits to policies under π, this mimics type t's equilibrium payoff). That is,

$$\tilde{w}_{s,t}(x, r) = u_t(s, x) + b + \delta \sum_{s'} p(s'|s, x)\left[\tilde{V}_t(s') + \frac{rb}{1 - \delta R}\right],$$

and note that $\tilde{w}_{s,t}(x, r)$ is continuous. For the restricted domain $r \in [\epsilon, 1]$, define $\underline{x}_{s,t}(r) = \text{argmin}_{x \in \text{supp}(\tilde{\pi}_{\kappa(s)}(\cdot|s))} \tilde{w}_{s,t}(x, r)$ and $\underline{w}_{s,t}(r) = \tilde{w}_{s,t}(\underline{x}_{s,t}(r), r)$, and note that $\underline{w}_{s,t}(r)$ is continuous (by the Maximum Theorem). Finally, for all $x \in \text{supp}(\tilde{\pi}_{\kappa(s)}(\cdot|s))$ and $r \in [\epsilon, 1]$, define $r_{s,t}(x, r)$ as the solution $r' \in (0, r]$ to $\underline{w}_{s,t}(r) = \tilde{w}_{s,t}(x, r')$. To show that $r_{s,t}(x, r)$ is well-defined, first note that $\tilde{w}_{s,t}(x, r)$ is strictly increasing in r because $\delta > 0$, so that $r_{s,t}(x, r) \leqslant r$ follows from the fact that $\underline{w}_{s,t}(r) \leqslant \tilde{w}_{s,t}(x, r)$. Second, note that

$$\underline{w}_{s,t}(r) - \tilde{w}_{s,t}(x, 0) \geqslant \underline{w}_{s,t}(\epsilon) - \tilde{w}_{s,t}(x, 0)$$

$$\geqslant \frac{\underline{u}}{1 - \delta} + b\left[1 + \frac{\delta\epsilon}{1 - \delta\epsilon}\right] - \left[\frac{\bar{u}}{1 - \delta} + b\right]$$

$$= \frac{\delta b\epsilon}{1 - \delta\epsilon} - \frac{\bar{u} - \underline{u}}{1 - \delta}$$

$$> 0,$$

where the final inequality follows from (5), so that $r_{s,t}(x, r) > 0$ for all $r \in [\epsilon, 1]$. Finally, note that $r_{s,t}(x, r)$ is continuous.

Towards defining the mapping f, we first define a collection $\{R'_{s,t}(R)\}_{s,t}$, where each $R'_{s,t}(R) \in [\epsilon, 1]$. Fix state s and type t. If $\int_x r_{s,t}(x, 1)\tilde{\pi}_{\kappa(s)}(dx|s) \leqslant R$, then we set

$$R'_{s,t}(R) = \int_x r_{s,t}(x, 1)\tilde{\pi}_{\kappa(s)}(dx|s).$$

To ensure that $R'_{s,t}(R)$ is well-defined in this case, we need to verify that $R'_{s,t}(R) \geqslant \epsilon$, for which it is sufficient to show that $r_{s,t}(x, 1) \geqslant \epsilon$ for all $x \in \text{supp}(\tilde{\pi}_{\kappa(s)}(\cdot|s))$. To see this, note that

$$\underline{w}_{s,t}(1) - \tilde{w}_{s,t}(x, \epsilon) \geqslant \frac{\underline{u}}{1 - \delta} + b\left[1 + \frac{\delta}{1 - \delta R}\right] - \left[\frac{\bar{u}}{1 - \delta} + b\left[1 + \frac{\delta\epsilon}{1 - \delta R}\right]\right]$$

$$\geqslant \delta b\left[\frac{1 - \epsilon}{1 - \delta\epsilon}\right] - \frac{\bar{u} - \underline{u}}{1 - \delta}$$

$$> 0,$$

yielding the desired contradiction, where the final inequality follows from (5). If instead $\int_x r_{s,t}(x, 1)\tilde{\pi}_{\kappa(s)}(dx|s) > R$, then we set $R'_{s,t}(R) = R$. Furthermore, in this case there exists $r^* \in [\epsilon, 1)$ such that

$$\int_x r_{s,t}(x, r^*)\tilde{\pi}_{\kappa(s)}(dx|s) = R.$$

To see this, note that by our previous results

$$\int_x r_{s,t}(x, \epsilon)\tilde{\pi}_{\kappa(s)}(dx|s) \leqslant \epsilon$$

$$\leqslant R,$$

so that the claim follows from the continuity of $r_{s,t}(x, r)$. Finally, we define the continuous mapping $f : [\epsilon, 1] \rightarrow [\epsilon, 1]$ such that $f(R) = \inf_{s,t} R'_{s,t}(R)$ for all $R \in [\epsilon, 1]$. This mapping has a fixed point $R^* = f(R^*)$ by Brouwer's fixed point theorem, and, for all states s and types t,

$$R'_{s,t}(R^*) \leqslant R^* \leqslant R'_{s,t}(R^*),$$

where the first inequality follows by construction of $R'_{s,t}(R^*)$ and the second inequality follows from the fact that R^* is a fixed point of f. Therefore, $R_{s,t}(R^*) = R^*$ for all s and t.

Finally, we can use the fixed point R^* of the mapping f to back out the voting strategy $\rho(s, t, x)$ for all states s, types t and policies $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$. Note that, by construction, $\int_x r_{s,t}(x, 1)\tilde{\pi}_{\kappa(s)}(dx|s) \geqslant R^*$ for all $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, so that there exists $r^* \in [\epsilon, 1]$ such that $\int_x r_{s,t}(x, r^*)\tilde{\pi}_{\kappa(s)}(dx|s) = R^*$. Therefore, we define $\rho(s, t, x) = r_{s,t}(x, r^*)$.

To verify the optimality of the policy strategy, note that, by construction, politician t is indifferent between all $x \in \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, which all yield payoff $\underline{w}_{s,t}(r^*)$. Furthermore, by choosing policy $x \notin \text{supp}\,(\tilde{\pi}_{\kappa(s)}(\cdot|s))$, the politician t can obtain a payoff of at most

$$\frac{\bar{u}}{1 - \delta} + b < \underline{w}_{s,t}(\epsilon)$$

$$\leqslant \underline{w}_{s,t}(r^*),$$

where the first inequality follows by (5) and the final inequality follows because, by construction, $r^* \geqslant \epsilon$. Therefore, there is no profitable deviation in s for politician t to a policy outside the support of $\tilde{\pi}_{\kappa(s)}(\cdot|s)$ □

Proof of Theorem 2 Fix a convergent Markov electoral equilibrium $\sigma = (\pi, \rho)$. A first claim is that the symmetry of policy strategies and the optimality of the voting strategy ρ imply that, for all states s, types t and t' and policies x, $V^F_{\kappa(s)}(s, t') = V^F_{\kappa(s)}(s, t)$, $V^I_{\kappa(s)}(s, t, x) = V^I_{\kappa(s)}(s, t', x)$, $V^C_{\kappa(s)}(s, t, x) = V^C_{\kappa(s)}(s, t', x)$ and

$$\rho(s, t, x)V^I_{\kappa(s)}(s, t, x) + (1 - \rho(s, t, x))V^C_{\kappa(s)}(s, t, x)$$
$$= \rho(s, t', x)V^I_{\kappa(s)}(s, t', x) + (1 - \rho(s, t', x))V^C_{\kappa(s)}(s, t', x).$$

Now suppose, towards a contradiction, that there exists a policy $x \in X(s)$ such that, for all types t,

$$u_{\kappa(s)}(s, x) + \delta\Big[\rho(s, t, x)V^I_{\kappa(s)}(s, t, x) + (1 - \rho(s, t, x))V^C_{\kappa(s)}(s, t, x)\Big]$$

$$> \int_{x'} \Big[u_{\kappa(s)}(s, x') + \delta\Big[\rho(s, t, x')V^I_{\kappa(s)}(s, t, x') + (1 - \rho(s, t, x'))V^C_{\kappa(s)}(s, t, x')\Big]\Big]\pi_t(dx'|s)$$

$$= V^F_{\kappa(s)}(s, t),$$

$$(6)$$

and notice that we have that $x \in Y(s)$ if and only if $\varphi(x) \in Y^c(s)$ also satisfies (6). Similarly, because the equilibrium is convergent, we have that $x \in Y(s)$ if and only if $\xi(s) \in Y^d(s)$ also satisfies (6). Correspondingly, in the sequel we assume that $x \in Y^c(s)$. If a politician of type $\kappa(s)$ commits to policy x in state s, it follows that

$$V^I_{\kappa(s)}(s, \kappa(s), x) - V^C_{\kappa(s)}(s, \kappa(s), x)$$

$$= p(s|s, x)\Big[u_{\kappa(s)}(s, x) + \delta V^I_{\kappa(s)}(s, \kappa(s), x) - V^F_{\kappa(s)}(s, \kappa(s))\Big]$$

$$\geqslant \delta p(s|s, x)(1 - \rho(s, \kappa(s), x))\Big[V^I_{\kappa(s)}(s, \kappa(s), x) - V^C_{\kappa(s)}(s, \kappa(s), x)\Big],$$

where the inequality, which is strict because $p(s|s, x) > 0$, follows from (6). Therefore, because $\delta p(s|s, x)(1 - \rho(s, \kappa(s), x)) < 1$, we have that

$$V^I_{\kappa(s)}(s, \kappa(s), x) > V^C_{\kappa(s)}(s, \kappa(s), x),$$

and hence $\rho(s, \kappa(s), x) = 1$.

To complete the proof, suppose that the equilibrium σ is reelection-balanced with ex ante reelection probability R^*. Therefore, the payoff to politician $\kappa(s)$ from choosing according to policy strategy $\pi_{\kappa(s)}(\cdot|s)$ in state s is

$$V^F_{\kappa(s)}(s, \kappa(s)) + b\Big[1 + \frac{\delta R^*}{1 - \delta R^*}\Big]. \tag{7}$$

If instead politician $\kappa(s)$ chooses deviating policy x in state s, her payoff is

$$u_{\kappa(s)}(s, x) + \delta V^I_{\kappa(s)}(s, \kappa(s), x) + b\Big[1 + \frac{\delta}{1 - \delta p(s|s, x)}\Big[\frac{1 - \delta p(s|s, x)R^*}{1 - \delta R^*}\Big]\Big],$$

which, using (6), is strictly higher than her equilibrium payoff from (7), yielding the desired contradiction. $\square$

Proof of Corollary 1 Fix a convergent and reelection-balanced Markov electoral equilibrium σ with reelection probability R^*, and consider the profile $\tilde{\pi}$ in the representative voting game defined such that $\tilde{\pi}_{\kappa(s)}(\cdot|s) = \pi_t(\cdot|s)$ for all s and arbitrary t. First, if either (*i*) the policy profile π is pure or (*ii*) politicians never commit to policies, i.e., $\pi_t(Y^c(s)|s) = 0$ for all s and t, then we have that, for all states s, types t and t' and policies x,

$$V_t^F(s, t') = \tilde{V}_t(s), \text{ and}$$

$$V_t^I(s, t', x) = p(s|s, x)\tilde{V}_t(s, x) + \sum_{s' \neq s} p(s'|s, x)\tilde{V}_t(s'), \tag{8}$$

where payoffs $\tilde{V}_t(s, x)$ and $\tilde{V}_t(s)$ are defined as in (2) Second, by Theorem 2, and invoking the optimality of the voting strategy ρ, we have that $\pi_t(\cdot|s)$ puts probability one on solutions to

$$\max_{x \in Y(s)} u_{\kappa(s)}(s, x) + \delta \max \left\{ V_{\kappa(s)}^I(s, t, x), V_{\kappa(s)}^C(s, t, x) \right\}. \tag{9}$$

Third, in both cases (i) and (ii), we have that $V_{\kappa(s)}^I(s, t, x) = V_{\kappa(s)}^C(s, t, x)$, so that any solution to (9) must also be a solution to

$$\max_{x \in Y(s)} u_{\kappa(s)}(s, x) + \delta V_{\kappa(s)}^I(s, t, x). \tag{10}$$

Finally, that $\tilde{\pi}_{\kappa(s)}(\cdot|s)$ must put probability one on solutions to (3) follows by substituting (8) into (10), so that the profile $\tilde{\pi}$ is a stationary Markov perfect equilibrium of the representative voting game. □

References

Bai JH, Lagunoff R (2011) On the Faustian dynamics of policy and political power. Rev Econ Stud 78(1):17–48

Banks J, Duggan J (2008) A dynamic model of democratic elections in multidimensional policy spaces. Q J Polit Sci 3:269–299

Bernhardt D, Dubey S, Hughson E (2004) Term limits and pork barrel politics. J Public Econ 88:2383–2422

Besley T, Coate S (1997) An economic model of representative democracy. Quart J Econ 112:85–114

Calvert R (1985) Robustness of the multi-dimensional voting model: candidate motivations, uncertainty and convergence. Am J Polit Sci 29:69–95

Downs A (1957) An economic theory of democracy. Harper and Bros, New York

Duggan J (2000) Repeated elections with asymmetric information. Econ Polit 12:109–136

Duggan J (2014) Majority voting over lotteries: conditions for existence of a decisive voter. Econ Bull 34:263–270

Duggan J, Forand JG (2018) Existence of Markov electoral equilibria. Unpublished paper

Duggan J, Forand JG (2019) Accountability via delegation in dynamic elections. Unpublished Paper

Federgruen A (1978) On n-person stochastic games with denumerable state space. Adv Appl Probab 10:452–471

Forand JG (2014) Two-party competition with persistent policies. J Econ Theory 152:64–91

Osborne M, Slivinski A (1996) A model of competition with citizen candidates. Q J Econ 111:65–96

Van Weelden R (2013) Candidates, credibility, and re-election incentives. Rev Econ Stud 80(4):1622–1651